

Program Equivalence for Algebraic Effects via Modalities



Cristina Matache
St Cross College
University of Oxford

A thesis submitted for the degree of
MSc in Computer Science
Trinity 2018

Abstract

This dissertation is concerned with the study of program equivalence and algebraic effects as they arise in the theory of programming languages. Algebraic effects represent impure behaviour in a functional programming language, such as input and output, exceptions, nondeterminism etc. all treated in a generic way. Program equivalence aims to identify which programs can be considered equal in some sense. This question has been studied for a long time but has only recently been extended to languages with algebraic effects, which are a newer development. Much work remains to be done in order to understand program equivalence in the presence of algebraic effects. In particular, there is no characterisation of contextual equivalence using a logic.

We define a logic whose formulas express properties of higher-order programs with algebraic effects. We then investigate three notions of program equivalence for algebraic effects: *logical equivalence* induced by the aforementioned logic, *applicative bisimilarity* and *contextual equivalence*. For the programming language used in this dissertation, we prove that they all coincide.

Therefore, the main novel contribution of the dissertation is defining the first logic for algebraic effects whose induced program equivalence coincides with contextual equivalence.

Acknowledgements

I would like to thank my supervisor, Sam Staton, for his guidance and patience without which this project would not have been possible. I am also grateful to Alex Simpson and Niels Voorneveld for insightful discussions about their work. Finally, I want to thank my parents for their support and for making this year in Oxford possible.

Contents

1	Introduction	9
1.1	Motivation	9
1.2	Contributions	11
1.3	Structure of the Dissertation	12
2	Background and Explanation of Problem	13
2.1	Program Equivalence	13
2.2	Algebraic Effects and Logical Properties	15
2.3	Program Equivalence for Algebraic Effects	16
2.4	PCF with Effects – EPCF	16
2.5	Operational Semantics of EPCF	19
2.6	Summary of Results about EPCF	22
2.7	Problem: Contextual vs. Logical Equivalence	25
2.8	Chapter Summary	27
3	Introducing the ECPS Language	29
3.1	A New Language – ECPS	29
3.2	Operational Semantics of ECPS	31
3.3	A Coinduction Proof Principle	34
3.4	Logical Relations	35
3.5	Chapter Summary	36
4	Comparison: ECPS vs. EPCF	37
4.1	CPS Translation	37
4.2	Typing and CPS Translation for Stacks	40
4.3	Correctness of the CPS Translation	40
4.4	ECPS Is More Expressive than EPCF	47
4.5	Chapter Summary	49
5	Applicative Bisimilarity for ECPS	51
5.1	Observations for ECPS	51
5.2	Applicative \mathfrak{B} -Bisimilarity	53
5.3	Applicative \mathfrak{B} -Bisimilarity is a Congruence	56
5.4	Howe’s Method	60
5.5	Chapter Summary	64

6	Logical Equivalence for ECPS	67
6.1	Two Logics for ECPS	67
6.2	Logical Equivalence	70
6.3	Logical Equivalence Coincides with Bisimilarity	71
6.4	Chapter Summary	74
7	Contextual Equivalence for ECPS	77
7.1	Contextual Equivalence Coinductively	77
7.2	Contextual Equivalence Coincides with Bisimilarity	78
7.3	Contextual Equivalence via Contexts	82
7.4	Chapter Summary	87
8	Conclusion	89
8.1	Summary	89
8.2	Comparison with Previous Work	90
8.3	Future Work	91
8.4	Personal Reflections	91
	Bibliography	93
	Appendices	96
A	Proofs about the CPS translation	97
B	Proofs about Applicative Bisimilarity	109
B.1	Howe’s Method	111
C	Proofs about Logical Equivalence	121
D	Proofs about Contextual Equivalence	125

List of Figures

2.1	Typing judgements for EPCF [SV18].	18
3.1	Typing judgements for ECPS.	30
4.1	CPS translation of EPCF into ECPS – first part.	39
4.2	CPS translation of EPCF into ECPS – continued.	39
5.1	Compatibility rules.	57
5.2	Compatible refinement rules.	61
6.1	Value formulas in the logic \mathcal{F}	68
6.2	Satisfaction relation \models for the logic \mathcal{F}	68
6.3	Translation from \mathcal{F} to \mathcal{V} and vice-versa.	72
7.1	Typing rules for contexts that accept a value.	84
7.2	Typing rules for contexts that accept a computation.	85

Chapter 1

Introduction

This dissertation is a theoretical study of program equivalence for a higher-order language with algebraic effects. In particular, we are interested in finding a logic of program properties that characterises contextual equivalence. This chapter reviews the history of program equivalence and explains why it is a hard but nevertheless interesting problem. At the same time, we outline the main research questions that led to the present work. Finally, we summarise the contributions of our work and the structure of the dissertation.

1.1 Motivation

Although undecidable in general, program equivalence is a fundamental problem both in the theory of programming languages and in formal verification. Programming language researchers are concerned with theoretical ways of reasoning about program equivalence based on formal semantics, like denotational or operational semantics. We adopt this point of view in the dissertation. Verification combines theory with building automated tools for checking program equivalence whenever possible.

Program equivalence seeks to establish when two programs are interchangeable, or when they behave the same, for some definition of behaviour. For the semantics of programming languages, this is useful because one can regard the meaning of a program as the equivalence class of programs that it belongs to. From a more practical perspective, checking program equivalence can determine whether a program implements a specification, which is itself given as a more familiar program. For example, establishing whether certain compiler optimisations are safe can be done in this way.

The first definition of program equivalence that springs to mind is that two programs are equivalent when they return the same result. In the case of a Turing-complete language, one would also like to take into account the possibility of divergence. Therefore, naively two programs are equivalent if they either return the same result or they both diverge.

This definition has two problems. Firstly, establishing divergence of programs is undecidable due to the halting problem, so program equivalence is in general hard to calculate. Working around this issue is mainly the focus of the formal verification community. Secondly, the above naive definition of program equivalence becomes unclear in the presence of higher-order functions. A higher-order function can receive as arguments and return as results other functions, rather than just ground data, such as natural numbers or booleans.

As an example, consider the following programs:

$$one = \lambda f. \lambda y. (f y)$$

$$two = \lambda f. \lambda y. f (f y)$$

$$two' = \lambda f. \lambda y. f (one f y)$$

where *one* and *two* represent the first two Church numerals. Morally, the functions *two* and *two'* should be equivalent in the λ -calculus, although they are not syntactically equal. Therefore the problem is: what does it mean for two higher-order functions to be the same?

Over the years, the programming languages community has proposed many answers to this question. At the beginning, they focussed on pure higher-order languages, with no side effects, such as the λ -calculus and PCF [Plo77]. PCF is a simply-typed extension of the λ -calculus with general recursion and a datatype of natural numbers and is therefore Turing-complete. These languages were chosen for their simplicity and because they have had well-understood formal semantics for a long time.

In this context, program equivalence based on denotational semantics was proposed: if two programs have the same denotation, they are equivalent. An alternative notion based on operational semantics is Morris-style contextual equivalence [Mor69]: two programs are equivalent if they have the same *observable* behaviour in all program contexts. In this dissertation, we are not concerned with denotational equality; contextual equivalence will be discussed more in Section 2.1. Notice that it improves on the naive definition of program equivalence because, to test the equality of functions, a context can provide them with arguments and observe their reduction behaviour.

To address some of the shortcomings of denotational equality and contextual equivalence, other notions of program equivalence that use operational semantics have been proposed. Examples include logical equivalence, that is, equivalence induced by satisfaction of formulas in a logic of program properties (e.g. [HM85]), applicative bisimilarity [Abr90] and logical relations [Tai67]. Logical equivalence is the main focus of this dissertation. Notably, it has strong connections with formal verification, where modal logics such as temporal logics [Pnu77] and the modal μ -calculus [Koz83] are used to specify and verify properties of programs. Applicative bisimilarity will be discussed in more detail in Section 2.1, but we will not be concerned with logical relations for program equivalence.

Given the wide variety of program equivalences mentioned so far, a natural question is comparing them: do they coincide, is one included in the other or are they altogether different? This question has been studied as new definitions arose. For example, Plotkin showed that for PCF denotational equality implies contextual equivalence.

Gradually, the questions identified so far have been posed for more complex languages than the λ -calculus or PCF. An interesting addition are impure operations such as input and output, exceptions, nondeterminism, state or continuations, whose semantics has been studied more recently. They are known as *computational effects*. In his influential work, Eugenio Moggi [Mog91] gave a general denotational semantics for computational effects using monads.

Subsequently, Plotkin and Power [PP01, PP02, PP03] initiated a program of research concerned with *algebraic effects*, a subset of computational effects whose behaviour can

be axiomatised by a set of equations. All the example of effects mentioned before are algebraic with the exception of continuations and exception handling. The purpose of algebraic effects is to give a unified, uniform treatment of the semantics of effects and of their combinations [HPP06]. The latter cannot easily be achieved using Moggi’s semantics.

In this context, a new question becomes apparent: when are two higher-order programs *exhibiting algebraic effects* equivalent? The natural starting point is to extend the existing notions of program equivalence to languages with algebraic effects, and once again compare the resulting relations. A lot of work has been done in this direction for specific effects, such as nondeterminism (e.g. [Las98]) and probabilistic choice (e.g. [CL14]). Ideally however, a notion of program equivalence should be applicable to *any* algebraic effect.

Several such notions of program equivalence for generic algebraic effects have been developed. Johann, Simpson and Voigtländer [JSV10] study contextual equivalence and a corresponding logical relation. Dal Lago, Gavazzo and Levy [LGL17a] are concerned with applicative bisimilarity. Plotkin and Pretnar [PP08] propose a logic for algebraic effects that is sound with respect to other notions of program equivalence, but not complete in general. Simpson and Voorneveld [SV18] propose a modal logic whose induced program equivalence coincides with applicative bisimilarity, but not with contextual equivalence.

Thus, the main question of the dissertation arises: can we find a logic that *characterises* contextual equivalence for a higher-order language with generic algebraic effects?

1.2 Contributions

The programming language we consider in this dissertation is named ECPS. It is a call-by-value continuation-passing variant of PCF with generic algebraic effects. It is a higher-order Turing-complete language.

Being a continuation-passing language means that functions receive an additional argument which specifies how the computation should proceed once the function terminates. This argument is called a *continuation*. Compared to the direct style of programming, continuation-passing style makes control flow and the order of evaluation explicit. Thanks to these features, continuation-passing languages are often used as intermediate languages inside compilers. So ECPS could be seen as a simple variant of an intermediate language.

In the dissertation, we are concerned with three notions of program equivalence for ECPS: logical equivalence, contextual equivalence and applicative bisimilarity, and with the relationship between them. Most importantly, we wish to define a logic whose induced program equivalence coincides with contextual equivalence.

Any notion of program equivalence needs to satisfy two key properties: being an equivalence relation and being compatible. Compatibility means that equivalent programs can be substituted for a variable in a program equation, thus allowing compositional reasoning about program equivalence. These requirements are both explained in more detail in Section 2.1.

The novel contributions of the dissertation are the following:

1. We study the relationship between ECPS and the language used in the work of Simpson and Voorneveld [SV18]. We provide a correct embedding of their language into ECPS (Theorem 4.3.1).

2. We develop applicative bisimilarity for ECPS and prove it is a compatible equivalence relation (Lemma 5.3.1 and Theorem 5.3.7).
3. We define a logic whose formulas express properties of ECPS programs. We prove that program equivalence induced by the logic coincides with applicative bisimilarity (Theorem 6.2.3). Therefore, logical equivalence is compatible.
4. We present two equivalent definitions of contextual equivalence for ECPS, which are both equivalence relations and compatible. We prove that contextual equivalence coincides with applicative bisimilarity (Theorem 7.2.2). This leads to the main result of the dissertation: logical equivalence coincides with contextual equivalence (Corollary 7.2.3).

1.3 Structure of the Dissertation

Chapter 2 starts with an informal introduction to program equivalence and algebraic effects. It then reviews the work of Simpson and Voorneveld [SV18] since it is the most closely related to this dissertation. Finally, it exemplifies the distinction between contextual equivalence and logical equivalence in the context of their development.

Chapter 3 introduces the language ECPS and its operational semantics. It then reviews two proof techniques, namely coinduction and logical relations.

Chapters 4 to 7 contain the novel technical content. Chapter 4 is concerned with justifying the use of the ECPS language. It gives a translation from the programming language used by Simpson and Voorneveld into ECPS and proves the translation correct.

The following three chapters study program equivalence. Chapter 5 defines applicative bisimilarity and proves its main properties. Chapter 6 introduces a logic for ECPS and proves that logical equivalence coincides with applicative bisimilarity. Finally, Chapter 7 develops contextual equivalence and proves it coincides with logical equivalence.

The last chapter reviews the material in the dissertation and surveys previous work. It then sketches several directions for future work and closes with some personal remarks.

Chapters 2 through 7 all end with an accessible summary of their most important points. For an overview of the dissertation, one can consult the “Chapter Summaries”. Almost all the mathematical proofs completed as part of the project appear in the dissertation. To facilitate reading, routine or overlong proofs appear in the appendices rather than in the main body of the text. Therefore, Chapters 4 to 7 all have a corresponding appendix.

Chapter 2

Background and Explanation of Problem

This chapter starts with an informal discussion of program equivalence and algebraic effects. It then reviews in some detail the work of Simpson and Voorneveld [SV18] on program equivalence for algebraic effects. The syntax and operational semantics of a programming language with algebraic effects named EPCF is introduced. Two forms of program equivalence are discussed: logical equivalence and applicative bisimilarity, and they are compared to contextual equivalence.

2.1 Program Equivalence

As Pitts observes [Pit11], for a notion of program equivalence, or equality, to be useful it should be a congruence, that is, satisfy two key properties: being an *equivalence relation* and being *compatible*. The former allows reasoning through a chain of equations in order to establish that two programs P_1 and P_n are equal:

$$P_1 \simeq P_2 \simeq \dots \simeq P_n.$$

Assuming that programs can take parameters, $P(x)$, compatibility means that we can substitute equivalent programs for a parameter in an equation:

$$P_1(x) \simeq P_2(x) \text{ and } Q_1 \simeq Q_2 \implies P_1(Q_1) \simeq P_2(Q_2).$$

This property is important because it allows us to reason compositionally about programs. In order to decide whether two programs are equivalent, it suffices to investigate whether its subphrases are equivalent. When the parameter is allowed to be a function or a process, rather than just ground data, like booleans or integers, compatibility becomes even more important but also harder to establish.

One of the most intuitive notions of program equivalence is Morris-style *contextual equivalence* [Mor69]. Two programs are contextually equivalent if and only if they have the same observable behaviour in all program contexts:

$$P_1 \simeq P_2 \iff \forall C. \mathfrak{D}(C[P_1]) = \mathfrak{D}(C[P_2]).$$

In the case of the untyped λ -calculus, a possible definition for the “observable behaviour” of program P , $\mathfrak{D}(P)$, is whether or not P terminates, written $P\Downarrow$. Therefore $\mathfrak{D}(C[P_1]) = \mathfrak{D}(C[P_2])$ becomes:

$$C[P_1]\Downarrow \iff C[P_2]\Downarrow.$$

However, contextual equivalence is difficult to establish for particular programs because of the quantification over all contexts. Therefore, Abramsky [Abr90] proposed another notion of equivalence for the untyped λ -calculus named *applicative bisimilarity*.

Bisimilarity was first defined by Milner [Mil80] for the process calculus CCS which models concurrency. The main idea is that two processes are bisimilar if whenever one of them can advance by one step, the other can perform a matching step, and the two resulting processes are again bisimilar. The circularity of this definition of bisimilarity suggests that it can be defined coinductively as the greatest relation with a certain property.

In the case of the λ -calculus, the steps that programs can take are β -reduction steps. Therefore, applicative similarity, the one-sided version of bisimilarity, is defined as:

The greatest relation \lesssim , such that $P_1 \lesssim P_2$ implies

$$\begin{aligned} P_1 \longrightarrow^* \lambda x.P'_1 &\implies \\ \exists P'_2 \text{ such that } P_2 \longrightarrow^* \lambda x.P'_2 &\text{ and for any value } v, P'_1[v/x] \lesssim P'_2[v/x]. \end{aligned}$$

Applicative bisimilarity is defined analogously.

For the λ -calculus, applicative bisimilarity coincides with contextual equivalence e.g. [Abr90, Pit11]. Therefore, checking whether two λ -terms are bisimilar is a sound and complete proof technique for establishing contextual equivalence, easier to use in practice.

Another approach to program equivalence is to define a logic \mathcal{L} whose formulas ϕ represent program properties. In this setting, two programs are equivalent if and only if they satisfy the same formulas in this logic:

$$P_1 \simeq P_2 \iff (\forall \phi \text{ in } \mathcal{L}. P_1 \models \phi \iff P_2 \models \phi).$$

An example of a logic that describes program properties is Hennessy-Milner logic [HM85], which concerns CCS processes. In fact, bisimilarity for CCS coincides with the equivalence induced by Hennessy-Milner logic.

In this dissertation, we are most interested in logical equivalence, so we will consider some informal examples of formulas.

Example 2.1.1. Consider a call-by-value simply-typed λ -calculus with natural numbers. A logical formula could be:

$$\phi = \{3\} \mapsto \{2\}.$$

A function f satisfies ϕ if, given argument 3, the expression $(f\ x)$ reduces to 2. Consider for example the following function:

$$f = \lambda n. \mathbf{if}\ n = 3 \mathbf{then}\ \mathbf{pred}\ n \mathbf{else}\ \mathbf{succ}\ n.$$

We can see that f indeed satisfies ϕ , but f does not satisfy $\{4\} \mapsto \{3\}$.

2.2 Algebraic Effects and Logical Properties

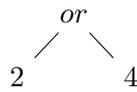
In general, programming languages that are purely functional do not provide “impure” operations such as input and output, nondeterministic or probabilistic choice, global state etc. These features are known as algebraic effects. To include them in a programming language, it suffices to add relevant operations to the language.

As an example consider nondeterministic choice. This can be implemented by adding an operation $or(-, -)$, where, in the term $or(t_1, t_2)$, an external agent chooses nondeterministically whether to execute term t_1 or t_2 .

Example 2.2.1. Consider the simply-typed λ -calculus from the previous example extended with or . Now there is more than one value that a term may reduce to. For example:

$$(g\ 3) \quad \text{where } g = \lambda n. or(\mathbf{pred}\ n, \mathbf{succ}\ n)$$

may reduce to either 2 or 4. The reduction behaviour of $(g\ 3)$ could be represented as a tree:



Recall the formula $\{3\} \mapsto \{2\}$ from the previous section. We can interpret it either as: “the function *always* returns 2” or “*may* return 2”. As a result we have two new formulas: $\phi_1 = \{3\} \mapsto \Box\{2\}$ and $\phi_2 = \{3\} \mapsto \Diamond\{2\}$. We can see that g satisfies the latter but not the former.

Example 2.2.2. As another example, consider a higher-order function:

$$h = \lambda f. \lambda n. or(f(\mathbf{pred}\ n), f(\mathbf{succ}\ n)) : (\mathbb{N} \rightarrow \mathbb{N}) \rightarrow (\mathbb{N} \rightarrow \mathbb{N}).$$

If we apply it to arguments g and 2, and then g and 4, the trees of $(h\ g\ 2)$ and $(h\ g\ 4)$ respectively are:



Now consider the formula:

$$\psi = (\{3\} \mapsto \Diamond\{2\}) \mapsto (\{2, 4\} \mapsto \Diamond\{2\}).$$

It says that, given a function f that satisfies $\{3\} \mapsto \Diamond\{2\}$, h returns another function which when given either 2 or 4 as argument *may* return 2. Function h satisfies ψ because in either case it may call $(f\ 3)$, which we know may return 2.

2.3 Program Equivalence for Algebraic Effects

The notions of program equivalence presented in Section 2.1 have been extended recently to programming languages with algebraic effects. Johann, Simpson and Voigtländer [JSV10] study contextual equivalence for a polymorphic language with recursion and generic effects. They characterise contextual equivalence using a logical relation and thus prove some of its fundamental properties.

Dal Lago, Gavazzo and Levy [LGL17a] give an abstract account of applicative bisimilarity for an untyped λ -calculus with generic algebraic effects. They show that applicative bisimilarity is included in contextual equivalence, but they note that this inclusion is strict in general.

Simpson and Voorneveld [SV18] consider a simply-typed programming language with recursion and generic algebraic effects. They propose a modal logic in which formulas expressing program behaviour are very similar in spirit to the example formulas we have seen so far. They also define applicative bisimilarity following [LGL17a]. The logical equivalence induced by the modal logic is then proved to coincide with applicative bisimilarity.

In the conclusion of their paper, Simpson and Voorneveld observe that logical equivalence is included in contextual equivalence. However, contextual equivalence equates more programs than logical equivalence does. Therefore, an open research direction is finding a logic that characterises contextual equivalence. This is the main goal of this dissertation.

2.4 PCF with Effects – EPCF

The programming language used in the work of Simpson and Voorneveld [SV18] is a call-by-value, simply-typed λ -calculus with recursion, a datatype of natural numbers and algebraic effects. Therefore, the language is a variant of Plotkin’s PCF [Pl077] extended with algebraic effects; in this work we will refer to it as EPCF.

In order to simplify their proofs, Simpson and Voorneveld formulate EPCF as fine-grained call-by-value [LPT03]. This means that there is a distinction between terms that are values and terms that are computations; they form separate syntactic categories. For example, $\lambda x:\mathbb{N}.S(x)$ and 3 are values because they cannot reduce, while $(\lambda x:\mathbb{N}.S(x))\ 3$ is a computation. Here $S(x)$ represents the successor of x . The fine-grained call-by-value formulation is equivalent to the usual call-by-value formulation.

Definition 2.4.1 (EPCF). *Types and environments:*

$$\tau, \rho := \mathbb{1} \mid \mathbb{N} \mid \rho \rightarrow \tau$$

$$\Gamma := \emptyset \mid \Gamma, x : \tau.$$

Values and computations are defined by the grammar:

$$V, W := \star \mid Z \mid S(V) \mid \lambda x:\tau.M \mid x$$

$$M, N := V\ W \mid \mathbf{return}\ V \mid \mathbf{let}\ M \Rightarrow x \mathbf{in}\ N \mid \mathbf{fix}\ V \mid \mathbf{case}\ V \mathbf{in}\ \{Z \Rightarrow M, S(x) \Rightarrow N\}.$$

The ground types are unit $\mathbb{1}$ and natural numbers \mathbb{N} . There is a countably infinite set of variables ranged over by x . The environment $\Gamma, x : \tau$ assumes that x does not appear in Γ .

Terms V, W represent values, and M, N represent computations, that is, terms which can be evaluated. The intuitive semantics of computations is the following: **return** V immediately returns the value V . The construct **let** $M \Rightarrow x$ **in** N is a sequencing operation: first it evaluates M , if this returns a value V , V is substituted for x in N , then $N[V/x]$ is evaluated. The computation **fix** V calculates the fixed point of the function V . The **case** V construct branches according to whether the natural number V is zero or a successor.

The language EPCF incorporates effects in a general way. Instead of specifying all the effect operations in the language, the definition of EPCF is parametrised by a set of effect operations Σ . The set Σ can be instantiated in turn for nondeterminism, probabilistic choice, global store, input and output etc. This is done in a series of examples at the end of the section.

Each operation $\sigma \in \Sigma$ has an arity which specifies what arguments the operation takes and what type the resulting computation has. The possible arities are:

$$\alpha^n \rightarrow \alpha \quad \mathbb{N} \times \alpha^n \rightarrow \alpha \quad \alpha^{\mathbb{N}} \rightarrow \alpha \quad \mathbb{N} \times \alpha^{\mathbb{N}} \rightarrow \alpha$$

where α can be regarded as a type variable. They should be interpreted as follows: $\sigma : \mathbb{N} \times \alpha^n \rightarrow \alpha$ is an operation that takes a natural number V and n computations of type α , M_0, M_1, \dots, M_{n-1} . The resulting computation $\sigma(V; M_0, M_1, \dots, M_{n-1})$ has type α . The expression $\alpha^{\mathbb{N}}$ represents a function from natural numbers to the type α .

Definition 2.4.1 (EPCF – continued). *Fix a set of effect operations Σ , with associated arities. The grammar of computations is extended as follows:*

$$M, N := \dots \mid \sigma(M_0, M_1, \dots, M_{n-1}) \mid \sigma(V; M_0, M_1, \dots, M_{n-1}) \mid \sigma(V) \mid \sigma(V; W).$$

The typing relations $\Gamma \vdash V : \tau$ and $\Gamma \vdash M : \tau$ are the least relations closed under the rules in Figure 2.1. The judgement $\Gamma \vdash V : \tau$ should be read as V has type τ in environment Γ .

Substitution of values for free variables inside values, $W[V/x]$, and inside computations, $M[V/x]$, is defined by recursion on the structure of W and M in a standard way. An example for effect operations is:

$$\sigma(W; M_0, M_1, \dots, M_{n-1})[V/x] = \sigma(W[V/x]; M_0[V/x], M_1[V/x], \dots, M_{n-1}[V/x]).$$

We use the notation $Val(\tau)$ for the set of closed values of type τ and $Comp(\tau)$ for the set of closed computations of type τ . We denote natural numbers $S^n(Z)$ by \bar{n} . Everywhere in the dissertation we consider terms up to α -conversion.

Below are examples of effect operations from [SV18] which will be used throughout the dissertation:

Example 2.4.2 (Pure functional computation). In this case, the language has no effects so the set Σ is empty.

Example 2.4.3 (Nondeterminism). There is one effect operation representing binary choice $or : \alpha^2 \rightarrow \alpha$ so $\Sigma = \{or\}$. It takes as arguments two computations of type α and chooses to run one of them. The choice is determined by an external agent.

Example 2.4.4 (Probabilistic choice). Define $\Sigma = \{p-or\}$ where $p-or : \alpha^2 \rightarrow \alpha$ is a binary choice operator. With probability 0.5 it executes the first computation, otherwise the second computation.

$$\begin{array}{c}
\frac{}{\Gamma, x : \tau \vdash x : \tau} \text{(VAR)} \quad \frac{}{\Gamma \vdash \star : \mathbb{1}} \text{(UNIT)} \quad \frac{}{\Gamma \vdash Z : \mathbb{N}} \text{(ZERO)} \quad \frac{\Gamma \vdash V : \mathbb{N}}{\Gamma \vdash S(V) : \mathbb{N}} \text{(SUCC)} \\
\\
\frac{\Gamma \vdash V : \tau}{\Gamma \vdash \mathbf{return} V : \tau} \text{(RET)} \quad \frac{\Gamma, x : \tau \vdash M : \rho}{\Gamma \vdash \lambda x : \tau. M : \tau \rightarrow \rho} \text{(LBD)} \quad \frac{\Gamma \vdash V : \tau \rightarrow \rho \quad \Gamma \vdash W : \tau}{\Gamma \vdash V W : \rho} \text{(APP)} \\
\\
\frac{\Gamma \vdash V : (\tau \rightarrow \rho) \rightarrow (\tau \rightarrow \rho)}{\Gamma \vdash \mathbf{fix} V : \tau \rightarrow \rho} \text{(FIX)} \quad \frac{\Gamma \vdash M : \tau \quad \Gamma, x : \tau \vdash N : \rho}{\Gamma \vdash \mathbf{let} M \Rightarrow x \mathbf{in} N : \rho} \text{(LET)} \\
\\
\frac{\Gamma \vdash V : \mathbb{N} \quad \Gamma \vdash M : \tau \quad \Gamma, x : \mathbb{N} \vdash N : \tau}{\mathbf{case} V \mathbf{in} \{Z \Rightarrow M, S(x) \Rightarrow N\} : \tau} \text{(CASE)} \\
\\
\frac{\sigma : \alpha^n \rightarrow \alpha \quad \Gamma \vdash M_i : \tau}{\Gamma \vdash \sigma(M_0, M_1, \dots, M_{n-1}) : \tau} \text{(OP1)} \quad \frac{\sigma : \alpha^{\mathbb{N}} \rightarrow \alpha \quad \Gamma \vdash V : \mathbb{N} \rightarrow \tau}{\Gamma \vdash \sigma(V) : \tau} \text{(OP2)} \\
\\
\frac{\sigma : \mathbb{N} \times \alpha^n \rightarrow \alpha \quad \Gamma \vdash V : \mathbb{N} \quad \Gamma \vdash M_i : \tau}{\Gamma \vdash \sigma(V; M_0, M_1, \dots, M_{n-1}) : \tau} \text{(OP3)} \\
\\
\frac{\sigma : \mathbb{N} \times \alpha^{\mathbb{N}} \rightarrow \alpha \quad \Gamma \vdash V : \mathbb{N} \quad \Gamma \vdash W : \mathbb{N} \rightarrow \tau}{\Gamma \vdash \sigma(V; W) : \tau} \text{(OP4)}
\end{array}$$

Figure 2.1: Typing judgements for EPCF [SV18].

Example 2.4.5 (Global store). Fix a finite set of locations \mathbb{L} for storing natural numbers. For each $l \in \mathbb{L}$, Σ contains the following operations:

$$\text{lookup}_l : \alpha^{\mathbb{N}} \rightarrow \alpha$$

$$\text{update}_l : \mathbb{N} \times \alpha \rightarrow \alpha.$$

The intuition is the following: the computation $\text{lookup}_l(V)$ looks up the number at location l then passes it to the function V ; $\text{update}_l(\bar{n}, M)$ writes n to location l then runs the computation M .

Example 2.4.6 (Input/output). In this case, $\Sigma = \{\text{read}, \text{write}\}$ where

$$\text{read} : \alpha^{\mathbb{N}} \rightarrow \alpha$$

$$\text{write} : \mathbb{N} \times \alpha \rightarrow \alpha.$$

The computation $\text{read}(V)$ reads a natural number from the input channel and passes it to the function V , which then executes. The computation $\text{write}(\bar{n}, M)$ outputs the number n then continues as M .

These operations seem very similar to those for global store. One difference is that updating location l with value n then immediately looking up the value will always yield \bar{n} . This is not the case for I/O operations write and read . There are no guarantees about the values on the output and input channels.

2.5 Operational Semantics of EPCF

Simpson and Voorneveld define an operational semantics for EPCF where computations evaluate to trees, following Plotkin and Power [PP01].

Definition 2.5.1. *The operational semantics uses evaluation stacks S to implement sequencing. They are defined as:*

$$S := \text{id} \mid S \circ (\mathbf{let} (-) \Rightarrow x \mathbf{in} M).$$

Define the operation of ‘filling in the hole’ of a stack with a closed computation as:

$$\text{id}\{N\} = N$$

$$(S \circ (\mathbf{let} (-) \Rightarrow x \mathbf{in} M))\{N\} = S\{\mathbf{let} N \Rightarrow x \mathbf{in} M\}.$$

Write $\text{Stack}(\tau, \rho)$ for the set of stacks S which when given a computation $N \in \text{Comp}(\tau)$, return a computation $S\{N\} \in \text{Comp}(\rho)$.

The operational semantics consists of two relations, one between closed computations, and one between configurations (S, M) , where $M \in \text{Comp}(\tau)$ and $S \in \text{Stack}(\tau, \rho)$.

$$(\lambda x:\tau.M) V \rightsquigarrow M[V/x]$$

$$\mathbf{fix} F \rightsquigarrow \mathbf{return} \lambda x:\tau.\mathbf{let} F (\lambda y:\tau.\mathbf{let} \mathbf{fix} F \Rightarrow z \mathbf{in} z y) \Rightarrow w \mathbf{in} w x$$

$$\mathbf{case} Z \mathbf{in} \{Z \Rightarrow M, S(x) \Rightarrow N\} \rightsquigarrow M$$

$$\mathbf{case} S(V) \mathbf{in} \{Z \Rightarrow M, S(x) \Rightarrow N\} \rightsquigarrow N[V/x]$$

$$\begin{aligned}
(S, \mathbf{let} \ N \Rightarrow x \ \mathbf{in} \ M) &\mapsto (S \circ \mathbf{let} \ (-) \Rightarrow x \ \mathbf{in} \ M, N) \\
(S \circ \mathbf{let} \ (-) \Rightarrow x \ \mathbf{in} \ M, \mathbf{return} \ V) &\mapsto (S, M[V/x]) \\
(S, M) &\mapsto (S, M') \text{ if } M \rightsquigarrow M'
\end{aligned}$$

Denote by \mapsto^* the reflexive-transitive closure of \mapsto .

The reduction rule for the fixed point $\mathbf{fix} \ F$ is somewhat complicated by the syntactic restrictions imposed by fine-grained call-by-value. Intuitively, we can think of $\mathbf{fix} \ F$ as reducing to the thunk of $F(\mathbf{fix} \ F)$.

By inspecting the reduction relation \mapsto we can see that it is deterministic. There are two ways \mapsto can get stuck. If $(S, M) \mapsto^* (id, \mathbf{return} \ V)$; in this case there is nothing left to do so the computation should terminate. Or if $(S, M) \mapsto^* (S', \sigma(\dots))$. In this case, an effect operation should take place and the execution should continue from S' with the computation chosen by the effect operation. It is also possible that \mapsto never terminates due to the presence of recursion.

This suggests that a computation of type τ should evaluate to an *effect tree* with leaves values of type τ . Denote the set of all such trees by $Trees(\tau)$. A possibly infinite tree in $Trees(\tau)$ can have:

- a leaf labelled by \perp , which signifies nontermination of \mapsto ;
- a leaf labelled by a value $V \in Val(\tau)$;
- a node labelled σ where $(\sigma : \alpha^n \rightarrow \alpha) \in \Sigma$; this has children t_0, t_1, \dots, t_{n-1} ;
- a node labelled σ where $(\sigma : \alpha^{\mathbb{N}} \rightarrow \alpha) \in \Sigma$; this has infinitely many children t_0, t_1, \dots , one for each natural number;
- a node labelled σ_m where $(\sigma : \mathbb{N} \times \alpha^n \rightarrow \alpha) \in \Sigma$; the label m is the natural number that σ takes as an argument; the node has children t_0, t_1, \dots, t_{n-1} ;
- a node labelled σ_m where $(\sigma : \mathbb{N} \times \alpha^{\mathbb{N}} \rightarrow \alpha) \in \Sigma$; this has children t_0, t_1, \dots .

Before defining the tree associated to each computation, we need to define a partial order on $Trees(\tau)$ as follows:

$$tr_1 \leq tr_2 \iff tr_1 \text{ can be obtained from } tr_2 \text{ by replacing some of its subtrees by } \perp.$$

This ordering endows $Trees(\tau)$ with an ω -CPO structure [Fio17] with least element \perp . This means that every increasing chain $tr_1 \leq tr_2 \leq \dots$ has a least upper bound $\bigsqcup_{n \in \mathbb{N}} tr_n$.

Definition 2.5.1 (Continued). *Define a family of functions*

$$|- , -|_{(-)} : Stack(\tau, \rho) \times Comp(\tau) \times \mathbb{N} \longrightarrow Trees(\rho).$$

The tree $|S, M|_n$ represents the unfolding of computation M for n steps starting in stack

S. The formal definition is:

$$|S, M|_0 = \perp$$

$$|S, M|_{n+1} = \begin{cases} V & \text{if } S = id \text{ and } M = \mathbf{return} V \\ |S', M'|_n & \text{if } (S, M) \mapsto (S', M') \\ \sigma(|S, M_0|_n, \dots, |S, M_{k-1}|_n) & \text{if } \sigma : \alpha^k \rightarrow \alpha \\ & \text{and } M = \sigma(M_0, M_1, \dots, M_{k-1}) \\ \sigma(|S, V \bar{0}|_n, |S, V \bar{1}|_n, \dots) & \text{if } \sigma : \alpha^{\mathbb{N}} \rightarrow \alpha \text{ and } M = \sigma(V) \\ \sigma_m(|S, M_0|_n, \dots, |S, M_{k-1}|_n) & \text{if } \sigma : \mathbb{N} \times \alpha^k \rightarrow \alpha \\ & \text{and } M = \sigma(\bar{m}; M_0, M_1, \dots, M_{k-1}) \\ \sigma_m(|S, V \bar{0}|_n, |S, V \bar{1}|_n, \dots) & \text{if } \sigma : \mathbb{N} \times \alpha^{\mathbb{N}} \rightarrow \alpha \text{ and } M = \sigma(\bar{m}; V) \\ \perp & \text{otherwise.} \end{cases}$$

From this definition we can see that $|S, M|_n \leq |S, M|_{n+1}$. Therefore, the effect tree associated with a computation is defined as:

$$|-| : Comp(\tau) \longrightarrow Trees(\tau)$$

$$|M| = \bigsqcup_{n \in \mathbb{N}} |id, M|_n.$$

We call $|M|$ a computation tree.

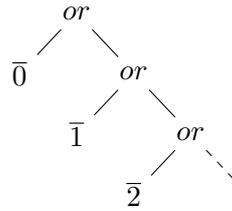
The intuitive interpretation of computation trees is that a path through the tree represents a potential execution path of the program. However, the operational semantics is not aware of this interpretation. In this sense, the effect operations are purely formal as in [PP01]. This will be illustrated in Example 2.5.5 below.

Example 2.5.2 (Pure functional computation). In this case, there are no effect operations so all computation trees are leaves.

Example 2.5.3 (Nondeterminism). Consider a computation that generates a natural number nondeterministically (from [SV18]):

$$?nat = \mathbf{let} \ \mathbf{fix} \ (\lambda f : \mathbb{1} \rightarrow \mathbb{N}. \mathit{or}(\lambda y : \mathbb{1}. Z, \lambda y : \mathbb{1}. \mathbf{let} \ f \ y \Rightarrow u \ \mathbf{in} \ S(u))) \Rightarrow w \ \mathbf{in} \ w \ \star.$$

Its computation tree is:

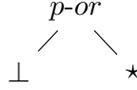


Thus, the notion of a tree whose nodes are nondeterministic choice points, discussed informally in Section 2.2, is formalised by a computation tree.

Example 2.5.4 (Probabilistic choice). Define the following computation whose execution never terminates:

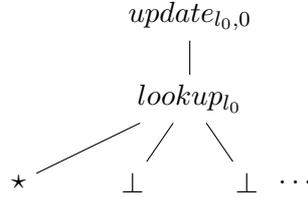
$$loop = \mathbf{let} \ (\mathbf{fix} \ \lambda f:\mathbb{N} \rightarrow \mathbb{1}.\mathbf{return} \ f) \Rightarrow g \ \mathbf{in} \ g \ Z.$$

The computation tree of $p\text{-or}(loop, \mathbf{return} \ \star) : \mathbb{1}$ is therefore:



Example 2.5.5 (Global store). Consider a location $l_0 \in \mathbb{L}$. The following computation writes $\bar{0}$ to location l_0 then immediately reads this value:

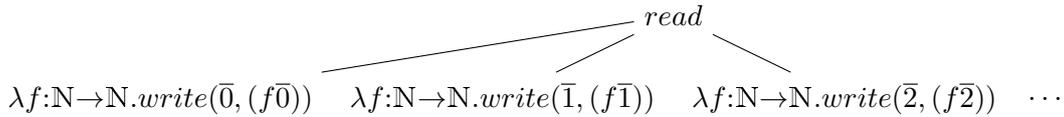
$$|update_{l_0}(\bar{0}; lookup_{l_0}(\lambda x:\mathbb{N}.\mathbf{case} \ x \ \mathbf{in} \ \{Z \Rightarrow \mathbf{return} \ \star, S(y) \Rightarrow loop\}))| =$$



According to the intuitive interpretation that we give to the *lookup* and *update* operations only the path that returns a \star can occur. However, the tree contains a path for each natural number that could be in l_0 because the $|-|$ function treats operation symbols as syntax without any interpretation. If we replaced *lookup* and *update* with the *read* and *write* operations from I/O then all paths would be relevant.

Example 2.5.6 (Input/output). The following computation reads a natural number from the input channel, then returns a function whose behaviour depends on this number:

$$read(\lambda x:\mathbb{N}.\mathbf{return} \ \lambda f:\mathbb{N} \rightarrow \mathbb{N}.\mathbf{write}(x; (f \ x))) \quad : (\mathbb{N} \rightarrow \mathbb{N}) \rightarrow \mathbb{N}$$



A path through the tree is not only a possible execution path of the computation, it also corresponds to an I/O trace.

2.6 Summary of Results about EPCF

The main contribution of Simpson's and Voorneveld's work [SV18] is that they define a modal logic whose formulas represent properties of EPCF programs. We will refer to it as EPCF logic. Moreover, they give a general analysis of the modalities involved.

In EPCF logic, there are two kinds of formulas all of which are attached an EPCF type. A formula $\phi : \tau$ describes a value of type τ , while $\Phi : \tau$ describes a computation of type τ . The definition of the logic starts from a set of basic formulas at each type which is then closed under negation, arbitrary conjunctions and disjunctions.

The basic formulas for values of type \mathbb{N} are:

$$\{n\} \text{ where } n \in \mathbb{N}.$$

A closed value $W : \mathbb{N}$ satisfies $\{n\}$, written $W \models \{n\}$, if and only if $W = \bar{n}$.

For values of function types $\tau \rightarrow \rho$, the basic formulas are:

$$\phi \mapsto \Phi$$

where ϕ is a value formula of type τ and Φ is a computation formula of type ρ . The satisfaction $W \models \phi \mapsto \Phi$ holds if and only if:

$$\forall V \models \phi. (W V) \models \Phi.$$

A value formula $\phi \mapsto \Phi$ tests the behaviour of a function when it is being applied. This behaviour is the fundamental property of a function, hence the choice of value formula is reasonable.

Finally, computation formulas make use of a set of *modalities* \mathcal{O} . The set \mathcal{O} contains sets of effect trees of type $\mathbb{1}$, that is, $\mathcal{O} \subseteq \mathcal{P}(\text{Trees}(\mathbb{1}))$. A basic computation formula of type τ is:

$$o\phi$$

where $o \in \mathcal{O}$ and ϕ is a value formula of type τ . We can see that o lifts a formula for values, ϕ , to a formula for computations. This is why o is named a *modality*.

For a tree $tr \in \text{Trees}(\tau)$ and a value formula $\phi : \tau$ denote by:

$$tr \models \phi$$

the tree in $\text{Trees}(\mathbb{1})$ obtained by replacing the leaves V of tr by \star if $V \models \phi$ and by \perp otherwise. We can now define satisfaction of computation formulas as:

$$M \models o\phi \iff |M| \models \phi \in o.$$

A computation formula $o\phi$ tests whether the possible return values of a computation satisfy ϕ and also tests the shape of the effect tree of the computation. These two pieces of information form the *observable behaviour* of a computation, in the sense of Section 2.1.

The definition of \mathcal{O} depends on the effects present in the language so we will look at the nondeterminism example. The other effects are treated in more detail in Chapter 5, in the context of ECPS.

Example 2.6.1 (Nondeterminism). Define $\mathcal{O} = \{\diamond, \square\}$ where:

$$\diamond = \{tr \in \text{Trees}(\mathbb{1}) \mid tr \text{ has some } \star \text{ leaf}\}$$

$$\square = \{tr \in \text{Trees}(\mathbb{1}) \mid tr \text{ has finite height and every leaf is a } \star\}.$$

The formula $\diamond\phi$ says that a computation may return a value satisfying ϕ , whereas $\square\phi$ asserts that it must return such a value. We can see that \square and \diamond which we discussed informally in Examples 2.2.1 and 2.2.2 are now defined as modalities.

In fact, the formulas used in those examples are valid formulas in EPCF logic:

$$\begin{aligned}\phi_1 &= \{3\} \mapsto \Box\{2\} \\ \phi_2 &= \{3\} \mapsto \Diamond\{2\} \\ \psi &= (\{3\} \mapsto \Diamond\{2\}) \mapsto ((\{2\} \vee \{4\}) \mapsto \Diamond\{2\})\end{aligned}$$

and the example functions:

$$\begin{aligned}g &= \lambda n:\mathbb{N}.or(\mathbf{pred} \ n, \ \mathbf{succ} \ n) \\ h &= \lambda f:\mathbb{N}\rightarrow\mathbb{N}. \lambda n:\mathbb{N}. or(f \ (\mathbf{pred} \ n), \ f \ (\mathbf{succ} \ n))\end{aligned}$$

are valid EPCF terms for suitable encodings of **succ** and **pred**. We can see that $(id, (g \ \bar{3}))$ may return either $\bar{2}$ or $\bar{4}$ so using the definition of logical satisfaction we indeed obtain:

$$g \not\models \phi_1 \quad \text{and} \quad g \models \phi_2.$$

Using the same definition we can also see that $h \models \psi$.

Simpson and Voorneveld [SV18] are concerned with two notions of program equivalence: applicative bisimilarity and logical equivalence induced by satisfaction in EPCF logic. Applicative bisimilarity is a family of relations defined between well-typed EPCF terms. The definition is technically involved so we omit it. However, some key features of applicative bisimilarity are:

- Two natural numbers are bisimilar if and only if they are equal.
- Two function values are bisimilar if and only if for all arguments they yield bisimilar computations. This condition is the same as in the informal explanation of similarity from Section 2.1.
- To specify when two computations are bisimilar, their effect trees are inspected. Roughly speaking, M and N are bisimilar if: the possible return values of N approximate those of M well enough to preserve the properties of the effect tree of M , and vice-versa. The set of modalities \mathcal{O} is used express properties of computation trees and to lift bisimilarity of values to a relation between trees.

Because the definition of bisimilarity depends on the set of modalities \mathcal{O} , it is named applicative \mathcal{O} -bisimilarity.

Definition 2.6.2. *Logical equivalence between well-typed EPCF terms is defined as:*

$$\begin{aligned}V \equiv_{EPCF} W &\iff \forall \phi : \tau. V \models \phi \iff W \models \phi \\ M \equiv_{EPCF} N &\iff \forall \Phi : \tau. M \models \Phi \iff N \models \Phi.\end{aligned}$$

It is easy to prove that logical equivalence and applicative bisimilarity for EPCF are equivalence relations. In order to prove they are *compatible*, Simpson and Voorneveld identify two sufficient conditions that the set of modalities \mathcal{O} should satisfy in general:

A computation that satisfies this formula is one such that: at least one of its possible return values is a function which, when given argument \star , *may* return *any* natural number. Note that for type $\mathbb{1}$, the only formulas are *true* and *false*, represented as the empty conjunction and disjunction respectively, and the only value is \star .

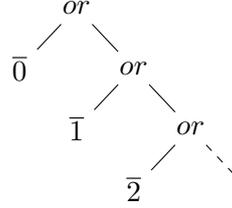
We can show that $M \models \Phi$. We need to check that $|M| \models true \mapsto \bigwedge_{n \in \mathbb{N}} \diamond\{n\} \in \diamond$. Computation M returns immediately so $|M|$ is just a leaf labelled $\lambda x:\mathbb{1}.?nat$. The set \diamond contains the trees with at least one \star leaf. So it is sufficient to check:

$$\lambda x:\mathbb{1}.?nat \models true \mapsto \bigwedge_{n \in \mathbb{N}} \diamond\{n\}$$

that is,

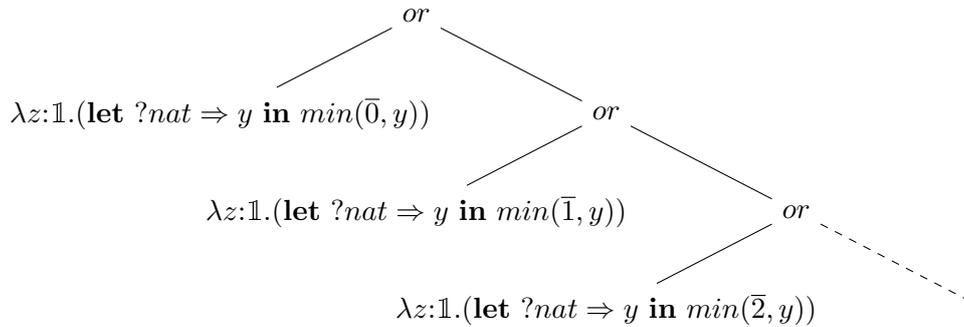
$$?nat \models \bigwedge_{n \in \mathbb{N}} \diamond\{n\}.$$

Recall that the computation tree of $?nat$ is:



It has a leaf labelled with each natural number, therefore $?nat \models \bigwedge_{n \in \mathbb{N}} \diamond\{n\}$ is true.

On the other hand, $P \not\models \Phi$. To see this note that $|P|$ is:



In order for $|P| \models true \mapsto \bigwedge_{n \in \mathbb{N}} \diamond\{n\} \in \diamond$ to be true we need to find some $m \in \mathbb{N}$ such that:

$$\lambda z:\mathbb{1}.(\mathbf{let} ?nat \Rightarrow y \mathbf{in} \mathit{min}(\bar{m}, y)) \models true \mapsto \bigwedge_{n \in \mathbb{N}} \diamond\{n\}$$

that is,

$$\mathbf{let} ?nat \Rightarrow y \mathbf{in} \mathit{min}(\bar{m}, y) \models \bigwedge_{n \in \mathbb{N}} \diamond\{n\}.$$

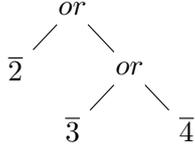
By contradiction, assume that such an m exists. The tree of $\mathbf{let} ?nat \Rightarrow y \mathbf{in} \mathit{min}(\bar{m}, y)$ has as leaves all the numbers from 0 to m , but none greater than m :

tree. For example the computation M where

$$F = \lambda f:\mathbb{N}\rightarrow\mathbb{N}.or(f \bar{1}, or(f \bar{2}, f \bar{3}))$$

$$M = F (\lambda n:\mathbb{N}.\mathbf{return} S(n))$$

has tree:



Simpson and Voorneveld [SV18] proposed a modal logic that expresses properties of EPCF programs, named EPCF logic (Section 2.6). For each effect, there is a set of *modalities* that express properties of computations which exhibit that effect. For nondeterminism these are \Box and \Diamond . For example, $\Phi_1 = \Box(\{2\} \vee \{3\} \vee \{4\})$ says that computation M *always* returns a result from the set $\{2, 3, 4\}$, whereas $\Phi_2 = \Diamond\{2\}$ says that M *may* return 2.

For functions, logical formulas have the form $\phi \mapsto \Phi$. This says that, if the argument of the function satisfies ϕ , then the resulting application satisfies Φ . For example, F satisfies the following property: $F \models (\{1\} \mapsto \Box\{2\}) \mapsto \Diamond\{2\}$.

Simpson and Voorneveld defined applicative bisimilarity for EPCF using modalities and proved it compatible. They showed that program equivalence induced by EPCF logic coincides with applicative bisimilarity but not with contextual equivalence. Section 2.7 described two EPCF programs exhibiting nondeterminism that are contextually equivalent but not logically equivalent. The problem is that a context can only test a function on a finite number of arguments, whereas a logical formula can test it on infinitely many arguments, using infinitary connectives. Thus, we have identified the main problem of the dissertation: finding a logic that characterises contextual equivalence for a higher-order language with algebraic effects.

Chapter 3

Introducing the ECPS Language

This chapter introduces the programming language ECPS and its operational semantics. ECPS will be used in the rest of the dissertation to study program equivalence. Moreover, a few results that will be used later are outlined: a coinduction proof principle and general intuitions about logical relations.

3.1 A New Language – ECPS

To make it easier to formulate a logic that characterises contextual equivalence, we introduce a new programming language ECPS. It is a variant of EPCF in which programs are written in continuation-passing style (CPS) [Rey93]. This means that functions carry an additional argument, namely the *continuation* to which they pass their result. The intuition is that a continuation specifies how the execution should proceed once a function has finished.

Given a fixed return type R , a continuation has type $\alpha \rightarrow R$. It is waiting for an argument of type α to produce a return value of type R . Consider for example a function that adds two natural numbers. Usually, it has type $\mathbf{nat} \rightarrow \mathbf{nat} \rightarrow \mathbf{nat}$. In continuation-passing style, this function would look like:

$$\begin{aligned} \mathbf{addc} &: \mathbf{nat} \rightarrow \mathbf{nat} \rightarrow (\mathbf{nat} \rightarrow R) \rightarrow R \\ \mathbf{addc} \ n \ m \ cont &= cont \ (n + m). \end{aligned}$$

Instead of directly returning the result $n + m$, the function \mathbf{addc} passes it to the continuation $cont$.

The key property of ECPS that allows the formulation of the new logic, which will be introduced in Section 6.1, is that functions do not have a return type. Once a function has been applied, the resulting computation is expected to run forever. In other words, the return type R of continuations is chosen to be \perp . Thus, programs no longer return values that we can observe. We might however observe termination, which is now treated as an effect, or other side effects such as output values.

Definition 3.1.1 (ECPS). *The types are defined by the following grammar:*

$$A, B := \neg(A_1, \dots, A_n) \mid \mathbf{nat} \mid \mathbf{unit}.$$

$$\begin{array}{c}
\frac{}{\Gamma, x : A \vdash x : A} \text{(VAR)} \quad \frac{\Gamma, x_1 : A_1, \dots, x_n : A_n \vdash t}{\Gamma \vdash \lambda(x_1, \dots, x_n) : (A_1, \dots, A_n).t : \neg(A_1, \dots, A_n)} \text{(LBD)} \\
\\
\frac{\Gamma \vdash v : \neg(A_1, \dots, A_n) \quad \Gamma \vdash w_1 : A_1 \quad \dots \quad \Gamma \vdash w_n : A_n}{\Gamma \vdash v(w_1, \dots, w_n)} \text{(APP)} \\
\\
\frac{}{\Gamma \vdash \mathbf{zero} : \mathbf{nat}} \text{(ZERO)} \quad \frac{\Gamma \vdash v : \mathbf{nat}}{\Gamma \vdash \mathbf{succ}(v) : \mathbf{nat}} \text{(SUCC)} \quad \frac{}{\Gamma \vdash \star : \mathbf{unit}} \text{(UNIT)} \\
\\
\frac{\Gamma \vdash v : \mathbf{nat} \quad \Gamma \vdash t \quad \Gamma, x : \mathbf{nat} \vdash s}{\Gamma \vdash \mathbf{case } v \text{ in } \{\mathbf{zero} \Rightarrow t, \mathbf{succ}(x) \Rightarrow s\}} \text{(CASE)} \\
\\
\frac{\Gamma, x : \neg(\vec{A}) \vdash v : \neg(\vec{A}) \quad \Gamma \vdash \vec{w} : \vec{A}}{\Gamma \vdash (\mu x.v)(\vec{w})} \text{(MU)} \\
\\
\frac{\Gamma, x : \mathbf{nat} \vdash t \quad \Gamma \vdash v : \mathbf{nat}}{\Gamma \vdash \sigma(v, x.t)} \sigma \in \Sigma \text{ (OP)} \quad \frac{}{\Gamma \vdash \downarrow} \text{(STOP)}
\end{array}$$

Figure 3.1: Typing judgements for ECPS.

Fix a set Σ of effect operations σ , each with arity $\mathbf{nat} \times \alpha^{\mathbf{nat}} \rightarrow \alpha$, where α stands for a computation. Values and computations are defined respectively as:

$$v, w := \mathbf{zero} \mid \mathbf{succ}(v) \mid \star \mid \lambda(x_1, \dots, x_n) : (A_1, \dots, A_n).t \mid x$$

$$t, u := v(w_1, \dots, w_n) \mid (\mu x.v)(\vec{w}) \mid \sigma(v, x.t) \mid \downarrow \mid \mathbf{case } v \text{ in } \{\mathbf{zero} \Rightarrow t, \mathbf{succ}(x) \Rightarrow u\}.$$

There are two typing relations, one for values $\Gamma \vdash v : \tau$, and one for computations which do not have a type, $\Gamma \vdash t$. These are the least relations closed under the rules in Figure 3.1.

ECPS is a fine-grained call-by-value language; it makes a distinction between values and computations. Ignoring effects, ECPS is in fact a fragment of Levy's Jump-With-Argument programming language [LL07].

The type $\neg(A_1, \dots, A_n)$ is the type of a function which takes n arguments of types A_1, \dots, A_n respectively and returns a computation. This computation is not expected to terminate so we can think of $\neg(A_1, \dots, A_n)$ as $(A_1, \dots, A_n) \rightarrow \perp$. We can also think of $k : \neg A$ as a continuation of type $A \rightarrow R$, where the return type of all continuations R has been set to \perp . The base types \mathbf{nat} and \mathbf{unit} are the same as in EPCF.

In $\mu x.v$, the expression v is constrained by the typing rules to be a function. Thus, intuitively $\mu x.v$ is a recursive definition of the function v , where x represents v and can appear free inside v . Computation $(\mu x.v)(\vec{w})$ arises by applying this function to \vec{w} .

Example 3.1.2. We can now implement the function `addc` in ECPS using continuations and recursion. The type of this function is:

$$\mathbf{addc} : \neg(\mathbf{nat}, \mathbf{nat}, \neg\mathbf{nat})$$

$$\begin{aligned}
\text{addc} &= \lambda(x, y, k) : (\text{nat}, \text{nat}, \neg\text{nat}). \\
&(\mu f. \lambda(x', y', k') : (\text{nat}, \text{nat}, \neg\text{nat}). \\
&\quad \text{case } x' \text{ in } \{\text{zero} \Rightarrow k' y', \\
&\quad\quad \text{succ}(x'') \Rightarrow \text{case } y' \text{ in} \\
&\quad\quad\quad \{\text{zero} \Rightarrow k' x', \\
&\quad\quad\quad\quad \text{succ}(y'') \Rightarrow f(x'', y'', \lambda z : \text{nat}. k' \text{succ}(\text{succ}(z)))\}) \\
&\quad) (x, y, k).
\end{aligned}$$

The behaviour of this function can be explained intuitively as follows: if x is zero then the result of the addition is y , so y is passed to the current continuation k . Similarly if y is zero. If x and y are both greater than zero, add $x - 1$ and $y - 1$ and pass the result to continuation $\lambda z : \text{nat}. k' \text{succ}(\text{succ}(z))$. This continuation adds two to $(x - 1) + (y - 1)$, then passes the result to the current continuation k . This explanation might become clearer if read in conjunction with the operational semantics from the next section.

ECPS does not have a **let** constructor for sequencing. Sequencing can be achieved instead by manipulating the continuation passed to a program.

As discussed above, termination in ECPS is an effect. Its associated effect operation is \downarrow , which does not take any arguments.

The different arities for effect operations from EPCF are conflated into the most general one: $\text{nat} \times \alpha^{\text{nat}} \rightarrow \alpha$. Therefore, operation σ takes as arguments a natural number v and a function from a natural number x to a computation t , and returns a computation $\sigma(v, x.t)$.

Substitution of values for free variables, $v[w/x]$ and $t[w/x]$, is defined in a standard way by recursion on the structure of v and t . We will use \bar{n} to denote the natural number $\text{succ}^n(\text{zero})$. Let (\vdash) be the set of closed computations and $(\vdash A)$ the set of closed values of type A .

3.2 Operational Semantics of ECPS

Next, we present the operational semantics of ECPS. It does not need to use stacks because any control flow is explicitly encoded inside a computation using continuations.

Definition 3.2.1. *The operational semantics is given by two families of relations on closed computation terms*

$$(\longrightarrow) \subseteq (\vdash) \times (\vdash)$$

$$(\xrightarrow{\sigma(v)}) \subseteq (\vdash) \times (\vdash)^{\mathbb{N}} \quad \text{for any } \sigma \in \Sigma \text{ and } \vdash v : \text{nat}$$

defined as:

$$\begin{aligned}
&\sigma(v, x.t) \xrightarrow{\sigma(v)} (t[\bar{n}/x])_{n \in \mathbb{N}} \\
&(\lambda(\vec{x}) : (\vec{A}).t) (\vec{w}) \longrightarrow t[\vec{w}/\vec{x}] \\
&(\mu x.v) (\vec{w}) \longrightarrow (v[(\lambda(\vec{y}) : (\vec{A}).(\mu x.v)(\vec{y}))/x]) (\vec{w}) \\
&\quad \text{case zero in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\} \longrightarrow s \\
&\quad \text{case succ}(v) \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\} \longrightarrow t[v/x].
\end{aligned}$$

Denote by \longrightarrow^* the reflexive-transitive closure of \longrightarrow .

We can see that the \longrightarrow reduction relation can only get stuck when encountering an effect operation $\sigma(\dots)$ or \downarrow . It might also be the case that \longrightarrow never terminates.

There are no reduction rules of any kind for \downarrow since it signifies termination. For $\sigma \in \Sigma$, the $\xrightarrow{\sigma(v)}$ reduction rule has on its right-hand-side a set of computations, one for each natural number. Therefore, repeated applications of this rule lead to the construction of an infinitely branching tree.

Given this observation we can define effect trees for ECPS. Denote the set of all effect trees by $Trees_\Sigma$. A tree in this set can have:

- a leaf labelled \perp , which signifies nontermination of \longrightarrow ;
- a leaf labelled \downarrow , which signifies termination;
- nodes labelled σ_n , where $\sigma \in \Sigma$ and $n \in \mathbb{N}$; such a node has an infinite number of children t_0, t_1, \dots

We can define a partial order on $Trees_\Sigma$ which makes it an ω -CPO. This is similar to the order on EPCF trees:

$$tr_1 \leq tr_2 \iff tr_1 \text{ can be obtained by replacing subtrees of } tr_2 \text{ by } \perp.$$

Using this order, we can give a domain theoretic definition of the tree associated to a closed computation. The tree might have infinite depth and width.

Definition 3.2.2 (Computation trees for ECPS). *Define a family of maps*

$$\llbracket - \rrbracket_{(-)} : (\vdash) \times \mathbb{N} \longrightarrow Trees_\Sigma$$

$$\llbracket t \rrbracket_0 = \perp$$

$$\llbracket t \rrbracket_{n+1} = \begin{cases} \llbracket s \rrbracket_n & \text{if } t \longrightarrow s \\ \sigma_m(\llbracket s[0/x] \rrbracket_n, \dots, \llbracket s[\bar{k}/x] \rrbracket_n, \dots) & \text{if } t \xrightarrow{\sigma(\bar{m})} (s[\bar{k}/x])_{k \in \mathbb{N}} \\ \downarrow & \text{if } t = \downarrow \\ \perp & \text{otherwise.} \end{cases}$$

We can see that $\llbracket t \rrbracket_n \leq \llbracket t \rrbracket_{n+1}$ so we can define $\llbracket - \rrbracket : (\vdash) \longrightarrow Trees_\Sigma$ as the least upper bound of the chain $\{\llbracket t_n \rrbracket\}_{n \in \mathbb{N}}$:

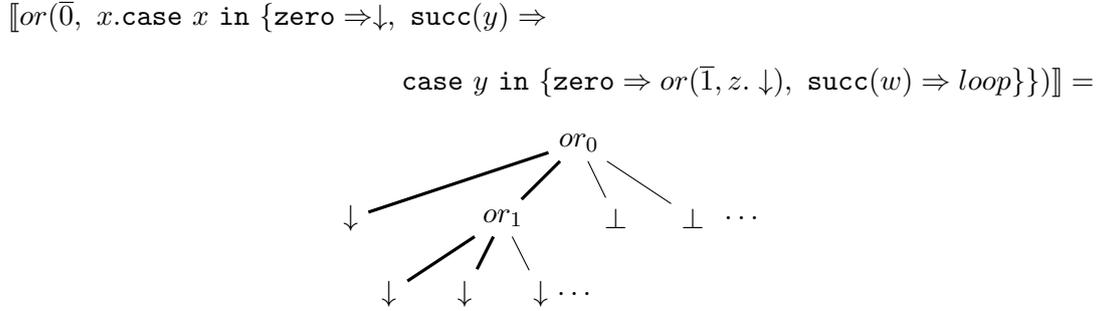
$$\llbracket t \rrbracket = \bigsqcup_{n \in \mathbb{N}} \llbracket t \rrbracket_n.$$

Example 3.2.3 (Pure functional computation). In this case, the signature Σ is empty and the only effect operation is \downarrow . Therefore, computation trees can only be leaves: \downarrow for a computation that terminates and \perp for one that does not. For example:

$$\llbracket loop \rrbracket = \llbracket (\mu f. \lambda x: \text{nat}. (f x)) \text{ zero} \rrbracket = \perp.$$

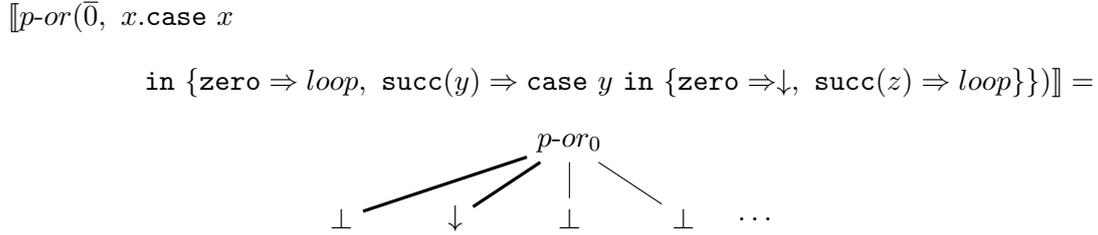
Example 3.2.4 (Nondeterminism). Define $\Sigma = \{or\}$. All effects have arity $\mathbf{nat} \times \alpha^{\mathbf{nat}} \rightarrow \alpha$. The intuitive interpretation of $or(v, x.t)$ is that it ignores v and performs a nondeterministic choice between $t[\bar{0}/x]$ and $t[\bar{1}/x]$.

For example, consider the computation tree of:



As far as the intuitive interpretation of this computation is concerned, the indices 0 and 1 are irrelevant, and only the paths highlighted in bold can occur. However, $\llbracket - \rrbracket$ treats or as uninterpreted syntax, hence the paths that can never occur are still present in the effect tree.

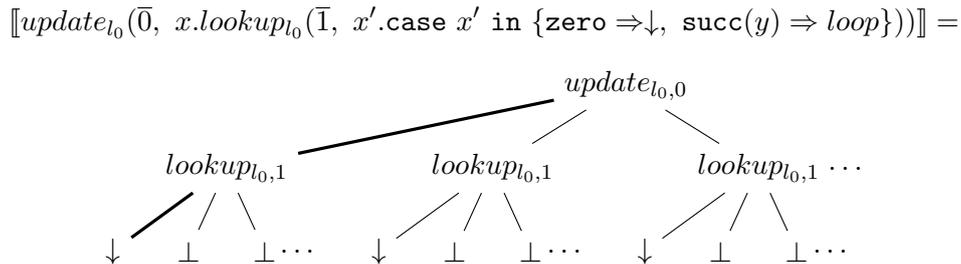
Example 3.2.5 (Probabilistic choice). Define $\Sigma = \{p-or\}$. Intuitively, the operation $p-or(v, x.t)$ chooses between $t[\bar{0}/x]$ and $t[\bar{1}/x]$ with probability 0.5. The rest of the branches can never occur. The following computation:



is analogous to the EPCF computation $p-or(\text{loop}, \mathbf{return} \star)$ from Example 2.5.4, irrespective of the index of $p-or$, here 0.

Example 3.2.6 (Global store). There is a finite set of locations \mathbb{L} that can store natural numbers and $\Sigma = \{lookup_l, update_l \mid l \in \mathbb{L}\}$. The intuitive interpretation of $lookup_l(v, x.t)$ is that it ignores v , it looks up the value at location l , if this is \bar{n} it continues with $t[\bar{n}/x]$. For $update_l(v, x.t)$ the intuition is: write the number v in location l then continue with the computation $t[\bar{0}/x]$.

The EPCF computation tree from Example 2.5.5 can be adapted here:

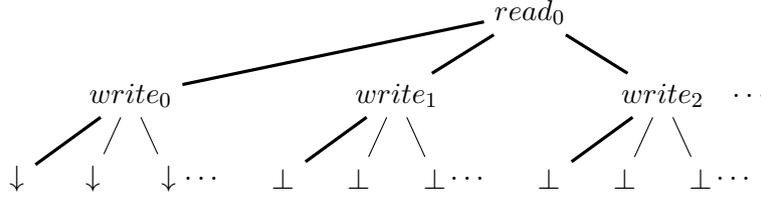


Only the path in bold can occur in the computation above.

Example 3.2.7 (Input/output). Define $\Sigma = \{read, write\}$. Intuitively, the computation $read(v, x.t)$ ignores v , accepts as input a number \bar{n} and continues with $t[\bar{n}/x]$. The computation $write(v, x.t)$ writes v to the output channel then continues with computation $t[\bar{0}/x]$.

Below is a computation that inputs a number then outputs it immediately. Only the paths in bold can occur:

$$\llbracket read(\bar{0}, x.write(x, x'.case\ x\ in\ \{zero\Rightarrow\downarrow,\ succ(y)\Rightarrow loop\})) \rrbracket =$$



3.3 A Coinduction Proof Principle

Coinduction is a proof techniques that will be used in the following chapter. We give an abstract overview of a coinduction proof principle using the basic notions of a category, functor, terminal object and coalgebra, all of which can be found in an introduction to category theory such as [AT18]. The following definition appears in [JR11]:

Definition 3.3.1. Let $T : \mathbf{Set} \Rightarrow \mathbf{Set}$ be a functor. Take two T -coalgebras $(X, a_X : X \rightarrow T(X))$ and $(Y, a_Y : Y \rightarrow T(Y))$. A T -bisimulation between (X, a_X) and (Y, a_Y) is a relation $\mathcal{R} \subseteq X \times Y$ for which there exists a T -coalgebra structure $g : \mathcal{R} \rightarrow T(\mathcal{R})$ such that the two projection functions $\pi_1 : \mathcal{R} \rightarrow X$ and $\pi_2 : \mathcal{R} \rightarrow Y$ are T -coalgebra morphisms:

$$\begin{array}{ccccc} X & \xleftarrow{\pi_1} & \mathcal{R} & \xrightarrow{\pi_2} & Y \\ \downarrow a_X & & \downarrow g & & \downarrow a_Y \\ T(X) & \xleftarrow{T(\pi_1)} & T(\mathcal{R}) & \xrightarrow{T(\pi_2)} & T(Y) \end{array}$$

Use the following notation:

$$x \cong y \iff \text{there exists a } T\text{-bisimulation } \mathcal{R} \subseteq X \times Y \text{ with } (x, y) \in \mathcal{R}.$$

We can now formulate the following coinduction proof principle:

Proposition 3.3.2 (From [JR11]). Consider the final T -coalgebra $(Z, c : Z \rightarrow T(Z))$, if it exists. Let $\alpha_X : X \rightarrow Z$ and $\alpha_Y : Y \rightarrow Z$ be coalgebra morphisms. For all $x \in X$ and $y \in Y$:

$$\text{if } x \cong y \text{ then } \alpha_X(x) = \alpha_Y(y).$$

Proof. It suffices to show that the following diagram commutes, where $\mathcal{R} \subseteq X \times Y$ is a bisimulation and $(x, y) \in \mathcal{R}$:

$$\begin{array}{ccccccccc}
 Z & \xleftarrow{\alpha_X} & X & \xleftarrow{\pi_1} & \mathcal{R} & \xrightarrow{\pi_2} & Y & \xrightarrow{\alpha_Y} & Z \\
 \downarrow c & & \downarrow a_X & & \downarrow g & & \downarrow a_Y & & \downarrow c \\
 T(Z) & \xleftarrow{T(\alpha_X)} & T(X) & \xleftarrow{T(\pi_1)} & T(\mathcal{R}) & \xrightarrow{T(\pi_2)} & T(Y) & \xrightarrow{T(\alpha_Y)} & T(Z)
 \end{array}$$

This is true because all the small squares commute. Then both $\alpha_X \circ \pi_1$ and $\alpha_Y \circ \pi_2$ are coalgebra morphisms into Z . By finality of Z , they must be equal. \square

3.4 Logical Relations

“Logical relations” is a proof technique that involves defining a family of relations by induction on the types of a programming language. A relation for type τ contains only pairs of terms of type τ . The relation is *logical* if, given related functions f_1 and f_2 of type $\rho \rightarrow \tau$ and related arguments $x_1 : \rho$ and $x_2 : \rho$, the terms $f_1 x_1$ and $f_2 x_2$ are related.

Logical predicates, the unary version of logical relations, have been used to prove strong normalisation of the simply-typed λ -calculus [Tai67], [GTL89, Chapter 6]. Other versions of logical relations have been used, for example, to characterise program equivalence [Ahm06] and to prove compiler correctness [BH09]. These are examples of syntactic logical relations, based on the operational semantics of a language. These are the kinds of relations used in the next chapter to prove that a continuation-passing translation of EPCF into ECPS is correct.

There are also logical relations based on denotational models (e.g. [Pit96, Fio17]). One of their applications is proving that a denotational model of a programming language is computationally adequate. Adequacy, means that denotational equality is a sound technique for establishing contextual equivalence of programs.

The particular flavour of logical relations we will need is *step-indexed biorthogonal* logical relations. As Jaber and Tabareau explain [JT11], biorthogonality allows us to define which terms should be related by specifying their interaction with program contexts.

A biorthogonal logical relation contains a collection of relations on values, \mathcal{V}_τ , for each type τ . There is a collection of relations on program contexts defined using the relations on values:

$$\mathcal{C}_\tau = \{(C_1, C_2) \mid \forall (v_1, v_2) \in \mathcal{V}_\tau. \mathfrak{D}(C_1[v_1], C_2[v_2])\}.$$

And finally, a collection of relations on terms defined using the relations on contexts:

$$\mathcal{T}_\tau = \{(t_1, t_2) \mid \forall (C_1, C_2) \in \mathcal{C}_\tau. \mathfrak{D}(C_1[t_1], C_2[t_2])\}.$$

The notation $\mathfrak{D}(C_1[t_1], C_2[t_2])$ stands for an observation about programs $C_1[t_1]$ and $C_2[t_2]$. The notion of observation is chosen on a case-by-case basis, but it usually involves the reduction behaviour of the terms under consideration.

Step-indexing was introduced to deal with programming languages with recursion. This approach, instead of defining a single relation for type τ , defines a family of relations for type τ indexed by natural numbers. Intuitively, terms in \mathcal{T}_τ^n are allowed to reduce at most n steps. The natural number indices help to break the vicious circle introduced by recursion in proofs about the logical relation.

Step-indexing and biorthogonality can be combined in a straightforward way as explained for example by Pitts [Pit10]. Because EPCF contains recursion, and because computations can only be evaluated in a stack, it is useful to use both step-indexing and biorthogonality when proving correctness of the translation from EPCF to ECPS. This will be explained in the next chapter.

3.5 Chapter Summary

This chapter introduced a new language ECPS in Sections 3.1 and 3.2. ECPS is a continuation-passing variant of EPCF which will be used in the rest of the dissertation to study program equivalence.

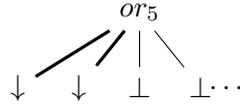
Being a continuation-passing language means that functions receive an additional argument, which specifies how the computation should proceed once the function has terminated. For example, the successor function in ECPS is:

$$f = \lambda(n, k):(\mathbf{nat}, \neg\mathbf{nat}).(k \text{ succ}(n)) : \neg(\mathbf{nat}, \neg\mathbf{nat}). \quad (3.5.1)$$

Here $k : \neg\mathbf{nat}$ is a *continuation*, a function that is waiting for a result of type \mathbf{nat} , but is not expected to return.

Because everything is written in continuation-passing style, ECPS computations do not usually return. Therefore, the operational semantics maps an ECPS computation to a tree whose nodes are effect operations and leaves are either \downarrow , which signifies the termination effect, or \perp for nontermination. For example, consider the tree of:

$$m = (\lambda(n, k):(\mathbf{nat}, \neg\mathbf{nat}). \text{or}(\bar{5}, y. \text{case } y \text{ in } \{\text{zero} \Rightarrow (k \text{ succ}(n)), \text{succ}(y') \Rightarrow \text{case } y' \text{ in } \{\text{zero} \Rightarrow (k \bar{0}), \text{succ}(y'') \Rightarrow \text{loop}\}\})) (\bar{3}, \lambda x:\mathbf{nat}. \downarrow)$$



In ECPS all effect operations have arity $\mathbf{nat} \times \alpha^{\mathbf{nat}} \rightarrow \alpha$, where α stands for a computation. Therefore, each node in a tree has a child for each natural number. In the case of *or*, the tree carries redundant information: only the paths in bold can occur during computation, and the index 5 is irrelevant.

Finally, two proof methods were reviewed: coinduction (Section 3.3) and logical relations (Section 3.4), both of which will be used in the next chapter.

Chapter 4

Comparison: ECPS vs. EPCF

This chapter presents a continuation-passing translation from the language PCF with effects (EPCF), introduced in Chapter 2, to its continuation-passing variant (ECPS), defined in Chapter 3. This translation is proved correct using logical relations and coinduction. The last section is an informal argument for why translating ECPS into EPCF is not possible. The contents of this chapter justify the choice of the ECPS language, so the following chapters are only concerned with ECPS.

4.1 CPS Translation

Preliminary ideas for the translation from EPCF to ECPS appear in a technical report by Lafont, Reus and Streicher [LRS93], but for simpler languages. In order to simplify the translation, we replace all the EPCF effect operations by operations with arity $\mathbb{N} \times \alpha^{\mathbb{N}} \rightarrow \alpha$, as follows:

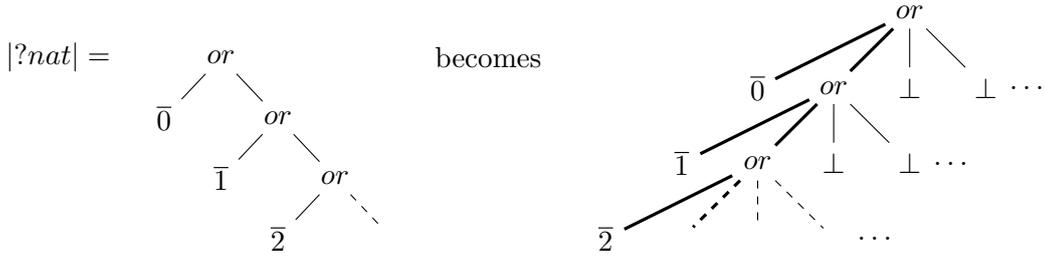
$$\begin{aligned}\sigma(M_0, \dots, M_{n-1}) &\text{ becomes } \sigma(Z; W) \\ \sigma(V; M_0, \dots, M_{n-1}) &\text{ becomes } \sigma(V; W) \\ \sigma(V) &\text{ becomes } \sigma(Z; V)\end{aligned}$$

where

$$\begin{aligned}W &= \lambda x:\mathbb{N}.\mathbf{case } x \mathbf{ in } \{Z \Rightarrow M_0, S(x_1) \Rightarrow \\ &\quad \mathbf{case } x_1 \mathbf{ in } \{Z \Rightarrow M_1, S(x_2) \Rightarrow \dots \\ &\quad \quad \mathbf{case } x_{n-1} \mathbf{ in } \{Z \Rightarrow M_{n-1}, S(x_n) \Rightarrow \mathit{loop}\} \dots\}\} \\ \mathit{loop} &= \mathbf{let } (\mathbf{fix } \lambda f:\mathbb{N} \rightarrow \tau.\mathbf{return } f) \Rightarrow x \mathbf{ in } x Z. \quad (\text{where } \tau \text{ is the type of } M_i)\end{aligned}$$

The function W chooses between the arguments M_0, \dots, M_{n-1} , just as σ used to do. This simplification means that EPCF computation trees carry more redundant information. For example:

Example 4.1.1 (Nondeterminism). The computation tree of $?nat$



where only the paths in bold can occur. The definitions of modalities in \mathcal{O} need to be adjusted to take this redundancy into account. For the effects considered in the dissertation, Scott-openness, decomposability and program equivalence are preserved.

The CPS translation for types is defined as follows:

$$\begin{aligned} \mathbb{N}^* &= \mathbf{nat} \\ \mathbb{1}^* &= \mathbf{unit} \\ (\rho \rightarrow \tau)^* &= \neg(\rho^*, \neg\tau^*). \end{aligned}$$

To translate contexts, assume that all variable names from EPCF appear in ECPS:

$$\begin{aligned} \emptyset^* &= \emptyset \\ (\Gamma, x : \tau)^* &= \Gamma^*, x : \tau^*. \end{aligned}$$

The CPS translation for values and computations is given in Figures 4.1 and 4.2. It relates values in context $\Gamma \vdash V : \tau$ to ECPS values $\Gamma^* \vdash V^* : \tau^*$. A function $F = \lambda x:\rho.M$ of type $\rho \rightarrow \tau$ is translated to $F^* = \lambda(x, k):(\rho^*, \neg\tau^*).M^* k$ of type $\neg(\rho^*, \neg\tau^*)$. The function F^* is waiting for an argument of type ρ^* and a continuation of type $\neg\tau^*$. The continuation is itself waiting for the result of computation M .

A computation $\Gamma \vdash M : \tau$ is translated to a function $\Gamma^* \vdash M^* : \neg\neg\tau^*$. This is because M^* is waiting for a continuation of type $\neg\tau^*$ to which it can pass its result.

Computation **return** $V : \tau$ becomes $\lambda k:\neg\tau^*.k V^*$. So value V^* is passed to the current continuation. Application $V W : \tau$ becomes $\lambda k:\neg\tau^*.V^* (W^*, k)$, which means that V^* is given arguments W^* and a continuation that is waiting for the result of the application. Sequencing **let** $M \Rightarrow x$ **in** $N : \tau$ is implemented as $\lambda k:\neg\tau^*.M^* \lambda x:\rho^*.(N^* k)$. First, M^* is evaluated and its result is passed to the continuation $\lambda x:\rho^*.(N^* k)$ which executes N^* and passes the result to k .

To translate effects, assume that the set of effect operations Σ is the same in EPCF and ECPS. The translation of $\sigma(V; W) : \tau$ is as expected, the function W^* is passed a continuation in addition to argument x .

The translation of **(fix F)** $: \tau \rightarrow \rho$ is complicated by the reduction behaviour of **fix F**. Recall that:

$$\mathbf{fix} F \rightsquigarrow \mathbf{return} \lambda x:\tau.\mathbf{let} F (\lambda y:\tau.\mathbf{let} \mathbf{fix} F \Rightarrow z \mathbf{in} z y) \Rightarrow w \mathbf{in} w x.$$

As a sanity check for the translation, it can be shown that

$$(\mathbf{fix} F)^* k \longrightarrow^* k (\lambda x:\tau.\mathbf{let} F (\lambda y:\tau.\mathbf{let} \mathbf{fix} F \Rightarrow z \mathbf{in} z y) \Rightarrow w \mathbf{in} w x)^*$$

$$\begin{aligned}
(\Gamma \vdash Z : \mathbb{N})^* &= \Gamma^* \vdash \mathbf{zero} : \mathbf{nat} \\
(\Gamma \vdash S(V) : \mathbb{N})^* &= \Gamma^* \vdash \mathbf{succ}(V^*) : \mathbf{nat} \\
(\Gamma \vdash x : \tau)^* &= \Gamma^* \vdash x : \tau^* \\
(\Gamma \vdash \star : \mathbb{1})^* &= \Gamma^* \vdash \star : \mathbf{unit} \\
(\Gamma \vdash \lambda x:\rho.M : \rho \rightarrow \tau)^* &= \Gamma^* \vdash \lambda(x, k):(\rho^*, \neg\tau^*).(M^* k) : \neg(\rho^*, \neg\tau^*) \\
(\Gamma \vdash V W : \tau)^* &= \Gamma^* \vdash \lambda k:\neg\tau^*.V^* (W^*, k) : \neg\neg\tau^* \\
(\Gamma \vdash \mathbf{return} V : \tau)^* &= \Gamma^* \vdash \lambda k:\neg\tau^*.(k V^*) : \neg\neg\tau^* \\
(\Gamma \vdash \mathbf{let} M \Rightarrow x \mathbf{in} N : \tau)^* &= \Gamma^* \vdash \lambda k:\neg\tau^*.M^* \lambda x:\rho^*.(N^* k) : \neg\neg\tau^* \\
(\Gamma \vdash \mathbf{case} V \mathbf{in} \{Z \Rightarrow M, S(x) \Rightarrow N\} : \tau)^* &= \\
\Gamma^* \vdash \lambda k:\neg\tau^*.\mathbf{case} V^* \mathbf{in} \{\mathbf{zero} \Rightarrow M^* k, \mathbf{succ}(x) \Rightarrow N^* k\} : \neg\neg\tau^* \\
(\Gamma \vdash \sigma(V; W) : \tau)^* &= \Gamma^* \vdash \lambda k:\neg\tau^*.\sigma(V^*, x.W^*(x, k)) : \neg\neg\tau^*
\end{aligned}$$

Figure 4.1: CPS translation of EPCF into ECPS – first part.

$$\begin{aligned}
(\Gamma \vdash \mathbf{fix} F : \tau \rightarrow \rho)^* &= \\
\Gamma^* \vdash \lambda k:\neg(\neg(\tau^*, \neg\rho^*)). & \\
(\mu x.\lambda c:\neg(\neg(\tau^*, \neg\rho^*)). & \\
c \lambda(x', k'):(\tau^*, \neg\rho^*). & \\
\lambda k'':\neg\rho^*.\lambda l:\neg(\neg(\tau^*, \neg\rho^*)).(F^* & (\lambda(y, l'):(\tau^*, \neg\rho^*). \\
\lambda p:\neg\rho^*.x (\lambda z:\neg(\tau^*, \neg\rho^*).(\lambda p':\neg\rho^*.z & (y, p')) \\
p) & \\
l', & \\
l) & \\
) & \\
\lambda w:\neg(\tau^*, \neg\rho^*).((\lambda l'':\neg\rho^*.w & (x', l'')) k'')) \\
k' & \\
) k & \\
: \neg\neg(\neg(\tau^*, \neg\rho^*)) &
\end{aligned}$$

Figure 4.2: CPS translation of EPCF into ECPS – continued.

which matches:

$$\begin{aligned} (\mathbf{return} \lambda x:\tau.\mathbf{let} F (\lambda y:\tau.\mathbf{let} \mathbf{fix} F \Rightarrow z \mathbf{in} z y) \Rightarrow w \mathbf{in} w x)^* k &\longrightarrow \\ k (\lambda x:\tau.\mathbf{let} F (\lambda y:\tau.\mathbf{let} \mathbf{fix} F \Rightarrow z \mathbf{in} z y) \Rightarrow w \mathbf{in} w x)^* &. \end{aligned}$$

4.2 Typing and CPS Translation for Stacks

In the CPS translation $(\mathbf{return} V)^* = \lambda k:\neg\tau^*.k V^*$, the continuation k plays a similar role to the stack $S \circ \mathbf{let} (-) \Rightarrow x \mathbf{in} M$, in the reduction:

$$(S \circ \mathbf{let} (-) \Rightarrow x \mathbf{in} M, \mathbf{return} V) \mapsto (S, M[V/x]).$$

This suggests that EPCF stacks can be translated to ECPS continuations. Since the CPS translation is typed, we allow stacks to have free variables and introduce typing judgements for stacks. The typing judgement $\Gamma \vdash S : \tau \Rightarrow \rho$ says that by substituting a *closed* computation M of type τ for the hole in S we obtain a computation $\Gamma \vdash S\{M\} : \rho$.

$$\frac{}{\Gamma \vdash id : \rho \Rightarrow \rho} \text{(SID)} \quad \frac{\Gamma, x : \tau \vdash M : \tau' \quad \Gamma \vdash S : \tau' \Rightarrow \rho}{\Gamma \vdash S \circ (\mathbf{let} (-) \Rightarrow x \mathbf{in} M) : \tau \Rightarrow \rho} \text{(SLET)}$$

The substitution of closed values for free variables can be extended to stacks:

$$id[V/y] = id$$

$$(S \circ \mathbf{let} (-) \Rightarrow x \mathbf{in} M) = S[V/y] \circ (\mathbf{let} (-) \Rightarrow x \mathbf{in} M[V/y]).$$

Now EPCF stacks $\Gamma \vdash S : \tau \Rightarrow \rho$ can be translated to ECPS function values $\Gamma^* \vdash S^* : \neg\tau^*$, which are in fact continuations:

$$(\Gamma \vdash id : \rho \Rightarrow \rho)^* = \Gamma^* \vdash \lambda x:\rho^*. \downarrow : \neg\rho^*$$

$$(\Gamma \vdash S \circ (\mathbf{let} (-) \Rightarrow x \mathbf{in} M) : \tau \Rightarrow \rho)^* = \Gamma^* \vdash (\lambda x:\tau^*.M^* S^*) : \neg\tau^*.$$

The empty stack id is equivalent to a continuation which given any value terminates. Here, we can see an important difference between EPCF and ECPS: the stack id returns the value given to it, while id^* discards the value so we cannot observe it. The translation of $S \circ (\mathbf{let} (-) \Rightarrow x \mathbf{in} M)$ is a continuation which, when given value x , executes M^* and passes the result to S^* .

4.3 Correctness of the CPS Translation

The goal of this section is to prove that EPCF computation trees preserve their shape when translated into ECPS. This is stated as:

Theorem 4.3.1 (The CPS translation is correct). *For any closed computation $M : \tau$ in EPCF and any stack $S : \tau \Rightarrow \rho$ the following holds:*

$$\llbracket M^* S^* \rrbracket = |S, M|[\downarrow / l_1, \downarrow / l_2, \dots].$$

That is, all the value leaves of $|S, M|$ are replaced by \downarrow , but the nodes and the \perp -leaves stay the same.

Example 4.3.2 (Probabilistic choice). As explained in Section 4.1, EPCF operation $p\text{-or}$ is replaced by its version with arity $\mathbb{N} \times \alpha^{\mathbb{N}} \rightarrow \alpha$. So computation $p\text{-or}(loop, \mathbf{return} \bar{3})$ becomes:

$p\text{-or}(Z; \lambda x:\mathbb{N}.\mathbf{case} x \mathbf{in} \{Z \Rightarrow loop, S(x') \Rightarrow \mathbf{case} x' \mathbf{in} \{Z \Rightarrow \mathbf{return} \bar{3}, S(x'') \Rightarrow loop\}\})$.

However, we will still use notation $p\text{-or}(loop, \mathbf{return} \bar{3})$ for convenience, even when the version with arity $\mathbb{N} \times \alpha^{\mathbb{N}} \rightarrow \alpha$ is meant. Its computation tree is:

$$|id, p\text{-or}(loop, \mathbf{return} \bar{3})| = \begin{array}{c} p\text{-or}_0 \\ \swarrow \quad \downarrow \quad \searrow \\ \perp \quad \bar{3} \quad \perp \quad \perp \dots \end{array}$$

According to Theorem 4.3.1, when translated to ECPS, this computation tree becomes:

$$\begin{aligned} \llbracket (p\text{-or}(loop, \mathbf{return} \bar{3}))^* id^* \rrbracket = & \\ \llbracket (\lambda k:\mathbf{nat}. p\text{-or}(\bar{0}, y.(\lambda(x, k'):(\mathbf{nat}, \neg\mathbf{nat}). & \\ & (\lambda k'':\mathbf{nat}. \mathbf{case} x \mathbf{in} \{\mathbf{zero} \Rightarrow loop^* k'', \\ & \mathbf{succ}(x') \Rightarrow (\lambda k''':\mathbf{nat}. \mathbf{case} x' \mathbf{in} \{\mathbf{zero} \Rightarrow (\lambda l:\mathbf{nat}. l \bar{3}) k''', \\ & \mathbf{succ}(x'') \Rightarrow loop^* k'''\}) \\ &) k''\}) \\ &) k' \\ &) (y, k) \\ &) \\ \rrbracket (\lambda x:\mathbf{nat}. \downarrow) \rrbracket = & \begin{array}{c} p\text{-or}_0 \\ \swarrow \quad \downarrow \quad \searrow \\ \perp \quad \downarrow \quad \perp \quad \perp \dots \end{array} \end{aligned}$$

In order to prove Theorem 4.3.1, we introduce a step-indexed biorthogonal logical relation, as discussed in Section 3.4. The purpose of the logical relation is to relate EPCF and ECPS computations that have similar reduction behaviour. Because EPCF computations reduce in a stack, it is useful to use biorthogonality, and have a relation between stacks and continuations as well. Step-indexing is needed to deal with the **fix** constructor.

The notion of observation $\mathfrak{D}((S, M), t)$ that the logical relation uses is a *step-indexed similarity* relation \mathcal{S} :

$$\mathcal{S}_{\tau, \rho}^n \subseteq (Stack(\tau, \rho) \times Comp(\tau)) \times (\vdash)$$

defined for each pair of EPCF types τ, ρ and each $n \in \mathbb{N}$. This is not related to the applicative similarity discussed before because it contains pairs of programs from two different languages. However, we will show that \mathcal{S} is related to the abstract notion of bisimulation from Section 3.3. First, we recall the definition of a step-indexed relation from [Pit10] and then define step-indexed similarity:

Definition 4.3.3. A *step-indexed relation* on the set X is an \mathbb{N} -indexed family of sets:

$$R = (R_n \mid n \in \mathbb{N}) \text{ satisfying } X \supseteq R_0 \supseteq R_1 \supseteq R_2 \supseteq \dots$$

Definition 4.3.4. The *step-indexed similarity relation* \mathcal{S} is the family of greatest relations such that $(S, M) \mathcal{S}_{\tau, \rho}^n t$ implies:

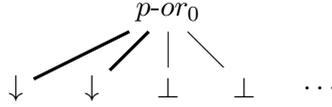
1. $\forall p < n. ((S, M) \rightsquigarrow^p (S', \sigma(V; W)) \implies t \longrightarrow^* \sigma(v, x.t')$ where $V = \bar{m}$ and $v = \bar{m}$ and $\forall l \in \mathbb{N}. (S', W \bar{l}) \mathcal{S}_{\tau', \rho}^{n-p} t'[\bar{l}/x]$.
2. $\forall p \leq n. ((S, M) \rightsquigarrow^p (id, \mathbf{return} V) \text{ for some } V : \rho \implies \tau \longrightarrow^* \downarrow)$.

The statement $(S, M) \mathcal{S}_{\tau, \rho}^n t$ means that the ECPS computation t simulates the reduction of the EPCF configuration (S, M) for n steps. The novel feature of this notion of observation is that it tracks effect operations, not only the termination behaviour of related programs, as compared to previous work on logical relations [Ahm06, BH09, Pit10, JT11].

Example 4.3.5 (Probabilistic choice). For example, the ECPS computation:

$$t = (\lambda k: \text{nat. } p\text{-or}(\bar{0}, x.\text{case } x \text{ in } \{\text{zero} \Rightarrow k \bar{2}, \text{succ}(x') \Rightarrow \text{case } x' \text{ in } \{\text{zero} \Rightarrow k \bar{3}, \text{succ}(x'') \Rightarrow \text{loop}\}\})) (\lambda x: \text{nat. } \downarrow)$$

with computation tree



and EPCF computation $M = p\text{-or}(\text{loop}, \mathbf{return} \bar{3})$ satisfy $(id, M) \mathcal{S}_{\mathbb{N}, \mathbb{N}}^n t$ for any $n > 0$.

Intuitively, as the index n increases, step-indexed similarity checks for more and more matching reduction steps. Therefore, \mathcal{S} respects the definition of a step-indexed relation. Below is a useful fact about step-indexed similarity which will be used later. The proof appears in Appendix A.

Lemma 4.3.6. *If $(S, M) \rightsquigarrow^k (S', M')$ and $t \longrightarrow^* t'$ where $k < n$ then:*

$$(S, M) \mathcal{S}_{\tau, \rho}^n t \iff (S', M') \mathcal{S}_{\tau, \rho}^{n-k} t'.$$

We can define a well-typed relation \mathcal{S} which encodes the fact that an ECPS computation t simulates configuration (S, M) for *any* number of reduction steps. Step-indexed similarity implies this new notion of similarity.

Definition 4.3.7. *Define similarity as the family \mathcal{S} of greatest relations indexed by EPCF types τ and ρ such that $(S, M) \mathcal{S}_{\tau, \rho} t$ implies:*

1. $(S, M) \rightsquigarrow^* (S', \sigma(V; W)) \implies t \longrightarrow^* \sigma(v, x.t')$ where $V = \bar{m}$ and $v = \bar{m}$ and $\forall l \in \mathbb{N}. (S', W \bar{l}) \mathcal{S}_{\tau', \rho} t'[\bar{l}/x]$.
2. $(S, M) \rightsquigarrow^* (id, \mathbf{return} V) \text{ for some } V : \rho \implies \tau \longrightarrow^* \downarrow$.

Lemma 4.3.8. *For any configuration (S, M) and closed computation t :*

$$(\forall n \in \mathbb{N}. (S, M) \mathcal{S}_{\tau, \rho}^n t) \implies (S, M) \mathcal{S}_{\tau, \rho} t.$$

Proof. The strategy is to show $\bigcap_{n \in \mathbb{N}} \mathcal{S}^n$ is a simulation relating (S, M) and t , and thus it is included in \mathcal{S} . The rest of the proof appears in Appendix A. \square

Now we are ready to define the logical relation. It is in fact a collection of relations $(\mathcal{R}_\tau^{v,n}, \mathcal{R}_\tau^{c,n}, \mathcal{R}_{\tau,\rho}^{s,n})$ for all EPCF types τ, ρ and $n \in \mathbb{N}$.

$$\mathcal{R}_\tau^{v,n} \subseteq \text{Val}(\tau) \times (\vdash \tau^*)$$

$$\mathcal{R}_\tau^{c,n} \subseteq \text{Comp}(\tau) \times (\vdash \neg\neg\tau^*)$$

$$\mathcal{R}_{\tau,\rho}^{s,n} \subseteq \text{Stack}(\tau, \rho) \times (\vdash \neg\tau^*).$$

It is defined by well-founded induction on the natural numbers and on EPCF types, following a tutorial paper by Pitts [Pit10] which deals with the untyped λ -calculus.

Definition 4.3.9 (Logical Relation).

$$\mathcal{R}_{\mathbb{N}}^{v,n} = \{(V, v) \mid V = \bar{m} \text{ and } v = \bar{m} \text{ for some } m \in \mathbb{N}\}$$

$$\mathcal{R}_1^{v,n} = \{(V, v) \mid V = \star \text{ and } v = \star\}$$

$$\mathcal{R}_{\tau \rightarrow \rho}^{v,n} = \{(V, v) \mid \forall p < n. \forall (V_1, v_1) \in \mathcal{R}_\tau^{v,p}. (V V_1, \lambda k: \neg\rho^*.v (v_1, k)) \in \mathcal{R}_\rho^{c,p}\}$$

$$\mathcal{R}_{\tau,\rho}^{s,n} = \{(S, k) \mid p \leq n. \forall (V, v) \in \mathcal{R}_\tau^{v,p}. (S, \mathbf{return} V) \mathcal{S}_{\tau,\rho}^p (k v)\}$$

$$\mathcal{R}_\tau^{c,n} = \{(M, v) \mid \forall p \leq n. \forall (S, k) \in \mathcal{R}_{\tau,\rho}^{s,p}. (S, M) \mathcal{S}_{\tau,\rho}^p (v k)\}.$$

These relations satisfy the definition of a step-indexed relation at each type, $\mathcal{R}_\tau^{n+1} \subseteq \mathcal{R}_\tau^n$, because, as the index of $X \mathcal{R}_\tau^n x$ increases, the pair (X, x) needs to satisfy more conditions.

The logical relation is defined for closed values, stacks and computations and makes use of the step-indexed similarity relation \mathcal{S}^n . As expected, related functions map related values to related computations. A stack and a continuation are related if, given any related values, they have a similar reduction behaviour. Finally, computations are related if they have similar reduction behaviour in all related stack-continuation pairs.

We can extend the logical relation to open terms and stacks by substituting in related values. The definition below generalises over all indices:

$$\overrightarrow{x_i : \tau_i} \vdash V \mathcal{R}_\rho^v v \iff \overrightarrow{x_i : \tau_i} \vdash V : \rho \text{ and } (\overrightarrow{x_i : \tau_i})^* \vdash v : \rho^* \text{ and}$$

$$(\forall n \in \mathbb{N}. \forall \overrightarrow{(V_i, v_i) : \tau_i}. ((\overrightarrow{(V_i, v_i)} \in \mathcal{R}_{\tau_i}^{v,n}) \implies (V[\overrightarrow{V_i/x_i}], v[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_\rho^{v,n}))$$

$$\overrightarrow{x_i : \tau_i} \vdash M \mathcal{R}_\rho^c v \iff \overrightarrow{x_i : \tau_i} \vdash M : \rho \text{ and } (\overrightarrow{x_i : \tau_i})^* \vdash v : \neg\neg\rho^* \text{ and}$$

$$(\forall n \in \mathbb{N}. \forall \overrightarrow{(V_i, v_i) : \tau_i}. ((\overrightarrow{(V_i, v_i)} \in \mathcal{R}_{\tau_i}^{v,n}) \implies (M[\overrightarrow{V_i/x_i}], v[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_\rho^{c,n}))$$

$$\overrightarrow{x_i : \tau_i} \vdash S \mathcal{R}_{\rho,\rho'}^s k \iff \overrightarrow{x_i : \tau_i} \vdash S : \rho \Rightarrow \rho' \text{ and } (\overrightarrow{x_i : \tau_i})^* \vdash k : \neg\rho^* \text{ and}$$

$$(\forall n \in \mathbb{N}. \forall \overrightarrow{(V_i, v_i) : \tau_i}. ((\overrightarrow{(V_i, v_i)} \in \mathcal{R}_{\tau_i}^{v,n}) \implies (S[\overrightarrow{V_i/x_i}], k[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\rho,\rho'}^{s,n})).$$

An important property of the logical relation is reflexivity. Since \mathcal{R} relates EPCF terms to ECPS terms the property is formulated as below:

Lemma 4.3.10 (Fundamental property of the logical relation). *For any value $\overline{x_i : \tau_i} \vdash V : \rho$, any computation $\overline{x_i : \tau_i} \vdash M : \rho$ and any stack $\overline{x_i : \tau_i} \vdash S : \rho \Rightarrow \rho'$ in EPCF:*

1. $\overline{x_i : \tau_i} \vdash V \mathcal{R}_\rho^v V^*$.
2. $\overline{x_i : \tau_i} \vdash M \mathcal{R}_\rho^c M^*$.
3. $\overline{x_i : \tau_i} \vdash S \mathcal{R}_{\rho, \rho'}^s S^*$.

Proof. The proof is by induction on the typing derivations of V , M and S . The most interesting case is $M = \mathbf{fix} W$. It uses the indices of the logical relation to reason about the number of times the fixed point of W is unfolded. All the cases appear in Appendix A. \square

The similarity relation \mathcal{S}^n expresses the fact that computation t simulates configuration (S, M) for n steps. Analogously, we can define a notion of similarity \mathcal{S}'^n which says that (S, M) simulates t for n steps. This will be a family of relations parametrised by $n \in \mathbb{N}$, and EPCF types τ and ρ :

$$\mathcal{S}_{\tau, \rho}^n \subseteq (\text{Stack}(\tau, \rho) \times \text{Comp}(\tau)) \times (\vdash).$$

Definition 4.3.11. *The step-indexed similarity relation \mathcal{S}' is the family of greatest relations such that $(S, M) \mathcal{S}'_{\tau, \rho} t$ implies:*

1. $\forall p < n. (t \xrightarrow{p} \sigma(v, x.t')) \implies (S, M) \rightsquigarrow^* (S', \sigma(V; W))$ where $V = \overline{m}$ and $v = \overline{m}$ and $\forall l \in \mathbb{N}. (S', W \bar{l}) \mathcal{S}'_{\tau, \rho}^{n-p} t'[\bar{l}/x]$.
2. $\forall p \leq n. (t \xrightarrow{p} \downarrow) \implies (S, M) \rightsquigarrow^* (\text{id}, \mathbf{return} V)$ for some $V : \rho$.

Using \mathcal{S}'^n , we can define a family of relations \mathcal{S}' indexed by types τ and ρ , such that $(S, M) \mathcal{S}' t$ says that (S, M) simulates t for any number of steps. The definition is similar to that of \mathcal{S} (Definition 4.3.7). The analogous of Lemma 4.3.8 can be proved for \mathcal{S}'^n and \mathcal{S}' .

Finally, we can define a logical relation \mathcal{R}' exactly as \mathcal{R} was defined but using \mathcal{S}' instead of \mathcal{S} . The logical relation \mathcal{R}' also has the fundamental property.

The next step is proving that the trees of a configuration and a computation which simulate each other are closely related. The proof is done by coinduction. We first discuss some properties of the sets of closed ECPS computations and EPCF stack-computation pairs.

Consider the functor:

$$T : \mathbf{Set} \implies \mathbf{Set} \text{ where } T(X) = \{\downarrow, \perp\} + (\Sigma \times \mathbb{N} \times X^{\mathbb{N}}).$$

Consider the following coalgebra for this functor: $(\text{Trees}_\Sigma, c : \text{Trees}_\Sigma \longrightarrow T(\text{Trees}_\Sigma))$ where

$$c(tr) = \begin{cases} \downarrow & \text{if } tr = \downarrow \\ \perp & \text{if } tr = \perp \\ (\sigma, n, \overrightarrow{tr'}) & \text{if } tr = \sigma_n(\overrightarrow{tr'}). \end{cases}$$

It is a standard result that the category of coalgebras and coalgebra morphisms for the functor T has a terminal object. This is shown for example by Jacobs [Jac16, Theorem 2.3.9]. Moreover, the terminal object is (Trees_Σ, c) . The proof of this is very similar to the proof of Proposition 2.3.5 from [Jac16].

Lemma 4.3.12. *The function $\llbracket - \rrbracket : (\vdash) \longrightarrow \mathit{Trees}_\Sigma$ is a coalgebra morphism in the category of coalgebras for the functor T .*

Proof. First, note that (\vdash) is indeed a coalgebra by considering the following function $a : (\vdash) \longrightarrow T(\vdash)$ on closed ECPS computations:

$$a(t) = \begin{cases} \downarrow & \text{if } t \longrightarrow^* \downarrow \\ \perp & \text{if } t \longrightarrow^\infty \\ (\sigma, k, \overrightarrow{t'[\bar{n}/x]}) & \text{if } t \longrightarrow^* \sigma(\bar{k}, x.t'). \end{cases}$$

By definition of \longrightarrow exactly one of the cases above will occur, so a is a well-defined function.

To prove that $\llbracket - \rrbracket$ is a coalgebra morphism, we use standard domain theoretic techniques. All the necessary results can be found for example in [Fio17]. The proof involves making the order \leq on Trees_Σ more precise, and defining an order \leq_T on $T(\mathit{Trees}_\Sigma)$ which makes it an ω -CPO. The full development can be found in Appendix A. \square

Define the following function indexed by EPCF types τ and ρ :

$$|- , -|_{(-)}^* : \mathit{Stack}(\tau, \rho) \times \mathit{Comp}(\tau) \times \mathbb{N} \longrightarrow \mathit{Trees}_\Sigma.$$

similarly to how $|- , -|_{(-)}$ was defined. Assuming all EPCF effect operations have been replaced with operations of arity $\mathbb{N} \times \alpha^{\mathbb{N}} \rightarrow \alpha$, the definition is:

$$|S, M|_0^* = \perp$$

$$|S, M|_{n+1}^* = \begin{cases} \downarrow & \text{if } S = \mathit{id} \text{ and } M = \mathbf{return} V \\ |S', M'|_n^* & \text{if } (S, M) \rightsquigarrow (S', M') \\ \sigma_k(|S, V \bar{0}|_n^*, |S, V \bar{1}|_n^*, \dots) & \text{if } \sigma : \mathbb{N} \times \alpha^{\mathbb{N}} \rightarrow \alpha \text{ and } M = \sigma(\bar{k}; V) \\ \perp & \text{otherwise.} \end{cases}$$

The tree $|S, M|_n^*$ is different from $|S, M|_n$ because all value leaves are replaced by \downarrow . We can see that $|S, M|_n^* \leq |S, M|_{n+1}^*$ in Trees_Σ so we can define:

$$|- , -|^* : \mathit{Stack}(\tau, \rho) \times \mathit{Comp}(\tau) \longrightarrow \mathit{Trees}_\Sigma$$

$$|S, M|^* = \bigsqcup_{n \in \mathbb{N}} |S, M|_n^*.$$

Define the set of all well-formed stack-computation pairs as:

$$\mathit{Stack} \times \mathit{Comp} = \{(S, M) \mid S \in \mathit{Stack}(\tau, \rho), M \in \mathit{Comp}(\tau) \text{ for some EPCF types } \tau, \rho\}$$

and extend the function $|- , -|^*$ to this set:

$$\beta : \mathit{Stack} \times \mathit{Comp} \longrightarrow \mathit{Trees}_\Sigma$$

$$\beta(S, M) = |S, M|^*.$$

Lemma 4.3.13. *The function $\beta : \mathit{Stack} \times \mathit{Comp} \longrightarrow \mathit{Trees}_\Sigma$ is a coalgebra morphism in the category of coalgebras for the functor T .*

Proof. To see that $Stack \times Comp$ is indeed a coalgebra, consider the following function ($Stack \times Comp, b : Stack \times Comp \longrightarrow T(Stack \times Comp)$):

$$b(S, M) = \begin{cases} \downarrow & \text{if } (S, M) \mapsto^* (id, \mathbf{return} V) \\ \perp & \text{if } (S, M) \mapsto^\infty \\ (\sigma, k, ((S', V\bar{0}), (S', V\bar{1}), \dots)) & \text{if } (S, M) \mapsto^* (S', \sigma(\bar{k}; V)). \end{cases}$$

By definition of \mapsto the cases above are exhaustive, and by determinacy only one of them can occur, so b is a well-defined function.

The proof that β is a coalgebra morphism is very similar to the proof for $\llbracket - \rrbracket$ so we omit it. It uses the ω -CPO structure of $Trees_\Sigma$ and $T(Trees_\Sigma)$ to show that the coalgebra morphism diagram commutes. \square

Proposition 4.3.14. *For any well-typed EPCF configuration (S, M) , where $S : \tau \Rightarrow \rho$, and any ECPS computation t :*

$$((S, M), t) \in \mathcal{S}_{\tau, \rho} \cap \mathcal{S}'_{\tau, \rho} \implies \llbracket t \rrbracket = |S, M|[\downarrow / l_1, \downarrow / l_2, \dots].$$

Proof. Apply the coinduction proof principle from Proposition 3.3.2, use Lemmas 4.3.12 and 4.3.13 and the fact that $(Trees_\Sigma, c : Trees_\Sigma \longrightarrow T(Trees_\Sigma))$ is a final coalgebra. The full proof appears in Appendix A. \square

Finally, we can prove the correctness of the CPS translation:

Theorem 4.3.1. *For any closed computation $M : \tau$ in EPCF and any stack $S : \tau \Rightarrow \rho$ the following holds:*

$$\llbracket M^* S^* \rrbracket = |S, M|[\downarrow / l_1, \downarrow / l_2, \dots].$$

That is, all the value leaves of $|S, M|$ are replaced by \downarrow , but the nodes and the \perp -leaves stay the same.

Proof. From the fundamental property of the logical relation (Lemma 4.3.10 and its analogue for \mathcal{R}') we know that for any closed configuration (S, M) , where $S : \tau \Rightarrow \rho$:

$$\forall n \in \mathbb{N}. ((S, M), (M^* S^*)) \in \mathcal{S}_{\tau, \rho}^n \cap \mathcal{S}'_{\tau, \rho}^n.$$

Then using Lemma 4.3.8 and its analogue for \mathcal{S}' we can deduce that:

$$((S, M), (M^* S^*)) \in \mathcal{S}_{\tau, \rho} \cap \mathcal{S}'_{\tau, \rho}.$$

Applying Proposition 4.3.14 we know that:

$$\llbracket M^* S^* \rrbracket = |S, M|[\downarrow / l_1, \downarrow / l_2, \dots].$$

\square

4.4 ECPS Is More Expressive than EPCF

From the point of view of the propositions-as-types correspondence, or the Curry-Howard isomorphism, the λ -calculus corresponds to intuitionistic logic [SU06, Chapter 6]. Griffin showed that a language with *control operators* can extend this correspondence to classical logic [Gri90].

One such control operator is `call-cc` (call-with-current-continuation) from the programming language Scheme. Griffin showed that `call-cc` can be assigned the type of Peirce's law, a proposition which is only provable classically:

$$((A \rightarrow B) \rightarrow A) \rightarrow A.$$

It is known that PCF and EPCF do not contain any control operators but we will show that ECPS does. In ECPS we can write a term which has the type of Peirce's law and the reduction behaviour of `call-cc`. In this sense, ECPS is more expressive than EPCF. Below is an informal explanation of `call-cc` and how it arises in ECPS.

One way of adding `call-cc` to PCF [RT99] is to add a new type $Cont(A)$ which stands for a continuation waiting for a term of type A . Such a continuation can be written as $\gamma x:A.E[x]$ where E is a PCF evaluation context and $E[x]$ is a PCF term. Then add the following term constructors to PCF:

$$\frac{}{\Gamma \vdash \text{callcc} : (Cont(A) \rightarrow A) \rightarrow A} \quad \frac{}{\Gamma \vdash \text{throw} : Cont(A) \rightarrow A \rightarrow B}$$

The reduction rules for these new constructors are:

$$E[\text{callcc } (\lambda x:Cont(A).M)] \longrightarrow E[M[x := (\gamma y:A.E[y])]] \quad (4.4.1)$$

$$E'[\text{throw } (\gamma y:A.E[y]) V] \longrightarrow E'[V]. \quad (4.4.2)$$

The first rule says that when `callcc` is encountered the current evaluation context E is bound to the continuation $x = \gamma y:A.E[y]$. Then evaluation of M proceeds normally. If M never throws, the control flow is not changed, the return value of `callcc` was M . If at some point M invokes `throw` $(\gamma y:A.E[y]) V$, the current evaluation context E' is abandoned. Instead, the context E is restored with value V . This looks as if the return value of `callcc` was V .

The type $Cont(A)$ can be interpreted as $A \rightarrow B$ for some type B . Thus `callcc` has type $((A \rightarrow B) \rightarrow A) \rightarrow A$. And `throw` has type $(A \rightarrow B) \rightarrow A \rightarrow B$. When the continuation $(\gamma y:A.E[y]) : Cont(A)$ is thrown, the current context E' that is waiting for type B is discarded. Therefore, B can be anything.

In ECPS, the PCF type $A \rightarrow B$ is interpreted as $\neg(A, \neg B)$. In other words, the implication $A \rightarrow B$ is expressed as $\neg(A \wedge \neg B)$. Peirce's law then becomes:

$$\neg(\neg(\neg(\mathbf{A} \wedge \neg \mathbf{B}) \wedge \neg \mathbf{A}) \wedge \neg \mathbf{A}).$$

In ECPS, a term of this type is:

$$\text{callcc}^* = \lambda(f, k):(\neg(\neg(\mathbf{A}, \neg \mathbf{B}), \neg \mathbf{A}), \neg \mathbf{A}). f (\lambda(y, k'):(A, \neg B).k y, k).$$

A term with the type of **throw**:

$$(A \rightarrow B) \rightarrow A \rightarrow B = \neg(\neg(A \wedge \neg B) \wedge \neg\neg(A \wedge \neg B))$$

is

$$\mathbf{throw}^* = \lambda(k, l):(\neg(A, \neg B), \neg\neg(A, \neg B)).(l\ k).$$

To illustrate the reduction behaviour of **callcc**^{*} and **throw**^{*} consider the following examples where $A = \mathbb{N}$:

$$f_1 = \lambda(x, k):(\neg(\mathbb{N}, \neg B), \neg\mathbb{N}). \mathbf{throw}^* (x, \lambda g:\neg(\mathbb{N}, \neg B).g (\bar{3}, \lambda w:B.loop)).$$

Here x is the analogous of x from $\lambda x:Cont(A).M$, equation 4.4.1. The function f_1 throws continuation x with value $V = \bar{3}$, in the context $\lambda w:B.loop$ analogous to E' , from equation 4.4.2.

Now consider the following computation, where $\lambda z:\mathbb{N}.\downarrow$ stands for the context E from equation 4.4.1:

$$\begin{aligned} & \mathbf{callcc}^* (f_1, \lambda z:\mathbb{N}.\downarrow) \\ & \longrightarrow f_1 (\lambda(y, k'):(\mathbb{N}, \neg B).(\lambda z:\mathbb{N}.\downarrow) y, \lambda z:\mathbb{N}.\downarrow) \\ & \longrightarrow \mathbf{throw}^* (\lambda(y, k'):(\mathbb{N}, \neg B).(\lambda z:\mathbb{N}.\downarrow) y, \lambda g:\neg(\mathbb{N}, \neg B).g (\bar{3}, \lambda w:B.loop)) \\ & \longrightarrow^2 (\lambda(y, k'):(\mathbb{N}, \neg B).(\lambda z:\mathbb{N}.\downarrow) y) (\bar{3}, \lambda w:B.loop) \\ & \longrightarrow (\lambda z:\mathbb{N}.\downarrow) \bar{3}. \end{aligned}$$

When **callcc**^{*} is called, the current continuation $\lambda z:\mathbb{N}.\downarrow$ is saved inside

$$\lambda(y, k'):(\mathbb{N}, \neg B).(\lambda z:\mathbb{N}.\downarrow) y.$$

Then when **throw**^{*} occurs, the now current continuation $\lambda w:B.loop$ is abandoned and the continuation $\lambda z:\mathbb{N}.\downarrow$ is restored with value $\bar{3}$.

As another example, consider a function f_2 which does not throw. It just invokes the continuation k that was passed to it:

$$f_2 = \lambda(x, k):(\neg(\mathbb{N}, \neg B), \neg\mathbb{N}). k \bar{3}$$

$$\begin{aligned} & \mathbf{callcc}^* (f_2, \lambda z:\mathbb{N}.\downarrow) \\ & \longrightarrow f_2 (\lambda(y, k'):(\mathbb{N}, \neg B).(\lambda z:\mathbb{N}.\downarrow) y, \lambda z:\mathbb{N}.\downarrow) \\ & \longrightarrow (\lambda z:\mathbb{N}.\downarrow) \bar{3}. \end{aligned}$$

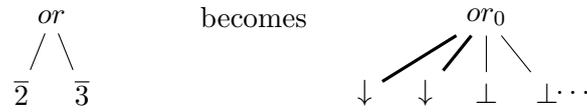
Here, the control flow is not changed by the use of **callcc**^{*}.

This section is not a full proof that ECPS is more expressive than EPCF. In particular, we have only shown examples that **callcc**^{*} has the desired behaviour. However, this is strong evidence to suggest that an embedding of ECPS in EPCF is not possible. This is why we only studied a translation of EPCF into ECPS in this chapter.

4.5 Chapter Summary

This chapter started the exposition of the novel technical content of the dissertation. In Section 4.1, we presented a continuation-passing translation from EPCF to ECPS. For example, an EPCF function V of type $\tau \rightarrow \mathbb{N}$ is translated to an ECPS function $V^* : \neg(\tau^*, \neg\mathbf{nat})$, where $\neg\mathbf{nat}$ is the type of a continuation waiting for the result of V . An EPCF computation $M : \tau$ is translated to a function $M^* = (\lambda k : \neg\tau^* \dots) : \neg\neg\tau^*$, where k is the continuation waiting for the result of M .

The correctness theorem of this translation (Theorem 4.3.1) implies that: the tree of M^* ($\lambda x : \neg\tau^* . \downarrow$) is obtained from the tree of EPCF computation M by replacing all value leaves with \downarrow . For example:



This means that the reduction behaviour of terms is preserved by the translation.

We could not prove Theorem 4.3.1 directly by induction on EPCF terms. To obtain a stronger induction hypothesis, we defined a *logical relation* (Definition 4.3.9) between EPCF terms and ECPS terms. Most importantly, two functions are related if and only if they send related arguments to related computations. Two computations are related when they simulate each other's behaviour, including effect operations. This is a custom notion of similarity introduced in Definition 4.3.4.

We then proved the fundamental property of the logical relation (Lemma 4.3.10): for any EPCF term T , the pair (T, T^*) is in the relation. Using this, and the coinduction proof principle from the previous chapter (Proposition 3.3.2), we proved the correctness of the translation.

Finally, we argued informally that ECPS is strictly more expressive than EPCF because it contains the control operator `call-cc`. Overall, this chapter showed that ECPS is a reasonable choice of language for studying program equivalence. Moreover, because ECPS contains more program contexts than EPCF, it becomes more likely that contextual equivalence equals applicative bisimilarity, which we will prove in Chapter 7, even though this is false for EPCF.

Chapter 5

Applicative Bisimilarity for ECPS

This chapter starts by defining observations for ECPS for all the running examples of effects. Using them, applicative \mathfrak{P} -bisimilarity is defined. Two sufficient conditions for bisimilarity to be compatible are identified: Scott-openness and a novel notion of decomposability. All the example observations are proved decomposable. The final section uses Howe's method to prove bisimilarity is indeed compatible. In the following chapters, applicative \mathfrak{P} -bisimilarity is compared with other notions of program equivalence.

5.1 Observations for ECPS

To define applicative bisimulation for ECPS we first fix a set of observations \mathfrak{P} , which contains subsets of $Trees_\Sigma$. The set \mathfrak{P} depends on the effects that are present in the language. It can be used to define various forms of program equivalence which check whether computation trees are in $P \in \mathfrak{P}$.

Observations $P \in \mathfrak{P}$ play a similar role to modalities from \mathcal{O} and to the observations defined by Johann, Simpson and Voigtländer [JSV10]. For the example effects considered so far, observations are defined as follows:

Example 5.1.1 (Pure functional computation). Define $\mathfrak{P} = \{\Downarrow\}$ where $\Downarrow = \{\downarrow\}$. There are no effect operations so the \Downarrow observation only checks for termination.

Example 5.1.2 (Nondeterminism). Define $\mathfrak{P} = \{Trees_\Sigma, \diamond, \square\}$ where:

$$\diamond = \{tr \in Trees_\Sigma \mid \text{at least one of the paths in } tr \text{ that can occur has a } \downarrow \text{ leaf}\}$$

$$\square = \{tr \in Trees_\Sigma \mid \text{the paths in } tr \text{ that can occur are all finite and finish with a } \downarrow\}.$$

The intuition is that, if $\llbracket t \rrbracket \in \diamond$, then computation t *may* terminate. Whereas, if $\llbracket t \rrbracket \in \square$, t *must* terminate. Notice that there is no condition on the value with which t terminates because ECPS computations do not have a return value.

As discussed in Section 3.2, every node in a computation tree has infinitely many children so some paths in the tree can never be executed. For example, the *or* operation always chooses between its first two children. The definitions of \diamond and \square take this into account.

The set of all trees $Trees_\Sigma$ is chosen to be an observation for technical reasons that will become clear in the next section. However, the fact that $Trees_\Sigma$ is an observation will not affect any notion of program equivalence because for all computations t , $\llbracket t \rrbracket \in Trees_\Sigma$.

Example 5.1.3 (Probabilistic choice). Define the set of observations as:

$$\mathfrak{P} = \{\mathbf{P}_{>q} \mid q \in \mathbb{Q}, 0 \leq q < 1\} \cup \{Trees_{\Sigma}\}.$$

Define $\mathbb{P} : Trees_{\Sigma} \rightarrow [0, 1]$ to be the least function, by the pointwise order, such that:

$$\mathbb{P}(tr) = \begin{cases} 1 & \text{if } tr = \downarrow \\ \frac{1}{2}\mathbb{P}(tr_0) + \frac{1}{2}\mathbb{P}(tr_1) & \text{if } tr = p\text{-or}(tr_0, tr_1). \end{cases}$$

Given functions $f_1, f_2 : Trees_{\Sigma} \rightarrow [0, 1]$ the pointwise order is defined as:

$$f_1 \leq f_2 \iff \forall tr \in Trees_{\Sigma}. f_1(tr) \leq f_2(tr).$$

Observations are defined as below:

$$\mathbf{P}_{>q} = \{tr \in Trees_{\Sigma} \mid \mathbb{P}(tr) > q\}.$$

This means that $tr \in \mathbf{P}_{>q}$ if the probability that tr terminates is greater than q . A p -or node chooses between its first two children with probability 0.5, so the probability that tree tr terminates is calculated over these choices. Notice that $\mathbb{P}(\perp) = 0$.

Example 5.1.4 (Global store). Define the set of states as the set of functions from storage locations to natural numbers: $State = \mathbb{L} \rightarrow \mathbb{N}$. The set \mathfrak{P} is defined as:

$$\mathfrak{P} = \{(s \mapsto r) \mid s, r \in State\} \cup \{Trees_{\Sigma}\}.$$

Define the execution of a tree starting in a state as the *least* partial function:

$$exec : Trees_{\Sigma} \times State \rightarrow \{\downarrow\} \times State$$

which satisfies

$$exec(tr, s) = \begin{cases} (\downarrow, s) & \text{if } tr = \downarrow \\ exec(tr_{s(l)}, s) & \text{if } tr = lookup_{l,n}(tr_0, tr_1, \dots) \\ & \text{and } exec(tr_{s(l)}, s) \text{ is defined} \\ exec(tr_0, s[l := n]) & \text{if } tr = update_{l,n}(tr_0, tr_1, \dots) \\ & \text{and } exec(tr_0, s[l := n]) \text{ is defined.} \end{cases}$$

Now define observations as:

$$(s \mapsto r) = \{tr \in Trees_{\Sigma} \mid exec(tr, s) \text{ is defined and } exec(tr, s) = (\downarrow, r)\}.$$

Notice that $exec(tr, s)$ is defined only when the execution of tr terminates. So $tr \in (s \mapsto r)$ only if the execution of tr started in state s terminates in state r .

Example 5.1.5 (Input/output). An I/O-trace is a finite word w over the alphabet

$$\{?n \mid n \in \mathbb{N}\} \cup \{!n \mid n \in \mathbb{N}\}.$$

Thus, a trace is a sequence of input and output operations, where $?n$ means that the number n was given as input to a *read* operation, and $!n$ means that the number n was output by a *write* operation. We can use them to define the set of observation:

$$\mathfrak{P} = \{\langle w \rangle \dots \mid w \text{ an I/O-trace}\}$$

where

$$\langle w \rangle \dots = \{tr \in \text{Trees}_\Sigma \mid \text{the execution of } tr \text{ produces I/O-trace } w\}.$$

To specify rigorously when “the execution of a tree produces an I/O trace”, we can define a relation between trees and I/O traces, by induction on traces. Denote this relation by \models .

$$\begin{aligned} tr \models \langle \epsilon \rangle \dots &\iff \text{true} \\ tr \models \langle (?n)w \rangle \dots &\iff tr = \text{read}_k(tr_0, tr_1, \dots) \text{ and } tr_n \models \langle w \rangle \dots \\ tr \models \langle (!n)w \rangle \dots &\iff tr = \text{write}_n(tr_0, tr_1, \dots) \text{ and } tr_0 \models \langle w \rangle \dots \end{aligned}$$

5.2 Applicative \mathfrak{P} -Bisimilarity

Definition 5.2.1 (Applicative \mathfrak{P} -simulation). *A collection of relations $\mathcal{R}_A^v \subseteq (\vdash A) \times (\vdash A)$ for each type A and $\mathcal{R}^c \subseteq (\vdash) \times (\vdash)$ is an applicative \mathfrak{P} -simulation if:*

1. $v \mathcal{R}_{\text{unit}}^v w \implies v = w = \star$.
2. $v \mathcal{R}_{\text{nat}}^v w \implies v = w$.
3. $s \mathcal{R}^c t \implies \forall P \in \mathfrak{P}. (\llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P)$.
4. $v \mathcal{R}_{\neg(A_1, \dots, A_n)}^v u \implies \forall \vdash w_1 : A_1, \dots, \vdash w_n : A_n. v(w_1, \dots, w_n) \mathcal{R}^c u(w_1, \dots, w_n)$.

Applicative \mathfrak{P} -similarity \lesssim is the union of all applicative \mathfrak{P} -simulations. Therefore, it is the greatest applicative \mathfrak{P} -simulation.

According to the definition above, unit values and natural number values are similar if and only if they are equal. The third clause says that t simulates s if the computation tree of t has all the properties of the computation tree of s . The properties are specified using the set of observations \mathfrak{P} . Notice that simulation for computations is not defined using simulations for values, since computations do not have a return value. The last clause compares the behaviour of functions for all possible arguments.

Definition 5.2.2 (Applicative \mathfrak{P} -bisimulation). *An applicative \mathfrak{P} -bisimulation is a symmetric \mathfrak{P} -simulation. Applicative \mathfrak{P} -bisimilarity \sim is the union of all applicative \mathfrak{P} -bisimulations. Therefore it is the greatest applicative \mathfrak{P} -bisimulation.*

Below are two properties of bisimilarity and similarity which will be used later. They are followed by two examples of how bisimilarity can be established.

Proposition 5.2.3. *Applicative \mathfrak{P} -bisimilarity coincides with the intersection between applicative \mathfrak{P} -similarity and its converse:*

$$(\sim) = (\lesssim) \cap (\lesssim)^{op}.$$

Proof. The proof is done using the definitions of similarity and bisimilarity. The \supseteq inclusion is shown by proving $(\lesssim) \cap (\lesssim)^{op}$ is a symmetric simulation. The full proof appears in Appendix B. \square

Lemma 5.2.4. *Similarity is preserved by the reduction relation, that is:*

$$\forall s, t. \vdash s \lesssim^c t \text{ and } s \longrightarrow^* s' \text{ and } t \longrightarrow^* t' \implies \vdash s' \lesssim^c t'.$$

Proof. By the definition of $\llbracket - \rrbracket$ we know that $\llbracket s \rrbracket = \llbracket s' \rrbracket$ and $\llbracket t \rrbracket = \llbracket t' \rrbracket$. So because $\vdash s \lesssim^c t$ we know that:

$$\forall P \in \mathfrak{P}. (\llbracket s \rrbracket = \llbracket s' \rrbracket \in P \implies \llbracket t \rrbracket = \llbracket t' \rrbracket \in P).$$

This is enough to establish that $\vdash s' \lesssim^c t'$. \square

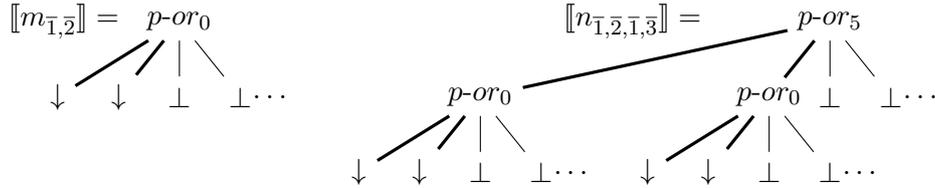
Example 5.2.5 (Probabilistic choice). Consider the following computations:

$$m_{\bar{1}, \bar{2}} = p\text{-or}(\bar{0}, y.\text{case } y \text{ in } \{\text{zero} \Rightarrow (\lambda x:\text{nat}. \downarrow) \bar{1}, \text{succ}(y') \Rightarrow \\ \text{case } y' \text{ in } \{\text{zero} \Rightarrow (\lambda x:\text{nat}. \downarrow) \bar{2}, \text{succ}(y'') \Rightarrow \text{loop}\}\}).$$

Computation $m_{\bar{1}, \bar{3}}$ is defined analogously to $m_{\bar{1}, \bar{2}}$ where the value $\bar{2}$ inside the computation is replaced by the value $\bar{3}$.

$$n_{\bar{1}, \bar{2}, \bar{1}, \bar{3}} = p\text{-or}(\bar{5}, y.\text{case } y \text{ in } \{\text{zero} \Rightarrow m_{\bar{1}, \bar{2}}, \text{succ}(y') \Rightarrow \\ \text{case } y' \text{ in } \{\text{zero} \Rightarrow m_{\bar{1}, \bar{3}}, \text{succ}(y'') \Rightarrow \text{loop}\}\}).$$

Their computation trees are:

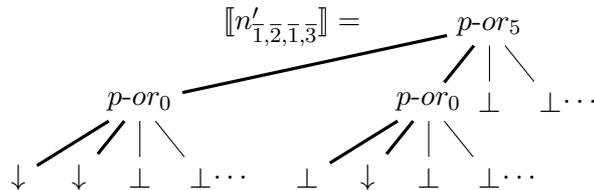


Both computations terminate with a \downarrow with probability 1 so

$$\forall P \in \mathfrak{P}. \llbracket m_{\bar{1}, \bar{2}} \rrbracket \in P \iff \llbracket n_{\bar{1}, \bar{2}, \bar{1}, \bar{3}} \rrbracket \in P$$

is true. Therefore $m_{\bar{1}, \bar{2}}$ and $n_{\bar{1}, \bar{2}, \bar{1}, \bar{3}}$ are bisimilar. Notice that the subscripts 0 and 5 on the $p\text{-or}$ nodes do not play any role in establishing bisimilarity.

However, if we consider computation $n'_{\bar{1}, \bar{2}, \bar{1}, \bar{3}}$ with tree:



it has a probability of 3/4 of terminating. Therefore, $\llbracket m_{\bar{1},\bar{2}} \rrbracket \in \mathbf{P}_{>0.9}$ but $\llbracket n'_{\bar{1},\bar{2},\bar{1},\bar{3}} \rrbracket \notin \mathbf{P}_{>0.9}$ so the two computations are not bisimilar.

Example 5.2.6 (Global store). Plotkin and Power [PP02] axiomatise the behaviour of the global store operations *lookup* and *update* using a set of program equations. We can show that these equations are in fact induced by applicative \mathfrak{B} -bisimilarity. For example, consider the following equation:

$$\forall loc, loc' \in \mathbb{L} \text{ where } loc \neq loc'. \forall n, n' \in \mathbb{N}. \forall (y : \mathbf{nat}, y' : \mathbf{nat} \vdash t).$$

$$l = \text{update}_{loc}(\bar{n}, y.\text{update}_{loc'}(\bar{n}', y'.t)) \sim^c \text{update}_{loc'}(\bar{n}', y'.\text{update}_{loc}(\bar{n}, y.t)) = r.$$

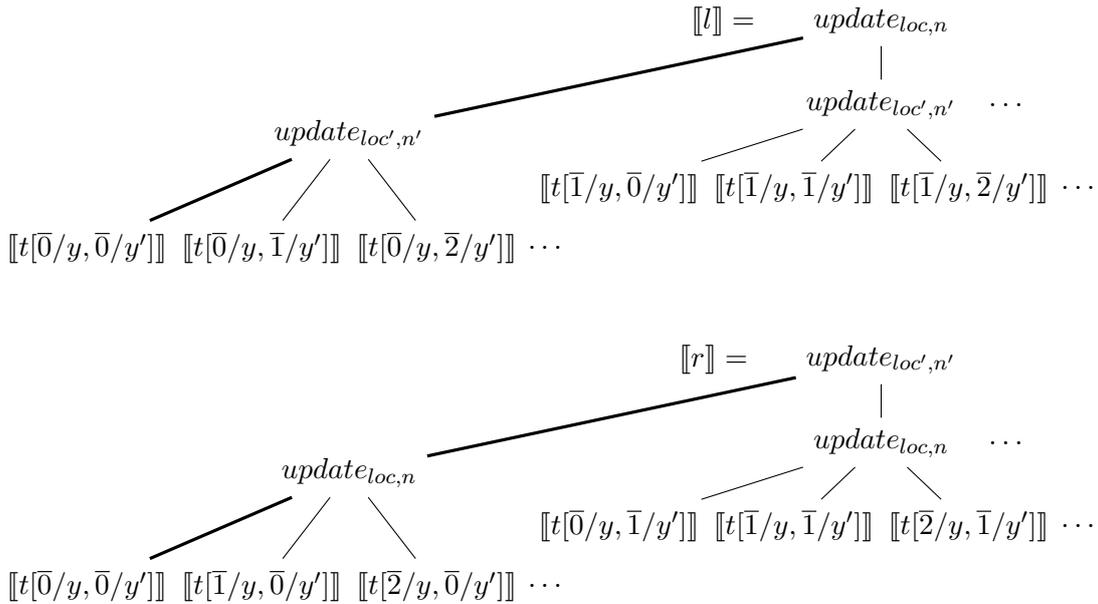
It says that writes to two different locations can be interchanged. Because we are dealing with computations, to prove bisimilarity it suffices to show:

$$\forall s_1, s_2 \in \text{State}. \llbracket l \rrbracket \in (s_1 \rightsquigarrow s_2) \iff \llbracket r \rrbracket \in (s_1 \rightsquigarrow s_2).$$

The observation $(s_1 \rightsquigarrow s_2)$ was defined using the partial function *exec* which formalises the execution of a computation tree. So in fact, we need to show that:

$$\begin{aligned} \forall s_1, s_2 \in \text{State}. \text{exec}(\llbracket l \rrbracket, s_1) \text{ is defined and } \text{exec}(\llbracket l \rrbracket, s_1) = (\downarrow, s_2) \\ \iff \text{exec}(\llbracket r \rrbracket, s_1) \text{ is defined and } \text{exec}(\llbracket r \rrbracket, s_1) = (\downarrow, s_2). \end{aligned} \quad (5.2.1)$$

The computation trees of *l* and *r* respectively are:



Using the trees and the definition of *exec*, we can deduce the following chain of equations,

which is enough to prove equation 5.2.1:

$$\begin{aligned}
exec(\llbracket l \rrbracket, s_1) &= exec(update_{loc', n'}(\overrightarrow{\llbracket t[\bar{0}/y, \bar{i}/y] \rrbracket}), s_1[loc := n]) \\
&= exec(\llbracket t[\bar{0}/y, \bar{0}/y'] \rrbracket, (s_1[loc = n])[loc' = n']) \\
&= exec(\llbracket t[\bar{0}/y, \bar{0}/y'] \rrbracket, s_1[loc = n, loc' = n']) && \text{(because } loc \neq loc') \\
&= exec(\llbracket t[\bar{0}/y, \bar{0}/y'] \rrbracket, (s_1[loc' = n'])[loc = n]) \\
&= exec(update_{loc, n}(\overrightarrow{\llbracket t[\bar{i}/y, \bar{0}/y] \rrbracket}), s_1[loc' = n']) = exec(\llbracket r \rrbracket, s_1) = (\downarrow, s_2).
\end{aligned}$$

The other six equations for global store from [PP02] can be proved analogously. Moreover, Plotkin and Power observe that adding more equations to this set of seven leads to inconsistency. Therefore, these are all the equations between computations for the global store effect.

5.3 Applicative \mathfrak{B} -Bisimilarity is a Congruence

This section discusses the two main properties required for \mathfrak{B} -bisimilarity to be a well-behaved program equivalence: being an equivalence relation and compatibility. The proof of the next lemma appears in Appendix B.

Lemma 5.3.1. *Applicative \mathfrak{B} -similarity is a preorder. Applicative \mathfrak{B} -bisimilarity is an equivalence relation.*

Applicative similarity is a *well-typed relation* on *closed* ECPS terms. Compatibility says we can substitute related programs for a variable inside related programs. Therefore, we need to talk about bisimilarity of programs with free variables. Bisimilarity can be extended to open terms in a standard way [LGL17a].

Definition 5.3.2 (Open extension). *Given a well-typed relation on closed terms, $\mathcal{R} = (\mathcal{R}_A^v, \mathcal{R}^c)$, the open extension of \mathcal{R} is $\mathcal{R}^o = (\mathcal{R}_A^{o,v}, \mathcal{R}^{o,c})$ where:*

$$\begin{aligned}
\overrightarrow{x_i : A_i} \vdash v \mathcal{R}_B^{o,v} w &\iff \forall \overrightarrow{u_i : A_i}. v[\overrightarrow{u_i/x_i}] \mathcal{R}_B^v w[\overrightarrow{u_i/x_i}] \\
\overrightarrow{x_i : A_i} \vdash s \mathcal{R}^{o,c} t &\iff \forall \overrightarrow{u_i : A_i}. s[\overrightarrow{u_i/x_i}] \mathcal{R}^c t[\overrightarrow{u_i/x_i}].
\end{aligned}$$

Definition 5.3.3 (Compatibility [LGL17a]). *A well-typed open relation $\mathcal{R} = (\mathcal{R}_A^v, \mathcal{R}^c)$ is compatible if it is closed under the rules in Figure 5.1. We define \mathcal{R} to be a pre-congruence if it is a compatible preorder, and a congruence if it is a compatible equivalence relation.*

The following lemma identifies some alternative compatibility rules which will be used in later proofs. Its proof can be found in Appendix B.

Lemma 5.3.4. *Consider a well-typed relation \mathcal{R} that is a preorder. The compatibility rules (COMP6), (COMP7), (COMP8) and (COMP10) from Figure 5.1 are equivalent to the conjunction of their single-premise versions. More explicitly:*

$$\begin{array}{c}
\frac{}{\Gamma \vdash x \mathcal{R}_A^v x} \text{(COMP1)} \quad \frac{}{\Gamma \vdash \star \mathcal{R}_{\text{unit}}^v \star} \text{(COMP2)} \\
\\
\frac{\Gamma, x_1 : A_1, \dots, x_n : A_n \vdash s \mathcal{R}^c t}{\Gamma \vdash \lambda \vec{x} : \vec{A}. s \mathcal{R}_{\neg(A_1, \dots, A_n)}^v \lambda \vec{x} : \vec{A}. t} \text{(COMP3)} \\
\\
\frac{}{\Gamma \vdash \text{zero} \mathcal{R}_{\text{nat}}^v \text{zero}} \text{(COMP4)} \quad \frac{\Gamma \vdash v \mathcal{R}_{\text{nat}}^v v'}{\Gamma \vdash \text{succ}(v) \mathcal{R}_{\text{nat}}^v \text{succ}(v')} \text{(COMP5)} \\
\\
\frac{\Gamma \vdash v \mathcal{R}_{\neg(A_1, \dots, A_n)}^v v' \quad \Gamma \vdash w_1 \mathcal{R}_{A_1}^v w'_1, \dots, \Gamma \vdash w_n \mathcal{R}_{A_n}^v w'_n}{\Gamma \vdash v(w_1, \dots, w_n) \mathcal{R}^c v'(w'_1, \dots, w'_n)} \text{(COMP6)} \\
\\
\frac{\Gamma, x : \neg(\vec{A}) \vdash v \mathcal{R}_{\neg(\vec{A})}^v v' \quad \Gamma \vdash w_i \mathcal{R}_{A_i}^v w'_i \text{ for each } i}{\Gamma \vdash (\mu x.v)(\vec{w}) \mathcal{R}^c (\mu x.v')(\vec{w}')} \text{(COMP7)} \\
\\
\frac{\Gamma \vdash v \mathcal{R}_{\text{nat}}^v v' \quad \Gamma, x : \text{nat} \vdash t \mathcal{R}^c t'}{\Gamma \vdash \sigma(v, x.t) \mathcal{R}^c \sigma(v', x.t')} \sigma \in \Sigma \text{ (COMP8)} \quad \frac{}{\Gamma \vdash \downarrow \mathcal{R}^c \downarrow} \text{(COMP9)} \\
\\
\frac{\Gamma \vdash v \mathcal{R}_{\text{nat}}^v v' \quad \Gamma \vdash s \mathcal{R}^c s' \quad \Gamma, x : \text{nat} \vdash t \mathcal{R}^c t'}{\Gamma \vdash \text{case } v \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\} \mathcal{R}^c \text{case } v' \text{ in } \{\text{zero} \Rightarrow s', \text{succ}(x) \Rightarrow t'\}} \\
\text{(COMP10)}
\end{array}$$

Figure 5.1: Compatibility rules.

- Rule (COMP6) is equivalent to the conjunction of the rules:

$$\frac{\Gamma \vdash v \mathcal{R}_{\neg(A_1, \dots, A_n)}^v v' \quad \Gamma \vdash \overline{w_j} : A_j}{\Gamma \vdash v(w_1, \dots, w_n) \mathcal{R}^c v'(w_1, \dots, w_n)} \text{(COMP6L)}$$

$$\frac{\Gamma \vdash v : \neg(\overline{A_j}) \quad \Gamma \vdash (\overline{w_{1,i-1}} : A_{1,i-1}) \quad \Gamma \vdash w_i \mathcal{R}_{A_i}^v w'_i \quad \Gamma \vdash (\overline{w_{i+1,n}} : A_{i+1,n})}{\Gamma \vdash v(w_1, \dots, w_i, \dots, w_n) \mathcal{R}^c v(w_1, \dots, w'_i, \dots, w_n)} \text{(COMP6R}_i \text{) for each } i = \overline{1, n}$$

- Rule (COMP7) is equivalent to the conjunction of the rules:

$$\frac{\Gamma, x : \neg(\overline{A_j}) \vdash v \mathcal{R}_{\neg(\overline{A_j})}^v v' \quad \Gamma \vdash \overline{w_j} : A_j}{\Gamma \vdash (\mu x.v)(\overline{w_j}) \mathcal{R}^c (\mu x.v')(\overline{w_j})} \text{(COMP7L)}$$

$$\frac{\Gamma, x : \neg(\overline{A_j}) \vdash v : \neg(\overline{A_j}) \quad \Gamma \vdash (\overline{w_{1,i-1}} : A_{1,i-1}) \quad \Gamma \vdash w_i \mathcal{R}_{A_i}^v w'_i \quad \Gamma \vdash (\overline{w_{i+1,n}} : A_{i+1,n})}{\Gamma \vdash (\mu x.v)(w_1, \dots, w_i, \dots, w_n) \mathcal{R}^c (\mu x.v)(w_1, \dots, w'_i, \dots, w_n)} \text{(COMP7R}_i \text{) for each } i = \overline{1, n}$$

- Rule (COMP8) is equivalent to the conjunction of the rules:

$$\frac{\Gamma \vdash v \mathcal{R}_{\text{nat}}^v v' \quad \Gamma, x : \text{nat} \vdash t}{\Gamma \vdash \sigma(v, x.t) \mathcal{R}^c \sigma(v', x.t)} \sigma \in \Sigma \text{(COMP8L)}$$

$$\frac{\Gamma \vdash v : \text{nat} \quad \Gamma, x : \text{nat} \vdash t \mathcal{R}^c t'}{\Gamma \vdash \sigma(v, x.t) \mathcal{R}^c \sigma(v, x.t')} \sigma \in \Sigma \text{(COMP8R)}$$

- Rule (COMP10) is equivalent to the conjunction of the rules:

$$\frac{\Gamma \vdash v \mathcal{R}_{\text{nat}}^v v' \quad \Gamma \vdash s \quad \Gamma, x : \text{nat} \vdash t}{\Gamma \vdash \text{case } v \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\} \mathcal{R}^c \text{case } v' \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\}} \text{(COMP10V)}$$

$$\frac{\Gamma \vdash v : \text{nat} \quad \Gamma \vdash s \mathcal{R}^c s' \quad \Gamma, x : \text{nat} \vdash t}{\Gamma \vdash \text{case } v \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\} \mathcal{R}^c \text{case } v \text{ in } \{\text{zero} \Rightarrow s', \text{succ}(x) \Rightarrow t\}} \text{(COMP10L)}$$

$$\frac{\Gamma \vdash v : \text{nat} \quad \Gamma \vdash s \quad \Gamma, x : \text{nat} \vdash t \mathcal{R}^c t'}{\Gamma \vdash \text{case } v \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\} \mathcal{R}^c \text{case } v \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t'\}} \text{(COMP10R)}$$

To prove \mathfrak{B} -bisimilarity is a congruence, we have identified two sufficient conditions that the set of observations \mathfrak{B} should satisfy. One of them is that every observation needs

to be Scott-open, as in the work of Simpson and Voorneveld [SV18]. The second condition is that \mathfrak{P} needs to be decomposable for a novel definition of decomposability.

Definition 5.3.5 (Scott-openness). *A set of trees X is Scott-open if:*

1. *It is upwards closed, that is: $tr \in X$ and $tr \leq tr'$ imply $tr' \in X$.*
2. *Whenever $tr_1 \leq tr_2 \leq \dots$ is an ascending chain with least upper bound $\bigsqcup tr_i \in X$, then $tr_j \in X$ for some j .*

Definition 5.3.6 (Decomposability). *The set of observations \mathfrak{P} is decomposable if for any $P \in \mathfrak{P}$ and for any $tr \in P$:*

$$tr = \sigma_n(\vec{tr}') \implies \exists \vec{P}' \in \mathfrak{P}. \vec{tr}' \in \vec{P}' \text{ and } \forall \vec{p}' \in \vec{P}'. \sigma_n(\vec{p}') \in P.$$

Decomposability says that, whenever a tree tr is part of an observation P , the children of tr 's root should themselves be part of some observations which fully capture the restrictions that P places on them. This is true for all examples of effects considered so far. Using the last two definitions we can state the main theorem of this chapter:

Theorem 5.3.7. *Given a decomposable set of Scott-open observations \mathfrak{P} :*

1. *The open extension of applicative \mathfrak{P} -similarity, \lesssim° , is compatible, and hence it is a precongruence.*
2. *The open extension of applicative \mathfrak{P} -bisimilarity, \sim° , is compatible, and hence it is a congruence.*

It is easy to check that all the running examples of observations are upwards closed. The proof that they satisfy the second condition in the definition of Scott-openness is the same as in [SV18]. It remains to check that \mathfrak{P} is decomposable:

Example 5.3.8 (Pure functional computation). The only observation is $\Downarrow = \{\downarrow\}$. There are no trees in \Downarrow whose root has children, so decomposability is satisfied.

Example 5.3.9 (Nondeterminism). Recall that $\mathfrak{P} = \{\text{Trees}_\Sigma, \diamond, \square\}$. For any $tr \in \text{Trees}_\Sigma$ we can choose each $P'_n = \text{Trees}_\Sigma$ to fulfil decomposability.

Consider $tr \in \diamond$. Either $tr = \downarrow$, in which case we are done, or $tr = or_n(tr'_0, tr'_1, tr'_2 \dots)$. It must be the case that either tr'_0 or tr'_1 have a reachable \downarrow -leaf. Without loss of generality, assume tr'_0 has a reachable \downarrow -leaf. Then we know $tr'_0 \in \diamond$ so we can choose $P'_0 = \diamond, P'_1 = \text{Trees}_\Sigma, P'_2 = \text{Trees}_\Sigma, \dots$. For any $\vec{p}' \in \vec{P}'$ we know $or_n(\vec{p}') \in \diamond$ because p'_0 has a reachable \downarrow -leaf.

The argument for $tr \in \square$ is analogous, if we choose $P'_0 = \square, P'_1 = \square, P'_2 = \text{Trees}_\Sigma, P'_3 = \text{Trees}_\Sigma, \dots$. This proof relies on the fact that Trees_Σ is an observation, which is why we decided to include it in \mathfrak{P} .

Example 5.3.10 (Probabilistic choice). Consider $tr = p-or_n(tr'_0, tr'_1, tr'_2, \dots) \in \mathbf{P}_{>q}$ for some $q \in \mathbb{Q}, 0 \leq q < 1$. Recall the definition of the partial function \mathbb{P} from Example 5.1.3. We know that:

$$\mathbb{P}(tr) = \frac{1}{2}\mathbb{P}(tr'_0) + \frac{1}{2}\mathbb{P}(tr'_1) > q. \quad (5.3.1)$$

Define:

$$q_0 = \frac{\mathbb{P}(tr'_0)}{\mathbb{P}(tr'_0) + \mathbb{P}(tr'_1)} \cdot 2q$$

$$q_1 = \frac{\mathbb{P}(tr'_1)}{\mathbb{P}(tr'_0) + \mathbb{P}(tr'_1)} \cdot 2q.$$

The two probabilities $\mathbb{P}(tr'_0)$ and $\mathbb{P}(tr'_1)$ are rational numbers because they are defined as a sum of rational numbers. The threshold q is rational by assumption, so q_1 and q_2 are rational.

From equation 5.3.1 we can deduce that:

$$1 \geq \mathbb{P}(tr'_0) > q_0$$

$$1 \geq \mathbb{P}(tr'_1) > q_1.$$

So we can choose $P'_0 = \mathbf{P}_{>q_0}$, $P'_1 = \mathbf{P}_{>q_1}$, $P'_2 = \text{Trees}_\Sigma$, $P'_3 = \text{Trees}_\Sigma, \dots$. Therefore, $\vec{tr}' \in \vec{P}'$ as required.

Consider some other subtrees $\vec{p}' \in \vec{P}'$. By the way we defined \vec{P}' it follows that:

$$\frac{1}{2}\mathbb{P}(p'_0) + \frac{1}{2}\mathbb{P}(p'_1) > \frac{1}{2}(q_0 + q_1) = q$$

so $p\text{-or}(\vec{p}') \in \mathbf{P}_{>q}$ as required.

Example 5.3.11 (Global store). Consider a tree $tr = \sigma_n(tr'_0, tr'_1, tr'_2, \dots) \in (s \mapsto r)$. It must be the case that $\text{exec}(tr, s) = (\downarrow, r)$.

If $\sigma_n = \text{lookup}_{i,n}$: because $\text{exec}(tr, s)$ is defined it must be the case that $\text{exec}(tr'_{s(l)}, s)$ is also defined and $\text{exec}(tr'_{s(l)}, s) = \text{exec}(tr, s) = (\downarrow, r)$ so we know that $tr'_{s(l)} \in (s \mapsto r)$. In the definition of decomposability, choose $P'_{s(l)} = (s \mapsto r)$ and $P'_{k \neq s(l)} = \text{Trees}_\Sigma$ and we are done.

If $\sigma_n = \text{update}_{i,n}$: because $\text{exec}(tr, s)$ is defined it must be the case that $\text{exec}(tr'_0, s[l := n])$ is also defined and $\text{exec}(tr'_0, s[l := n]) = \text{exec}(tr, s) = (\downarrow, r)$. Therefore $tr'_0 \in (s[l := n] \mapsto r)$. We can choose $P'_0 = (s[l := n] \mapsto r)$ and $P'_{k \neq 0} = \text{Trees}_\Sigma$ and we are done.

Example 5.3.12 (Input/output). Consider a tree $tr = \sigma_n(tr'_0, tr'_1, tr'_2, \dots) \in \langle w \rangle \dots$. If $w = \epsilon$, then decomposability is immediately satisfied by choosing $P'_k = \langle \epsilon \rangle \dots$. Assume $w \neq \epsilon$.

If $\sigma_n = \text{read}_n$, it must be the case that $w = (?k)w'$ and $tr'_k \models \langle w' \rangle \dots$. We can choose $P'_k = \langle w' \rangle \dots$ and $P'_{m \neq k} = \langle \epsilon \rangle \dots$ and we are done.

If $\sigma_n = \text{write}_n$, then $w = (!n)w'$ and $tr'_0 \models \langle w' \rangle \dots$. Choose $P'_0 = \langle w' \rangle \dots$ and $P'_{k \neq 0} = \langle \epsilon \rangle \dots$ and we are done.

5.4 Howe's Method

To prove Theorem 5.3.7 we will use a method originally due to Howe [How96]. The strategy is to define a relation $\lesssim^{\mathcal{H}}$ named the Howe extension of \lesssim , prove that it is compatible, and then prove that it coincides with \lesssim° .

$$\begin{array}{c}
\frac{}{\Gamma \vdash x \widehat{\mathcal{R}}_A^v x} \text{(C1)} \quad \frac{}{\Gamma \vdash \star \widehat{\mathcal{R}}_{\text{unit}}^v \star} \text{(C2)} \quad \frac{\Gamma, x_1 : A_1, \dots, x_n : A_n \vdash s \mathcal{R}^c t}{\Gamma \vdash \lambda \vec{x} : \vec{A}. s \widehat{\mathcal{R}}_{-(A_1, \dots, A_n)}^v \lambda \vec{x} : \vec{A}. t} \text{(C3)} \\
\\
\frac{}{\Gamma \vdash \text{zero} \widehat{\mathcal{R}}_{\text{nat}}^v \text{zero}} \text{(C4)} \quad \frac{\Gamma \vdash v \mathcal{R}_{\text{nat}}^v v'}{\Gamma \vdash \text{succ}(v) \widehat{\mathcal{R}}_{\text{nat}}^v \text{succ}(v')} \text{(C5)} \\
\\
\frac{\Gamma \vdash v \mathcal{R}_{-(A_1, \dots, A_n)}^v v' \quad \Gamma \vdash w_1 \mathcal{R}_{A_1}^v w'_1, \dots, \Gamma \vdash w_n \mathcal{R}_{A_n}^v w'_n}{\Gamma \vdash v(w_1, \dots, w_n) \widehat{\mathcal{R}}^c v'(w'_1, \dots, w'_n)} \text{(C6)} \\
\\
\frac{\Gamma, x : \neg(\vec{A}) \vdash v \mathcal{R}_{-\vec{A}}^v v' \quad \Gamma \vdash w_i \mathcal{R}_{A_i}^v w'_i \text{ for each } i}{\Gamma \vdash (\mu x.v)(\vec{w}) \widehat{\mathcal{R}}^c (\mu x.v')(\vec{w}')} \text{(C7)} \\
\\
\frac{\Gamma \vdash v \mathcal{R}_{\text{nat}}^v v' \quad \Gamma, x : \text{nat} \vdash t \mathcal{R}^c t'}{\Gamma \vdash \sigma(v, x.t) \widehat{\mathcal{R}}^c \sigma(v', x.t')} \sigma \in \Sigma \text{(C8)} \quad \frac{}{\Gamma \vdash \downarrow \widehat{\mathcal{R}}^c \downarrow} \text{(C9)} \\
\\
\frac{\Gamma \vdash v \mathcal{R}_{\text{nat}}^v v' \quad \Gamma \vdash s \mathcal{R}^c s' \quad \Gamma, x : \text{nat} \vdash t \mathcal{R}^c t'}{\Gamma \vdash \text{case } v \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\} \widehat{\mathcal{R}}^c \text{case } v' \text{ in } \{\text{zero} \Rightarrow s', \text{succ}(x) \Rightarrow t'\}} \text{(C10)}
\end{array}$$

Figure 5.2: Compatible refinement rules.

The justification that Pitts [Pit11] gives for the use of Howe's method in the case of the untyped λ -calculus applies to ECPS as well. A direct proof that \lesssim° is compatible is problematic because it requires proving a substitutivity property of \lesssim° which is very close to compatibility. Howe's method avoids this problem because $\lesssim^{\mathcal{H}}$ is compatible by construction, and can be proved substitutive.

The proof of compatibility of applicative \mathfrak{B} -bisimilarity follows a similar structure to that for applicative bisimilarity for EPCF, found in [SV17, Appendix] and [SV18, Section 6]. We present all the proofs in detail, filling in gaps in Simpson's and Voorneveld's presentation, and adapting to the setting of ECPS.

Definition 5.4.1 (Compatible refinement). *Given a well-typed open relation \mathcal{R} its compatible refinement $\widehat{\mathcal{R}}$ is inductively defined by the rules in Figure 5.2.*

Definition 5.4.2 (Howe extension). *Given a well-typed closed relation \mathcal{R} , we define its Howe extension $\mathcal{R}^{\mathcal{H}}$ to be the least relation \mathcal{S} such that $\mathcal{S} = \mathcal{R}^\circ \circ \widehat{\mathcal{S}}$.*

It has been observed by Levy ([Lev06, Proposition 5.4]) that the equation above determines a unique relation. The Howe extension can equivalently be defined inductively as the least relation closed under the rules:

$$\frac{\Gamma \vdash s \widehat{\mathcal{R}}^c t \quad \Gamma \vdash t \mathcal{R}^{\circ, c} r}{\Gamma \vdash s \mathcal{R}^{\mathcal{H}, c} r} \text{(HC)} \quad \frac{\Gamma \vdash v \widehat{\mathcal{R}}_A^v w \quad \Gamma \vdash w \mathcal{R}_A^{\circ, v} u}{\Gamma \vdash v \mathcal{R}_A^{\mathcal{H}, v} u} \text{(HV)}$$

This is shown in [LGL17b].

Below are two lemmas about the open extension and the Howe extension of a relation. Their proofs appear in Appendix B.1.

Lemma 5.4.3 (From [SV17, Appendix]). *Given a well-typed relation \mathcal{R} on closed terms that is reflexive:*

1. *The Howe extension of \mathcal{R} , $\mathcal{R}^{\mathcal{H}}$, is compatible and hence reflexive.*
2. $\mathcal{R}^{\circ} \subseteq \mathcal{R}^{\mathcal{H}}$.

Lemma 5.4.4 (From [SV17, Appendix]). *Given a well-typed relation \mathcal{R} on closed terms that is transitive:*

$$\mathcal{R}^{\circ} \circ \mathcal{R}^{\mathcal{H}} \subseteq \mathcal{R}^{\mathcal{H}}.$$

Lemma 5.4.5 (Substitutivity). *Given a well-typed relation \mathcal{R} on closed terms that is transitive, its Howe extension satisfies the following two value-substitutivity properties:*

1. $\overrightarrow{x_i : A_i}, y : B \vdash s \mathcal{R}^{\mathcal{H},c} t$ and $\overrightarrow{x_i : A_i} \vdash v \mathcal{R}_B^{\mathcal{H},v} w \implies \overrightarrow{x_i : A_i} \vdash s[v/y] \mathcal{R}^{\mathcal{H},c} t[w/y]$.
2. $\overrightarrow{x_i : A_i}, y : B \vdash u \mathcal{R}_C^{\mathcal{H},v} u'$ and $\overrightarrow{x_i : A_i} \vdash v \mathcal{R}_B^{\mathcal{H},v} w \implies \overrightarrow{x_i : A_i} \vdash u[v/y] \mathcal{R}_C^{\mathcal{H},v} u'[w/y]$.

Proof. The proof is done by induction on the structure of s and u . It can be found in Appendix B.1. \square

The following lemma will help prove that $\lesssim^{\mathcal{H}}$ restricted to closed terms is a simulation. Its proof appears in Appendix B.1.

Lemma 5.4.6. *Consider a well-typed closed relation \leq that is a \mathfrak{P} -simulation. For any closed values v and w :*

$$\vdash v \leq_{\text{nat}}^{\mathcal{H},v} w \implies v = w.$$

Using the domain theoretic definition of ECPS computation trees (Definition 3.2.2) we can state the Key Lemma which will help us prove $\lesssim^{\mathcal{H}}$ is a simulation. Recall that $\llbracket s \rrbracket_n$ is the tree resulting from n steps of evaluation of s .

Lemma 5.4.7 (Key Lemma). *Consider a decomposable set of Scott-open observations \mathfrak{P} . Consider a well-typed closed relation \leq that is a preorder and a \mathfrak{P} -simulation. For any closed computations s and t , $\vdash s \leq^{\mathcal{H},c} t$ implies:*

$$\forall n \in \mathbb{N}. \forall P \in \mathfrak{P}. \llbracket s \rrbracket_n \in P \implies \llbracket t \rrbracket \in P.$$

Proof. The proof is done by induction on $n \in \mathbb{N}$. It uses the fact that observations are upwards closed and \mathfrak{P} is decomposable. The full proof appears in Appendix B.1. \square

Proposition 5.4.8. *Consider a decomposable set of Scott-open observations \mathfrak{P} . Consider a well-typed closed relation \leq that is a preorder and a \mathfrak{P} -simulation. The Howe extension of \leq , $\leq^{\mathcal{H}}$, restricted to closed terms is an applicative simulation.*

Proof. We need to verify that all four conditions in the definition of applicative simulation are satisfied by $\leq^{\mathcal{H}}$ for closed terms.

1. $\vdash v \leq_{\text{unit}}^{\mathcal{H},v} w \implies v = w = \star$. Assume $\vdash v \leq_{\text{unit}}^{\mathcal{H},v} w$. The only closed value of type `unit` is \star so $v = w = \star$.
2. $\vdash v \leq_{\text{nat}}^{\mathcal{H},v} w \implies v = w$. This is Lemma 5.4.6.
3. $\vdash s \leq^{\mathcal{H},c} t \implies \forall P \in \mathfrak{P}. (\llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P)$. Assume that $\vdash s \leq^{\mathcal{H},c} t$ and $\llbracket s \rrbracket \in P$ for some $P \in \mathfrak{P}$. From the Key Lemma (Lemma 5.4.7) we know that:

$$\forall n \in \mathbb{N}. \forall P \in \mathfrak{P}. \llbracket s \rrbracket_n \in P \implies \llbracket t \rrbracket \in P. \quad (5.4.1)$$

We know $\llbracket s \rrbracket = \bigsqcup_{m \in \mathbb{N}} \llbracket s \rrbracket_m \in P$ and that $\{\llbracket s \rrbracket_m\}_{m \in \mathbb{N}}$ is an ascending chain. By Scott-openness of P , there exists $j \in \mathbb{N}$ such that $\llbracket s \rrbracket_j \in P$.

Therefore, by equation 5.4.1 we have the desired result: $\llbracket t \rrbracket \in P$.

4. $\vdash v \leq_{\neg(A_1, \dots, A_n)}^{\mathcal{H},v} w \implies \forall \vdash u_1 : A_1, \dots, \vdash u_n : A_n. v(u_1, \dots, u_n) \leq^{\mathcal{H},c} w(u_1, \dots, u_n)$. Assume $\vdash v \leq_{\neg(A_1, \dots, A_n)}^{\mathcal{H},v} w$. By Lemma 5.4.3 we know $\leq^{\mathcal{H}}$ is compatible and reflexive. Therefore, for all $i, u_i \leq_{A_i}^{\mathcal{H},v} u_i$ so by compatibility $v(u_1, \dots, u_n) \leq^{\mathcal{H},c} w(u_1, \dots, u_n)$. □

The following lemmas will be used to prove that similarity and bisimilarity are compatible. Their proofs can be found in Appendix B.1.

Lemma 5.4.9. *Given a well-typed open relation \mathcal{R} that is reflexive and has the two substitutivity properties from Lemma 5.4.5, and a well-typed closed relation \mathcal{S} then:*

if \mathcal{R} restricted to closed terms is included in \mathcal{S} then $\mathcal{R} \subseteq \mathcal{S}^\circ$.

Lemma 5.4.10. *Given a \mathfrak{P} -simulation \mathcal{R} , its reflexive-transitive closure, \mathcal{R}^* is also a \mathfrak{P} -simulation.*

Lemma 5.4.11. *Given a well-typed compatible relation \mathcal{R} , its reflexive-transitive closure \mathcal{R}^* is also compatible.*

Lemma 5.4.12 (From [Las98]). *Given a well-typed closed relation \mathcal{R} the following holds:*

if \mathcal{R}° is reflexive and symmetric, then $\mathcal{R}^{\mathcal{H}^}$ is symmetric.*

Where S^* denotes the reflexive-transitive closure of a relation \mathcal{S} .

Finally, we can prove similarity and bisimilarity are compatible. We first recall the formal statement of this:

Theorem 5.3.7. *Given a decomposable set of Scott-open observations \mathfrak{P} :*

1. *The open extension of applicative \mathfrak{P} -similarity, \lesssim° , is compatible, and hence it is a precongruence.*
2. *The open extension of applicative \mathfrak{P} -bisimilarity, \sim° , is compatible, and hence it is a congruence.*

Proof. This proof has the same structure as the proof of Theorem 3 from [SV18]. Here, we present significant details that were missing.

1. We need to prove that the open extension of applicative \mathfrak{P} -similarity, \simeq° , is compatible. We know that \simeq is a preorder (Lemma 5.3.1) and a simulation. Therefore, we can apply Proposition 5.4.8 to deduce that the restriction of $\simeq^{\mathcal{H}}$ to closed terms is a simulation, so it is included in the greatest simulation, \simeq .

Since \simeq is transitive we know from Lemma 5.4.5 that $\simeq^{\mathcal{H}}$ has the two substitution properties. Because \simeq is reflexive, we know from Lemma 5.4.3 that $\simeq^{\mathcal{H}}$ is reflexive.

We can use all these to apply Lemma 5.4.9 for $\simeq^{\mathcal{H}}$ and \simeq to deduce $\simeq^{\mathcal{H}} \subseteq \simeq^\circ$.

From Lemma 5.4.3 we already know that $\simeq^\circ \subseteq \simeq^{\mathcal{H}}$ and that $\simeq^{\mathcal{H}}$ is compatible. Therefore, the open extension of applicative \mathfrak{P} -similarity equals the Howe extension, $\simeq^\circ = \simeq^{\mathcal{H}}$, and is compatible.

Since \simeq is a preorder, \simeq° is also a preorder, so it is a precongruence.

2. We need to prove that \sim° is compatible. The relation \sim is also a \mathfrak{P} -simulation. From Lemma 5.3.1 we know \sim is an equivalence relation, hence a preorder. Therefore we can use Proposition 5.4.8 to deduce that $\sim^{\mathcal{H}}$ restricted to closed terms is a simulation.

From Lemma 5.4.10, we obtain that $\sim^{\mathcal{H}^*}$ restricted to closed terms is a simulation, because restricting to closed terms and taking the reflexive-transitive closure are commutative operations.

Because \sim is reflexive and symmetric it follows that \sim° is also reflexive and symmetric. We can then apply Lemma 5.4.12 for \sim to deduce that $\sim^{\mathcal{H}^*}$ is symmetric. Therefore, $\sim^{\mathcal{H}^*}$ restricted to closed terms is also symmetric.

As a result, we know that $\sim^{\mathcal{H}^*}$ restricted to closed terms is a bisimulation, so it is included in the greatest bisimulation, \sim .

Now we would like to use Lemma 5.4.9 to deduce $\sim^{\mathcal{H}^*} \subseteq \sim^\circ$. We know \sim is transitive so we can apply Lemma 5.4.5 to deduce that $\sim^{\mathcal{H}}$ has the required substitutivity properties. Then it is easy to prove using transitivity and reflexivity that $\sim^{\mathcal{H}^*}$ also has the required substitutivity properties. Moreover, $\sim^{\mathcal{H}^*}$ is reflexive by definition. So we can apply Lemma 5.4.9.

Because \sim is reflexive we know from Lemma 5.4.3 that $\sim^\circ \subseteq \sim^{\mathcal{H}}$ and that $\sim^{\mathcal{H}}$ is compatible. By definition of the reflexive transitive closure we know $\sim^{\mathcal{H}} \subseteq \sim^{\mathcal{H}^*}$.

Therefore we know $\sim^{\mathcal{H}^*} \subseteq \sim^\circ$ and $\sim^\circ \subseteq \sim^{\mathcal{H}} \subseteq \sim^{\mathcal{H}^*}$. So $\sim^\circ = \sim^{\mathcal{H}^*}$.

Because $\sim^{\mathcal{H}}$ is compatible, from Lemma 5.4.11, $\sim^{\mathcal{H}^*}$ is also compatible. Therefore, \sim° is compatible. Since \sim is an equivalence relation, \sim° is also an equivalence relation, so it is a congruence.

□

5.5 Chapter Summary

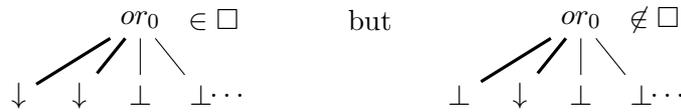
This chapter started by defining a set \mathfrak{P} of *observations*, for each of the following effects: nondeterminism, probabilistic choice, global store and I/O. In ECPS, we can observe termination and some effect operations, both of which are encoded by computation trees.

Therefore, each observation is a set of trees. For example, for nondeterminism the observations are:

$$\diamond = \{\text{trees in which at least one execution path that can occur ends in } \downarrow\}$$

$$\square = \{\text{trees in which all paths that can occur have finite height and end in } \downarrow\}.$$

We used observations to define applicative bisimilarity for ECPS (Definition 5.2.2). Two values of ground type are bisimilar when they are equal. Two functions are bisimilar if, for all arguments, they yield bisimilar computations. Finally, two computations are bisimilar if they belong to exactly the same elements of \mathfrak{P} , in other words, they have the same observable behaviour. For example, two computations with the trees below are not bisimilar:



Next, we defined compatibility formally (Definition 5.3.3). The definition says that two related programs can be substituted in related contexts to yield another pair of related programs. We identified two sufficient conditions that the set of observations \mathfrak{P} should satisfy in order for bisimilarity to be compatible. These are Scott-openness, which has to do with the topology of the elements of \mathfrak{P} , and *decomposability* (Definition 5.3.6). This definition of decomposability is novel. It states that, for each tree tr in an observation $P \in \mathfrak{P}$, there exist observations that characterise the restrictions that P places on the children of tr .

The main result of this chapter is Theorem 5.3.7, which states that, given a decomposable set \mathfrak{P} of Scott-open observations, bisimilarity is compatible. The proof of this uses Howe's method (Section 5.4) and required significant effort. The idea is to define a new relation named the Howe extension of bisimilarity, which is compatible by construction, and prove it equal to bisimilarity.

Chapter 6

Logical Equivalence for ECPS

This chapter introduces the logic \mathcal{F} whose formulas express properties of ECPS terms. The logic is defined using the set of observations \mathfrak{P} . We prove that program equivalence induced by \mathcal{F} is an equivalence relation and coincides with applicative \mathfrak{P} -bisimilarity. Therefore, using the main theorem of Chapter 5, we can deduce that \mathcal{F} -logical equivalence is compatible. The results in this chapter will be used in Chapter 7 to show that the logic \mathcal{F} characterises contextual equivalence, which is our main goal.

6.1 Two Logics for ECPS

Recall the set of observations \mathfrak{P} , defined in Section 5.1, which contains sets of ECPS effect trees. Each $P \in \mathfrak{P}$ specifies the shape of computation trees for a particular effect. Using \mathfrak{P} , we define two slightly different logics for ECPS named \mathcal{V} and \mathcal{F} respectively.

In the logic \mathcal{V} values appear inside logical formulas, whereas this is not the case in \mathcal{F} . Both logics make a distinction between value formulas and computation formulas. Value formulas are always associated an ECPS type.

Definition 6.1.1 (Logic \mathcal{F}). *The value formulas of the logic \mathcal{F} are constructed from basic formulas $\phi = \{n\}$ and $(\phi_1, \dots, \phi_n) \mapsto P$, where P is an observation from \mathfrak{P} , according to the rules in Figure 6.1. In these rules, A stands for an ECPS type. The computation formulas are the elements of \mathfrak{P} .*

The satisfaction relation \models relates a closed value $\vdash v : A$ to a value formula $\phi : A$ of the same type, or a closed computation t to an observation P . The definition of \models appears in Figure 6.2. Intuitively, $v \models \phi$ means that the program v has property ϕ .

Let \mathcal{F}^+ be the fragment of \mathcal{F} without negation.

Definition 6.1.2 (Logic \mathcal{V}). *The logic \mathcal{V} is the same as \mathcal{F} except that the (VAL) rule is replaced by:*

$$\frac{\vdash w_1 : A_1 \dots \vdash w_n : A_n}{(w_1, \dots, w_n) \mapsto P : \neg(A_1, \dots, A_n)} P \in \mathfrak{P} \text{ (VAL')}$$

$$v \models (w_1, \dots, w_n) \mapsto P \iff \llbracket v(w_1 \dots w_n) \rrbracket \in P.$$

That is, formulas of function type are now constructed using ECPS values.

Let \mathcal{V}^+ be the fragment of \mathcal{V} without negation.

$$\begin{array}{c}
\frac{n \in \mathbb{N}}{\{n\} : \mathbf{nat}} \text{(NAT)} \quad \frac{\phi_1 : A_1 \dots \phi_n : A_n}{(\phi_1, \dots, \phi_n) \mapsto P : \neg(A_1, \dots, A_n)} P \in \mathfrak{P} \text{ (VAL)} \\
\\
\frac{(\phi_i : A)_{i \in I}}{\bigvee_{i \in I} \phi_i : A} \text{(DISJ)} \quad \frac{(\phi_i : A)_{i \in I}}{\bigwedge_{i \in I} \phi_i : A} \text{(CONJ)} \quad \frac{\phi : A}{\neg \phi : A} \text{(NEG)}
\end{array}$$

Figure 6.1: Value formulas in the logic \mathcal{F} .

$$\begin{array}{l}
v \models \{n\} \iff v = \bar{n} \\
v \models (\phi_1, \dots, \phi_n) \mapsto P \iff \text{for all closed values } w_1, \dots, w_n \text{ such that } w_i \models \phi_i \\
\text{then } \llbracket v(w_1 \dots w_n) \rrbracket \in P \\
v \models \bigvee_{i \in I} \phi_i \iff \text{there exists } j \in I \text{ such that } v \models \phi_j \\
v \models \bigwedge_{i \in I} \phi_i \iff \text{for all } j \in I, v \models \phi_j \\
v \models \neg \phi \iff \text{it is false that } v \models \phi \\
t \models P \iff \llbracket t \rrbracket \in P
\end{array}$$

Figure 6.2: Satisfaction relation \models for the logic \mathcal{F} .

Notice that a computation formula is just one of the observations $P \in \mathfrak{P}$. For example, it can be \square , $\mathbf{P}_{>0.5}$, $(s \rightsquigarrow r)$ depending on the effects present in the language. Therefore, $t \models P$ tests the shape of the computation tree of t without looking at its possible return values. This is consistent with the fact that ECPS computations are not expected to return.

However, this is unlike computation formulas in EPCF logic, $o\phi$, which test whether return values satisfy ϕ (Section 2.6). Therefore, P is no longer a modality in the same sense as o because it does not lift formulas. This is why we called P an *observation*. In the logic \mathcal{F} , it can still be argued that P is a modality in the traditional sense because it takes the value formula ϕ to another value formula $\phi \mapsto P$.

In both \mathcal{F} and \mathcal{V} the value formulas of type \mathbf{nat} are obtained from the natural numbers, arbitrary conjunctions and disjunctions, and negation. However, the basic formulas of function type are different. In \mathcal{V} :

$$v \models (\lambda x : \mathbf{nat}. \downarrow) \mapsto \diamond$$

means that $v : \neg(\neg(\mathbf{nat}))$ may terminate when given argument $\lambda x : \mathbf{nat}. \downarrow$. In \mathcal{F} , an analogous statement is:

$$v \models ((\bigvee_{n \in \mathbb{N}} \{n\}) \mapsto \square) \mapsto \diamond.$$

This says that v may terminate when given as argument a function that satisfies $\psi =$

$(\bigvee_{n \in \mathbb{N}} \{n\}) \mapsto \square$. Notice that $(\lambda x:\mathbf{nat}. \downarrow)$ indeed satisfies ψ .

There is no need to include logical connectives at the level on computation formulas because they can be encoded in value formulas. For example $\phi \mapsto \bigwedge_{i \in I} P_i$ can be expressed as:

$$\bigwedge_{i \in I} (\phi \mapsto P_i).$$

The statement $t \models \bigwedge_{i \in I} P_i$ can instead be expressed as:

$$\lambda x:\mathbf{unit}.t \models \bigwedge_{i \in I} (\mathit{true} \mapsto P_i).$$

The indexing set I in $\bigwedge_{i \in I}$ may be uncountable. However, the sets of values and computations are countable. Since logical formulas are interpreted over values and computations, all conjunctions and disjunctions are semantically equivalent to countable ones.

The following example compares logical formulas for EPCF with \mathcal{F} -formulas:

Example 6.1.3 (Nondeterminism). Recall the EPCF logic formula $\phi_2 = \{3\} \mapsto \diamond\{2\}$ from Examples 2.2.1 and 2.6.1 and the function:

$$g = \lambda n:\mathbb{N}.\mathit{or}(\mathbf{pred} \ n, \ \mathbf{succ} \ n).$$

We have previously established that $g \models_{EPCF} \phi_2$ because, when given argument $\bar{3}$ in the empty stack id , $(id, g \ \bar{3})$ may return $\bar{2}$.

Now consider the CPS translation (Section 4.1) of the function g into ECPS, named g^* . Apart from the natural number argument that g receives, g^* also receives a continuation k to which it passes its result.

$$\begin{aligned} g^* &= \lambda(n, k):(\mathbf{nat}, \neg\mathbf{nat}). \\ & \quad (\lambda k':\neg\mathbf{nat}.\mathit{or}(\bar{0}, x.(\lambda(y, k'')):(\mathbf{nat}, \neg\mathbf{nat}). \\ & \quad \quad (\lambda k''':\neg\mathbf{nat}.\mathit{case} \ y \ \mathbf{in} \ \{\mathbf{zero} \Rightarrow (\mathbf{pred} \ n)^* \ k''', \\ & \quad \quad \quad \mathbf{succ}(y') \Rightarrow \mathit{case} \ y' \ \mathbf{in} \ \{\mathbf{zero} \Rightarrow (\mathbf{succ} \ n)^* \ k''', \\ & \quad \quad \quad \quad \mathbf{succ}(y'') \Rightarrow \mathit{loop}^* \ k'''\}) \\ & \quad \quad \quad) \ k'' \\ & \quad \quad \quad) \ (x, k') \\ & \quad \quad \quad) \\ & \quad \quad) \ k. \end{aligned}$$

The tree of $(g^* \ (\bar{3}, k))$ is:

$$\llbracket g^* \ (\bar{3}, k) \rrbracket = \begin{array}{c} \mathit{or}0 \\ / \quad \backslash \\ \llbracket k \ \bar{2} \rrbracket \quad \llbracket k \ \bar{4} \rrbracket \quad \perp \ \dots \end{array} \quad \text{in particular } \llbracket g^* \ (\bar{3}, id^*) \rrbracket = \llbracket g^* \ (\bar{3}, \lambda x:\mathbf{nat}. \downarrow) \rrbracket = \begin{array}{c} \mathit{or}0 \\ / \quad \backslash \\ \downarrow \quad \downarrow \quad \perp \ \dots \end{array}$$

The formula ϕ_2 could be translated to the \mathcal{F} formula:

$$\phi_2^* = ((\{3\}, \{2\} \mapsto \diamond) \mapsto \diamond) \wedge ((\{3\}, \{2\} \mapsto \square) \mapsto \diamond).$$

Intuitively, $((\{3\}, \{2\} \mapsto \diamond) \mapsto \diamond)$ says that, when g^* is given as arguments $\bar{3}$ and a continuation k that satisfies $\{2\} \mapsto \diamond$, the computation may eventually terminate. So if g indeed returns $\bar{2}$ this formula will be satisfied. Similarly for $((\{3\}, \{2\} \mapsto \square) \mapsto \diamond)$. Therefore, we can see that $g \models_{EPCF} \phi_2$ implies $g^* \models_{\mathcal{F}} \phi_2^*$.

The conjunction in ϕ_2^* is over all nontrivial $P \in \mathfrak{P}$ that appear in $\{2\} \mapsto P$, namely \diamond and \square .

6.2 Logical Equivalence

The logics \mathcal{F} and \mathcal{V} induce a notion of program equivalence defined below:

Definition 6.2.1 (Logical preorder and equivalence). *Consider a fragment \mathcal{L} of one of the logics \mathcal{F} or \mathcal{V} . For any closed values $\vdash v_1 : A$ and $\vdash v_2 : A$:*

$$v_1 \sqsubseteq_{\mathcal{L}} v_2 \iff \forall \phi : A \text{ in } \mathcal{L}. (v_1 \models_{\mathcal{L}} \phi \implies v_2 \models_{\mathcal{L}} \phi).$$

And for any closed computations $\vdash s_1$ and $\vdash s_2$:

$$s_1 \sqsubseteq_{\mathcal{L}} s_2 \iff \forall P \text{ in } \mathcal{L}. (s_1 \models_{\mathcal{L}} P \implies s_2 \models_{\mathcal{L}} P).$$

Two terms (values or computations) are logically equivalent, $t_1 \equiv_{\mathcal{L}} t_2$, if $t_1 \sqsubseteq_{\mathcal{L}} t_2$ and $t_2 \sqsubseteq_{\mathcal{L}} t_1$.

The definition of logical equivalence provides a convenient way of proving that two programs are *not* equivalent: we just need to find a formula that one of them does not satisfy but the other does. For example:

Example 6.2.2 (Probabilistic choice). Consider the following ECPS functions, where $m_{\bar{1},\bar{2}}$, $n_{\bar{1},\bar{2},\bar{1},\bar{3}}$ and $n'_{\bar{1},\bar{2},\bar{1},\bar{3}}$ are defined as in Example 5.2.5:

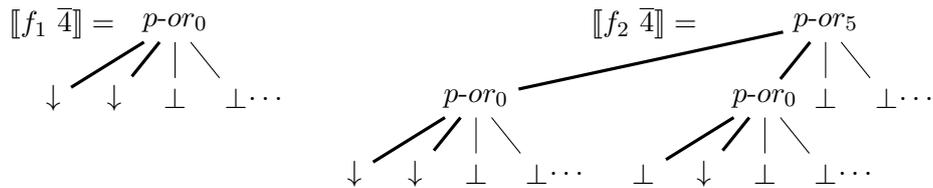
$$f_1 = \lambda x:\text{nat}.m_{\bar{1},\bar{2}}$$

$$f_2 = \lambda x:\text{nat}.\text{case } x \text{ in } \{\text{zero} \Rightarrow n_{\bar{1},\bar{2},\bar{1},\bar{3}}, \text{succ}(y) \Rightarrow n'_{\bar{1},\bar{2},\bar{1},\bar{3}}\}.$$

Consider the \mathcal{F} -formula:

$$\phi = \{4\} \mapsto \mathbf{P}_{>0.9}.$$

Formula ϕ distinguishes between these two functions because: $f_1 \models_{\mathcal{F}} \phi$ but $f_2 \not\models_{\mathcal{F}} \phi$. This can be seen by looking at their computation trees:



Therefore, f_1 and f_2 are not \mathcal{F} -logically equivalent. However, we can see that $f_1 \models_{\mathcal{F}} \{0\} \mapsto \mathbf{P}_{>0.9}$ and $f_2 \models_{\mathcal{F}} \{0\} \mapsto \mathbf{P}_{>0.9}$.

It is the program equivalence induced by \mathcal{F} , rather than \mathcal{V} , that we are mostly interested in. This is because \mathcal{F} enforces a natural separation between ECPS terms and program properties. Using \mathcal{F} , we can specify program properties without knowing the syntax of the programming language.

It can be easily seen that $\sqsubseteq_{\mathcal{F}^+}$ is a preorder and $\equiv_{\mathcal{F}}$ is an equivalence relation. Compatibility, the property that related programs can be substituted for variables in related contexts, is established in the main theorem of this chapter. The proof of this theorem appears at the end of the next section.

Theorem 6.2.3. *Given a decomposable set \mathfrak{P} of Scott-open observations:*

1. *Applicative \mathfrak{P} -similarity, \lesssim , coincides with the logical preorder induced by the logic \mathcal{F}^+ , $\sqsubseteq_{\mathcal{F}^+}$. Hence, the open extension of the \mathcal{F}^+ -logical preorder $\sqsubseteq_{\mathcal{F}^+}^\circ$ is compatible.*
2. *Applicative \mathfrak{P} -bisimilarity, \sim , coincides with the logical equivalence induced by the logic \mathcal{F} , $\equiv_{\mathcal{F}}$. Hence, the open extension of \mathcal{F} -logical equivalence $\equiv_{\mathcal{F}}^\circ$ is compatible.*

This theorem is important because, when combined with the result of the next chapter, it shows that the logic \mathcal{F} characterises contextual equivalence. This answers the main question asked in the introduction.

6.3 Logical Equivalence Coincides with Bisimilarity

The aim of this section is to show that program equivalence induced by the logic \mathcal{F} coincides with applicative \mathfrak{P} -bisimilarity, defined in the previous chapter. First, we show that this is the case for the logic \mathcal{V} . The proof appears in Appendix C.

Proposition 6.3.1. *Given a decomposable set \mathfrak{P} of Scott-open observations:*

1. *Applicative \mathfrak{P} -similarity, \lesssim , coincides with the logical preorder induced by the logic \mathcal{V}^+ , $\sqsubseteq_{\mathcal{V}^+}$. Therefore, the open extension of $\sqsubseteq_{\mathcal{V}^+}$ is compatible.*
2. *Applicative \mathfrak{P} -bisimilarity, \sim , coincides with the logical equivalence induced by the logic \mathcal{V} , $\equiv_{\mathcal{V}}$. Therefore, the open extension of $\equiv_{\mathcal{V}}$ is compatible.*

Next, we show that the logics \mathcal{F} and \mathcal{V} are in fact equivalent. This is done by translating \mathcal{F} -formulas into \mathcal{V} , and vice-versa, and proving that the satisfaction relation is preserved. Define a translation from \mathcal{F} to \mathcal{V} , $(-)^{\flat}$, and a translation from \mathcal{V} to \mathcal{F} , $(-)^{\sharp}$. The definition appears in Figure 6.3. It makes use of the following lemma, which is similar to a lemma for EPCF proved in [SV18]:

Lemma 6.3.2 (Characteristic formula). *For any fragment \mathcal{L} of \mathcal{F} or \mathcal{V} closed under countable conjunction it is true that*

for any closed value v there exists a formula $\chi_v \in \mathcal{L}$ such that:

$$u \models_{\mathcal{L}} \chi_v \iff v \sqsubseteq_{\mathcal{L}} u.$$

Proof. For each closed value u such that $v \not\sqsubseteq_{\mathcal{L}} u$ we can choose a formula ϕ_u such that $v \models_{\mathcal{L}} \phi_u$ but $u \not\models_{\mathcal{L}} \phi_u$. Define χ_v to be:

$$\chi_v = \bigwedge_{\{u \mid v \not\sqsubseteq_{\mathcal{L}} u\}} \phi_u.$$

We can see that $u \not\models_{\mathcal{L}} \chi_v \iff v \not\sqsubseteq_{\mathcal{L}} u$, which is what we need. □

Theorem 6.3.3. *Given a decomposable set \mathfrak{P} of Scott-open observations, the logics \mathcal{F}^+ and \mathcal{V}^+ are equi-expressive.*

$$\begin{aligned}
((\phi_1, \dots, \phi_n) \mapsto P)^b &= \bigwedge \{(w_1, \dots, w_n) \mapsto P \mid w_1 \models_{\mathcal{V}} \phi_1^b, \dots, w_n \models_{\mathcal{V}} \phi_n^b\} \\
((w_1, \dots, w_n) \mapsto P)^\sharp &= (\chi_{w_1}, \dots, \chi_{w_n}) \mapsto P \\
\{n\}^b &= \{n\} & \{n\}^\sharp &= \{n\} \\
P^b &= P & P^\sharp &= P \\
(\bigvee_{i \in I} \phi_i)^b &= \bigvee_{i \in I} \phi_i^b & (\bigvee_{i \in I} \phi_i)^\sharp &= \bigvee_{i \in I} \phi_i^\sharp \\
(\bigwedge_{i \in I} \phi_i)^b &= \bigwedge_{i \in I} \phi_i^b & (\bigwedge_{i \in I} \phi_i)^\sharp &= \bigwedge_{i \in I} \phi_i^\sharp \\
(\neg \phi)^b &= \neg \phi^b & (\neg \phi)^\sharp &= \neg \phi^\sharp
\end{aligned}$$

The formula χ_{w_i} is the characteristic formula of w_i in the logic \mathcal{F} , from Lemma 6.3.2.

Figure 6.3: Translation from \mathcal{F} to \mathcal{V} and vice-versa.

1. For any type A , for any formula ϕ in \mathcal{F}^+ , $\phi : A$ implies that for any value $\vdash v : A$:

$$v \models_{\mathcal{F}^+} \phi \iff v \models_{\mathcal{V}^+} \phi^b.$$

For any $P \in \mathfrak{P}$ and any computation $\vdash t$:

$$t \models_{\mathcal{F}^+} P \iff t \models_{\mathcal{V}^+} P^b.$$

2. For any type A , for any formula ϕ in \mathcal{V}^+ , $\phi : A$ implies that for any value $\vdash v : A$:

$$v \models_{\mathcal{V}^+} \phi \iff v \models_{\mathcal{F}^+} \phi^\sharp.$$

For any $P \in \mathfrak{P}$ and any computation $\vdash t$:

$$t \models_{\mathcal{V}^+} P \iff t \models_{\mathcal{F}^+} P^\sharp.$$

Proof. Statement 1. For computation formulas the result is immediate because they do not change when translated.

For value formulas we prove the following property:

$$\Phi(\phi, A) = (\phi : A \implies (\forall \vdash v : A. v \models_{\mathcal{F}^+} \phi \iff v \models_{\mathcal{V}^+} \phi^b))$$

by induction on the rules in Figure 6.1, which specify when $\phi : A$ is well-formed.

In the case (NAT), $\phi = \{n\}$. The equivalence holds because $\{n\}^b = \{n\}$ and the satisfaction relation does not change with the translation.

The cases for the logical connectives follow from the induction hypothesis.

In the case (VAL), $\phi = (\phi_1, \dots, \phi_n) \mapsto P$. Let $v \models_{\mathcal{F}^+} \phi$ and consider some arbitrary $w_1 \models_{\mathcal{V}^+} \phi_1^b, \dots, w_n \models_{\mathcal{V}^+} \phi_n^b$. By the induction hypothesis we know $w_1 \models_{\mathcal{F}^+} \phi_1, \dots,$

$w_n \models_{\mathcal{F}^+} \phi_n$. So by assumption $\llbracket v(w_1, \dots, w_n) \rrbracket \in P$. Therefore it is true that, in \mathcal{V}^+ , v satisfies $(w_1, \dots, w_n) \mapsto P$ so in general $v \models_{\mathcal{V}^+} \phi^b$.

For the reverse implication let $v \models_{\mathcal{V}^+} \phi^b$ and consider some arbitrary $w_1 \models_{\mathcal{F}^+} \phi_1, \dots, w_n \models_{\mathcal{F}^+} \phi_n$. By the induction hypothesis $w_1 \models_{\mathcal{V}^+} \phi_1^b, \dots, w_n \models_{\mathcal{V}^+} \phi_n^b$, so by assumption $\llbracket v(w_1, \dots, w_n) \rrbracket \in P$. Therefore $v \models_{\mathcal{F}^+} \phi$ as required.

Statement 2. For computation formulas $P^\sharp = P$ so the equivalence holds.

For value formulas proceed by induction on the type A . If $A = \mathbf{nat}$, then the formulas ϕ and ϕ^\sharp represent the same set of natural numbers. Therefore, $v \models_{\mathcal{V}^+} \phi$ is equivalent to $v \models_{\mathcal{F}^+} \phi^\sharp$. For $A = \mathbf{unit}$ the only formulas are *true* and *false* so the equivalence holds trivially. For $A = \neg(B_1, \dots, B_n)$ the induction hypothesis is, for each B_i :

For any formula ϕ' in \mathcal{V}^+ , $\phi' : B_i$ implies that for any value $\vdash v : B_i$:

$$v \models_{\mathcal{V}^+} \phi' \iff v \models_{\mathcal{F}^+} \phi'^\sharp.$$

We do an additional induction on ϕ :

Case $\phi = (w_1, \dots, w_n) \mapsto P : \neg(B_1, \dots, B_n)$. Assume $v \models_{\mathcal{V}^+} \phi$, that is $\llbracket v(w_1, \dots, w_n) \rrbracket \in P$. We need to prove that for any $u_1 : B_1, \dots, u_n : B_n$ such that $u_i \models_{\mathcal{F}^+} \chi_{w_i}$ for all i , we have $\llbracket v(u_1, \dots, u_n) \rrbracket \in P$.

By the definition of χ_{w_i} we know that $w_i \sqsubseteq_{\mathcal{F}^+} u_i$. We can show $w_i \sqsubseteq_{\mathcal{V}^+} u_i$ as follows: consider an arbitrary $\psi : B_i$ such that $w_i \models_{\mathcal{V}^+} \psi$. Then by the induction hypothesis for the type B_i we know $w_i \models_{\mathcal{F}^+} \psi^\sharp$. Hence deduce $u_i \models_{\mathcal{F}^+} \psi^\sharp$ from $w_i \sqsubseteq_{\mathcal{F}^+} u_i$. Again from the induction hypothesis for B_i , we have $u_i \models_{\mathcal{V}^+} \psi$, as required.

Now that we have $w_i \sqsubseteq_{\mathcal{V}^+} u_i$ we can use compatibility of $\sqsubseteq_{\mathcal{V}^+}$, Proposition 6.3.1, and reflexivity to deduce:

$$v(w_1, \dots, w_n) \sqsubseteq_{\mathcal{V}^+} v(u_1, \dots, u_n).$$

So from $\llbracket v(w_1, \dots, w_n) \rrbracket \in P$ we get the desired result $\llbracket v(u_1, \dots, u_n) \rrbracket \in P$.

For the reverse implication assume $v \models_{\mathcal{F}^+} (\chi_{w_1}, \dots, \chi_{w_n}) \mapsto P$. We need to prove that $\llbracket v(w_1, \dots, w_n) \rrbracket \in P$. This follows from the fact that $w_i \models_{\mathcal{F}^+} \chi_{w_i}$ because $\sqsubseteq_{\mathcal{F}^+}$ is reflexive.

Case $\phi = \bigvee_{i \in I} \varphi_i : \neg(B_1, \dots, B_n)$. From the type of ϕ we know that for all i , $\varphi_i : \neg(B_1, \dots, B_n)$. This means that the induction hypothesis for φ_i gives us:

For any value $\vdash v : \neg(B_1, \dots, B_n)$:

$$v \models_{\mathcal{V}^+} \varphi_i \iff v \models_{\mathcal{F}^+} \varphi_i^\sharp.$$

Assume $v \models_{\mathcal{V}^+} \bigvee_{i \in I} \varphi_i$. There exists $j \in I$ such that $v \models_{\mathcal{V}^+} \varphi_j$. By the induction hypothesis for φ_j we have that $v \models_{\mathcal{F}^+} \varphi_j^\sharp$. So $v \models_{\mathcal{F}^+} \phi^\sharp$.

The reverse implication is analogous.

Case $\phi = \bigwedge_{i \in I} \varphi_i : \neg(B_1, \dots, B_n)$. Analogous to the previous case. \square

Theorem 6.3.4. *Given a decomposable set \mathfrak{P} of Scott-open observations, the logics \mathcal{F} and \mathcal{V} are equi-expressive.*

Proof. We need to prove the same statements as in Theorem 6.3.3, where \mathcal{V}^+ is replaced by \mathcal{V} and \mathcal{F}^+ is replaced by \mathcal{F} . The proof is very similar and the differences are pointed out in Appendix C. \square

Using all the results in this section we can finally prove Theorem 6.2.3, which says that \mathcal{F} -logical equivalence coincides with bisimilarity.

Proof of Theorem 6.2.3. From Theorems 6.3.3 and 6.3.4 we can deduce that:

$$\begin{aligned} v \sqsubseteq_{\mathcal{V}^+} u &\iff v \sqsubseteq_{\mathcal{F}^+} u \\ v \equiv_{\mathcal{V}} u &\iff v \equiv_{\mathcal{F}} u \end{aligned}$$

and similarly for computations.

Then by Proposition 6.3.1 we have $(\preceq) = (\sqsubseteq_{\mathcal{F}^+})$ and $(\sim) = (\equiv_{\mathcal{F}})$. From Theorem 5.3.7 we know \preceq° and \sim° are compatible, so this is also the case for $\sqsubseteq_{\mathcal{F}^+}^\circ$ and $\equiv_{\mathcal{F}}^\circ$. \square

Notice that the proofs of Theorems 6.3.3 and 6.3.4 make use of compatibility of $\sqsubseteq_{\mathcal{V}^+}$ and $\equiv_{\mathcal{V}}$, which was established via Howe's method. So a direct proof of Theorem 6.2.3 would require us to prove a compatibility property of $\sqsubseteq_{\mathcal{F}^+}$ first. As we have seen in the previous chapter, proofs of compatibility are laborious. Therefore, the method of going through the logic \mathcal{V} to prove Theorem 6.2.3 is justified.

6.4 Chapter Summary

This chapter defined a logic \mathcal{F} in which each formula expresses a property of an ECPS program (Definition 6.1.1). Formulas which concern computations are elements of the set of observations \mathfrak{P} . Formulas for function values have the form $(\phi_1, \phi_2, \dots, \phi_n) \mapsto P$. They assert that, if the arguments x_1, \dots, x_n of a function satisfy ϕ_1, \dots, ϕ_n respectively, the resulting computation is in $P \in \mathfrak{P}$.

Recall the successor function from equation 3.5.1:

$$f = \lambda(n, k):(\mathbf{nat}, \neg\mathbf{nat}).(k \text{ succ}(n)) : \neg(\mathbf{nat}, \neg\mathbf{nat}).$$

We can see that it satisfies the following formula:

$$\phi = (\{2\}, \{3\} \mapsto \diamond) \mapsto \diamond.$$

This says that, given argument $\bar{2}$ and a continuation k which may terminate for input $\bar{3}$, the body of f may terminate. Consider the formula $\phi' = \{2\} \mapsto \diamond\{3\}$, which describes a program in direct style that may return $\bar{3}$. This is similar to the formulas discussed in Example 2.2.1, but it is not a valid \mathcal{F} -formula. Formula ϕ can be viewed as a translation of ϕ' describing a program in continuation-passing style instead.

The goal of this chapter was to prove that program equivalence induced by the logic \mathcal{F} coincides with applicative bisimilarity (Theorem 6.2.3). Hence, according to Theorem 5.3.7

from the previous chapter, for a decomposable set \mathfrak{B} of Scott-open observations, \mathcal{F} -logical equivalence is compatible.

To prove this, we defined a logic \mathcal{V} which is similar to \mathcal{F} , but in which ECPS values can appear inside formulas. From this point of view \mathcal{V} is not satisfactory as a specification language because we need knowledge of the syntax of ECPS to express properties of programs.

Nevertheless, it is relatively straightforward to prove that \mathcal{V} -logical equivalence coincides with applicative bisimilarity (Proposition 6.3.1), but we do not know any proof of a similar result for \mathcal{F} -logical equivalence. Instead, we proved that \mathcal{F} and \mathcal{V} are expressive, using translations between the two logics (Theorem 6.3.4). Thus, we obtained a proof of Theorem 6.2.3. This theorem will be used in the next chapter to prove the main result of the dissertation: that \mathcal{F} -logical equivalence coincides with contextual equivalence.

Chapter 7

Contextual Equivalence for ECPS

This chapter defines contextual equivalence for ECPS coinductively and proves that it is compatible and an equivalence relation. Then, contextual equivalence is proved to coincide with applicative bisimilarity for ECPS (Theorem 7.2.2). As a result, they both coincide with logical equivalence induced by the logic \mathcal{F} (Corollary 7.2.3). Therefore, the logic \mathcal{F} characterises contextual equivalence. Establishing this was the main goal of the dissertation. Finally, we present an alternative definition of contextual equivalence using program contexts, and prove it equivalent with the coinductive definition (Theorem 7.3.7).

7.1 Contextual Equivalence Coinductively

This section presents a coinductive definition of contextual preorder and equivalence, initially proposed by Lassen [Las98] and Gordon [Gor98]. Contextual preorder is defined as the greatest compatible and adequate relation, for a suitable definition of adequacy.

The advantage of this definition is that we do not need to deal with contexts explicitly. This is important in the case of ECPS where contexts are duplicated because of the distinction between values and computations, as Section 7.3 will show.

Definition 7.1.1 (Adequacy). *Consider a set of observations \mathfrak{P} and a well-typed relation on possibly open terms $\mathcal{R} = (\mathcal{R}_A^v, \mathcal{R}^c)$, where \mathcal{R}^c relates computations. The relation \mathcal{R} is \mathfrak{P} -adequate if:*

$$\forall s, t. \vdash s \mathcal{R}^c t \implies \forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P.$$

The relation \mathcal{R} is \mathfrak{P} -biadequate if:

$$\forall s, t. \vdash s \mathcal{R}^c t \implies \forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \iff \llbracket t \rrbracket \in P.$$

The definition of adequacy is motivated by the fact that, in ECPS, the observable behaviour of a program s , in the sense of Section 2.1, is whether $s \in P$, where $P \in \mathfrak{P}$. So adequacy checks that t simulates the observable behaviour of s .

Definition 7.1.2 (Contextual preorder). *Let $\mathbb{C}\mathbb{A}$ be the set of well-typed relations on possibly open terms that are both compatible and \mathfrak{P} -adequate. Define the contextual preorder \sqsubseteq_{ctx} to be $\bigcup \mathbb{C}\mathbb{A}$.*

The next proposition establishes that contextual preorder is a precongruence. This will help prove that contextual equivalence is a congruence, hence a well-behaved notion of program equivalence.

Proposition 7.1.3. *The contextual preorder \sqsubseteq_{ctx} is a preorder, and is moreover compatible and \mathfrak{P} -adequate. Thus, it is the greatest compatible and \mathfrak{P} -adequate preorder.*

Proof. The proof follows the structure of the proof of Proposition 4 from [LGL17b]. To prove reflexivity, we show that the open identity relation \mathcal{I} is in \mathbb{CA} . To show transitivity it suffices to show that the composition of relations in \mathbb{CA} is itself in \mathbb{CA} . The relation \sqsubseteq_{ctx} is shown compatible using the definition of \mathbb{CA} . The complete proof of this proposition can be found in Appendix D. \square

Definition 7.1.4 (Contextual equivalence). *Let \mathbb{CAS} be the set of well-typed relations on possibly open terms that are both compatible and \mathfrak{P} -biadequate. Define contextual equivalence \equiv_{ctx} to be $\bigcup \mathbb{CAS}$.*

Proposition 7.1.5. *Contextual equivalence is the intersection of the contextual preorder with its converse:*

$$(\equiv_{ctx}) = (\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}.$$

The proof of the above proposition appears in Appendix D. This relationship between contextual equivalence and preorder in ECPS is expected. It also holds in the case of the untyped λ -calculus with generic effects, as shown by [LGL17a]. Finally, we can prove contextual equivalence is a congruence.

Proposition 7.1.6. *Contextual equivalence \equiv_{ctx} is an equivalence relation, and is moreover compatible and \mathfrak{P} -biadequate. Thus, it is the greatest compatible and \mathfrak{P} -biadequate equivalence relation.*

Proof. From Proposition 7.1.5 we know $(\equiv_{ctx}) = (\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}$. We have proved \sqsubseteq_{ctx} is a preorder, so $(\sqsubseteq_{ctx})^{op}$ is also a preorder. The intersection of two preorders is another preorder so \equiv_{ctx} is a preorder. Moreover, it is symmetric because it is the intersection of a relation and its converse. Thus \equiv_{ctx} is an equivalence relation. In the proof of Proposition 7.1.5 we have shown $(\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}$ is compatible and biadequate, so \equiv_{ctx} is as well. \square

7.2 Contextual Equivalence Coincides with Bisimilarity

This section shows that the coinductive notion of contextual equivalence (Definition 7.1.4) coincides with applicative bisimilarity. Using the results from the previous chapter about logical equivalence, we can in fact deduce that all notions of program equivalence for ECPS considered so far are the same. Thus, we have defined a logic \mathcal{F} whose induced program equivalence characterises contextual equivalence. This is the main contribution of this dissertation.

To obtain this result, the set of observations \mathfrak{P} needs to satisfy one more condition apart from decomposability and Scott-openness, named *consistency*.

Definition 7.2.1 (Consistency). *A set of observations \mathfrak{P} is consistent if there exists at least one observation $P_0 \in \mathfrak{P}$ such that:*

1. $P_0 \neq \text{Trees}_\Sigma$ and
2. *there exists at least one computation t_0 such that $\llbracket t_0 \rrbracket \in P_0$.*

This definition says that \mathfrak{P} contains a non-trivial observation P_0 which contains at least one computation tree $\llbracket t_0 \rrbracket$. If this were not the case, then contextual equivalence would equate all terms, including natural numbers.

On the other hand, applicative bisimilarity would equate all computations, but not all natural numbers because $v \sim_{\text{nat}}^p w \iff v = w$. Hence, it would not be the case that applicative bisimilarity coincides with contextual equivalence.

However, program equivalence induced by a set of observations \mathfrak{P} which does not satisfy consistency is not meaningful because not enough programs are distinguished. In particular, all computations would be equivalent to *loop*, the program that runs forever. Indeed, for all example of effects considered so far \mathfrak{P} is consistent. Therefore consistency is a reasonable assumption.

Theorem 7.2.2. *Consider a decomposable set of Scott-open observations \mathfrak{P} which is consistent. Then:*

1. *The open extension of applicative \mathfrak{P} -similarity, \lesssim° , coincides with the contextual preorder, \sqsubseteq_{ctx} .*
2. *The open extension of applicative \mathfrak{P} -bisimilarity, \sim° , coincides with contextual equivalence, \equiv_{ctx} .*

Proof. **We first show** $(\lesssim^\circ) = (\sqsubseteq_{ctx})$. We have shown in Theorem 5.3.7 that \lesssim° is compatible under the current assumptions. Consider $\vdash s \lesssim^\circ t$. Then $\vdash s \lesssim t$ so by the definition of simulation we know that:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P.$$

Therefore, \lesssim° is adequate. Being both compatible and adequate, \lesssim° is included in \sqsubseteq_{ctx} .

Now we need to show $(\sqsubseteq_{ctx}) \subseteq (\lesssim^\circ)$. We first show that \sqsubseteq_{ctx} restricted to closed terms is included in \lesssim , then extend this to open terms. To do this, we show \sqsubseteq_{ctx} restricted to closed terms is a simulation by checking it satisfies the four conditions in the definition of simulation (Definition 5.2.1).

We will concentrate on the case for natural numbers because it is the most interesting. It makes use of the existence of computation t_0 and observation P_0 from the definition of consistency. The complete proof can be found in Appendix D.

2. Assume $\vdash v (\sqsubseteq_{ctx})_{\text{nat}}^p u$. Consider the computation:

$$\text{loop} = (\mu f. \lambda x: \text{nat}. f x)(\bar{0}).$$

This computation leads to an infinite chain of reductions:

$$\begin{aligned} (\mu f. \lambda x: \text{nat}. f x) \bar{0} &\longrightarrow (\lambda x: \text{nat}. (\lambda y: \text{nat}. (\mu f. \lambda x: \text{nat}. f x) y) x) \bar{0} \longrightarrow^2 \\ &(\mu f. \lambda x: \text{nat}. f x) \bar{0} \longrightarrow^* \end{aligned}$$

so $\llbracket \text{loop} \rrbracket = \perp$.

Since v and u are closed values there exist m and n in \mathbb{N} such that $u = \bar{n}$ and $v = \bar{m}$. Consider the following computation:

$$\begin{aligned} & \text{if } x = \bar{n} \text{ then } t_0 \text{ else } \text{loop} = \\ & \quad \text{case } x \text{ in } \{ \text{zero} \Rightarrow \text{loop}, \text{succ}(x_1) \Rightarrow \text{case} \dots \\ & \quad \quad x_n \text{ in } \{ \text{zero} \Rightarrow t_0, \text{succ}(x_{n+1}) \Rightarrow \text{loop} \} \dots \} \end{aligned}$$

which evaluates to t_0 if $x = \bar{n}$ or loops otherwise. We know by consistency that $\llbracket t_0 \rrbracket \in P_0 \neq \text{Trees}_\Sigma$. It must be the case that $\perp \notin P_0$ because otherwise by upwards-closure of P_0 we would obtain $P_0 = \text{Trees}_\Sigma$.

By compatibility and reflexivity of \sqsubseteq_{ctx} we know that:

$$\vdash (\text{if } \bar{n} = \bar{n} \text{ then } t_0 \text{ else } \text{loop}) (\sqsubseteq_{ctx})^c (\text{if } \bar{m} = \bar{n} \text{ then } t_0 \text{ else } \text{loop}). \quad (7.2.1)$$

Suppose by contradiction that $n \neq m$. Then:

$$\begin{aligned} \llbracket \text{if } \bar{n} = \bar{n} \text{ then } t_0 \text{ else } \text{loop} \rrbracket &= \llbracket t_0 \rrbracket \in P_0 \\ \llbracket \text{if } \bar{m} = \bar{n} \text{ then } t_0 \text{ else } \text{loop} \rrbracket &= \llbracket \text{loop} \rrbracket = \perp \notin P_0. \end{aligned}$$

But this contradicts equation 7.2.1 because \sqsubseteq_{ctx} is adequate. Therefore, $n = m$ and $v = u$ as required.

3. By adequacy of \sqsubseteq_{ctx} .
4. By compatibility of \sqsubseteq_{ctx} .

So we have established $(\sqsubseteq_{ctx}) \subseteq (\preceq)$ for closed terms. Now we need to prove $(\sqsubseteq_{ctx}) \subseteq (\preceq^\circ)$ in general. To do this, we consider computations and each type of value separately. Again, we concentrate on the case for natural numbers. The cases for computations and function values are proved using compatibility of \sqsubseteq_{ctx} . We give the proof of the computation case as an example. The proof for function values can be found in Appendix D.

If $\overrightarrow{x_j : A_j} \vdash v (\sqsubseteq_{ctx})_{\text{nat}}^v w$ then for any $\overrightarrow{u_j : A_j}$ there exist natural numbers m and k such that $v[\overrightarrow{u_j/x_j}] = \bar{m}$ and $w[\overrightarrow{u_j/x_j}] = \bar{k}$. We want to show $m = k$ and therefore:

$$\forall (\overrightarrow{u_j : A_j}). v[\overrightarrow{u_j/x_j}] = w[\overrightarrow{u_j/x_j}].$$

This would imply by the definition of simulation and of open extension that $\overrightarrow{x_i : A_i} \vdash v \preceq^\circ w$, as required. To do this consider the computation:

$$\begin{aligned} eq(y, v) &= (\mu f. \lambda(y_0, x_0). (\text{nat}, \text{nat}). \\ & \quad \text{case } y_0 \text{ in } \{ \text{zero} \Rightarrow \text{case } x_0 \text{ in } \{ \text{zero} \Rightarrow t_0, \text{succ}(x_1) \Rightarrow \text{loop} \} \\ & \quad \quad , \text{succ}(y_1) \Rightarrow \text{case } x_0 \text{ in } \{ \text{zero} \Rightarrow \text{loop}, \text{succ}(x_1) \Rightarrow f(y_1, x_1) \} \}) (y, v). \end{aligned}$$

If $y = v$ then $\llbracket eq(y, v) \rrbracket = \llbracket t_0 \rrbracket$, otherwise $\llbracket eq(y, v) \rrbracket = \perp$. Using context weakening and compatibility and reflexivity of \sqsubseteq_{ctx} deduce that:

$$\overrightarrow{x_j : A_j}, y : \mathbf{nat} \vdash eq(y, v) (\sqsubseteq_{ctx})^c eq(y, w)$$

and then again by compatibility

$$\vdash \lambda(\overrightarrow{x_j}, y) : (\overrightarrow{A_j}, \mathbf{nat}).eq(y, v) (\sqsubseteq_{ctx})_{-(\overrightarrow{A_j}, \mathbf{nat})}^v \lambda(\overrightarrow{x_j}, y) : (\overrightarrow{A_j}, \mathbf{nat}).eq(y, w).$$

From $(\sqsubseteq_{ctx}) \subseteq (\lesssim)$ for closed terms we can now establish that:

$$\vdash \lambda(\overrightarrow{x_j}, y) : (\overrightarrow{A_j}, \mathbf{nat}).eq(y, v) \lesssim_{-(\overrightarrow{A_j}, \mathbf{nat})}^v \lambda(\overrightarrow{x_j}, y) : (\overrightarrow{A_j}, \mathbf{nat}).eq(y, w)$$

so from the definition of \lesssim and the fact that reduction preserves similarity (Lemma 5.2.4) we know that:

$$\forall (\vdash \overrightarrow{u_j} : \overrightarrow{A_j}). \forall (\overline{n} : \mathbf{nat}). \vdash eq(\overline{n}, v[\overrightarrow{u_j}/\overrightarrow{x_j}]) \lesssim^c eq(\overline{n}, w[\overrightarrow{u_j}/\overrightarrow{x_j}]). \quad (7.2.2)$$

Assume by contradiction that there exists $\overrightarrow{u_j} : \overrightarrow{A_j}$ such that $v[\overrightarrow{u_j}/\overrightarrow{x_j}] \neq w[\overrightarrow{u_j}/\overrightarrow{x_j}]$. Choose $\overline{n} = v[\overrightarrow{u_j}/\overrightarrow{x_j}]$. Then:

$$\llbracket eq(\overline{n}, v[\overrightarrow{u_j}/\overrightarrow{x_j}]) \rrbracket = \llbracket t_0 \rrbracket \in P_0$$

$$\llbracket eq(\overline{n}, w[\overrightarrow{u_j}/\overrightarrow{x_j}]) \rrbracket = \perp \notin P_0$$

but this contradicts equation 7.2.2 because of the definition of \lesssim . Therefore, it must be the case that $\forall \vdash \overrightarrow{u_j} : \overrightarrow{A_j}. v[\overrightarrow{u_j}/\overrightarrow{x_j}] = w[\overrightarrow{u_j}/\overrightarrow{x_j}]$, which is what he had to prove.

If $\overrightarrow{x_i} : \overrightarrow{A_i} \vdash s (\sqsubseteq_{ctx})^c t$ then by compatibility of \sqsubseteq_{ctx} we know that:

$$\vdash \lambda(\overrightarrow{x_i}) : \overrightarrow{A_i}.s (\sqsubseteq_{ctx})_{-(\overrightarrow{A_i})}^v \lambda(\overrightarrow{x_i}) : \overrightarrow{A_i}.t.$$

So using $(\sqsubseteq_{ctx}) \subseteq (\lesssim)$ for closed terms we have:

$$\vdash \lambda(\overrightarrow{x_i}) : \overrightarrow{A_i}.s \lesssim_{-(\overrightarrow{A_i})}^v \lambda(\overrightarrow{x_i}) : \overrightarrow{A_i}.t.$$

By the definition of \lesssim for function values we can deduce:

$$\forall (\vdash \overrightarrow{v_i} : \overrightarrow{A_i}). \vdash (\lambda(\overrightarrow{x_i}) : \overrightarrow{A_i}.s) (\overrightarrow{v_i}) \lesssim^c (\lambda(\overrightarrow{x_i}) : \overrightarrow{A_i}.t) (\overrightarrow{v_i}).$$

From Lemma 5.2.4 we know that reduction preserves similarity so:

$$\forall (\vdash \overrightarrow{v_i} : \overrightarrow{A_i}). \vdash s[\overrightarrow{v_i}/\overrightarrow{x_i}] \lesssim^c t[\overrightarrow{v_i}/\overrightarrow{x_i}]$$

which by the definition of open extensions means that:

$$\overrightarrow{x_i} : \overrightarrow{A_i} \vdash s \lesssim^{o,c} t.$$

Now show that $(\sim^\circ) = (\equiv_{ctx})$. This is done using $(\preceq^\circ) = (\sqsubseteq_{ctx})$ and the facts that $(\equiv_{ctx}) = (\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}$ (Proposition 7.1.5) and $(\sim) = (\preceq) \cap (\preceq^{op})$ (Proposition 5.2.3). The full proof can be found in Appendix D. \square

Corollary 7.2.3. *For a decomposable set of Scott-open observations \mathfrak{P} that is consistent, the logic \mathcal{F} characterises contextual equivalence for ECPS. That is:*

1. *The open extension of \mathcal{F} -logical preorder coincides with the contextual preorder:*
 $(\sqsubseteq_{\mathcal{F}^+}^\circ) = (\sqsubseteq_{ctx})$.
2. *The open extension of \mathcal{F} -logical equivalence coincides with contextual equivalence:*
 $(\equiv_{\mathcal{F}}^\circ) = (\equiv_{ctx})$.

Hence, applicative bisimilarity, logical equivalence and contextual equivalence all coincide.

Proof. Recall Theorem 6.2.3 which says that, for a decomposable set of Scott-open observations \mathfrak{P} , $(\preceq) = (\sqsubseteq_{\mathcal{F}^+})$ and $(\sim) = (\equiv_{\mathcal{F}})$. From Theorem 7.2.2 we obtain the desired result:

$$\begin{aligned} (\sqsubseteq_{\mathcal{F}^+}^\circ) &= (\sqsubseteq_{ctx}) = (\preceq^\circ) \\ (\equiv_{\mathcal{F}}^\circ) &= (\equiv_{ctx}) = (\sim^\circ). \end{aligned}$$

\square

7.3 Contextual Equivalence via Contexts

For completeness, we give a more familiar definition of contextual equivalence using program contexts, following Crary and Harper [CH07] and Pitts [Pit11]. This way of defining contextual equivalence implements the intuition that equal programs behave the same in all contexts. This definition is proved to be the same as the coinductive one (Definition 7.1.4), a fact that is not immediately apparent and which justifies the use of the coinductive definition.

Informally, a context is an ECPS program with a hole. Because ECPS makes a distinction between values and computations, contexts are divided according to whether they accept and produce computations or values. This leads to a duplication of contexts which makes contextual equivalence tedious to work with. This is why we preferred working with the coinductive definition of contextual equivalence to establish the main results of this chapter.

Definition 7.3.1. *Program contexts for ECPS are defined by the following grammar:*

$$\begin{aligned} C_v^\circ &:= [-]^\circ \mid \text{succ}(C_v^\circ) \mid \lambda \vec{x}_i : \vec{A}_i . C_c^\circ \\ C_c^\circ &:= C_v^\circ(\vec{w}_i) \mid v(w_1, \dots, C_v^\circ, \dots, w_n) \mid (\mu x . C_v^\circ)(\vec{w}_i) \mid (\mu x . v)(w_1, \dots, C_v^\circ, \dots, w_n) \mid \\ &\quad \sigma(C_v^\circ, x.t) \mid \sigma(v, x.C_c^\circ) \mid \text{case } C_v^\circ \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\} \mid \\ &\quad \text{case } v \text{ in } \{\text{zero} \Rightarrow C_c^\circ, \text{succ}(x) \Rightarrow t\} \mid \text{case } v \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow C_c^\circ\} \end{aligned}$$

$$\begin{aligned}
C_c^c &:= [-]^c \mid C_v^c(\vec{w}_i) \mid v(w_1, \dots, C_v^c, \dots, w_n) \mid (\mu x. C_v^c)(\vec{w}_i) \mid (\mu x.v)(w_1, \dots, C_v^c, \dots, w_n) \mid \\
&\quad \sigma(v, x.C_c^c) \mid \text{case } v \text{ in } \{\text{zero} \Rightarrow C_c^c, \text{succ}(x) \Rightarrow t\} \mid \\
&\quad \text{case } v \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow C_c^c\} \\
C_v^c &:= \lambda \vec{x}_i. \vec{A}_i. C_c^c.
\end{aligned}$$

In the context $v(w_1, \dots, C_v^v, \dots, w_n)$, C_v^v can appear in any of the positions 1 to n , for any $n \in \mathbb{N}$. Similarly for the other expressions which take multiple arguments.

The notation C_c^v stands for a context whose hole can only be filled with a value, and the resulting term is a computation; C_v^v , C_c^c and C_v^c should be read analogously.

Filling the whole of a context with a value or computation, as appropriate, is defined by recursion on the structure of contexts. The definition is standard so we present only a few cases:

$$\begin{aligned}
[-]^v[u] &= u \\
(\lambda(\vec{x}_i).(\vec{A}_i).C_c^v)[u] &= \lambda(\vec{x}_i).(\vec{A}_i).C_c^v[u] \\
((\mu x.C_v^v)(\vec{w}_i))[u] &= (\mu x.C_v^v[u])(\vec{w}_i) \\
&\dots
\end{aligned}$$

Note that contexts can bind free variables in the term that fills their hole.

There are no contexts of the form $\sigma(C_v^c, x.t)$, $\text{case } C_v^c \dots$ or $\text{succ}(C_v^c)$. In these cases $C_v^c[t]$ needs to be a value of type nat , for some computation t . But the values of type nat are either variables x or natural numbers \bar{n} , which do not contain any computation. Hence, they cannot be obtained from C_v^c .

We write $C[C']$ for composition of contexts. This means replacing the hole, $[-]$ in C with the context C' . We can now define typing judgements for contexts:

Definition 7.3.2. *The typing relation $C_c^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)$ asserts that, given a value $\Gamma' \vdash u : B$, $C_c^v[u]$ is a well-formed computation in the environment Γ .*

Similarly, define typing relations:

$$\begin{aligned}
C_v^v &: (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash A) \\
C_c^c &: (\Gamma' \vdash) \Rightarrow (\Gamma \vdash) \\
C_v^c &: (\Gamma' \vdash) \Rightarrow (\Gamma \vdash A).
\end{aligned}$$

These relations are the least relations closed under the rules in Figures 7.1 and 7.2.

Lemma 7.3.3. *Context instantiation and context composition yield well-typed contexts:*

1a. *If $C_v^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash A)$ and $\Gamma' \vdash u : B$ then $\Gamma \vdash C_v^v[u] : A$.*

And the analogous statements for contexts C_c^v , C_c^c , C_v^c .

$$\begin{array}{c}
\frac{}{[-]^v : (\Gamma \vdash A) \Rightarrow (\Gamma \vdash A)} \text{(VV-ID)} \quad \frac{C_v^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash \mathbf{nat})}{\text{succ}(C_v^v) : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash \mathbf{nat})} \text{(VV-NAT)} \\
\\
\frac{C_c^v : (\Gamma' \vdash B) \Rightarrow (\Gamma, \overrightarrow{x_i : A_i} \vdash)}{\lambda \overrightarrow{x_i : A_i}. C_c^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash \neg(\overrightarrow{A_i}))} \text{(VV-LBD)} \quad \frac{C_v^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash \neg(\overrightarrow{A_i})) \quad \Gamma \vdash \overrightarrow{w_i} : \overrightarrow{A_i}}{C_v^v(\overrightarrow{w_i}) : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)} \text{(VC-APPL)} \\
\\
\frac{\Gamma \vdash v : \neg(\overrightarrow{A_j}) \quad \Gamma \vdash (w_1, \dots, w_{i-1}) : (A_1, \dots, A_{i-1}) \quad C_v^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash A_i) \quad \Gamma \vdash (w_{i+1}, \dots, w_n) : (A_{i+1}, \dots, A_n)}{v(w_1, \dots, w_{i-1}, C_v^v, w_{i+1}, \dots, w_n) : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)} \text{(VC-APPR}_i\text{)} \\
\\
\frac{C_v^v : (\Gamma' \vdash B) \Rightarrow (\Gamma, x : \neg(\overrightarrow{A_i}) \vdash) \quad \Gamma \vdash \overrightarrow{w_i} : \overrightarrow{A_i}}{(\mu x. C_v^v)(\overrightarrow{w_i}) : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)} \text{(VC-MUL)} \\
\\
\frac{\Gamma, x : \neg(\overrightarrow{A_j}) \vdash v : \neg(\overrightarrow{A_j}) \quad \Gamma \vdash (w_1, \dots, w_{i-1}) : (A_1, \dots, A_{i-1}) \quad C_v^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash A_i) \quad \Gamma \vdash (w_{i+1}, \dots, w_n) : (A_{i+1}, \dots, A_n)}{(\mu x. v)(w_1, \dots, w_{i-1}, C_v^v, w_{i+1}, \dots, w_n) : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)} \text{(VC-MUR}_i\text{)} \\
\\
\frac{C_v^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash \mathbf{nat}) \quad \Gamma, x : \mathbf{nat} \vdash t}{\sigma(C_v^v, x.t) : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)} \text{(VC-OPL)} \\
\\
\frac{\Gamma \vdash v : \mathbf{nat} \quad C_c^v : (\Gamma' \vdash B) \Rightarrow (\Gamma, x : \mathbf{nat} \vdash)}{\sigma(v, x.C_c^v) : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)} \text{(VC-OPR)} \\
\\
\frac{C_v^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash \mathbf{nat}) \quad \Gamma \vdash s \quad \Gamma, x : \mathbf{nat} \vdash t}{\text{case } C_v^v \text{ in } \{\mathbf{zero} \Rightarrow s, \text{succ}(x) \Rightarrow t\} : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)} \text{(VC-CASEV)} \\
\\
\frac{\Gamma \vdash v : \mathbf{nat} \quad C_c^v : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash) \quad \Gamma, x : \mathbf{nat} \vdash t}{\text{case } v \text{ in } \{\mathbf{zero} \Rightarrow C_c^v, \text{succ}(x) \Rightarrow t\} : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)} \text{(VC-CASEL)} \\
\\
\frac{\Gamma \vdash v : \mathbf{nat} \quad \Gamma \vdash s \quad C_c^v : (\Gamma' \vdash B) \Rightarrow (\Gamma, x : \mathbf{nat} \vdash)}{\text{case } v \text{ in } \{\mathbf{zero} \Rightarrow s, \text{succ}(x) \Rightarrow C_c^v\} : (\Gamma' \vdash B) \Rightarrow (\Gamma \vdash)} \text{(VC-CASER)}
\end{array}$$

Figure 7.1: Typing rules for contexts that accept a value.

$$\begin{array}{c}
\frac{}{[-]^c : (\Gamma \vdash) \Rightarrow (\Gamma \vdash)} \text{(CC-ID)} \quad \frac{C_v^c : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash \neg(\vec{A}_i)) \quad \Gamma \vdash \vec{w}_i : \vec{A}_i}{C_v^c(\vec{w}_i) : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash)} \text{(CC-APPL)} \\
\\
\frac{\Gamma \vdash v : \neg(\vec{A}_j) \quad \Gamma \vdash (w_1, \dots, w_{i-1}) : (A_1, \dots, A_{i-1}) \quad C_v^c : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash A_i) \quad \Gamma \vdash (w_{i+1}, \dots, w_n) : (A_{i+1}, \dots, A_n)}{v(w_1, \dots, w_{i-1}, C_v^c, w_{i+1}, \dots, w_n) : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash)} \text{(CC-APPR}_i\text{)} \\
\\
\frac{C_v^c : (\Gamma' \vdash) \Rightarrow (\Gamma, x : \neg(\vec{A}_i) \vdash) \quad \Gamma \vdash \vec{w}_i : \vec{A}_i}{(\mu x. C_v^c)(\vec{w}_i) : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash)} \text{(CC-MUL)} \\
\\
\frac{\Gamma, x : \neg(\vec{A}_j) \vdash v : \neg(\vec{A}_j) \quad \Gamma \vdash (w_1, \dots, w_{i-1}) : (A_1, \dots, A_{i-1}) \quad C_v^c : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash A_i) \quad \Gamma \vdash (w_{i+1}, \dots, w_n) : (A_{i+1}, \dots, A_n)}{(\mu x. v)(w_1, \dots, w_{i-1}, C_v^c, w_{i+1}, \dots, w_n) : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash)} \text{(CC-MUR}_i\text{)} \\
\\
\frac{\Gamma \vdash v : \text{nat} \quad C_c^c : (\Gamma' \vdash) \Rightarrow (\Gamma, x : \text{nat} \vdash)}{\sigma(v, x. C_c^c) : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash)} \text{(CC-OP)} \\
\\
\frac{\Gamma \vdash v : \text{nat} \quad C_c^c : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash) \quad \Gamma, x : \text{nat} \vdash t}{\text{case } v \text{ in } \{\text{zero} \Rightarrow C_c^c, \text{succ}(x) \Rightarrow t\} : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash)} \text{(CC-CASEL)} \\
\\
\frac{\Gamma \vdash v : \text{nat} \quad \Gamma \vdash s \quad C_c^c : (\Gamma' \vdash) \Rightarrow (\Gamma, x : \text{nat} \vdash)}{\text{case } v \text{ in } \{\text{zero} \Rightarrow s, \text{succ}(x) \Rightarrow C_c^c\} : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash)} \text{(CC-CASER)} \\
\\
\frac{C_c^c : (\Gamma' \vdash) \Rightarrow (\Gamma, \overrightarrow{x_i} : \vec{A}_i \vdash)}{\lambda \overrightarrow{x_i} : \vec{A}_i. C_c^c : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash \neg(\vec{A}_i))} \text{(CV-LBD)}
\end{array}$$

Figure 7.2: Typing rules for contexts that accept a computation.

2a. If $C_v^c : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash A)$ and $C_c^v : (\Gamma'' \vdash B) \Rightarrow (\Gamma' \vdash)$ then

$$C_v^c[C_c^v] : (\Gamma'' \vdash B) \Rightarrow (\Gamma \vdash A).$$

And the analogous statements for all valid combinations of contexts C and C' .

Proof. In both cases proceed by induction on the typing derivation of the context C .

1. The two base cases (VV-ID) and (CC-ID) follow by assumption. All the other cases are solved by applying the induction hypothesis then the appropriate typing rule for terms.
2. Again, the base cases follow by assumption. In the other cases apply the induction hypothesis then apply the context typing rule that matches C .

□

Contextual equivalence checks whether two terms have the same behaviour in all program contexts that yield closed computations. Only these contexts are used because we can only observe the behaviour of closed computations. The set of observations \mathfrak{P} encodes the observable behaviour of programs, so contextual equivalence makes use of it.

Definition 7.3.4. *Contextual preorder and equivalence are well-typed relations on possibly open terms defined as follows:*

1. Given values $\Gamma \vdash v, u : A$, they are in the contextual preorder, $\Gamma \vdash v (\leq_{ctx})_A^v u$, if and only if:

$$\forall C_c^v : (\Gamma \vdash A) \Rightarrow (\emptyset \vdash). \forall P \in \mathfrak{P}. \llbracket C_c^v[v] \rrbracket \in P \implies \llbracket C_c^v[u] \rrbracket \in P.$$

2. Given two computations $\Gamma \vdash s, t$, they are in the contextual preorder, $\Gamma \vdash s (\leq_{ctx})^c t$, if and only if:

$$\forall C_c^c : (\Gamma \vdash) \Rightarrow (\emptyset \vdash). \forall P \in \mathfrak{P}. \llbracket C_c^c[s] \rrbracket \in P \implies \llbracket C_c^c[t] \rrbracket \in P.$$

Two values v and u are contextually equivalent, $\Gamma \vdash v (=_{ctx})_A^v u$, if $\Gamma \vdash v (\leq_{ctx})_A^v u$ and $\Gamma \vdash u (\leq_{ctx})_A^v v$. And similarly for computations. Therefore:

$$=_{ctx} = (\leq_{ctx}) \cap (\leq_{ctx})^{op}.$$

Next, we prove that the definition above coincides with the coinductive definition of contextual equivalence from the previous section (Definition 7.1.4). The proofs of the following two lemmas can be found in Appendix D.

Lemma 7.3.5. *Contextual preorder defined with contexts, \leq_{ctx} , is a compatible and adequate preorder. Hence, it is included in contextual preorder defined coinductively, \sqsubseteq_{ctx} .*

Lemma 7.3.6. *Contextual preorder defined coinductively, \sqsubseteq_{ctx} , is closed under program contexts, that is:*

1. If $\Gamma' \vdash v (\sqsubseteq_{ctx})_A^v u$ and $C_v^v : (\Gamma' \vdash A) \Rightarrow (\Gamma \vdash B)$ then $\Gamma \vdash C_v^v[v] (\sqsubseteq_{ctx})_B^v C_v^v[u]$.

And the analogous statement for C_c^v .

2. If $\Gamma' \vdash s (\sqsubseteq_{ctx})^c t$ and $C_c^c : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash)$ then $\Gamma \vdash C_c^c[s] (\sqsubseteq_{ctx})^c C_c^c[t]$.

And the analogous statement for C_v^c .

Theorem 7.3.7. *Contextual preorder defined using program contexts, \leq_{ctx} , coincides with contextual preorder defined coinductively, \sqsubseteq_{ctx} . Moreover, $(=_{ctx}) = (\equiv_{ctx})$.*

Proof. We have already shown in Lemma 7.3.5 that $(\leq_{ctx}) \subseteq (\sqsubseteq_{ctx})$. So it only remains to show the inverse inclusion.

Consider $\Gamma \vdash v (\sqsubseteq_{ctx})_A^v u$. Then by Lemma 7.3.6 we know that:

$$\forall C_c^v : (\Gamma \vdash A) \Rightarrow (\emptyset \vdash). \emptyset \vdash C_c^v[v] (\sqsubseteq_{ctx})^c C_c^v[u].$$

By adequacy of \sqsubseteq_{ctx} we can deduce that:

$$\forall C_c^v : (\Gamma \vdash A) \Rightarrow (\emptyset \vdash). \forall P \in \mathfrak{P}. \llbracket C_c^v[v] \rrbracket \in P \implies \llbracket C_c^v[u] \rrbracket \in P$$

which is equivalent to:

$$\Gamma \vdash v (\leq_{ctx})_A^v u.$$

A similar reasoning can be applied for computations. So $(\sqsubseteq_{ctx}) \subseteq (\leq_{ctx})$ and hence $(\sqsubseteq_{ctx}) = (\leq_{ctx})$.

We know from Proposition 7.1.5 that $(\equiv_{ctx}) = (\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}$. By definition of $=_{ctx}$ we know $(=_{ctx}) = (\leq_{ctx}) \cap (\leq_{ctx})^{op}$. So since $(\sqsubseteq_{ctx})^{op} = (\leq_{ctx})^{op}$ as well, we know that:

$$(=_{ctx}) = (\equiv_{ctx}).$$

□

7.4 Chapter Summary

Section 7.1 defined contextual equivalence for ECPS coinductively, as the greatest compatible relation that is adequate. Adequacy means that related computations are part of exactly the same observations from \mathfrak{P} . Thus, contextual equivalence is compatible by construction.

We then identified a condition on \mathfrak{P} named *consistency* (Definition 7.2.1), which requires contextual equivalence to distinguish between at least two computations. This is a reasonable condition; otherwise, all computations would be identified with the divergent one.

Given a decomposable set \mathfrak{P} of Scott-open observations which is consistent, we proved that contextual equivalence coincides with applicative bisimilarity (Theorem 7.2.2). The proof proceeds by showing the equality of the two relations for closed terms, and then extends this to open terms. Consistency is used in the proof to construct ECPS computations which distinguish between two different natural numbers. The construction of these terms required some ingenuity.

Theorems 7.2.2 and 6.2.3 allowed us to deduce the main result of the dissertation: that under the assumptions of Scott-openness, decomposability and consistency of \mathfrak{P} , \mathcal{F} -logical equivalence characterises contextual equivalence (Corollary 7.2.3).

The coinductive definition of contextual equivalence (Definition 7.1.4) might seem un-intuitive, so we also presented the standard definition (Definition 7.3.4). This says that programs are related if they have the same *observable* behaviour in all program contexts. In ECPS the observable behaviour is encoded by the set \mathfrak{P} . So two programs are contextually equivalent if, when substituted in an arbitrary context yielding closed computations, the resulting computations are part of the same observations $P \in \mathfrak{P}$.

This definition of contextual equivalence is difficult to work with because of the universal quantification over a large number of contexts. Therefore, we preferred working with the coinductive definition in order to establish the most important results in this chapter.

Chapter 8

Conclusion

This chapter summarises the motivation and outcomes of the dissertation. It provides a comparison with previous work and outlines two directions for further research. The chapter concludes with an assessment of my personal development throughout the project.

8.1 Summary

Program equivalence, establishing when two programs are interchangeable, is one of the most important problems in the theory of programming languages. As explained in the introduction, the case of higher-order functions is especially difficult. Many definitions of program equivalence for increasingly complex higher-order languages have been proposed over the years. This variety of approaches has led researchers to investigate the relationships between different definitions of program equivalence.

This dissertation studied program equivalence for a higher-order language with algebraic effects. Algebraic effects are an approach to giving uniform formal semantics to impure operations such as input and output, probabilistic choice, nondeterminism etc. not usually found in purely functional languages. This method was discovered relatively recently [PP01, PP02, PP03, HPP06] so work still needs to be done to understand program equivalence in the presence of generic algebraic effects.

One open question we identified was: can we formulate a logic of program properties that characterises contextual equivalence for a higher-order language with generic algebraic effects? The language we chose to study is named ECPS (Section 3.1). It is a call-by-value extension of the simply-typed λ -calculus with algebraic effects, natural numbers and general recursion. Moreover, ECPS is a continuation-passing language in which programs are not expected to return.

We answered the above question positively in the context of ECPS by developing the logic \mathcal{F} , defined in Chapter 6. Thus, we obtained the first logic whose induced program equivalence coincides with contextual equivalence for algebraic effects (Corollary 7.2.3).

The starting point of our work was a paper by Simpson and Voorneveld [SV18]. They proposed a modal logic that characterises applicative bisimilarity for the EPCF language, but not contextual equivalence. EPCF (Section 2.4) is a variant of ECPS in which programs are written in direct style.

To justify the use of ECPS and to link our work to previous research, we first inves-

tigated the relationship between ECPS and EPCF, in Chapter 4. We proved that EPCF can be embedded into ECPS via a continuation-passing translation which preserves computation trees (Theorem 4.3.1). Then we argued informally that ECPS is more expressive than EPCF.

The remaining chapters were concerned with three forms of program equivalence for ECPS and the relationship between them: applicative bisimilarity, logical equivalence and contextual equivalence. In Chapter 5, we defined applicative bisimilarity for ECPS and proved it compatible (Theorem 5.3.7). Compatibility is a fundamental property that a meaningful program equivalence should satisfy. It says that related programs can be substituted for a variable on opposite sides of a program equation.

Chapter 6 defined the logic \mathcal{F} , which expresses properties of ECPS programs. We proved that \mathcal{F} -logical equivalence coincides with applicative bisimilarity, and is therefore compatible (Theorem 6.2.3). In Chapter 7, we developed contextual equivalence and proved it compatible. Then, we showed that applicative bisimilarity coincides with contextual equivalence (Theorem 7.2.2). This helped us prove the main result of the dissertation: that \mathcal{F} -logical equivalence coincides with contextual equivalence (Corollary 7.2.3).

8.2 Comparison with Previous Work

Other logics of program properties for effects have been proposed. For example, the *evaluation logic* of Pitts [Pit91] concerns general computational effects. It uses two built-in modalities \Box and \Diamond to talk about evaluation of programs. Observations $P \in \mathfrak{P}$ from the logic \mathcal{F} play a similar role to these modalities. The difference is that the contents of the set of observations \mathfrak{P} depend on the effects present in the language, rather than being built into the logic. Moreover, program equivalence induced by evaluation logic is not compared with other operational notions of program equivalence.

Plotkin and Pretnar [PP08] propose a *logic for algebraic effects* which is shown to be sound for establishing different kinds of program equivalence, but not complete in general. According to Pnueli’s classification [Pnu77], this logic for algebraic effects is an exogenous logic because computations are allowed to appear inside formulas. In contrast, the logic \mathcal{F} is endogenous because a formula concerns only one computation.

Another difference is that in Plotkin’s and Pretnar’s logic there is a modality for each algebraic operation. In the logic \mathcal{F} , we instead adopt the view that each observation should express a behavioural property of programs. Thus, although observations depend on the effects in the language, they do not depend on the syntax of these effects. So we obtain a greater separation between the logic and the syntax of the programming language.

Simpson’s and Voorneveld’s work [SV18] is the most closely related to ours. They propose the EPCF logic (Section 2.6), in which modalities specify shapes of computation trees, and show it characterises applicative bisimilarity. Therefore, modalities play a similar role to the observations from \mathcal{F} .

The difference is that EPCF formulas, and applicative bisimilarity, check for the return values of programs. Because ECPS is a continuation-passing language, programs do not return, so neither \mathcal{F} -formulas nor applicative bisimilarity can check return values. As explained in Chapter 4, EPCF is in fact a fragment of ECPS. Therefore, contextual equivalence in ECPS is more restrictive since there are more program contexts. These

differences provide an insight into why in ECPS we were able to obtain the relationship:

$$\text{applicative bisimilarity} = \mathcal{F}\text{-logical equivalence} = \text{contextual equivalence}.$$

Whereas in EPCF the situation is:

$$\text{applicative bisimilarity} = \text{EPCF-logical equivalence} \neq \text{contextual equivalence}.$$

Compared to modal logics used to specify program properties in practice, such as LTL or CTL, the logic \mathcal{F} is not directly suitable for verification because of its infinitary connectives. However, \mathcal{F} achieves its goal of expressing behavioural properties of higher-order programs with algebraic effects. As a result, it can be used as the starting point for designing higher-level logics for algebraic effects that would be suitable for verification.

8.3 Future Work

The examples of effects considered in this work were nondeterminism, probabilistic choice, global store and I/O. A next step would be to consider local store, which is also an algebraic effect [PP02]. Relevant questions here are integrating local store into the framework for generic algebraic effects provided by the logic \mathcal{F} , and investigating whether the current relationship between different forms of program equivalence still holds.

Operational notions of program equivalence for higher-order languages with local store have been studied by Pitts and Stark [PS98], who obtained a characterisation of contextual equivalence using a logical relation. Yoshida, Honda and Berger [YHB08] propose an extension of Hoare logic to reason about local store. They prove that this logic characterises contextual equivalence. However, local store remains a challenging effect. We anticipate that the notion of computation tree, on which observations from \mathcal{F} are based, would need to be changed to account for local store.

Finally, one could study a fourth notion of program equivalence for ECPS called *normal-form bisimilarity* [LL07]. This has not been extended to algebraic effects before. One difference between normal-form bisimilarity and applicative bisimilarity is that the latter only relates closed terms, whereas the former also considers open terms.

The definition of normal-form bisimulation is related to game semantics [AM99], which models open programs as strategies in a two-player game. Similarly, the operational semantics of ECPS can be extended to open terms. A satisfaction relation between open terms and logical formulas can be defined based on this operational semantics and the observations in \mathfrak{B} . The question here is how does normal-form bisimilarity compare to the program equivalence induced by game satisfaction? and furthermore, how do they both compare to applicative bisimilarity? We conjecture that normal-form bisimilarity is closer to game satisfaction than applicative bisimilarity, but leave these questions for future investigation.

8.4 Personal Reflections

From a technical perspective the project presented many challenges. I learnt about algebraic effects and became familiar with proof techniques such as logical relations, coinduction and Howe's method. I completed a significant number of fairly long inductive proofs which required care and organisation.

The proof of correctness of the CPS translation (Theorem 4.3.1) was more difficult than expected. Two approaches I tried initially failed: an inductive proof using domain theoretic techniques and a coinductive proof. I finally settled on the combination between logical relations and coinduction which proved successful.

The *Categories, Proofs and Processes* course [AT18] I took this year helped me understand some of the abstract material needed for the project, such as coinduction, while the *Principles of Programming Languages* lectures I attended provided useful background on continuations. Other useful courses were *Formal Verification* and *Automata, Logic and Games*, where I learnt about modal logics used in verification to express program properties, such as LTL, CTL and the modal μ -calculus.

Overall, my ability to understand new theoretical material and to carry out complex proofs has improved. Considering this and the novel theoretical contribution that the dissertation makes, I can say that the project has been a success.

Bibliography

- [Abr90] S. Abramsky. The lazy λ -calculus. In D.A. Turner, editor, *Research Topics in Functional Programming*, chapter 4, pages 65–117. Addison Wesley, 1990.
- [Ahm06] A.J. Ahmed. Step-Indexed Syntactic Logical Relations for Recursive and Quantified Types. In *ESOP*, 2006.
- [AM99] S. Abramsky and G. McCusker. Game Semantics. In *Computational Logic*. Springer Berlin Heidelberg, 1999.
- [AT18] S. Abramsky and N. Tzevelekos. Introduction to Categories and Categorical Logic. Lecture notes for the MSc in Computer Science, University of Oxford, 2017/18.
- [BH09] N. Benton and C.-K. Hur. Biorthogonality, step-indexing and compiler correctness. In *ICFP*, 2009.
- [CH07] K. Crary and R. Harper. Syntactic Logical Relations for Polymorphic and Recursive Types. *Electr. Notes Theor. Comput. Sci.*, 172:259–299, 2007.
- [CL14] R. Crubillé and U. Dal Lago. On Probabilistic Applicative Bisimulation and Call-by-Value λ -Calculi. In *ESOP*, 2014.
- [Fio17] M. Fiore. Denotational Semantics. Lecture notes for Part II of the Computer Science Tripos, University of Cambridge, 2016/17.
- [Gor98] A.D. Gordon. Operational Equivalences for Untyped and Polymorphic Object Calculi. In A.D. Gordon and A.M. Pitts, editors, *Higher Order Operational Techniques in Semantics*, pages 9–54. Cambridge University Press, 1998.
- [Gri90] T. Griffin. A Formulae-as-Types Notion of Control. In *POPL*, 1990.
- [GTL89] J.-Y. Girard, P. Taylor, and Y. Lafont. *Proofs and Types*. Cambridge University Press, 1989.
- [HM85] M. Hennessy and R. Milner. Algebraic Laws for Nondeterminism and Concurrency. *J. ACM*, 32(1):137–161, 1985.
- [How96] D.J. Howe. Proving Congruence of Bisimulation in Functional Programming Languages. *Inf. Comput.*, 124(2):103–112, 1996.
- [HPP06] M. Hyland, G.D. Plotkin, and J. Power. Combining effects: Sum and tensor. *Theor. Comput. Sci.*, 357(1-3):70–99, 2006.

- [Jac16] B. Jacobs. *Introduction to Coalgebra: Towards Mathematics of States and Observation*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016.
- [JR11] B. Jacobs and J. Rutten. An Introduction to (co)algebra and (co)induction. In D. Sangiorgi and J. Rutten, editors, *Advanced Topics in Bisimulation and Coinduction*, chapter 2, pages 38–99. Cambridge University Press, 2011.
- [JSV10] P. Johann, A. Simpson, and J. Voigtländer. A Generic Operational Metatheory for Algebraic Effects. In *LICS*, 2010.
- [JT11] G. Jaber and N. Tabareau. The Journey of Biorthogonal Logical Relations to the Realm of Assembly Code. In *LOLA*, 2011.
- [Koz83] D. Kozen. Results on the Propositional μ -Calculus. *Theor. Comput. Sci.*, 27:333–354, 1983.
- [Las98] S.B. Lassen. *Relational Reasoning about Functions and Nondeterminism*. PhD thesis, University of Aarhus, BRICS, December 1998.
- [Lev06] P.B. Levy. Infinitary Howe’s Method. *Electr. Notes Theor. Comput. Sci.*, 164(1):85–104, 2006.
- [LGL17a] U. Dal Lago, F. Gavazzo, and P.B. Levy. Effectful applicative bisimilarity: Monads, relators, and Howe’s method. In *LICS*, 2017.
- [LGL17b] U. Dal Lago, F. Gavazzo, and P.B. Levy. Effectful Applicative Bisimilarity: Monads, Relators, and Howe’s Method (Long Version). *CoRR*, abs/1704.04647, 2017.
- [LL07] S.B. Lassen and P.B. Levy. Typed Normal Form Bisimulation. In *CSL*, 2007.
- [LPT03] P.B. Levy, J. Power, and H. Thielecke. Modelling environments in call-by-value programming languages. *Inf. Comput.*, 185(2):182 – 210, 2003.
- [LRS93] Y. Lafont, B. Reus, and T. Streicher. Continuations Semantics or Expressing Implication by Negation. Technical Report 9321, Ludwig-Maximilians-Universität, München, 1993.
- [Mil80] R. Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer, Berlin, Heidelberg, 1980.
- [Mog91] E. Moggi. Notions of Computation and Monads. *Inf. Comput.*, 93(1):55–92, 1991.
- [Mor69] J. Morris. *Lambda Calculus Models of Programming Languages*. PhD thesis, MIT, 1969.
- [Pit91] A.M. Pitts. Evaluation Logic. In *IVth Higher Order Workshop*, 1991.
- [Pit96] A.M. Pitts. Relational Properties of Domains. *Inf. Comput.*, 127(2):66–90, 1996.

- [Pit10] A.M. Pitts. Step-Indexed Biorthogonality: a Tutorial Example. In *Modelling, Controlling and Reasoning About State*, Dagstuhl Seminar Proceedings, 2010.
- [Pit11] A.M. Pitts. Howe’s method for higher-order languages. In D. Sangiorgi and J. Rutten, editors, *Advanced Topics in Bisimulation and Coinduction*, chapter 5, pages 197–232. Cambridge University Press, 2011.
- [Plo77] G.D. Plotkin. LCF Considered as a Programming Language. *Theor. Comput. Sci.*, 5(3):223–255, 1977.
- [Pnu77] A. Pnueli. The Temporal Logic of Programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pages 46–57, 1977.
- [PP01] G.D. Plotkin and J. Power. Adequacy for Algebraic Effects. In *FOSSACS*, 2001.
- [PP02] G.D. Plotkin and J. Power. Notions of Computation Determine Monads. In *FOSSACS*, 2002.
- [PP03] G.D. Plotkin and J. Power. Algebraic Operations and Generic Effects. *Appl. Categ. Structures*, 11(1):69–94, 2003.
- [PP08] G.D. Plotkin and M. Pretnar. A Logic for Algebraic Effects. In *LICS*, 2008.
- [PS98] A.M. Pitts and I.D.B. Stark. Operational reasoning for functions with local state. In A.D. Gordon and A.M. Pitts, editors, *Higher Order Operational Techniques in Semantics*, pages 227–273. Cambridge University Press, 1998.
- [Rey93] J.C. Reynolds. The Discoveries of Continuations. *Lisp and Symbolic Computation*, 6(3-4):233–248, 1993.
- [RT99] J.G. Riecke and H. Thielecke. Typed Exeptions and Continuations Cannot Macro-Express Each Other. In *ICALP*, 1999.
- [SU06] M. Sørensen and P. Urzyczyn. *Lectures on the Curry-Howard Isomorphism*, volume 149 of *Studies in Logic and the Foundations of Mathematics*. Elsevier Science Inc., 2006.
- [SV17] A. Simpson and N. Voorneveld. Behavioural equivalence via modalities for algebraic effects, 2017. Unpublished manuscript.
- [SV18] A. Simpson and N. Voorneveld. Behavioural equivalence via modalities for algebraic effects. In *ESOP*, 2018.
- [Tai67] W.W. Tait. Intensional Interpretations of Functionals of Finite Type I. *J. Symb. Log.*, 32(2):198–212, 1967.
- [YHB08] N. Yoshida, K. Honda, and M. Berger. Logical Reasoning for Higher-Order Functions with Local State. *Logical Methods in Computer Science*, 4(4), 2008.

Appendix A

Proofs about the CPS translation

Lemma 4.3.6. *If $(S, M) \rightsquigarrow^k (S', M')$ and $t \longrightarrow^* t'$ where $k < n$ then:*

$$(S, M) \mathcal{S}_{\tau, \rho}^n t \iff (S', M') \mathcal{S}_{\tau, \rho}^{n-k} t'.$$

Proof. Assume $(S, M) \mathcal{S}_{\tau, \rho}^n t$. To show $(S', M') \mathcal{S}_{\tau, \rho}^{n-k} t'$ it suffices to show that the two conditions in Definition 4.3.4 are satisfied for $((S', M'), t')$.

Assume that for some $p < n - k$, $(S', M') \longrightarrow^p (S'', \sigma(V; W))$. Then we know $(S, M) \longrightarrow^{p+k} (S'', \sigma(V; W))$ and $p + k < n$. So we can deduce $t \longrightarrow^* t' \longrightarrow^* \sigma(v, x.t')$ where $V = \bar{l}$ and $v = \bar{l}$, and $\forall l \in \mathbb{N}$. $(S'', W \bar{l}) \mathcal{S}_{\tau'', \rho}^{n-k-p} t'[\bar{l}/x]$, as required.

Assume that for some $p \leq n - k$, $(S', M') \longrightarrow^p (id, \mathbf{return} V)$. Then $(S, M) \longrightarrow^{p+k} (id, \mathbf{return} V)$. So from the initial assumption we know $t \longrightarrow^* t' \longrightarrow^* \downarrow$, as required.

The reverse implication is proved similarly. \square

Lemma 4.3.8. *For any configuration (S, M) and closed computation t :*

$$(\forall n \in \mathbb{N}. (S, M) \mathcal{S}_{\tau, \rho}^n t) \implies (S, M) \mathcal{S}_{\tau, \rho} t.$$

Proof. Assume $\forall n \in \mathbb{N}. (S, M) \mathcal{S}_{\tau, \rho}^n t$.

To prove $(S, M) \mathcal{S}_{\tau, \rho} t$, it suffices to find a simulation relation respecting the conditions in Definition 4.3.7 which contains the pair $((S, M), t)$. Because \mathcal{S} is the greatest simulation, this relation will be contained in \mathcal{S} .

Let $\bigcap_{m \in \mathbb{N}} \mathcal{S}^m$ be the candidate simulation. Check the conditions in Definition 4.3.7:

1. Assume $(S, M) \rightsquigarrow^* (S', \sigma(V; W))$. Then there exists $p \in \mathbb{N}$ such that $(S, M) \rightsquigarrow^p (S', \sigma(V; W))$. By assumption we know $\forall m > 0. (S, M) \mathcal{S}_{\tau, \rho}^{p+m} t$. So by the first condition in Definition 4.3.4 we know that $t \longrightarrow^* \sigma(v, x.t')$ where $V = \bar{l}$ and $v = \bar{l}$, and that $\forall l \in \mathbb{N}. (S', W \bar{l}) \bigcap_{m > 0} \mathcal{S}_{\tau', \rho}^m t'[\bar{l}/x]$. For $m = 0$, $(S', W \bar{l}) \mathcal{S}_{\tau', \rho}^0 t'[\bar{l}/x]$ can be easily seen to hold. So we have $\forall l \in \mathbb{N}. (S', W \bar{l}) \bigcap_{m \in \mathbb{N}} \mathcal{S}_{\tau', \rho}^m t'[\bar{l}/x]$ as required.
2. Assume $(S, M) \rightsquigarrow^p (id, \mathbf{return} V)$ for some $p \in \mathbb{N}$. Then from $(S, M) \mathcal{S}_{\tau, \rho}^p t$ we have $t \longrightarrow^* \downarrow$ as required.

\square

Lemma 4.3.10 (Fundamental property of the logical relation). *For any value $\overline{x_i : \tau_i} \vdash V : \rho$, any computation $\overline{x_i : \tau_i} \vdash M : \rho$ and any stack $\overline{x_i : \tau_i} \vdash S : \rho \Rightarrow \rho'$ in EPCF:*

1. $\overrightarrow{x_i : \tau_i} \vdash V \mathcal{R}_\rho^v V^*$.
2. $\overrightarrow{x_i : \tau_i} \vdash M \mathcal{R}_\rho^c M^*$.
3. $\overrightarrow{x_i : \tau_i} \vdash S \mathcal{R}_{\rho, \rho'}^s S^*$.

Proof. The proof is by induction on the typing derivations of V , M and S . All cases are presented below, starting with **fix** which is the most interesting.

Case (fix), $M = \mathbf{fix} W$. Assume $\overrightarrow{x_i : \tau_i} \vdash W \mathcal{R}_{(\rho_1 \rightarrow \rho_2) \rightarrow (\rho_1 \rightarrow \rho_2)}^v W^*$. We need to prove:

$$\forall n \in \mathbb{N}. \forall \overrightarrow{(V_i, v_i) : \tau_i}. ((V_i, v_i) \in \mathcal{R}_{\tau_i}^{v, n}) \implies (\mathbf{fix} W[\overrightarrow{V_i/x_i}], (\mathbf{fix} W)^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2}^{c, n}.$$

This is proved by induction on n .

Base case $n = 0$. To prove $(\mathbf{fix} W[\overrightarrow{V_i/x_i}], (\mathbf{fix} W)^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2}^{c, 0}$ we need to check that:

$$\forall (S, k) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2, \rho}^{s, 0}. (S, \mathbf{fix} W[\overrightarrow{V_i/x_i}]) \mathcal{S}_{\rho_1 \rightarrow \rho_2, \rho}^0 ((\mathbf{fix} W)^*[\overrightarrow{v_i/x_i}] k).$$

The first condition in the definition of step-indexed similarity holds because there is no $p < 0$. The second condition holds because $(S, \mathbf{fix} W[\overrightarrow{V_i/x_i}]) \longrightarrow^0 (id, \mathbf{return} V)$ is false.

Induction step $n > 0$. The induction hypothesis is:

$$\forall m < n. \forall \overrightarrow{(V_i, v_i) : \tau_i}. ((V_i, v_i) \in \mathcal{R}_{\tau_i}^{v, m}) \implies (\mathbf{fix} W[\overrightarrow{V_i/x_i}], (\mathbf{fix} W)^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2}^{c, m}.$$

Consider some arbitrary $\overrightarrow{(V_i, v_i) \in \mathcal{R}_{\tau_i}^{v, n}}$. By the definition of a step-indexed relation we know $\forall m < n. \mathcal{R}_{\tau_i}^{v, n} \subseteq \mathcal{R}_{\tau_i}^{v, m}$.

We need to prove that:

$$\forall p \leq n. \forall (S, k) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2, \rho}^{s, p}. (S, \mathbf{fix} W[\overrightarrow{V_i/x_i}]) \mathcal{S}_{\rho_1 \rightarrow \rho_2, \rho}^p ((\mathbf{fix} W)^*[\overrightarrow{v_i/x_i}] k).$$

For $p < n$ this follows immediately from the induction hypothesis if we choose $m = n - 1$.

For the case $p = n$ let:

$$E = \lambda y : \rho_1. \mathbf{let} \mathbf{fix} W[\overrightarrow{V_i/x_i}] \Rightarrow z \mathbf{in} z y$$

$$F = \lambda x : \rho_1. \mathbf{let} W[\overrightarrow{V_i/x_i}] E \Rightarrow w \mathbf{in} w x.$$

From the operational semantics and the CPS translation we can see that:

$$(S, \mathbf{fix} W[\overrightarrow{V_i/x_i}]) \rightsquigarrow (S, \mathbf{return} F)$$

$$(\mathbf{fix} W)^*[\overrightarrow{v_i/x_i}] k \longrightarrow^* k F^*.$$

Therefore, using Lemma 4.3.6 it suffices to prove:

$$(S, \mathbf{return} F) \mathcal{S}_{\rho_1 \rightarrow \rho_2, \rho}^{n-1} k F^*.$$

We have assumed $(S, k) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2, \rho}^{5, n}$ so by definition of \mathcal{R} we know:

$$\forall (V, v) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2}^{v, n-1}. (S, \mathbf{return} V) \mathcal{S}_{\rho_1 \rightarrow \rho_2, \rho}^{n-1} k v.$$

So it is enough to show $(F, F^*) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2}^{v, n-1}$, that is:

$$\begin{aligned} \forall p_1 < n - 1. \forall (V_1, v_1) \in \mathcal{R}_{\rho_1}^{v, p_1}. \forall p_2 \leq p_1. \forall (S_2, k_2) \in \mathcal{R}_{\rho_2, \rho'}^{5, p_2}. \\ (S_2, F V_1) \mathcal{S}_{\rho_2, \rho'}^{p_2} (\lambda k_1: \neg \rho_2^*. F^* (v_1, k_1)) k_2. \end{aligned}$$

From the operational semantics we can deduce:

$$\begin{aligned} (S_1, F V_1) \mapsto^2 (S_2 \circ (\mathbf{let} (-) \Rightarrow w \mathbf{in} wV_1), W[\overrightarrow{V_i/x_i}] E) \\ (\lambda k_1: \neg \rho_2^*. F^* (v_1, k_1)) k_2 \longrightarrow^* W^*[\overrightarrow{v_i/x_i}] (E^*, \lambda w: (\rho_1 \rightarrow \rho_2)^*. ((\lambda l'': \neg \rho_2^*. w (v_1, l'')) k_2)). \end{aligned}$$

For $p_2 < 2$, we immediately have that $(S_2, F V_1) \mathcal{S}_{\rho_2, \rho'}^{p_2} (\lambda k_1: \neg \rho_2^*. F^* (v_1, k_1)) k_2$ because in the definition of \mathcal{S}^{p_2} , the premises of both implications are false.

If $p_2 \geq 2$, then using Lemma 4.3.6 it suffices to show:

$$\begin{aligned} (S_2 \circ (\mathbf{let} (-) \Rightarrow w \mathbf{in} wV_1), W[\overrightarrow{V_i/x_i}] E) \mathcal{S}_{\rho_1 \rightarrow \rho_2, \rho''}^{p_2-2} \\ W^*[\overrightarrow{v_i/x_i}] (E^*, \lambda w: (\rho_1 \rightarrow \rho_2)^*. ((\lambda l'': \neg \rho_2^*. w (v_1, l'')) k_2)). \end{aligned}$$

From the induction hypothesis for W we know that:

$$(W[\overrightarrow{V_i/x_i}], W^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{(\rho_1 \rightarrow \rho_2) \rightarrow (\rho_1 \rightarrow \rho_2)}^{v, p_2-1}. \quad (\text{A.1})$$

We would like to use this to prove the previous statement. Unpacking the definition of $\mathcal{R}_{(\rho_1 \rightarrow \rho_2) \rightarrow (\rho_1 \rightarrow \rho_2)}^{v, p_2-1}$ we see that if we prove:

1. $(E, E^*) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2}^{v, p_2-2}$.
2. $(S_2 \circ (\mathbf{let} (-) \Rightarrow w \mathbf{in} wV_1), \lambda w: (\rho_1 \rightarrow \rho_2)^*. ((\lambda l'': \neg \rho_2^*. w (v_1, l'')) k_2)) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2, \rho''}^{5, p_2-2}$.

we can then use equation A.1 to obtain the desired result.

We first prove the second statement, which is equivalent to:

$$\begin{aligned} \forall p_3 \leq p_2 - 2. \forall (V_3, v_3) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2}^{v, p_3}. \\ (S_2 \circ (\mathbf{let} (-) \Rightarrow w \mathbf{in} wV_1), V_3) \mathcal{S}_{\rho_1 \rightarrow \rho_2, \rho''}^{p_3} (\lambda w: (\rho_1 \rightarrow \rho_2)^*. ((\lambda l'': \neg \rho_2^*. w (v_1, l'')) k_2)) v_3. \end{aligned}$$

From the operational semantics:

$$\begin{aligned} (S_2 \circ (\mathbf{let} (-) \Rightarrow w \mathbf{in} wV_1), V_3) \mapsto (S_2, V_3 V_1) \\ (\lambda w: (\rho_1 \rightarrow \rho_2)^*. ((\lambda l'': \neg \rho_2^*. w (v_1, l'')) k_2)) v_3 \longrightarrow^* v_3 (v_1, k_2). \end{aligned}$$

Therefore, it is enough to show:

$$(S_2, V_3 V_1) \mathcal{S}_{\rho_2, \rho''}^{p_3-1} v_3 (v_1, k_2). \quad (\text{A.2})$$

We have assumed $(V_3, v_3) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2}^{v, p_3}$ and $(V_1, v_1) \in \mathcal{R}_{\rho_1}^{v, p_1} \subseteq \mathcal{R}_{\rho_1}^{v, p_3-1}$, so by definition of \mathcal{R}^v we know:

$$(V_3 \ V_1, \lambda k: \neg \rho_2^*. v_3 \ (v_1, k)) \in \mathcal{R}_{\rho_2}^{c, p_3-1}.$$

We have also assumed $(S_2, k_2) \in \mathcal{R}_{\rho_2, \rho'}^{s, p_2} \subseteq \mathcal{R}_{\rho_2, \rho'}^{s, p_3-1}$ so by definition of \mathcal{R}^c we know:

$$(S_2, V_3 \ V_1) \mathcal{S}_{\rho_2, \rho''}^{p_3-1} (\lambda k: \neg \rho_2^*. v_3 \ (v_1, k)) \ k_2.$$

So from Lemma 4.3.6 we obtain the desired result, equation A.2.

Now prove the first statement, 1: this is equivalent to proving

$$\forall p_4 < p_2 - 2. \ \forall (V_4, v_4) \in \mathcal{R}_{\rho_1}^{v, p_4}. \ \forall p_5 \leq p_4. \ \forall (S_5, k_5) \in \mathcal{R}_{\rho_2, \sigma}^{s, p_5}.$$

$$(S_5, E \ V_4) \mathcal{S}_{\rho_2, \sigma}^{p_5} (\lambda k: \neg \rho_2^*. E^* \ (v_4, k)) \ k_5.$$

Using the operational semantics we have:

$$\begin{aligned} (S_5, E \ V_4) &\mapsto^2 (S_5 \circ (\mathbf{let} \ (-) \Rightarrow z \ \mathbf{in} \ zV_4), \mathbf{fix} \ W \overrightarrow{[V_i/x_i]}) \\ &(\lambda k: \neg \rho_2^*. E^* \ (v_4, k)) \ k_5 \longrightarrow^* (\mathbf{fix} \ W)^* \overrightarrow{[v_i/x_i]} \ \lambda z: (\rho_1 \rightarrow \rho_2)^*. ((\lambda p': \neg \rho_2^*. z \ (v_4, p')) \ k_5). \end{aligned}$$

So it is enough to show:

$$\begin{aligned} (S_5 \circ (\mathbf{let} \ (-) \Rightarrow z \ \mathbf{in} \ zV_4), \mathbf{fix} \ W \overrightarrow{[V_i/x_i]}) &\mathcal{S}_{\rho_1 \rightarrow \rho_2, \sigma}^{p_5-2} \\ &(\mathbf{fix} \ W)^* \overrightarrow{[v_i/x_i]} \ \lambda z: (\rho_1 \rightarrow \rho_2)^*. ((\lambda p': \neg \rho_2^*. z \ (v_4, p')) \ k_5). \quad (\text{A.3}) \end{aligned}$$

From the induction hypothesis on natural numbers we know that:

$$(\mathbf{fix} \ W \overrightarrow{[V_i/x_i]}, (\mathbf{fix} \ W)^* \overrightarrow{[v_i/x_i]}) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2}^{c, p_5-2}.$$

Similarly to the proof for S_2 and k_2 , statement 2, we can show that:

$$(S_5 \circ (\mathbf{let} \ (-) \Rightarrow z \ \mathbf{in} \ zV_4), \lambda z: (\rho_1 \rightarrow \rho_2)^*. ((\lambda p': \neg \rho_2^*. z \ (v_4, p')) \ k_5)) \in \mathcal{R}_{\rho_1 \rightarrow \rho_2, \sigma}^{s, p_5-2}.$$

Combining the last two observations and using the definition of \mathcal{R}^c we obtain the desired result, equation A.3.

Case (op), $M = \sigma(V; W)$ for some $\sigma \in \Sigma$. The induction hypothesis is:

$$\overrightarrow{x_i : \sigma_i} \vdash V \ \mathcal{R}_{\mathbb{N}}^v \ V^* \quad \text{and} \quad \overrightarrow{x_i : \sigma_i} \vdash W \ \mathcal{R}_{\mathbb{N} \rightarrow \tau}^v \ W^*.$$

We need to prove;

$$\begin{aligned} \forall n \in \mathbb{N}. \ \forall \overrightarrow{(V_i, v_i) : \sigma_i}. \ (\forall i. \ (V_i, v_i) \in \mathcal{R}_{\sigma_i}^{v, n}) &\implies \\ &(\sigma(V; W) \overrightarrow{[V_i/x_i]}, \lambda k: \neg \tau^*. \sigma(V^*; x.W^*(x, k)) \overrightarrow{[v_i/x_i]}) \in \mathcal{R}_{\tau}^{c, n}. \end{aligned}$$

Consider $p \leq n$ and $(S, k) \in \mathcal{R}_{\tau, \rho}^{s, p}$. It suffices to prove:

$$(S, \sigma(V; W) \overrightarrow{[V_i/x_i]}) \mathcal{S}_{\tau, \rho}^p (\lambda k: \neg \tau^*. \sigma(V^*; x.W^*(x, k)) \overrightarrow{[v_i/x_i]}) \ k.$$

We prove this by checking the conditions in the definition of \mathcal{S}^p . If $p = 0$ the first condition is satisfied immediately because of the strict inequality. If $p > 0$ we can choose $p' = 0$ such that $(S, \sigma(V; W)[\overrightarrow{V_i/x_i}]) \rightsquigarrow^{p'} (S, \sigma(V; W)[\overrightarrow{V_i/x_i}])$. We see that $(\lambda k: \neg \tau^*. \sigma(V^*; x.W^*(x, k))[\overrightarrow{v_i/x_i}]) k \longrightarrow^* \sigma(V^*; x.W^*(x, k))[\overrightarrow{v_i/x_i}]$ as required. By the induction hypothesis for V we know that $V[\overrightarrow{V_i/x_i}] = \overline{m}$ and $V^*[\overrightarrow{v_i/x_i}] = \overline{m}$ for some $m \in \mathbb{N}$. It remains to show that:

$$\forall l \in \mathbb{N}. (S, W[\overrightarrow{V_i/x_i} \bar{l}]) \mathcal{S}_{\tau, \rho}^p W^*[\overrightarrow{v_i/x_i} (\bar{l}, k)].$$

This follows from the induction hypothesis for W , using Lemma 4.3.6.

The second condition in the definition of \mathcal{S}^p is satisfied because the premise of the implication is false. $(S, \sigma(V; W)[\overrightarrow{V_i/x_i}])$ cannot reduce anymore according to the \rightsquigarrow relation, and is not of the form $(S, \mathbf{return} V')$.

Case (ret), $M = \mathbf{return} W$. From the induction hypothesis we know:

$$\forall n \in \mathbb{N}. \forall (V_i, v_i) : \sigma_i. (\forall i. (V_i, v_i) \in \mathcal{R}_{\sigma_i}^{v, n}) \implies (W[\overrightarrow{V_i/x_i}], W^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\tau}^{v, n}.$$

We need to prove that:

$$\begin{aligned} \forall n \in \mathbb{N}. \forall (V_i, v_i) : \sigma_i. (\forall i. (V_i, v_i) \in \mathcal{R}_{\sigma_i}^{v, n}) \implies \\ (\mathbf{return} W[\overrightarrow{V_i/x_i}], (\mathbf{return} W)^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\tau}^{c, n}. \end{aligned}$$

So it suffices to show:

$$\begin{aligned} \forall n \in \mathbb{N}. (W[\overrightarrow{V_i/x_i}], W^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\tau}^{v, n} \implies \\ (\mathbf{return} W[\overrightarrow{V_i/x_i}], (\mathbf{return} W)^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\tau}^{c, n}. \end{aligned}$$

Consider $p \leq n$ and $(S, k) \in \mathcal{R}_{\tau, \rho}^{s, p}$. We need to show that:

$$(S, \mathbf{return} W[\overrightarrow{V_i/x_i}]) \mathcal{S}_{\tau, \rho}^p (\lambda k: \neg \tau^*. k W^*[\overrightarrow{v_i/x_i}]) k.$$

Using Lemma 4.3.6 it is enough to show $(S, \mathbf{return} W[\overrightarrow{V_i/x_i}]) \mathcal{S}_{\tau, \rho}^p k W^*[\overrightarrow{v_i/x_i}]$. This follows from $(S, k) \in \mathcal{R}_{\tau, \rho}^{s, p}$ and $(W[\overrightarrow{V_i/x_i}], W^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\tau}^{v, n} \subseteq \mathcal{R}_{\tau}^{v, p}$ by definition of \mathcal{R}^s .

Case (let), $M = \mathbf{let} N_1 \Rightarrow y \mathbf{in} N_2$. The induction hypothesis is:

$$\overline{x_i} : \sigma_i \vdash N_1 \mathcal{R}_{\tau_1}^c N_1^* \quad \text{and} \quad \overline{x_i} : \sigma_i, y : \tau_1 \vdash N_2 \mathcal{R}_{\tau_2}^c N_2^*.$$

We need to prove:

$$\begin{aligned} \forall n \in \mathbb{N}. \forall (V_i, v_i) : \sigma_i. (\forall i. (V_i, v_i) \in \mathcal{R}_{\sigma_i}^{v, n}) \implies \\ (\mathbf{let} N_1[\overrightarrow{V_i/x_i}] \Rightarrow y \mathbf{in} N_2[\overrightarrow{V_i/x_i}], (\mathbf{let} N_1 \Rightarrow y \mathbf{in} N_2)^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\tau_2}^{c, n}. \end{aligned}$$

Consider $p \leq n$ and $(S, k) \in \mathcal{R}_{\tau_2, \rho}^{s, p}$. For $p \geq 1$ it is enough to show:

$$(S \circ \mathbf{let} (-) \Rightarrow y \mathbf{in} N_2[\overrightarrow{V_i/x_i}], N_1) \mathcal{S}_{\tau_2, \rho}^{p-1} N_1^*[\overrightarrow{v_i/x_i}] (\lambda y: \tau_1^*. N_2^*[\overrightarrow{v_i/x_i}] k).$$

By induction hypothesis for N_1 it is enough to show:

$$(S \circ \mathbf{let} (-) \Rightarrow y \mathbf{in} N_2[\overrightarrow{V_i/x_i}], \lambda y: \tau_1^*. N_2^*[\overrightarrow{v_i/x_i}] k) \in \mathcal{R}_{\tau_1, \rho}^{s, p-1}.$$

Consider $p_1 \leq p$ and $(W_1, w_1) \in \mathcal{R}_{\tau_1}^{v, p_1}$. For $p_1 \geq 1$ it is enough to show:

$$(S, N_2[\overrightarrow{V_i/x_i}, W_1/y]) \mathcal{S}_{\tau_2, \rho}^{p_1-1} N_2^*[\overrightarrow{v_i/x_i}, w_1/y] k.$$

This follows from the induction hypothesis for N_2 .

Case (app), $M = W U$. If W has type $\tau_1 \rightarrow \tau_2$ and U has type τ_1 , for each $n \in \mathbb{N}$, we need to prove:

$$\forall p \leq n. \forall (S, k) \in \mathcal{R}_{\tau_2, \rho}^{s, p}. (S, (W U)[\overrightarrow{V_i/x_i}]) \mathcal{S}_{\tau_2, \rho}^p (\lambda k: \neg \tau_2^*. W^*[\overrightarrow{v_i/x_i}] (U^*[\overrightarrow{v_i/x_i}], k)) k.$$

This follows from $(W[\overrightarrow{V_i/x_i}], W^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\tau_1 \rightarrow \tau_2, \rho}^{v, p+1}$ and $(U[\overrightarrow{V_i/x_i}], U^*[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\tau_1}^{v, p}$, which we know from the induction hypothesis.

Case (case), $M = \mathbf{case} V \mathbf{in} \{Z \Rightarrow N_1, S(y) \Rightarrow N_2\}$. We need to prove for each $n \in \mathbb{N}$:

$$\begin{aligned} \forall p \leq n. \forall (S, k) \in \mathcal{R}_{\tau, \rho}^{s, p}. (S, (\mathbf{case} V \mathbf{in} \{Z \Rightarrow N_1, S(y) \Rightarrow N_2\})[\overrightarrow{V_i/x_i}]) \mathcal{S}_{\tau, \rho}^p \\ (\lambda k: \neg \tau^*. \mathbf{case} V^* \mathbf{in} \{\mathbf{zero} \Rightarrow N_1^* k, \mathbf{succ}(y) \Rightarrow N_2^* k\})[\overrightarrow{v_i/x_i}] k. \end{aligned}$$

We proceed by a case split on whether $V[\overrightarrow{V_i/x_i}] = \bar{0}$ or not. Then the result follows from the induction hypothesis for N_1 or N_2 , respectively.

Case (sid), $S = id$. In this case we need to prove:

$$\forall n \in \mathbb{N}. \forall p \leq n. \forall (V, v) \in \mathcal{R}_{\rho}^{v, p}. (id, \mathbf{return} V) \mathcal{S}_{\rho, \rho}^p (\lambda x: \neg \rho^*. \downarrow) v.$$

We show that the conditions in the definition of \mathcal{S}^p hold. The first condition holds because $(id, \mathbf{return} V) \not\rightarrow$ and $(id, \mathbf{return} V)$ is not of the form $(S, \sigma(V; W))$. For the second condition we can choose $p' = 0 \leq p$ such that $(id, \mathbf{return} V) \rightarrow^{p'} (id, \mathbf{return} V)$ and we have $(\lambda x: \neg \rho^*. \downarrow) v \rightarrow \downarrow$, as required.

Case (slet), $S = S' \circ \mathbf{let} (-) \Rightarrow y \mathbf{in} M$. The induction hypothesis is:

$$\overrightarrow{x_i: \sigma_i} \vdash S' \mathcal{R}_{\tau_2, \rho}^s S'^* \quad \text{and} \quad \overrightarrow{x_i: \sigma_i}, y: \tau_1 \vdash M \mathcal{R}_{\tau_2}^c M^*.$$

We need to prove:

$$\forall n \in \mathbb{N}. \forall p \leq n. \forall (V, v) \in \mathcal{R}_{\tau_1}^{v, p}.$$

$$(S'[\overrightarrow{V_i/x_i}] \circ \mathbf{let} (-) \Rightarrow y \mathbf{in} M[\overrightarrow{V_i/x_i}], \mathbf{return} V) \mathcal{S}_{\tau_1, \rho}^p (\lambda y: \tau_1^*. (M^* S'^*)[\overrightarrow{v_i/x_i}]) v.$$

If $p > 0$, from Lemma 4.3.6, it is enough to show that:

$$(S'[\overrightarrow{V_i/x_i}], M[\overrightarrow{V_i/x_i}, V/y]) \mathcal{S}_{\tau_2, \rho}^{p-1} M^*[\overrightarrow{v_i/x_i}, v/y] S'^*[\overrightarrow{v_i/x_i}].$$

The result follows from the induction hypothesis for M and S' .

Case (lbd), $V = \lambda y:\tau_1.M$. The result follows from the induction hypothesis for M , unpacking the definition of $\mathcal{R}_{\tau_1 \rightarrow \tau_2}^v$.

Case (var), $V = y$. We know by assumption that $\overrightarrow{x_i} : \sigma_i \vdash y : \tau$. We need to prove:

$$\forall n \in \mathbb{N}. (\forall (V_i, v_i) \in \mathcal{R}_{\sigma_i}^{v,n}) \implies (y[\overrightarrow{V_i/x_i}], y[\overrightarrow{v_i/x_i}]) \in \mathcal{R}_{\tau}^{v,n}.$$

This is true because $y[\overrightarrow{V_i/x_i}] = V_j$, $y[\overrightarrow{v_i/x_i}] = v_j$ and $\tau = \sigma_j$ for some j .

Case (zero), $V = Z$. By definition of $\mathcal{R}_{\mathbb{N}}^{v,m}$ we know that $(Z, Z) \in \mathcal{R}_{\mathbb{N}}^{v,m}$ for any $m \in \mathbb{N}$.

Case (unit), $V = \star$. Analogous to the previous case.

Case (succ), $V = S(W)$. We need to prove that:

$$\forall n \in \mathbb{N}. (\forall (V_i, v_i) \in \mathcal{R}_{\sigma_i}^{v,n}) \implies (S(W[\overrightarrow{V_i/x_i}]), \text{succ}(W^*[\overrightarrow{v_i/x_i}])) \in \mathcal{R}_{\mathbb{N}}^{v,n}.$$

From the induction hypothesis for W we know that $W[\overrightarrow{V_i/x_i}] = \overline{m}$ and $W^*[\overrightarrow{v_i/x_i}] = \overline{m}$ for some $m \in \mathbb{N}$. So $S(W[\overrightarrow{V_i/x_i}]) = \overline{m+1}$ and $\text{succ}(W^*[\overrightarrow{v_i/x_i}]) = \overline{m+1}$, which gives us the required result. \square

Lemma 4.3.12. *The function $\llbracket - \rrbracket : (\vdash) \longrightarrow \text{Trees}_{\Sigma}$ is a coalgebra morphism in the category of coalgebras for the functor T .*

Proof. We will use the definitions of the functor T and coalgebra morphisms c and a from Section 4.3.

First, we describe the ω -CPO structure on Trees_{Σ} . The partial order \leq satisfies the following:

$$\forall tr \in \text{Trees}_{\Sigma}. \perp \leq tr \text{ and } (tr \leq \perp \iff tr = \perp) \tag{A.4}$$

$$\forall tr \in \text{Trees}_{\Sigma}. (\downarrow \leq tr \iff tr = \downarrow) \text{ and } (tr \leq \downarrow \iff tr = \perp \text{ or } tr = \downarrow) \tag{A.5}$$

$$\forall \sigma, \sigma' \in \Sigma, n, n' \in \mathbb{N}, \overrightarrow{tr'_k}, \overrightarrow{tr''_k} \in \text{Trees}_{\Sigma}.$$

$$\sigma_n(\overrightarrow{tr'_k}) \leq \sigma'_{n'}(\overrightarrow{tr''_k}) \text{ iff } \sigma = \sigma' \text{ and } n = n' \text{ and } \forall k \in \mathbb{N}. tr_k \leq tr'_k. \tag{A.6}$$

From these properties we see that a chain $tr_0 \leq tr_1 \leq tr_2 \leq \dots$ can have one of the following three forms:

1. $\forall n \in \mathbb{N}. tr_n = \perp$. In this case define the least upper bound of the chain to be $\bigsqcup_{n \in \mathbb{N}} tr_n = \perp$.
2. $\exists k \in \mathbb{N}$ such that $\forall n \geq k. tr_n = \downarrow$ and $\forall n < k. tr_n = \perp$. So we define the least upper bound to be $\bigsqcup_{n \in \mathbb{N}} tr_n = \downarrow$.
3. $\exists k \in \mathbb{N}$ such that $\forall n \geq k. tr_n = \sigma_i(\overrightarrow{tr'_j})$ and $\forall n < k. tr_n = \perp$. By property A.6 of \leq , σ and i need to stay the same for all n . So we can define the least upper bound as: $\bigsqcup_{n \in \mathbb{N}} tr_n = \sigma_i(\bigsqcup_{n \geq k} tr_1^n, \bigsqcup_{n \geq k} tr_2^n, \dots)$.

Now we can define a partial order, \leq_T , on $T(\text{Trees}_\Sigma) = \{\downarrow, \perp\} + \Sigma \times \mathbb{N} \times (\text{Trees}_\Sigma)^\mathbb{N}$ which makes it into an ω -CPO:

$$\forall x \in T(\text{Trees}_\Sigma). \perp \leq_T x \text{ and } (x \leq_T \perp \iff x = \perp)$$

$$\forall x \in T(\text{Trees}_\Sigma). (\downarrow \leq_T x \iff x = \downarrow) \text{ and } (x \leq_T \downarrow \iff x = \perp \text{ or } x = \downarrow)$$

$$\forall \sigma, \sigma' \in \Sigma, n, n' \in \mathbb{N}, \vec{tr}_k, \vec{tr}'_k \in \text{Trees}_\Sigma.$$

$$(\sigma, n, (\vec{tr}_k)) \leq_T (\sigma', n', (\vec{tr}'_k)) \text{ iff } \sigma = \sigma' \text{ and } n = n' \text{ and } \forall k \in \mathbb{N}. tr_k \leq tr'_k.$$

Given this definition, a chain $x_0 \leq_T x_1 \leq_T x_2 \leq_T \dots$ can only have one of the three forms below:

1. $\forall n \in \mathbb{N}. x_n = \perp$. Define the least upped bound as: $\bigsqcup_{n \in \mathbb{N}} x_n = \perp$.
2. $\exists k \in \mathbb{N}$ such that $\forall n \geq k. x_n = \downarrow$ and $\forall n < k. x_n = \perp$. Define the least upped bound as: $\bigsqcup_{n \in \mathbb{N}} x_n = \downarrow$.
3. $\exists k \in \mathbb{N}$ such that $\forall n \geq k. x_n = (\sigma, i, (\vec{tr}_j^n))$ and $\forall n < k. x_n = \perp$. Again it must be the case that σ and i are the same for all $n \geq k$ so we define the least upper bound as: $\bigsqcup_{n \in \mathbb{N}} x_n = (\sigma, i, (\bigsqcup_{n \geq k} tr_1^n, \bigsqcup_{n \geq k} tr_2^n, \dots))$.

To prove $\llbracket - \rrbracket$ is a coalgebra morphism, we must show that the following diagram commutes:

$$\begin{array}{ccc} (\vdash) & \xrightarrow{\llbracket - \rrbracket} & \text{Trees}_\Sigma \\ \downarrow a & & \downarrow c \\ \{\downarrow, \perp\} + \Sigma \times \mathbb{N} \times (\vdash)^\mathbb{N} & \xrightarrow{T(\llbracket - \rrbracket)} & \{\downarrow, \perp\} + \Sigma \times \mathbb{N} \times (\text{Trees}_\Sigma)^\mathbb{N} \end{array}$$

First show that for any chain $tr_0 \leq tr_1 \leq tr_2 \leq \dots$ in Trees_Σ :

$$c\left(\bigsqcup_{n \in \mathbb{N}} tr_n\right) = \bigsqcup_{n \in \mathbb{N}} c(tr_n).$$

We do a case split on the structure of the chain following the cases that we identified previously:

1. $\forall n \in \mathbb{N}. tr_n = \perp$. Then $c(\bigsqcup_{n \in \mathbb{N}} tr_n) = \perp$ and $c(tr_n) = \perp$ for all n so we are done.
2. $\exists k \in \mathbb{N}$ such that $\forall n \geq k. tr_n = \downarrow$ and $\forall n < k. tr_n = \perp$. Then $c(\bigsqcup_{n \in \mathbb{N}} tr_n) = c(\downarrow) = \downarrow$ and $\bigsqcup_{n \in \mathbb{N}} c(tr_n) = \bigsqcup_{n \geq k} \downarrow = \downarrow$.
3. $\exists k \in \mathbb{N}$ such that $\forall n \geq k. tr_n = \sigma_i(\vec{tr}_j^n)$ and $\forall n < k. tr_n = \perp$. In this case:

$$c\left(\bigsqcup_{n \in \mathbb{N}} tr_n\right) = c\left(\bigsqcup_{n \geq k} tr_n\right) = c\left(\sigma_i\left(\bigsqcup_{n \geq k} tr_1^n, \bigsqcup_{n \geq k} tr_2^n, \dots\right)\right) = \left(\sigma, i, \left(\bigsqcup_{n \geq k} tr_1^n, \bigsqcup_{n \geq k} tr_2^n, \dots\right)\right)$$

$$\bigsqcup_{n \in \mathbb{N}} c(tr_n) = \bigsqcup_{n \geq k} c(tr_n) = \bigsqcup_{n \geq k} \left(\sigma, i, (tr_1^n, tr_2^n, \dots)\right) = \left(\sigma, i, \left(\bigsqcup_{n \geq k} tr_1^n, \bigsqcup_{n \geq k} tr_2^n, \dots\right)\right).$$

so the two sides are equal as required.

Now we can deduce that:

$$c(\llbracket t \rrbracket) = c\left(\bigsqcup_{n \in \mathbb{N}} \llbracket t \rrbracket_n\right) = \bigsqcup_{n \in \mathbb{N}} c(\llbracket t \rrbracket_n). \quad (\text{A.7})$$

On the other side of the diagram we have $T(\llbracket - \rrbracket)(a(t))$ which by definition of T is equal to:

$$T(\llbracket - \rrbracket)(a(t)) = \begin{cases} \downarrow & \text{if } a(t) = \downarrow \\ \perp & \text{if } a(t) = \perp \\ (\sigma, k, \overrightarrow{\llbracket t'[\bar{i}/x] \rrbracket}) & \text{if } a(t) = (\sigma, k, \overrightarrow{t'[\bar{i}/x]}). \end{cases}$$

By definition of a the side conditions can be expressed in terms of the reduction behaviour of t :

$$T(\llbracket - \rrbracket)(a(t)) = \begin{cases} \downarrow & \text{if } t \longrightarrow^* \downarrow \\ \perp & \text{if } t \longrightarrow^\infty \\ (\sigma, k, \overrightarrow{\bigsqcup_{n \in \mathbb{N}} \llbracket t'[\bar{i}/x] \rrbracket_n}) & \text{if } t \longrightarrow^* \sigma(\bar{k}, x.t'). \end{cases}$$

By the definition of least upper bound in $T(\text{Trees}_\Sigma)$ we can rewrite the last case as:

$$T(\llbracket - \rrbracket)(a(t)) = \begin{cases} \downarrow & \text{if } t \longrightarrow^* \downarrow \\ \perp & \text{if } t \longrightarrow^\infty \\ \bigsqcup_{n \in \mathbb{N}} (\sigma, k, \overrightarrow{\llbracket t'[\bar{i}/x] \rrbracket_n}) & \text{if } t \longrightarrow^* \sigma(\bar{k}, x.t'). \end{cases} \quad (\text{A.8})$$

To see that equations A.7 and A.8 are in fact equal we proceed by a case analysis on the reduction of t . By definition of \longrightarrow the following cases are exhaustive and only one of them can occur:

1. $t \longrightarrow^* \downarrow$. There exists m such that $t \longrightarrow^m \downarrow$. By definition of $\llbracket - \rrbracket$ we can deduce that $\forall n \geq m + 1. \llbracket t \rrbracket_n = \downarrow$ so $c(\llbracket t \rrbracket_n) = \downarrow$. And $\forall n < m + 1. \llbracket t \rrbracket_n = \perp$. Therefore we have:

$$\bigsqcup_{n \in \mathbb{N}} c(\llbracket t \rrbracket_n) = \downarrow.$$

as required.

2. $t \longrightarrow^\infty$. For each $n \in \mathbb{N}$ there exists t_n such that $t \longrightarrow^n t_n$. Therefore $\forall n \in \mathbb{N}. \llbracket t \rrbracket_n = \llbracket t_n \rrbracket_0 = \perp$ so by the definition of least upper bound in $T(\text{Trees}_\Sigma)$:

$$\bigsqcup_{n \in \mathbb{N}} c(\llbracket t \rrbracket_n) = \perp.$$

3. $t \longrightarrow^* \sigma(\bar{k}, x.t')$. There exists $m \in \mathbb{N}$ such that $t \longrightarrow^m \sigma(\bar{k}, x.t')$. By definition of $\llbracket - \rrbracket_n$ we can deduce that:

$$\forall n \geq m + 1. \llbracket t \rrbracket_n = \sigma_k(\llbracket t'[\bar{0}/x] \rrbracket_{n-m-1}, \llbracket t'[\bar{1}/x] \rrbracket_{n-m-1}, \dots)$$

$$\forall n < m + 1. \llbracket t \rrbracket_n = \perp.$$

From this we know that:

$$\bigsqcup_{n \in \mathbb{N}} c(\llbracket t \rrbracket_n) = \bigsqcup_{n \geq m+1} c(\llbracket t \rrbracket_n) = \bigsqcup_{n \geq m+1} (\sigma, k, \overrightarrow{\llbracket t'[\bar{i}/x] \rrbracket_{n-m-1}}) = T(\llbracket - \rrbracket)(a(t)).$$

□

Proposition 4.3.14. *For any well-typed EPCF configuration (S, M) , where $S : \tau \Rightarrow \rho$, and any ECPS computation t :*

$$((S, M), t) \in \mathcal{S}_{\tau, \rho} \cap \mathcal{S}'_{\tau, \rho} \implies \llbracket t \rrbracket = |S, M|[\downarrow / l_1, \downarrow / l_2, \dots].$$

Proof. By coinduction. Define the relation $\mathcal{B} \subseteq (\text{Stack} \times \text{Comp}) \times (\vdash)$ as:

$$\mathcal{B} = \bigcup_{\tau, \rho} (\mathcal{S}_{\tau, \rho} \cap \mathcal{S}'_{\tau, \rho}).$$

We will show that \mathcal{B} is a bisimulation in the abstract sense of Definition 3.3.1. Consider the following morphism of type $\mathcal{B} \longrightarrow T(\mathcal{B})$:

$$r((S, M), t) = \begin{cases} \downarrow & \text{if } (S, M) \mapsto^* (id, \mathbf{return} V) \\ \perp & \text{if } (S, M) \mapsto^\infty \\ (\sigma, n, \overrightarrow{((S', W\bar{l}), t'[\bar{l}/x])}) & \text{if } (S, M) \mapsto^* (S', \sigma(\bar{n}; W)), \text{ which implies} \\ & \text{from definition of } \mathcal{S} \text{ that } t \longrightarrow^* \sigma(\bar{n}, x.t'). \end{cases}$$

This morphism makes the following diagram commute, so it makes \mathcal{B} into a bisimulation:

$$\begin{array}{ccccc} \text{Stack} \times \text{Comp} & \xleftarrow{\pi_1} & \mathcal{B} & \xrightarrow{\pi_2} & (\vdash) \\ \downarrow b & & \downarrow r & & \downarrow a \\ T(\text{Stack} \times \text{Comp}) & \xleftarrow{T(\pi_1)} & T(\mathcal{B}) & \xrightarrow{T(\pi_2)} & T(\vdash) \end{array}$$

By definition of b and r we have:

$$\forall ((S, M), t) \in \mathcal{B}. b(\pi_1((S, M), t)) = T(\pi_1)(r((S, M), t)).$$

For the second equation: we know that for any $((S, M), t) \in \mathcal{B}$:

$$a(\pi_2((S, M), t)) = a(t) = \begin{cases} \downarrow & \text{if } t \longrightarrow^* \downarrow \\ \perp & \text{if } t \longrightarrow^\infty \\ (\sigma, k, \overrightarrow{t'[\bar{n}/x]}) & \text{if } t \longrightarrow^* \sigma(\bar{k}, x.t'). \end{cases} \quad (\text{A.9})$$

$$T(\pi_2)(r((S, M), t)) = \begin{cases} \downarrow & \text{if } (S, M) \mapsto^* (id, \mathbf{return} V) \\ \perp & \text{if } (S, M) \mapsto^\infty \\ (\sigma, n, \overrightarrow{t'[\bar{l}/x]}) & \text{if } (S, M) \mapsto^* (S', \sigma(\bar{n}; W)), \text{ which implies} \\ & \text{from definition of } \mathcal{S} \text{ that } t \longrightarrow^* \sigma(\bar{n}, x.t'). \end{cases} \quad (\text{A.10})$$

Expressions A.9 and A.10 can be proved equal by using the definitions of \mathcal{S} and \mathcal{S}' . If $(S, M) \mapsto^* (S', \sigma(\bar{n}; W))$ then we know by definition of \mathcal{S} that $t \longrightarrow^* \sigma(\bar{n}, x.t')$. If $(S, M) \mapsto^* (id, \mathbf{return} V)$ then again by definition of \mathcal{S} we have $t \longrightarrow^* \downarrow$. If $(S, M) \mapsto^\infty$, because the reduction relation \mapsto is deterministic we know that $(S, M) \not\mapsto^* (id, \mathbf{return} V)$ and $(S, M) \not\mapsto^* (S', \sigma(\bar{n}; W))$ for any $n \in \mathbb{N}$. Therefore, by definition of \mathcal{S}' , $t \not\mapsto^* \downarrow$ and

$t \not\rightarrow^* \sigma(v, x.t')$. So by the definition of \longrightarrow it must be the case that $t \longrightarrow^\infty$, as required. We can prove the reverse implication by making an assumption about the reduction of t and proceeding analogously.

By assumption, we know $((S, M), t) \in \mathcal{B}$. As discussed in Section 4.3, $Trees_\Sigma$ is a final coalgebra. From Lemmas 4.3.12 and 4.3.13 we know that $\llbracket - \rrbracket$ and β are coalgebra morphisms into it. Using the fact that \mathcal{B} is a bisimulation, we can apply the coinduction proof principle, Proposition 3.3.2, to deduce:

$$\llbracket t \rrbracket = \beta(S, M) = |S, M|^* = |S, M|[\downarrow / l_1, \downarrow / l_2, \dots].$$

as required. □

Appendix B

Proofs about Applicative Bisimilarity

Proposition 5.2.3. *Applicative \mathfrak{B} -bisimilarity coincides with the intersection between applicative \mathfrak{B} -similarity and its converse:*

$$(\sim) = (\lesssim) \cap (\lesssim)^{op}.$$

Proof. We prove each inclusion in turn.

“ \subseteq ”. Bisimilarity is a bisimulation, hence also a simulation so $(\sim) \subseteq (\lesssim)$. Therefore $(\sim)^{op} \subseteq (\lesssim)^{op}$. But bisimilarity is by definition symmetric so $(\sim) = (\sim)^{op}$ so we can deduce that $(\sim) \subseteq (\lesssim)^{op}$. Hence, $(\sim) \subseteq (\lesssim) \cap (\lesssim)^{op}$ as required.

“ \supseteq ”. The strategy is to check that $(\lesssim) \cap (\lesssim)^{op}$ is a symmetric simulation. Consider $(s, t) \in (\lesssim) \cap (\lesssim)^{op}$. Then $s \lesssim t$ and $t \lesssim s$, which implies $t \lesssim^{op} s$ and $t \lesssim s$. Therefore, $(t, s) \in (\lesssim) \cap (\lesssim)^{op}$ and $(\lesssim) \cap (\lesssim)^{op}$ is symmetric.

The relation $(\lesssim) \cap (\lesssim)^{op}$ satisfies the first three conditions in the definition of simulation because we know \lesssim is a simulation. For the fourth condition assume:

$$(v, u) \in ((\lesssim) \cap (\lesssim)^{op})_{-(A_i)}^v.$$

Then $v \lesssim_{-(A_i)}^v w$ and $w \lesssim_{-(A_i)}^v v$ and since \lesssim is a simulation we know that:

$$\forall \vdash \overrightarrow{u_i} : A_i. v(\overrightarrow{u_i}) \lesssim^c w(\overrightarrow{u_i})$$

$$\forall \vdash \overrightarrow{u_i} : A_i. w(\overrightarrow{u_i}) \lesssim^c v(\overrightarrow{u_i}).$$

Therefore:

$$\forall \vdash \overrightarrow{u_i} : A_i. (v(\overrightarrow{u_i}), w(\overrightarrow{u_i})) \in ((\lesssim) \cap (\lesssim)^{op})^c.$$

So we have shown that $(\lesssim) \cap (\lesssim)^{op}$ is a symmetric simulation and it is therefore included in the union of all symmetric simulations \sim . \square

Lemma 5.3.1. *Applicative \mathfrak{B} -similarity is a preorder. Applicative \mathfrak{B} -bisimilarity is an equivalence relation.*

Proof. Prove that similarity is reflexive and transitive.

Reflexivity. Let \mathcal{I}_A^v and \mathcal{I}^c be the identity relations. We will show $(\mathcal{I}_A^v, \mathcal{I}^c) \subseteq \lesssim$ so \lesssim is reflexive. For this it suffices to show $(\mathcal{I}_A^v, \mathcal{I}^c)$ is a simulation. Conditions 1 and 2 are satisfied by definition. For condition 3 assume $s \mathcal{I}^c t$. Then it must be the case that $s = t$ so $\llbracket s \rrbracket = \llbracket t \rrbracket$. Therefore $\forall P \in \mathfrak{P}. (\llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P)$. For condition 4 assume $v \mathcal{I}_{\neg(A_1, \dots, A_n)}^v u$. Then $v = u$ so the result follows from the definition of \mathcal{I}^c .

Transitivity. First assume that there exist some closed computations such that $r \lesssim s \lesssim t$. Then we know from the definition of similarity and transitivity of implication that:

$$\forall P \in \mathfrak{P}. \llbracket r \rrbracket \in P \implies \llbracket t \rrbracket \in P.$$

Since similarity is the greatest relation with property 3, it must be the case that $r \lesssim t$.

Assume $u \lesssim v \lesssim w$. Then there exist simulations \mathcal{R} and \mathcal{S} such that $u \mathcal{R}_A^v v$ and $v \mathcal{S}_A^v w$. Proceed by a case distinction on A .

If $A = \mathbf{nat}$ or $A = \mathbf{unit}$, then it must be the case that $u = v = w$. Let $\mathcal{T}_A^v = \{(u, w)\}$, $\mathcal{T}_{B \neq A}^v = \emptyset$, $\mathcal{T}^c = \emptyset$. The relation $(\mathcal{T}_A^v, \mathcal{T}_{B \neq A}^v, \mathcal{T}^c)$ is a simulation because $u = w$, and is therefore included in \lesssim .

If $A = \neg(A_1, \dots, A_n)$, consider the candidate simulation $\mathcal{T}_A^v = \{(u, w)\}$, $\mathcal{T}_{B \neq A}^v = \emptyset$, $\mathcal{T}^c = \{(u(w_1, \dots, w_n), w(w_1, \dots, w_n)) \mid \vdash w_1 : A_1, \dots, \vdash w_n : A_n\}$. Condition 4 is satisfied by definition of \mathcal{T}^c . Condition 3 is satisfied because $u \mathcal{R}_A^v v$ and $v \mathcal{S}_A^v w$ imply:

$$\forall \vdash w_1 : A_1, \dots, \vdash w_n : A_n. \forall P \in \mathfrak{P}. \llbracket u(w_1, \dots, w_n) \rrbracket \in P \implies \llbracket w(w_1, \dots, w_n) \rrbracket \in P.$$

From Proposition 5.2.3 we know that:

$$(\sim) = (\lesssim) \cap (\lesssim)^{op}.$$

The relation $(\lesssim) \cap (\lesssim)^{op}$ is reflexive and transitive because similarity is, and it is by definition symmetric. So \sim is an equivalence relation. \square

Lemma 5.3.4. *Consider a well-typed relation \mathcal{R} that is a preorder. The compatibility rules (COMP6), (COMP7), (COMP8) and (COMP10) are equivalent to the conjunction of their single-premise versions.*

Proof. As an example, consider the case of $((\text{COMP7}) \iff (\text{COMP7L}) \text{ and } (\text{COMP7R}_i))$. The proof for the other cases is analogous.

Assume that (COMP7) is true and that

$$\Gamma, x : \neg(\vec{A}) \vdash v \mathcal{R}_{\neg(\vec{A})}^v v'$$

$$\Gamma \vdash w_i \mathcal{R}_{A_i}^v w'_i \text{ for each } i.$$

By reflexivity of \mathcal{R} we know $\Gamma \vdash w_i \mathcal{R}_{A_i}^v w_i$ for each i . So we can apply rule (COMP7) to deduce the conclusion of rule (COMP7L).

Also by reflexivity, we know $\Gamma, x : \neg(\vec{A}) \vdash v \mathcal{R}_{\neg(\vec{A})}^v v$ and $\Gamma \vdash w_j \mathcal{R}_{A_j}^v w_j$ for each $j \neq i$. Now apply rule (COMP7) to deduce the conclusion of rule (COMP7R_i).

For the reverse implication assume that the rules (COMP7L) and (COMP7R_{*i*}) are true and that:

$$\Gamma, x : \neg(\vec{A}) \vdash v \mathcal{R}_{\neg(\vec{A})}^v v'$$

$$\Gamma \vdash w_i \mathcal{R}_{A_i}^v w'_i \text{ for each } i.$$

Because \mathcal{R} is well-typed we know that all the values considered above are well-typed. Therefore, we can apply (COMP7L) and (COMP7R_{*i*}) one by one to obtain:

$$\Gamma \vdash (\mu x.v)(\vec{w}) \mathcal{R}^c (\mu x.v')(\vec{w})$$

$$\Gamma \vdash (\mu x.v')(w_1, w_2, \dots, w_n) \mathcal{R}^c (\mu x.v')(w'_1, w_2, \dots, w_n)$$

...

$$\Gamma \vdash (\mu x.v')(w'_1, \dots, w'_i, \dots, w'_{n-1}, w_n) \mathcal{R}^c (\mu x.v')(w'_1, \dots, w'_i, \dots, w'_{n-1}, w'_n)$$

so by transitivity of \mathcal{R} we have $\Gamma \vdash (\mu x.v)(\vec{w}) \mathcal{R}^c (\mu x.v')(\vec{w}')$ as required. \square

B.1 Howe's Method

Lemma 5.4.3 (From [SV17, Appendix]). *Given a well-typed relation \mathcal{R} on closed terms that is reflexive:*

1. *The Howe extension of \mathcal{R} , $\mathcal{R}^{\mathcal{H}}$, is compatible and hence reflexive.*
2. $\mathcal{R}^\circ \subseteq \mathcal{R}^{\mathcal{H}}$.

Proof. 1. From the definitions of compatibility and compatible refinement we see that $\widehat{\mathcal{R}^{\mathcal{H}}} \subseteq \widehat{\mathcal{R}^{\mathcal{H}}}$ implies that $\mathcal{R}^{\mathcal{H}}$ is compatible.

We know \mathcal{R} is reflexive. Therefore \mathcal{R}° is reflexive. Let Id be the identity well-typed open relation. Then:

$$\widehat{\mathcal{R}^{\mathcal{H}}} = Id \circ \widehat{\mathcal{R}^{\mathcal{H}}} \subseteq \mathcal{R}^\circ \circ \widehat{\mathcal{R}^{\mathcal{H}}} = \mathcal{R}^{\mathcal{H}}.$$

To show $\mathcal{R}^{\mathcal{H}}$ is reflexive, that is, for all terms s , $s \mathcal{R}^{\mathcal{H}} s$ we can proceed by induction on s using the fact that $\mathcal{R}^{\mathcal{H}}$ is compatible.

2. Now we know that $\mathcal{R}^{\mathcal{H}}$ is reflexive. By the definition of compatible refinement we can easily see that this implies $\widehat{\mathcal{R}^{\mathcal{H}}}$ is reflexive. Therefore:

$$\mathcal{R}^\circ = \mathcal{R}^\circ \circ Id \subseteq \mathcal{R}^\circ \circ \widehat{\mathcal{R}^{\mathcal{H}}} = \mathcal{R}^{\mathcal{H}}.$$

\square

Lemma 5.4.4 (From [SV17, Appendix]). *Given a well-typed relation \mathcal{R} on closed terms that is transitive:*

$$\mathcal{R}^\circ \circ \mathcal{R}^{\mathcal{H}} \subseteq \mathcal{R}^{\mathcal{H}}.$$

Proof. Consider three terms s , t and r , either values or computations, such that $s \mathcal{R}^{\mathcal{H}} t \mathcal{R}^{\circ} r$. By definition of $\mathcal{R}^{\mathcal{H}}$ this means there exists a term t' such that:

$$s \widehat{\mathcal{R}^{\mathcal{H}}} t' \mathcal{R}^{\circ} t \mathcal{R}^{\circ} r.$$

Since \mathcal{R} is transitive, \mathcal{R}° is also transitive. Therefore:

$$s \widehat{\mathcal{R}^{\mathcal{H}}} t' \mathcal{R}^{\circ} r$$

so by definition of $\mathcal{R}^{\mathcal{H}}$ we have $s \mathcal{R}^{\mathcal{H}} r$. \square

Lemma 5.4.5 (Substitutivity). *Given a well-typed relation \mathcal{R} on closed terms that is transitive, its Howe extension satisfies the following two value-substitutivity properties:*

1. $\overrightarrow{x_i : A_i}, y : B \vdash s \mathcal{R}^{\mathcal{H},c} t$ and $\overrightarrow{x_i : A_i} \vdash v \mathcal{R}_B^{\mathcal{H},v} w \implies \overrightarrow{x_i : A_i} \vdash s[v/y] \mathcal{R}^{\mathcal{H},c} t[w/y]$.
2. $\overrightarrow{x_i : A_i}, y : B \vdash u \mathcal{R}_C^{\mathcal{H},v} u'$ and $\overrightarrow{x_i : A_i} \vdash v \mathcal{R}_B^{\mathcal{H},v} w \implies \overrightarrow{x_i : A_i} \vdash u[v/y] \mathcal{R}_C^{\mathcal{H},v} u'[w/y]$.

Proof. We prove the two statements by induction on s and u :

If s is a computation then $\overrightarrow{x_i : A_i}, y : B \vdash s \mathcal{R}^{\mathcal{H},c} t$ was derived using rule (HC) so it must be the case that:

$$\overrightarrow{x_i : A_i}, y : B \vdash s \widehat{\mathcal{R}^{\mathcal{H}}} t' \tag{B.1}$$

$$\overrightarrow{x_i : A_i}, y : B \vdash t' \mathcal{R}^{\circ,c} t. \tag{B.2}$$

By the definition of open extension we know from equation B.2 that:

$$\overrightarrow{x_i : A_i} \vdash t'[w/y] \mathcal{R}^{\circ,c} t[w/y].$$

If we can prove

$$\overrightarrow{x_i : A_i} \vdash s[v/y] \widehat{\mathcal{R}^{\mathcal{H}}} t'[w/y]$$

then we could use rule (HC) to deduce the desired result, $\overrightarrow{x_i : A_i} \vdash s[v/y] \mathcal{R}^{\mathcal{H},c} t[w/y]$.

Case $s = (\mu z.u)(\overrightarrow{w_i})$. It must be the case that equation B.1 was obtained using rule (C7) so we know that:

$$\begin{aligned} t' &= (\mu z.u')(\overrightarrow{w_i}) \\ \overrightarrow{x_i : A_i}, y : B, z : \neg(\overrightarrow{C}) &\vdash u \mathcal{R}_{\neg(\overrightarrow{C})}^{\mathcal{H},v} u' \\ \overrightarrow{x_i : A_i}, y : B &\vdash w_i \mathcal{R}_{C_i}^{\mathcal{H},v} w'_i \text{ for each } i. \end{aligned}$$

By induction hypothesis for u and $\overrightarrow{w_i}$ we can deduce:

$$\begin{aligned} \overrightarrow{x_i : A_i}, z : \neg(\overrightarrow{C}) &\vdash u[v/y] \mathcal{R}_{\neg(\overrightarrow{C})}^{\mathcal{H},v} u'[w/y] \\ \overrightarrow{x_i : A_i} &\vdash w_i[v/y] \mathcal{R}_{C_i}^{\mathcal{H},v} w'_i[w/y] \text{ for each } i \end{aligned}$$

and then apply rule (C7) to get $\overrightarrow{x_i : A_i} \vdash s[v/y] \widehat{\mathcal{R}^{\mathcal{H}}} t'[w/y]$.

Case $s = \downarrow$. In this case equation B.1 was obtained from rule (C9), so $t' = \downarrow$ and $s[v/y] = t'[w/y] = \downarrow$. We can then deduce $\overrightarrow{x_i : A_i} \vdash s[v/y] \widehat{\mathcal{R}}^{\mathcal{H}^c} t'[w/y]$ by rule (C9).

Cases $s = u(\overrightarrow{w_i})$, $s = \sigma(u, x.s')$, $s = \text{case } u \text{ in } \{\text{zero} \Rightarrow s', \text{succ}(z) \Rightarrow s''\}$. Analogous to the case $s = (\mu z.u)(\overrightarrow{w_i})$.

If u is a value then $\overrightarrow{x_i : A_i}, y : B \vdash u \mathcal{R}_C^{\mathcal{H},v} u'$ was derived using rule (HV) so it must be the case that:

$$\overrightarrow{x_i : A_i}, y : B \vdash u \widehat{\mathcal{R}}_C^{\mathcal{H},v} u'' \quad (\text{B.3})$$

$$\overrightarrow{x_i : A_i}, y : B \vdash u'' \mathcal{R}_C^{\circ,v} u'. \quad (\text{B.4})$$

From equation B.4 by the definition of open extension we have:

$$\overrightarrow{x_i : A_i} \vdash u''[w/y] \mathcal{R}_C^{\circ,v} u'[w/y] \quad (\text{B.5})$$

so using rule (HV) it suffices to prove:

$$\overrightarrow{x_i : A_i} \vdash u[v/y] \widehat{\mathcal{R}}_C^{\mathcal{H},v} u''[w/y]. \quad (\text{B.6})$$

Case $u = \lambda \overrightarrow{z_j} : \overrightarrow{C_j}.r$. It must be the case that equation B.3 was obtained by rule (C3) so:

$$u'' = \lambda \overrightarrow{z_j} : \overrightarrow{C_j}.r'$$

$$\overrightarrow{x_i : A_i}, y : B, \overrightarrow{z_j} : \overrightarrow{C_j} \vdash r \mathcal{R}^{\mathcal{H},c} r'.$$

By induction hypothesis for r and applying rule (C3) we can deduce:

$$\overrightarrow{x_i : A_i} \vdash \lambda \overrightarrow{z_j} : \overrightarrow{C_j}.r[v/y] \widehat{\mathcal{R}}_{-(\overrightarrow{C_j})}^{\mathcal{H},v} \lambda \overrightarrow{z_j} : \overrightarrow{C_j}.r'[w/y].$$

Case $u = \text{succ}(u')$. Analogous to the previous case.

Cases $u = \text{zero}$ and $u = \star$. Analogous to the case $s = \downarrow$.

Case $u = z$. If $z \neq y$ then $z = x_j$ for some j . Then equation B.3 must have been the conclusion of rule (C1) so $u'' = z$. By rule (C1) we have:

$$\overrightarrow{x_i : A_i} \vdash x_j \widehat{\mathcal{R}}_C^{\mathcal{H},v} x_j$$

which is what we had to prove.

If $z = y$ then instead of going through equation B.6 we will prove directly $\overrightarrow{x_i : A_i} \vdash u[v/y] \mathcal{R}_C^{\mathcal{H},v} u'[w/y]$, that is:

$$\overrightarrow{x_i : A_i} \vdash v \mathcal{R}_C^{\mathcal{H},v} u'[w/y].$$

Because \mathcal{R} is transitive, we can apply Lemma 5.4.4 to obtain: $\mathcal{R}^\circ \circ \mathcal{R}^{\mathcal{H}} \subseteq \mathcal{R}^{\mathcal{H}}$. We already know equation B.5:

$$\overrightarrow{x_i : A_i} \vdash u''[w/y] \mathcal{R}_C^{\circ,v} u'[w/y]$$

so we just need to show:

$$\overrightarrow{x_i : A_i} \vdash v \mathcal{R}_C^{\mathcal{H},v} u''[w/y].$$

Equation B.3 must be the conclusion of rule (C1) so $u'' = y$. Then we are left to prove:

$$\overrightarrow{x_i : A_i} \vdash v \mathcal{R}_C^{\mathcal{H},v} w$$

which we know by the initial assumption. \square

Lemma 5.4.6. *Consider a well-typed closed relation \leq that is a \mathfrak{B} -simulation. For any closed values v and w :*

$$\vdash v \leq_{\text{nat}}^{\mathcal{H},v} w \implies v = w.$$

Proof. Since v is a closed value of type nat there exists $n \in \mathbb{N}$ such that $v = \bar{n}$. The equation $\vdash \bar{n} \leq_{\text{nat}}^{\mathcal{H},v} w$ must have been derived by rule (HV) so:

$$\vdash \bar{n} \leq_{\text{nat}}^{\widehat{\mathcal{H}}^v} u \quad \text{and} \quad \vdash u \leq_{\text{nat}}^{o,v} w.$$

But u and w are closed so $u \leq_{\text{nat}}^v w$. Since \leq is a simulation it follows that $u = w$. So we know:

$$\vdash \bar{n} \leq_{\text{nat}}^{\widehat{\mathcal{H}}^v} w. \tag{B.7}$$

We prove by induction on n that:

$$\forall \vdash w : \text{nat}. \vdash \bar{n} \leq_{\text{nat}}^{\mathcal{H},v} w \implies \bar{n} = w.$$

Base case, $n = 0$. From $\vdash \bar{n} \leq_{\text{nat}}^{\mathcal{H},v} w$ we deduce equation B.7, which must be the conclusion of rule (C4). Therefore $\bar{n} = w = \text{zero}$.

Induction step. Assume $\overline{n+1} \leq_{\text{nat}}^{\mathcal{H},v} w'$ for some arbitrary $\vdash w' : \text{nat}$. Then equation B.7 becomes:

$$\vdash \text{succ}(\bar{n}) \leq_{\text{nat}}^{\widehat{\mathcal{H}}^v} w'.$$

This must be the conclusion of rule (C5) so $w' = \text{succ}(w'')$ and $\bar{n} \leq_{\text{nat}}^{\mathcal{H},v} w''$.

We can instantiate the induction hypothesis with the last equation to obtain $w'' = \bar{n}$. This means that $\overline{n+1} = \text{succ}(\bar{n}) = \text{succ}(w'') = w'$ as required. \square

Lemma 5.4.7 (Key Lemma). *Consider a decomposable set of Scott-open observations \mathfrak{B} . Consider a well-typed closed relation \leq that is a preorder and a \mathfrak{B} -simulation. For any closed computations s and t , $\vdash s \leq^{\mathcal{H},c} t$ implies:*

$$\forall n \in \mathbb{N}. \forall P \in \mathfrak{B}. \llbracket s \rrbracket_n \in P \implies \llbracket t \rrbracket \in P.$$

Proof. We prove by induction on $n \in \mathbb{N}$ that:

$$\forall n \in \mathbb{N}. \forall \text{computations } s', t'. \vdash s' \leq^{\mathcal{H},c} t' \implies (\forall P \in \mathfrak{B}. \llbracket s' \rrbracket_n \in P \implies \llbracket t' \rrbracket \in P).$$

Base case, $n = 0$. By definition $\llbracket s' \rrbracket_0 = \perp$. Assume that $\llbracket s' \rrbracket_0 \in P$. Then by upwards closure of P it follows that $\text{Tree}_\Sigma \subseteq P$, so we also have $\llbracket t' \rrbracket \in P$, as required.

Induction step. The induction hypothesis is:

$\forall k < n + 1 \in \mathbb{N}. \forall \text{ computations } s', t'.$

$$\vdash s' \leq^{\mathcal{H},c} t' \implies (\forall P \in \mathfrak{P}. \llbracket s' \rrbracket_k \in P \implies \llbracket t' \rrbracket \in P).$$

Assume that $\vdash s \leq^{\mathcal{H},c} t$. This must be the conclusion of rule (HC) so there exists r such that:

$$\vdash s \leq^{\widehat{\mathcal{H}}^c} r \quad \text{and} \quad \vdash r \leq^{\circ,c} t.$$

But r and t are closed terms so we in fact know $r \leq^c t$. Because \leq is a simulation it follows that:

$$\forall P' \in \mathfrak{P}. \llbracket r \rrbracket \in P' \implies \llbracket t \rrbracket \in P'.$$

Therefore, it suffices to show:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket_{n+1} \in P \implies \llbracket r \rrbracket \in P.$$

To do this we proceed by a case split on the structure of s .

Case $s = v(w_1, \dots, w_n)$. Because s is a well-typed closed computation, it must be the case that $s = (\lambda \vec{x}_i : \vec{A}_i. s') (w_1, \dots, w_n)$.

The equation $\vdash s \leq^{\widehat{\mathcal{H}}^c} r$ must have been obtained by rule (C3) so it must be the case that:

$$r = (\lambda \vec{x}_i : \vec{A}_i. r') (w'_1, \dots, w'_n) \tag{B.8}$$

$$\vdash \lambda \vec{x}_i : \vec{A}_i. s' \leq^{\mathcal{H},v}_{-(\vec{A}_i)} \lambda \vec{x}_i : \vec{A}_i. r' \tag{B.9}$$

$$\vdash w_i \leq^{\mathcal{H},v}_{A_i} w'_i \text{ for each } i. \tag{B.10}$$

By definition of $\llbracket - \rrbracket_{(-)}$ we know $\llbracket s \rrbracket_{n+1} = \llbracket (\lambda \vec{x}_i : \vec{A}_i. s') (w_1, \dots, w_n) \rrbracket_{n+1} = \llbracket s'[\vec{w}_i/\vec{x}_i] \rrbracket_n$, and similarly for r . So we only need to prove:

$$\forall P \in \mathfrak{P}. \llbracket s'[\vec{w}_i/\vec{x}_i] \rrbracket_n \in P \implies \llbracket r'[\vec{w}'_i/\vec{x}_i] \rrbracket \in P.$$

Equation B.9 must have been obtained by rule (HV) so there exists p such that:

$$\vdash \lambda \vec{x}_i : \vec{A}_i. s' \leq^{\widehat{\mathcal{H}}^v}_{-(\vec{A}_i)} p \tag{B.11}$$

$$\vdash p \leq^{\circ,v}_{-(\vec{A}_i)} \lambda \vec{x}_i : \vec{A}_i. r'. \tag{B.12}$$

Equation B.11 must be the conclusion of rule (C3) so:

$$p = \lambda \vec{x}_i : \vec{A}_i. p' \quad \text{and} \quad \vec{x}_i : \vec{A}_i \vdash s' \leq^{\mathcal{H},c} p'.$$

Because \leq is transitive we can apply Lemma 5.4.5 to obtain the substitutivity property for $\leq^{\mathcal{H}}$. Using this and equation B.10 we can deduce:

$$\vdash s'[\vec{w}_i/\vec{x}_i] \leq^{\mathcal{H},c} p'[\vec{w}'_i/\vec{x}_i].$$

Apply the induction hypothesis for this to obtain:

$$\forall P \in \mathfrak{P}. \llbracket s'[\overrightarrow{w_i/x_i}] \rrbracket_n \in P \implies \llbracket p'[\overrightarrow{w_i/x_i}] \rrbracket \in P.$$

We know that $\vdash \lambda \vec{x}_i : \vec{A}_i. p' \leq_{-(\vec{A}_i)}^v \lambda \vec{x}_i : \vec{A}_i. r'$ from equation B.12. By the definition of simulation we then have:

$$\forall P \in \mathfrak{P}. \llbracket p'[\overrightarrow{w_i/x_i}] \rrbracket \in P \implies \llbracket r'[\overrightarrow{w_i/x_i}] \rrbracket \in P.$$

From here we obtain the required result:

$$\forall P \in \mathfrak{P}. \llbracket s'[\overrightarrow{w_i/x_i}] \rrbracket_n \in P \implies \llbracket r'[\overrightarrow{w_i/x_i}] \rrbracket \in P.$$

Case $s = \sigma(v, x.s')$. In this case $\vdash s \leq^{\mathcal{H},c} r$ is the conclusion of rule (C8) so we know:

$$\begin{aligned} r &= \sigma(v', x.r') \\ \vdash v &\leq_{\text{nat}}^{\mathcal{H},v} v' \\ x : \text{nat} \vdash s' &\leq^{\mathcal{H},c} r'. \end{aligned}$$

Using $\vdash v \leq_{\text{nat}}^{\mathcal{H},v} v'$ and Lemma 5.4.6 deduce $v = v' = \bar{k}$.

By reflexivity of $\leq^{\mathcal{H}}$ (Lemma 5.4.3) we have $\vdash \bar{m} \leq_{\text{nat}}^{\mathcal{H},v} \bar{m}$. Using $x : \text{nat} \vdash s' \leq^{\mathcal{H},c} r'$ we obtain by substitutivity (Lemma 5.4.5) that:

$$\forall m \in \mathbb{N}. \vdash s'[\bar{m}/x] \leq^{\mathcal{H},c} r'[\bar{m}/x].$$

Applying the induction hypothesis to this we we obtain:

$$\forall m \in \mathbb{N}. \forall P' \in \mathfrak{P}. \llbracket s'[\bar{m}/x] \rrbracket_n \in P' \implies \llbracket r'[\bar{m}/x] \rrbracket \in P'. \quad (\text{B.13})$$

By definition of $\llbracket - \rrbracket_{(-)}$ it suffices to prove:

$$\forall P \in \mathfrak{P}. \sigma_k(\llbracket s'[\bar{0}/x] \rrbracket_n, \dots, \llbracket s'[\bar{k}/x] \rrbracket_n, \dots) \in P \implies \sigma_k(\llbracket r'[\bar{0}/x] \rrbracket, \dots, \llbracket r'[\bar{k}/x] \rrbracket, \dots) \in P.$$

Assume $\sigma_k(\llbracket s'[\bar{0}/x] \rrbracket_n, \dots, \llbracket s'[\bar{k}/x] \rrbracket_n, \dots) \in P$. Then by decomposability of \mathfrak{P} , Definition 5.3.6, we know there exist observations \vec{P}'_m such that:

$$\begin{aligned} \forall \vec{p}'_m \in \vec{P}'_m. \sigma_k(\vec{p}'_m) &\in P \\ \text{for each } m \in \mathbb{N} \quad \llbracket s'[\bar{m}/x] \rrbracket_n &\in P'_m. \end{aligned}$$

By equation B.13 we can deduce that for all $m \in \mathbb{N}$, $\llbracket r'[\bar{m}/x] \rrbracket \in P'_m$. So we have

$$\sigma_k(\llbracket r'[\bar{0}/x] \rrbracket, \dots, \llbracket r'[\bar{k}/x] \rrbracket, \dots) \in P$$

as required.

Case $s = (\mu x.v)(\vec{w}_i)$. In this case $s \leq^{\widehat{\mathcal{H}}^c} r$ is the conclusion of rule (C7) so:

$$r = (\mu x.v')(\vec{w}'_i) \quad (\text{B.14})$$

$$x : \neg(\vec{A}_i) \vdash v \leq_{\neg(\vec{A}_i)}^{\mathcal{H}, \mathfrak{v}} v' \quad (\text{B.15})$$

$$w_i \leq_{A_i}^{\mathcal{H}, \mathfrak{v}} w'_i \text{ for each } i. \quad (\text{B.16})$$

By definition of $\llbracket - \rrbracket_{(-)}$ it suffices to prove:

$$\forall P \in \mathfrak{P}. \llbracket v[\lambda \vec{y} : \vec{A}_i. (\mu x.v)(\vec{y})/x] (\vec{w}_i) \rrbracket_n \in P \implies \llbracket v'[\lambda \vec{y} : \vec{A}_i. (\mu x.v')(\vec{y})/x] (\vec{w}'_i) \rrbracket \in P.$$

By equation B.15 we can deduce using context weakening that:

$$\vec{y} : \vec{A}_i, x : \neg(\vec{A}_i) \vdash v \leq_{\neg(\vec{A}_i)}^{\mathcal{H}, \mathfrak{v}} v'.$$

Because \leq is reflexive, we can apply Lemma 5.4.3 to deduce $\leq^{\mathcal{H}}$ is reflexive. Therefore:

$$\vec{y} : \vec{A}_i \vdash y_i \leq_{A_i}^{\mathcal{H}, \mathfrak{v}} y_i.$$

From the last two equations and from rule (C7) we can deduce:

$$\vec{y} : \vec{A}_i \vdash (\mu x.v)(\vec{y}) \leq^{\widehat{\mathcal{H}}^c} (\mu x.v')(\vec{y}).$$

By reflexivity of \leq we know:

$$\vec{y} : \vec{A}_i \vdash (\mu x.v')(\vec{y}) \leq^{\circ, c} (\mu x.v')(\vec{y}).$$

Using the last two equations and rule (HC) we obtain:

$$\vec{y} : \vec{A}_i \vdash (\mu x.v)(\vec{y}) \leq^{\mathcal{H}, c} (\mu x.v')(\vec{y}).$$

From this, using rule (C3), we have:

$$\vdash \lambda \vec{y} : \vec{A}_i. (\mu x.v)(\vec{y}) \leq^{\widehat{\mathcal{H}}^{\mathfrak{v}}} \lambda \vec{y} : \vec{A}_i. (\mu x.v')(\vec{y}).$$

By reflexivity of \leq we have:

$$\vdash \lambda \vec{y} : \vec{A}_i. (\mu x.v')(\vec{y}) \leq_{\neg(\vec{A}_i)}^{\circ, \mathfrak{v}} \lambda \vec{y} : \vec{A}_i. (\mu x.v')(\vec{y}).$$

From the last two equations, by rule (HV), we obtain:

$$\vdash \lambda \vec{y} : \vec{A}_i. (\mu x.v)(\vec{y}) \leq_{\neg(\vec{A}_i)}^{\mathcal{H}, \mathfrak{v}} \lambda \vec{y} : \vec{A}_i. (\mu x.v')(\vec{y}).$$

Using this last equation and equation B.15: $x : \neg(\vec{A}_i) \vdash v \leq_{\neg(\vec{A}_i)}^{\mathcal{H}, \mathfrak{v}} v'$, we obtain by substitutivity for $\leq^{\mathcal{H}}$, Lemma 5.4.5, that:

$$\vdash v[(\lambda \vec{y} : \vec{A}_i. (\mu x.v)(\vec{y}))/x] \leq_{\neg(\vec{A}_i)}^{\mathcal{H}, \mathfrak{v}} v'[(\lambda \vec{y} : \vec{A}_i. (\mu x.v')(\vec{y}))/x].$$

From Lemma 5.4.3 we can deduce $\leq^{\mathcal{H}}$ is compatible. Using the last equation and equation B.16, $\vdash w_i \leq_{A_i}^{\mathcal{H}, \mathfrak{v}} w'_i$ for each i , we obtain by compatibility:

$$\vdash v[(\lambda \vec{y} : \vec{A}_i. (\mu x.v)(\vec{y}))/x] (\vec{w}_i) \leq^{\mathcal{H}, c} v'[(\lambda \vec{y} : \vec{A}_i. (\mu x.v')(\vec{y}))/x] (\vec{w}'_i).$$

By applying the induction hypothesis to this we obtain the desired result:

$$\forall P \in \mathfrak{P}. \llbracket v[\lambda \vec{y} : \vec{A}_i. (\mu x.v)(\vec{y})/x] (\vec{w}_i) \rrbracket_n \in P \implies \llbracket v'[\lambda \vec{y} : \vec{A}_i. (\mu x.v')(\vec{y})/x] (\vec{w}'_i) \rrbracket \in P.$$

Case $s = \text{case } v \text{ in } \{\text{zero} \Rightarrow s', \text{succ}(y) \Rightarrow s''\}$. In this case $\vdash s \leq^{\widehat{\mathcal{H}}^c} r$ is the conclusion of rule (C10) so we know:

$$\begin{aligned} r &= \text{case } v' \text{ in } \{\text{zero} \Rightarrow r', \text{succ}(y) \Rightarrow r''\} \\ \vdash v &\leq_{\text{nat}}^{\mathcal{H},v} v' \\ \vdash s' &\leq^{\mathcal{H},c} r' \\ y : \text{nat} &\vdash s'' \leq^{\mathcal{H},c} r''. \end{aligned}$$

Using $\vdash v \leq_{\text{nat}}^{\mathcal{H},v} v'$ we can apply Lemma 5.4.6 to deduce $v = v'$.

If $v = v' = \text{zero}$ then it suffices to prove:

$$\forall P \in \mathfrak{P}. \llbracket s' \rrbracket_n \in P \implies \llbracket r' \rrbracket \in P.$$

We can deduce this using the induction hypothesis for s' and r' and $\vdash s' \leq^{\mathcal{H},c} r'$.

If $v = v' = \text{succ}(w)$ where $\vdash w : \text{nat}$ then it suffices to prove:

$$\forall P \in \mathfrak{P}. \llbracket s''[w/y] \rrbracket_n \in P \implies \llbracket r''[w/y] \rrbracket \in P.$$

By reflexivity of $\leq^{\mathcal{H}}$ (Lemma 5.4.3) we have $\vdash w \leq_{\text{nat}}^{\mathcal{H},v} w$. We can use this, $y : \text{nat} \vdash s'' \leq^{\mathcal{H},c} r''$ and substitutivity for $\leq^{\mathcal{H}}$ (Lemma 5.4.5), to deduce:

$$\vdash s''[w/y] \leq^{\mathcal{H},c} r''[w/y].$$

We then get the desired result from the induction hypothesis for $s''[w/y]$ and $r''[w/y]$.

Case $s = \downarrow$. In this case $\vdash s \leq^{\widehat{\mathcal{H}}^c} r$ is the conclusion of rule (C9) so $r = \downarrow$. Then $\llbracket s \rrbracket_{n+1} = \llbracket r \rrbracket = \downarrow$. Therefore we have the required result:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket_{n+1} \in P \implies \llbracket r \rrbracket \in P.$$

□

Lemma 5.4.9. *Given a well-typed open relation \mathcal{R} that is reflexive and has the two substitutivity properties from Lemma 5.4.5, and a well-typed closed relation \mathcal{S} then:*

if \mathcal{R} restricted to closed terms is included in \mathcal{S} then $\mathcal{R} \subseteq \mathcal{S}^\circ$.

Proof. Consider terms s and t , values or computations, such that $\overrightarrow{x_i : A_i} \vdash s \mathcal{R} t$. By reflexivity of \mathcal{R} we know that for any values $\vdash \overrightarrow{v_i : A_i}$ we have $\vdash \overrightarrow{v_i} \mathcal{R}^\circ \overrightarrow{v_i}$. Therefore we can apply the substitutivity property of \mathcal{R} to obtain:

$$\forall \vdash \overrightarrow{v_i : A_i}. \vdash s[\overrightarrow{v_i/x_i}] \mathcal{R} t[\overrightarrow{v_i/x_i}].$$

From here we can deduce by assumption that:

$$\forall \vdash \overrightarrow{v_i : A_i}. \vdash s[\overrightarrow{v_i/x_i}] \mathcal{S} t[\overrightarrow{v_i/x_i}]$$

so by the definition of open extension we have:

$$\overrightarrow{x_i : A_i} \vdash s \mathcal{S}^\circ t.$$

□

Lemma 5.4.10. *Given a \mathfrak{P} -simulation \mathcal{R} , its reflexive-transitive closure, \mathcal{R}^* is also a \mathfrak{P} -simulation.*

Proof. We check all the conditions in the definition of \mathfrak{P} -simulation in turn:

1. $\vdash v \mathcal{R}_{\text{unit}}^{*,\flat} w \implies v = w = \star$. Assume $\vdash v \mathcal{R}_{\text{unit}}^{*,\flat} w$. The only closed value of type `unit` is \star so $v = w = \star$.
2. $\vdash v \mathcal{R}_{\text{nat}}^{*,\flat} w \implies v = w$. Assume $v \mathcal{R}_{\text{nat}}^{*,\flat} w$. Then by the definition of reflexive-transitive closure there must exist a chain of values v_1, \dots, v_k such that $v_1 = v$ and $v_k = w$ and $v_{i-1} \mathcal{R}_{\text{nat}}^{\flat} v_i$ for each $i > 1$. If $k = 1$ then $v = w$. If $k > 1$, we can use the fact that \mathcal{R} is a simulation to deduce $v_{i-1} = v_i$ for each $i > 1$. So by transitivity $v = w$.
3. $\vdash s \mathcal{R}^{*,\flat} t \implies \forall P \in \mathfrak{P}. (\llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P)$. Assume $s \mathcal{R}^{*,\flat} t$. There exists a chain of computations s_1, \dots, s_k such that $s_1 = s$ and $s_k = t$ and $s_{i-1} \mathcal{R}^{\flat} s_i$ for each $i > 1$. If $k = 1$ then $s = t$ so we have $\forall P \in \mathfrak{P}. (\llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P)$ as required. If $k > 1$ then for each $i > 1$:

$$\forall P \in \mathfrak{P}. (\llbracket s_{i-1} \rrbracket \in P \implies \llbracket s_i \rrbracket \in P).$$

From here we can deduce the desired result by transitivity of implication.

4. $\vdash v \mathcal{R}_{-(A_1, \dots, A_n)}^{*,\flat} w \implies \forall \vdash u_1 : A_1, \dots, \vdash u_n : A_n. v(u_1, \dots, u_n) \mathcal{R}^{*,\flat} w(u_1, \dots, u_n)$. Similar to the previous two cases. It uses the fact that $v \mathcal{R} u \mathcal{R} w$ implies $v \mathcal{R}^* w$.

□

Lemma 5.4.11. *Given a well-typed compatible relation \mathcal{R} , its reflexive-transitive closure \mathcal{R}^* is also compatible.*

Proof. The proof is similar to the proof of Lemma 5.4.10 in that it expands $s \mathcal{R}^* t$ into a chain of terms related by \mathcal{R} . It also uses the fact that \mathcal{R} compatible implies \mathcal{R} is reflexive, Lemma 5.4.3. □

Lemma 5.4.12 (From [Las98]). *Given a well-typed closed relation \mathcal{R} the following holds:*

if \mathcal{R}° is reflexive and symmetric, then $\mathcal{R}^{\mathcal{H}^}$ is symmetric.*

Where S^* denotes the reflexive-transitive closure of a relation \mathcal{S} .

Proof. By examining the compatible refinement rules we can observe that for any relation \mathcal{S} :

$$\widehat{\mathcal{S}^{op}} = \widehat{\mathcal{S}}^{op}. \tag{B.17}$$

Since \mathcal{R}° is reflexive, \mathcal{R} is also reflexive. Therefore we can apply Lemma 5.4.3 to deduce:

$$\mathcal{R}^{\circ} \subseteq \mathcal{R}^{\mathcal{H}} \tag{B.18}$$

and $\mathcal{R}^{\mathcal{H}}$ compatible.

Because $\mathcal{R}^{\mathcal{H}}$ is compatible, $\mathcal{R}^{\mathcal{H}^*}$ is also compatible using Lemma 5.4.11. By definition of compatibility and compatible refinement we see that $\mathcal{R}^{\mathcal{H}^*}$ compatible implies:

$$\widehat{\mathcal{R}^{\mathcal{H}^*}} \subseteq \mathcal{R}^{\mathcal{H}^*}. \quad (\text{B.19})$$

Using the fact that \mathcal{R}° is symmetric, equations B.17, B.19 and B.18, and the fact that taking the reflexive-transitive closure and the converse of a relation are commutative operations we obtain:

$$\begin{aligned} \mathcal{R}^{\circ} \circ \widehat{\mathcal{R}^{\mathcal{H}^*op}} &= \mathcal{R}^{\circop} \circ \widehat{\mathcal{R}^{\mathcal{H}^*op}} = \mathcal{R}^{\circop} \circ \widehat{\mathcal{R}^{\mathcal{H}^*}op} \subseteq \\ &\mathcal{R}^{\circop} \circ \mathcal{R}^{\mathcal{H}^*op} \subseteq \mathcal{R}^{\mathcal{H}op} \circ \mathcal{R}^{\mathcal{H}^*op} = \mathcal{R}^{\mathcal{H}op} \circ \mathcal{R}^{\mathcal{H}op*} = \mathcal{R}^{\mathcal{H}op*} = \mathcal{R}^{\mathcal{H}^*op}. \end{aligned}$$

This means that $\mathcal{R}^{\mathcal{H}^*op}$ is a solution \mathcal{S} to the inequation $\mathcal{R}^{\circ} \circ \widehat{\mathcal{S}} \subseteq \mathcal{S}$. So the relation $\mathcal{R}^{\mathcal{H}^*op}$ is closed under the rules (HC) and (HV). But $\mathcal{R}^{\mathcal{H}}$ is the least relation closed under those rules. Therefore:

$$\mathcal{R}^{\mathcal{H}} \subseteq \mathcal{R}^{\mathcal{H}^*op}.$$

Consider some terms s and t , values or computations, such that $s \mathcal{R}^{\mathcal{H}^*} t$. Then there exists a sequence of terms s_1, \dots, s_n such that $s_1 = s$ and $s_n = t$ and $s_1 \mathcal{R}^{\mathcal{H}} s_2 \mathcal{R}^{\mathcal{H}} \dots \mathcal{R}^{\mathcal{H}} s_n$. Therefore $s_1 \mathcal{R}^{\mathcal{H}^*op} s_2 \mathcal{R}^{\mathcal{H}^*op} \dots \mathcal{R}^{\mathcal{H}^*op} s_n$.

From here we can deduce $s \mathcal{R}^{\mathcal{H}^*op} t$, which means $t \mathcal{R}^{\mathcal{H}^*} s$. So $\mathcal{R}^{\mathcal{H}^*}$ is symmetric as required. \square

Appendix C

Proofs about Logical Equivalence

Proposition 6.3.1. *Given a decomposable set \mathfrak{P} of Scott-open observations:*

1. *Applicative \mathfrak{P} -similarity, \lesssim , coincides with the logical preorder induced by the logic \mathcal{V}^+ , $\sqsubseteq_{\mathcal{V}^+}$. Therefore, the open extension of $\sqsubseteq_{\mathcal{V}^+}$ is compatible.*
2. *Applicative \mathfrak{P} -bisimilarity, \sim , coincides with the logical equivalence induced by the logic \mathcal{V} , $\equiv_{\mathcal{V}}$. Therefore, the open extension of $\equiv_{\mathcal{V}}$ is compatible.*

Proof. 1. Consider two arbitrary closed computations s and t . Since similarity is the greatest simulation:

$$\begin{aligned} s \lesssim^c t & \text{ iff } \forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P \\ s \sqsubseteq_{\mathcal{V}^+} t & \text{ iff } \forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P. \end{aligned}$$

So we can see that $s \lesssim^c t \iff s \sqsubseteq_{\mathcal{V}^+} t$ as required.

For values v and u assume first that $v \lesssim_A^v u$. We need to prove $v \sqsubseteq_{\mathcal{V}^+} u$, that is:

$$\forall \phi : A \in \mathcal{V}^+. (v \models_{\mathcal{V}^+} \phi \implies u \models_{\mathcal{V}^+} \phi).$$

We proceed by a case distinction on the type A .

Case $A = \text{unit}$. The only formulas of this type are *true*, the empty conjunction, and *false*, the empty disjunction. All values satisfy *true* so in this case we are done. No values satisfy *false* so the implication above holds trivially.

Case $A = \text{nat}$. From the definition of simulation we know that $v = u$. We continue by induction on the formula ϕ .

If $\phi = \{m\}$, assume $v \models_{\mathcal{V}^+} \{m\}$. By definition of satisfaction this means $v = \bar{m}$ so $u = \bar{m}$. Therefore $u \models_{\mathcal{V}^+} \{m\}$ as required.

If $\phi = \bigwedge_{i \in I} \phi_i$ or $\phi = \bigvee_{i \in I} \phi_i$, the result follows from the induction hypothesis for ϕ_i .

Case $A = \neg(B_1, \dots, B_n)$. From $v \lesssim_A^v u$ we know that:

$$\forall \vdash w_1 : A_1, \dots, \vdash w_n : A_n. \forall P \in \mathfrak{P}. \llbracket v(\vec{w}_i) \rrbracket \in P \implies \llbracket u(\vec{w}_i) \rrbracket \in P.$$

We proceed by induction on ϕ .

If $\phi = (w_1, \dots, w_n) \mapsto P$, assume $v \models_{\mathcal{V}^+} (w_1, \dots, w_n) \mapsto P$, which means $\llbracket v(\vec{w}_i) \rrbracket \in P$. Therefore, by assumption $\llbracket u(\vec{w}_i) \rrbracket \in P$ so $u \models_{\mathcal{V}^+} (w_1, \dots, w_n) \mapsto P$ as required.

If $\phi = \bigwedge_{i \in I} \phi_i$ **or** $\phi = \bigvee_{i \in I} \phi_i$, the result follows from the induction hypothesis.

Now assume that for values v and u , $v \sqsubseteq_{\mathcal{V}^+} u$. To show $v \lesssim_A^v u$ we proceed by a case distinction on A .

Case $A = \mathbf{unit}$. The only closed value of type \mathbf{unit} is \star so $v = u = \star$. Since \lesssim is the greatest simulation, this is enough to establish $v \lesssim_{\mathbf{unit}}^v u$.

Case $A = \mathbf{nat}$. Since v is closed $v = \bar{n}$ for some $n \in \mathbb{N}$. Therefore $v \models_{\mathcal{V}^+} \{n\}$. Since $v \sqsubseteq_{\mathcal{V}^+} u$ we have that $u \models_{\mathcal{V}^+} \{n\}$ so $u = \bar{n} = v$ as required.

Case $A = \neg(B_1, \dots, B_n)$. We need to prove that:

$$\forall \vdash w_1 : A_1, \dots, \vdash w_n : A_n. \forall P \in \mathfrak{P}. \llbracket v(\vec{w}_i) \rrbracket \in P \implies \llbracket u(\vec{w}_i) \rrbracket \in P.$$

Assume $\llbracket v(\vec{w}_i) \rrbracket \in P$. Then $v \models_{\mathcal{V}^+} (w_1, \dots, w_n) \mapsto P$ so $w \models_{\mathcal{V}^+} (w_1, \dots, w_n) \mapsto P$. Therefore $\llbracket w(\vec{w}_i) \rrbracket \in P$ as required.

So we have proved $v \lesssim_A^v u \iff v \sqsubseteq_{\mathcal{V}^+} u$.

2. For computations we know that $s \sim^c t$ if and only if:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \iff \llbracket t \rrbracket \in P.$$

So we can see this is equivalent to $s \equiv_{\mathcal{V}} t$.

For values, assume $v \equiv_{\mathcal{V}} u$, that is:

$$\forall \phi : A \in \mathcal{V}. (v \models_{\mathcal{V}} \phi \iff u \models_{\mathcal{V}} \phi). \quad (\text{C.1})$$

We need to prove $v \sim_A^v u$. We proceed by a case split on the type A . The proof is the same as in point 1, except that in the case $A = \neg(B_1, \dots, B_n)$ we need to prove an equivalence. This is done by using the fact that equation C.1 is now an equivalence.

Now assume $v \sim_A^v u$. We need to prove $v \equiv_{\mathcal{V}} u$, that is:

$$\forall \phi : A \in \mathcal{V}. (v \models_{\mathcal{V}} \phi \iff u \models_{\mathcal{V}} \phi).$$

As in point 1, we proceed by a case distinction on the type A .

In case $A = \mathbf{unit}$ we have new formulas apart from *true* and *false* due to the addition of negation. However, they are all semantically equivalent to either *true* or *false*, so we are done.

In cases $A = \mathbf{nat}$ and $A = \neg(B_1, \dots, B_n)$ proceed by induction on $\phi : A$ as before. The \implies direction of the proof is the same. The \iff direction can be obtained from the previous one by choosing $\phi = \neg\phi$. The inductions have an additional case:

If $\phi = \neg\phi'$. By the induction hypothesis we know:

$$v \models_{\mathcal{V}} \phi' \iff u \models_{\mathcal{V}} \phi'$$

which is equivalent to

$$v \not\models_{\mathcal{V}} \phi' \iff u \not\models_{\mathcal{V}} \phi'$$

which is in turn equivalent to

$$v \models_{\mathcal{V}} \neg\phi' \iff u \models_{\mathcal{V}} \neg\phi'.$$

This is what we had to prove. □

Theorem 6.3.4. *Given a decomposable set \mathfrak{P} of Scott-open observations, the logics \mathcal{F} and \mathcal{V} are equi-expressive.*

Proof. We need to prove the same statements as in Theorem 6.3.3, where \mathcal{V}^+ is replaced by \mathcal{V} and \mathcal{F}^+ is replaced by \mathcal{F} . We will point out where the proofs need to be modified.

Statement 1. The proof for computation formulas $P \in \mathfrak{P}$ remains the same because these formulas do not change when adding negation.

For values, we prove by induction on the derivation of $\phi : A$ the following property:

$$\Phi(\phi, A) = (\phi : A \implies (\forall \vdash v : A. v \models_{\mathcal{F}} \phi \iff v \models_{\mathcal{V}} \phi^b)).$$

All the cases from the proof of Theorem 6.3.3 stay the same, but now we have an additional case (NEG).

In this case, $\phi = \neg\phi'$. Assume $\phi : A$. Then we know $\phi' : A$ so we can apply the induction hypothesis to get:

$$\forall \vdash v : A. v \models_{\mathcal{F}} \phi' \iff v \models_{\mathcal{V}} \phi'^b.$$

This is equivalent to:

$$\forall \vdash v : A. v \models_{\mathcal{F}} \neg\phi' \iff v \models_{\mathcal{V}} \neg\phi'^b$$

which is what we had to prove.

Statement 2. For computation formulas the equivalence is proved the same as in Theorem 6.3.3.

For value formulas we proceed by induction on the type A , as in the proof of Theorem 6.3.3. The case $A = \mathbf{nat}$ stays the same. In the $A = \mathbf{unit}$ case, we now have more formulas than *true* and *false* because of the addition of negation. However, all these new formulas are semantically equivalent to *true* and *false*, so their semantics does not change when translated. Therefore, the equivalence we need to prove is true. In the case $A = \neg(B_1, \dots, B_n)$ we do an induction on the formula ϕ .

We observe that $(\sqsubseteq_{\mathcal{F}}) = (\equiv_{\mathcal{F}})$ because \mathcal{F} contains negation. The proof of this goes as follows:

$$\forall \phi \in \mathcal{F}. v \models_{\mathcal{F}} \phi \implies u \models_{\mathcal{F}} \phi$$

implies that

$$\forall \phi' \in \mathcal{F}. v \models_{\mathcal{F}} \neg \phi' \implies u \models_{\mathcal{F}} \neg \phi'$$

so

$$\forall \phi' \in \mathcal{F}. u \models_{\mathcal{F}} \phi' \implies v \models_{\mathcal{F}} \phi'.$$

Using $(\sqsubseteq_{\mathcal{F}}) = (\equiv_{\mathcal{F}})$, we see that $\sqsubseteq_{\mathcal{F}}$ is compatible. Therefore, the proof of the case $\phi = (w_1, \dots, w_n) \mapsto P$ is the same as in Theorem 6.3.3. The cases $\phi = \bigvee_{i \in I} \varphi_i$ and $\phi = \bigwedge_{i \in I} \varphi_i$ remain unchanged.

There is one new case, namely $\phi = \neg \phi'$. By the induction hypothesis for ϕ' we know that:

$$\forall \vdash v : \neg(B_1, \dots, B_n). v \models_{\mathcal{V}} \phi' \iff v \models_{\mathcal{F}} \phi'^{\#}.$$

This is equivalent to:

$$\forall \vdash v : \neg(B_1, \dots, B_n). v \models_{\mathcal{V}} \neg \phi' \iff v \models_{\mathcal{F}} \neg \phi'^{\#}$$

which is what we had to prove. □

Appendix D

Proofs about Contextual Equivalence

Proposition 7.1.3. *The contextual preorder \sqsubseteq_{ctx} is a preorder, and is moreover compatible and \mathfrak{P} -adequate. Thus, it is the greatest compatible and \mathfrak{P} -adequate preorder.*

Proof. First prove that $\sqsubseteq_{ctx} = \bigcup \mathbb{CA}$ is a preorder. To prove reflexivity, we show that the open identity relation, \mathcal{I} , is in \mathbb{CA} . From the compatibility rules we can see \mathcal{I} is compatible. Given $\vdash s \mathcal{I} s$ it follows that $(\forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket s \rrbracket \in P)$ so \mathcal{I} is adequate. Hence, $\mathcal{I} \in \mathbb{CA}$, as required.

To show transitivity it suffices to show that the composition of relations in \mathbb{CA} is itself in \mathbb{CA} . Consider two relations \mathcal{R} and \mathcal{S} which are compatible and adequate.

We can show that $\mathcal{S} \circ \mathcal{R}$ is adequate. Consider $\vdash s \mathcal{R}^c t \mathcal{S}^c r$. Since \mathcal{R} and \mathcal{S} are adequate we know that:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P$$

$$\forall P \in \mathfrak{P}. \llbracket t \rrbracket \in P \implies \llbracket r \rrbracket \in P.$$

Therefore:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket r \rrbracket \in P$$

which means $\mathcal{S} \circ \mathcal{R}$ is adequate.

To prove $\mathcal{S} \circ \mathcal{R}$ is compatible we check each of the rules in the definition of compatibility (Definition 5.3.3) in turn. Since \mathcal{R} and \mathcal{S} are compatible we know $\Gamma \vdash x \mathcal{R}_A^v x$ and $\Gamma \vdash x \mathcal{S}_A^v x$ so $\Gamma \vdash x (\mathcal{S} \circ \mathcal{R})_A^v x$. Therefore, $\mathcal{S} \circ \mathcal{R}$ satisfies (COMP1). Rules (COMP2), (COMP4) and (COMP9) are proved similarly.

Consider $\Gamma, \overrightarrow{x_i} : \overrightarrow{A_i} \vdash s \mathcal{R}^c s' \mathcal{S}^c s''$. Then by compatibility of \mathcal{R} and \mathcal{S} we know that:

$$\Gamma \vdash \lambda(\overrightarrow{x_i}) : (\overrightarrow{A_i}).s \mathcal{R}_{-(\overrightarrow{A_i})}^v \lambda(\overrightarrow{x_i}) : (\overrightarrow{A_i}).s' \mathcal{S}_{-(\overrightarrow{A_i})}^v \lambda(\overrightarrow{x_i}) : (\overrightarrow{A_i}).s''.$$

So

$$\Gamma \vdash \lambda(\overrightarrow{x_i}) : (\overrightarrow{A_i}).s (\mathcal{S} \circ \mathcal{R})_{-(\overrightarrow{A_i})}^v \lambda(\overrightarrow{x_i}) : (\overrightarrow{A_i}).s''.$$

Therefore $\mathcal{S} \circ \mathcal{R}$ satisfies (COMP3). Rule (COMP5) is proved similarly.

Consider $\Gamma, x : \neg(\vec{A}_i) \vdash v \mathcal{R}_{\neg(\vec{A}_i)}^v v' \mathcal{S}_{\neg(\vec{A}_i)}^v v''$ and $\Gamma \vdash w_i \mathcal{R}_{A_i}^v w'_i \mathcal{S}_{A_i}^v w''_i$ for each i . By compatibility of \mathcal{R} and \mathcal{S} we infer that:

$$\Gamma \vdash (\mu x.v)(\vec{w}_i) \mathcal{R}^c (\mu x.v)(\vec{w}'_i) \mathcal{S}^c (\mu x.v)(\vec{w}''_i)$$

so

$$\Gamma \vdash (\mu x.v)(\vec{w}_i) (\mathcal{S} \circ \mathcal{R})^c (\mu x.v)(\vec{w}''_i).$$

Therefore, $\mathcal{S} \circ \mathcal{R}$ satisfies (COMP7). Proving rules (COMP6), (COMP8) and (COMP10) are satisfied is similar.

Now we prove that \sqsubseteq_{ctx} is compatible. We have already shown \sqsubseteq_{ctx} is reflexive so rules (COMP1), (COMP2), (COMP4) and (COMP9) are satisfied.

Suppose $\Gamma, \vec{x}_i : \vec{A}_i \vdash s (\sqsubseteq_{ctx})^c t$. Then there exists a relation $\mathcal{R} \in \mathbb{CA}$ such that $\Gamma, \vec{x}_i : \vec{A}_i \vdash s \mathcal{R}^c t$. By compatibility of \mathcal{R} we know that:

$$\Gamma \vdash \lambda \vec{x}_i : \vec{A}_i . s \mathcal{R}_{\neg(\vec{A}_i)}^v \lambda \vec{x}_i : \vec{A}_i . t$$

so since $\mathcal{R} \subseteq (\sqsubseteq_{ctx})$ we have:

$$\Gamma \vdash \lambda \vec{x}_i : \vec{A}_i . s (\sqsubseteq_{ctx})_{\neg(\vec{A}_i)}^v \lambda \vec{x}_i : \vec{A}_i . t.$$

Therefore, \sqsubseteq_{ctx} satisfies rule (COMP3). Similarly, we can prove \sqsubseteq_{ctx} satisfies (COMP5).

Since we have proved \sqsubseteq_{ctx} is a preorder, we can apply Lemma 5.3.4 to deduce that the compatibility clauses (COMP6), (COMP7), (COMP8) and (COMP10) are equivalent to their single-premise versions. We then use the same reasoning as before to show that these single-premise rules are satisfied.

Now show that \sqsubseteq_{ctx} is adequate. Consider $\vdash s (\sqsubseteq_{ctx})^c t$. Then there exists a relation $\mathcal{R} \in \mathbb{CA}$ such that $\vdash s \mathcal{R}^c t$. So by adequacy of \mathcal{R} we have:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P$$

as required.

Therefore, we have shown that \sqsubseteq_{ctx} is a compatible adequate preorder so we are done. \square

Proposition 7.1.5. *Contextual equivalence is the intersection of the contextual preorder with its converse:*

$$(\equiv_{ctx}) = (\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}.$$

Proof. We prove each inclusion in turn.

“ \subseteq ”. Consider two terms, values or computations, $s \equiv_{ctx} t$. By definition of \equiv_{ctx} there exists a compatible and biadequate relation \mathcal{R} such that $s \mathcal{R} t$. Since \mathcal{R} is biadequate, it is also adequate, so $\mathcal{R} \in \mathbb{CA}$. Thus, $\mathcal{R} \subseteq (\sqsubseteq_{ctx})$ and $s \sqsubseteq_{ctx} t$.

Notice that $(\sqsubseteq_{ctx})^{op} = \bigcup \{ \mathcal{S}^{op} \mid \mathcal{S} \in \mathbb{CA} \}$. If a relation \mathcal{S} is compatible then we see by the definition of compatibility that \mathcal{S}^{op} is also compatible. So $(\sqsubseteq_{ctx})^{op}$ is the union of all relations \mathcal{S} that are compatible and have the property:

$$\forall s', t'. \vdash s' \mathcal{R}^c t' \implies \forall P \in \mathfrak{P}. \llbracket t' \rrbracket \in P \implies \llbracket s' \rrbracket \in P.$$

Because \mathcal{R} is biadequate it has the above property. So $\mathcal{R} \subseteq (\sqsubseteq_{ctx})^{op}$. Therefore, $s (\sqsubseteq_{ctx})^{op} t$. We have shown that $(s, t) \in (\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}$ as required.

“ \supseteq ”. From Proposition 7.1.3 we know \sqsubseteq_{ctx} is compatible so $(\sqsubseteq_{ctx})^{op}$ is also compatible. From the definition of compatibility we can see that the intersection of two compatible relations is also compatible so $(\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}$ is compatible.

Let s and t be two closed computations such that $(s, t) \in (\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}$. Then by adequacy of \sqsubseteq_{ctx} (Proposition 7.1.3) we have:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P$$

$$\forall P \in \mathfrak{P}. \llbracket t \rrbracket \in P \implies \llbracket s \rrbracket \in P$$

which means that $(\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op}$ is biadequate.

Therefore $(\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op} \in \mathbb{CAS}$ so $(\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op} \subseteq (\equiv_{ctx})$. \square

Theorem 7.2.2. *Consider a decomposable set of Scott-open observations \mathfrak{P} that is consistent. Then:*

1. *The open extension of applicative \mathfrak{P} -similarity, \lesssim° , coincides with the contextual preorder, \sqsubseteq_{ctx} .*
2. *The open extension of applicative \mathfrak{P} -bisimilarity, \sim° , coincides with contextual equivalence, \equiv_{ctx} .*

Proof. **We first show** $(\lesssim^\circ) = (\sqsubseteq_{ctx})$. We have shown in Section 7.2 that \lesssim° is included in \sqsubseteq_{ctx} .

Now we need to show $(\sqsubseteq_{ctx}) \subseteq (\lesssim^\circ)$. We first show that \sqsubseteq_{ctx} restricted to closed terms is included in \lesssim , then extend this to open terms. To do this, we show \sqsubseteq_{ctx} restricted to closed terms is a simulation by checking it satisfies the four conditions in the definition of simulation.

1. Assume $\vdash v (\sqsubseteq_{ctx})_{\mathbf{unit}}^v u$. The only closed value of type **unit** is \star so $v = u = \star$ as required.
2. Assume $\vdash v (\sqsubseteq_{ctx})_{\mathbf{nat}}^v u$. This is shown in Section 7.2.
3. Assume $\vdash s (\sqsubseteq_{ctx})^c t$. Because \sqsubseteq_{ctx} is adequate (Proposition 7.1.3) we have the desired result:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P.$$

4. Assume $\vdash v (\sqsubseteq_{ctx})_{-(A_1, \dots, A_n)}^v u$. Consider arbitrary values $\vdash w_1 : A_1, \dots, \vdash w_n : A_n$. Because \sqsubseteq_{ctx} is a preorder it is reflexive so $w_i (\sqsubseteq_{ctx})_{A_i}^v w_i$ for each i . We know \sqsubseteq_{ctx} is compatible so we obtain:

$$v(w_1, \dots, w_n) (\sqsubseteq_{ctx})^c u(w_1, \dots, w_n)$$

as required.

So we have established $(\sqsubseteq_{ctx}) \subseteq (\lesssim)$ for closed terms. Now we need to prove $(\sqsubseteq_{ctx}) \subseteq (\lesssim^\circ)$ in general. To do this, we consider computations and each type of value separately. The cases for natural numbers and computations are shown in Section 7.2 so we only show the cases for type **unit** and function values.

If $\overrightarrow{x_j : A_j} \vdash v (\sqsubseteq_{ctx})_{\mathbf{unit}}^v w$ then each of v and w can either be \star or a variable x_i . In any case, because \star is the only closed value of type \mathbf{unit} , we know that:

$$\forall (\vdash \overrightarrow{u_i : A_i}). v[\overrightarrow{u_i/x_i}] = w[\overrightarrow{u_i/x_i}] = \star.$$

This means that:

$$\forall (\vdash \overrightarrow{u_i : A_i}). \vdash v[\overrightarrow{u_i/x_i}] \lesssim_{\mathbf{unit}}^v w[\overrightarrow{u_i/x_i}]$$

so by the definition of open extension we have that $\overrightarrow{x_j : A_j} \vdash v \lesssim_{\mathbf{unit}}^{\circ, v} w$ as required.

If $\overrightarrow{x_j : A_j} \vdash v (\sqsubseteq_{ctx})_{-(B_1, \dots, B_n)}^v w$ then by compatibility and reflexivity of \sqsubseteq_{ctx} we know that

$$\forall (\vdash \overrightarrow{u_i : B_i}). \overrightarrow{x_j : A_j} \vdash v(\overrightarrow{u_i}) (\sqsubseteq_{ctx})^c w(\overrightarrow{u_i})$$

and again by compatibility:

$$\forall (\vdash \overrightarrow{u_i : B_i}). \vdash \lambda(\overrightarrow{x_j}):(A_j).v(\overrightarrow{u_i}) (\sqsubseteq_{ctx})_{-(A_j)}^v \lambda(\overrightarrow{x_j}):(A_j).w(\overrightarrow{u_i}).$$

Using $(\sqsubseteq_{ctx}) \subseteq (\lesssim)$ we can deduce:

$$\forall (\vdash \overrightarrow{u_i : B_i}). \vdash \lambda(\overrightarrow{x_j}):(A_j).v(\overrightarrow{u_i}) \lesssim_{-(A_j)}^v \lambda(\overrightarrow{x_j}):(A_j).w(\overrightarrow{u_i})$$

so by the definition of \lesssim we know that

$$\forall (\vdash \overrightarrow{p_j : A_j}). \forall (\vdash \overrightarrow{u_i : B_i}). \vdash (\lambda(\overrightarrow{x_j}):(A_j).v(\overrightarrow{u_i}))(\overrightarrow{p_j}) \lesssim^c (\lambda(\overrightarrow{x_j}):(A_j).w(\overrightarrow{u_i}))(\overrightarrow{p_j}).$$

From Lemma 5.2.4 we know reduction preserves similarity so:

$$\forall (\vdash \overrightarrow{p_j : A_j}). \forall (\vdash \overrightarrow{u_i : B_i}). \vdash v[\overrightarrow{p_j/x_j}] (\overrightarrow{u_i}) \lesssim^c w[\overrightarrow{p_j/x_j}] (\overrightarrow{u_i}).$$

By the definition of similarity for function values this means that:

$$\forall (\vdash \overrightarrow{p_j : A_j}). \vdash v[\overrightarrow{p_j/x_j}] \lesssim_{-(B_i)}^v w[\overrightarrow{p_j/x_j}]$$

so by the definition of open extension

$$\overrightarrow{x_j : A_j} \vdash v \lesssim_{-(B_i)}^{\circ, v} w$$

as required.

Now show that $(\sim^\circ) = (\equiv_{ctx})$. We have shown that $(\lesssim^\circ) = (\sqsubseteq_{ctx})$ so:

$$(\lesssim^{op})^\circ = (\lesssim^\circ)^{op} = (\sqsubseteq_{ctx})^{op}$$

because taking the converse of a relation and its open extension are commutative operations. From Proposition 7.1.5 we know that:

$$(\equiv_{ctx}) = (\sqsubseteq_{ctx}) \cap (\sqsubseteq_{ctx})^{op} = (\lesssim^\circ) \cap (\lesssim^{op})^\circ.$$

We can show that $(\lesssim^\circ) \cap (\lesssim^{op})^\circ = ((\lesssim) \cap (\lesssim^{op}))^\circ$. The equation:

$$\overrightarrow{x_i : A_i} \vdash s \lesssim^\circ t \quad \text{and} \quad \overrightarrow{x_i : A_i} \vdash s (\lesssim^{op})^\circ t$$

is equivalent to

$$\forall (\vdash \overrightarrow{u_i : A_i}). \vdash s[\overrightarrow{u_i/x_i}] \lesssim t[\overrightarrow{u_i/x_i}] \text{ and } \vdash s[\overrightarrow{u_i/x_i}] \lesssim^{op} t[\overrightarrow{u_i/x_i}]$$

which in turn is equivalent to $(s, t) \in ((\lesssim) \cap (\lesssim^{op}))^\circ$, as required.

Therefore, we know that:

$$(\equiv_{ctx}) = ((\lesssim) \cap (\lesssim^{op}))^\circ$$

and by Proposition 5.2.3 we know that:

$$(\equiv_{ctx}) = ((\lesssim) \cap (\lesssim^{op}))^\circ = (\sim)^\circ$$

which is what we had to prove. \square

Lemma 7.3.5. *Contextual preorder defined with contexts, \leq_{ctx} , is a compatible and adequate preorder. Hence, it is included in contextual preorder defined coinductively, \sqsubseteq_{ctx} .*

Proof. First prove \leq_{ctx} is a preorder. Given terms $\Gamma \vdash s = t$ we know that for any context C , $\llbracket C[s] \rrbracket = \llbracket C[t] \rrbracket$. Therefore $\Gamma \vdash s \leq_{ctx} t$, so \leq_{ctx} is reflexive. Because implication is transitive we can see that \leq_{ctx} is also transitive.

To prove \leq_{ctx} is adequate consider closed computations $\emptyset \vdash s (\leq_{ctx})^\epsilon t$. Then by definition of \leq_{ctx} , choosing $C_c^\epsilon = [-]^\epsilon$, we know that:

$$\forall P \in \mathfrak{P}. \llbracket s \rrbracket \in P \implies \llbracket t \rrbracket \in P$$

as required.

To prove \leq_{ctx} is compatible show that it satisfies each compatibility rule. Rules (COMP1), (COMP2), (COMP4) and (COMP9) are satisfied by reflexivity.

For rule (COMP3) assume $\Gamma, \overrightarrow{x_i : A_i} \vdash s (\leq_{ctx})^\epsilon t$. Consider an arbitrary context $C_c^\nu : (\Gamma \vdash \neg(\overrightarrow{A_i})) \Rightarrow (\emptyset \vdash)$. Now consider the context:

$$C_c^\epsilon = C_c^\nu[\lambda \overrightarrow{x_i} : \overrightarrow{A_i}. [-]^\epsilon] : (\Gamma, \overrightarrow{x_i : A_i} \vdash) \Rightarrow (\emptyset \vdash).$$

Instantiate the assumption that s and t are in the contextual preorder with C_c^ϵ to deduce that:

$$\forall P \in \mathfrak{P}. \llbracket C_c^\nu[\lambda \overrightarrow{x_i} : \overrightarrow{A_i}. [s]^\epsilon] \rrbracket \in P \implies \llbracket C_c^\nu[\lambda \overrightarrow{x_i} : \overrightarrow{A_i}. [t]^\epsilon] \rrbracket \in P.$$

Then we know $\Gamma \vdash \lambda \overrightarrow{x_i} : \overrightarrow{A_i}. s (\leq_{ctx})_{\neg(\overrightarrow{A_i})}^\nu \lambda \overrightarrow{x_i} : \overrightarrow{A_i}. t$ which is what we had to prove. Rule (COMP5) can be proved similarly choosing:

$$C_c^\nu = C_c^\nu[\text{succ}([-]^\nu)].$$

Using the fact that \leq_{ctx} is a preorder we can apply Lemma 5.3.4 to replace the four compatibility rules that we still need to prove with their single-premise versions. Proving these single-premise rules hold is analogous to proving (COMP3). As an example, we prove rule (COMP7L).

Assume $\Gamma, x : \neg(\overrightarrow{A_i}) \vdash v (\leq_{ctx})_{\neg(\overrightarrow{A_i})}^\nu v'$ and $\Gamma \vdash \overrightarrow{w_i : A_i}$. Consider an arbitrary context $C_c^\epsilon : (\Gamma \vdash) \Rightarrow (\emptyset \vdash)$ and the context:

$$C_c^\nu = C_c^\epsilon[(\mu x. [-]^\nu)(\overrightarrow{w_i})].$$

Using the context typing rules (VV-ID), (VC-MUL) and Lemma 7.3.3 we can deduce that:

$$C_c^v : (\Gamma, x : \neg(\vec{A}_i) \vdash) \Rightarrow (\emptyset \vdash).$$

Therefore we can apply the assumption about v and v' to get:

$$\forall P \in \mathfrak{P}. \llbracket C_c^v[(\mu x.[v]^v)(\vec{w}_i)] \rrbracket \in P \implies \llbracket C_c^v[(\mu x.[v']^v)(\vec{w}_i)] \rrbracket \in P$$

which means $\Gamma \vdash (\mu x.v)(\vec{w}_i) (\leq_{ctx})^c (\mu x.v')(\vec{w}_i)$ as required.

We know that \sqsubseteq_{ctx} is the greatest compatible and adequate relation (Lemma 7.1.3) and we have shown \leq_{ctx} is compatible and adequate. Therefore $(\leq_{ctx}) \subseteq (\sqsubseteq_{ctx})$. \square

Lemma 7.3.6. *Contextual preorder defined coinductively, \sqsubseteq_{ctx} , is closed under program contexts, that is:*

1. *If $\Gamma' \vdash v (\sqsubseteq_{ctx})_A^v u$ and $C_v^v : (\Gamma' \vdash A) \Rightarrow (\Gamma \vdash B)$ then $\Gamma \vdash C_v^v[v] (\sqsubseteq_{ctx})_B^v C_v^v[u]$.*

And the analogous statement for C_c^v .

2. *If $\Gamma' \vdash s (\sqsubseteq_{ctx})^c t$ and $C_c^c : (\Gamma' \vdash) \Rightarrow (\Gamma \vdash)$ then $\Gamma \vdash C_c^c[s] (\sqsubseteq_{ctx})^c C_c^c[t]$.*

And the analogous statement for C_v^c .

Proof. By induction on the typing derivation of C . The two base cases (VV-ID) and (CC-ID) follow from the assumptions $\Gamma' \vdash v (\sqsubseteq_{ctx})_A^v u$ and $\Gamma' \vdash s (\sqsubseteq_{ctx})^c t$ respectively.

In the case (VV-LBD):

$$C_v^v = \lambda \vec{x}_i : \vec{A}_i. \vec{C}_c^v : (\Gamma' \vdash A) \Rightarrow (\Gamma \vdash \neg(\vec{A}_i))$$

$$\text{where } C_c^v : (\Gamma' \vdash A) \Rightarrow (\Gamma, x : \neg(\vec{A}_i) \vdash).$$

Assume $\Gamma' \vdash v (\sqsubseteq_{ctx})_A^v u$. Then by induction hypothesis for C_c^v we know that:

$$\Gamma, x : \neg(\vec{A}_i) \vdash C_c^v[v] (\sqsubseteq_{ctx})^c C_c^v[u]$$

so by compatibility of \sqsubseteq_{ctx} , rule (COMP3) we can deduce that:

$$\Gamma \vdash \lambda \vec{x}_i : \vec{A}_i. \vec{C}_c^v[v] (\sqsubseteq_{ctx})_{\neg(\vec{A}_i)}^c \lambda \vec{x}_i : \vec{A}_i. \vec{C}_c^v[u]$$

which is what we had to prove. Cases (CV-LBD) and (VV-NAT) are analogous.

For the remaining cases, we use the fact that \sqsubseteq_{ctx} is a preorder, so the single-premise compatibility rules from Lemma 5.3.4 hold. The proof then proceeds similarly to the proof of (VV-LBD): apply the induction hypothesis then use one of the single-premise compatibility rules. \square