

University of Southampton Research Repository  
ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON  
FACULTY OF PHYSICAL AND APPLIED SCIENCES  
Electronics and Computer Science

**Behavioural Properties and Dynamic Software Update for Concurrent  
Programs**

by

**Gabrielle Anderson**

Thesis for the degree of Doctor of Philosophy

April 2013



UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES  
Electronics and Computer Science

Doctor of Philosophy

BEHAVIOURAL PROPERTIES AND DYNAMIC SOFTWARE UPDATE FOR  
CONCURRENT PROGRAMS

by **Gabrielle Anderson**

Software maintenance is a major part of the development cycle. The traditional methodology for rolling out an update to existing programs is to shut down the system, modify the binary, and restart the program. Downtime has significant disadvantages. In response to such concerns, researchers and practitioners have investigated how to perform update on running programs whilst maintaining various desired properties. In a multi-threaded setting this is further complicated by the interleaving of different threads' actions. In this thesis we investigate how to prove that safety and liveness are preserved when updating a program. We present two possible approaches; the main intuition behind each of these is to find quiescent points where updates are safe. The first approach requires global synchronisation, and is more generally applicable, but can delay updates indefinitely. The second restricts the class of programs that can be updated, but permits update without global synchronisation, and guarantees application of update. We provide full proofs of all relevant properties.



# Contents

<b>Declaration of Authorship</b>	<b>xi</b>
<b>Acknowledgements</b>	<b>xiii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Safety and Liveness of Concurrent Programs . . . . .	3
1.2 Dynamic Software Update . . . . .	4
<b>2 Background Material and Related Work</b>	<b>5</b>
2.1 Type and Effect Systems . . . . .	5
2.2 Session Typing . . . . .	8
2.3 Dynamic Software Update . . . . .	14
2.3.1 Update Design Decisions . . . . .	14
2.3.1.1 Updating Code . . . . .	15
2.3.1.2 Updating Data . . . . .	16
2.3.1.3 Update Linking . . . . .	17
2.3.1.4 Update Timing . . . . .	17
2.3.2 Safety . . . . .	19
2.3.2.1 Type Safety . . . . .	19
2.3.2.2 Transaction Safety . . . . .	20
2.3.2.3 Update Validity . . . . .	21
2.3.2.4 Assurance Safety By Testing and Verification . . . . .	22
2.3.2.5 No Safety . . . . .	22
2.3.3 Other Considerations . . . . .	23
<b>3 Safety and Liveness Of Concurrent Programs</b>	<b>25</b>
3.1 Motivation: Message Passing Programs . . . . .	26
3.1.1 Communication Errors . . . . .	26
3.1.2 Deadlock . . . . .	28
3.2 Language Definitions . . . . .	29
3.3 Behavioural Abstraction . . . . .	31
3.3.1 Parallel and Local Session Types . . . . .	31
3.3.2 Typing Rules . . . . .	33
3.3.3 Semantics of Effects . . . . .	33
3.4 General Formulation . . . . .	35
3.4.1 Key Lemmas . . . . .	37
3.4.1.1 Subject Reduction . . . . .	37
3.4.1.2 Liveness . . . . .	40

3.4.2	Key Theorems . . . . .	44
3.5	Specific Formulations . . . . .	46
3.5.1	Blocking Message Passing . . . . .	46
3.5.1.1	Safe Program Points . . . . .	47
3.5.1.2	Formal Properties . . . . .	51
3.5.2	Non-Blocking Message Passing . . . . .	56
3.6	Conclusions . . . . .	61
<b>4</b>	<b>Dynamic Software Update</b>	<b>63</b>
4.1	Motivation: Message Passing Programs . . . . .	64
4.2	Update Calculus . . . . .	67
4.2.1	Language Additions . . . . .	67
4.2.2	Updates . . . . .	70
4.2.3	Behavioural Abstraction . . . . .	73
4.3	General Formulation . . . . .	74
4.3.1	Global Update Compatibility . . . . .	74
4.3.1.1	Applying Updates . . . . .	76
4.3.1.2	Introducing Updates . . . . .	76
4.3.2	Key Properties . . . . .	76
4.4	Conclusions . . . . .	79
<b>5</b>	<b>Specific Formulations for DSU</b>	<b>81</b>
5.1	Global Typability . . . . .	82
5.1.1	Formalisation . . . . .	82
5.1.2	Properties and Proofs . . . . .	85
5.1.3	Summary . . . . .	88
5.2	Local Update . . . . .	89
5.2.1	Properties and Proofs . . . . .	94
5.2.2	Summary . . . . .	98
5.3	Conclusions . . . . .	99
<b>6</b>	<b>Conclusions and Future Work</b>	<b>101</b>
6.1	Summary . . . . .	101
6.2	Future Work . . . . .	102
<b>A</b>	<b>General Formulation</b>	<b>107</b>
A.1	Environment Extension Lemma . . . . .	107
A.2	Substitution Lemma . . . . .	111
A.3	Recursive Variable Substitution Lemma . . . . .	117
A.4	Update Equivalence Lemma . . . . .	122
A.5	Expression Update Consistency Lemma . . . . .	125
A.6	Thread Update Consistency Lemma . . . . .	130
A.7	Expression Safety Lemma . . . . .	131
A.8	Safety Theorem . . . . .	134
A.9	Value Effect Lemma . . . . .	135
A.10	General Compatibility Well Definedness Lemma . . . . .	136
A.11	Expression Subject Reduction Lemma . . . . .	138
A.12	Process Subject Reduction Lemma . . . . .	145

A.13 Empty Effect Expression Liveness Lemma . . . . .	147
A.14 Active Effect Expression Liveness Lemma . . . . .	151
A.15 Empty Effect Thread Liveness Lemma . . . . .	158
A.16 Active Effect Thread Liveness Lemma . . . . .	159
<b>B Blocking Message Passing</b>	<b>161</b>
B.1 Separate Channels Lemma . . . . .	161
B.2 Send Action Equivalence Lemma . . . . .	164
B.3 Receive Action Equivalence Lemma . . . . .	166
B.4 Valid States Safety Lemma . . . . .	167
B.5 Valid States Liveness Lemma . . . . .	169
<b>C Non Blocking Message Passing</b>	<b>171</b>
C.1 Non Blocking Send Action Equivalence Lemma . . . . .	171
C.2 Non Blocking Receive Action Equivalence Lemma . . . . .	173
C.3 Valid States Safety Lemma . . . . .	174
C.4 Valid States Liveness Lemma . . . . .	177
<b>D Global Typability and Local Update</b>	<b>181</b>
D.1 Participants Relation as Strict Partial Order . . . . .	182
D.2 Equivalent Projection Give Equivalent Orderings . . . . .	183
D.3 Global Session Type Updatability Lemma . . . . .	184



# List of Figures

2.1	Lambda Calculus . . . . .	5
2.2	Types and Effects . . . . .	6
2.3	Side Effect Analysis . . . . .	6
2.4	Message Passing Language and Resource Syntaxes . . . . .	8
2.5	Message Passing Semantics . . . . .	9
2.6	Types and Effects for Message Passing Language . . . . .	9
2.7	Type and Effect Analysis for Message Passing Language . . . . .	11
2.8	Complementary Relation . . . . .	12
2.9	Global Session Types . . . . .	13
2.10	Global Session Type Projection Function . . . . .	13
3.1	Message Passing Example . . . . .	25
3.2	Safe Blocking Semantics Reduction . . . . .	27
3.3	Non Blocking Semantics Gives Rise to Communication Error . . . . .	28
3.4	Language . . . . .	29
3.5	Operational Semantics . . . . .	30
3.6	Blocking Message Passing Calculus Embedding . . . . .	31
3.7	Types, and Local and Parallel Effects . . . . .	32
3.8	Typing Rules . . . . .	34
3.9	Effect Semantics . . . . .	35
3.10	Label Abstraction . . . . .	35
3.11	Global Compatibility . . . . .	36
3.12	Global Session Types Projection . . . . .	46
3.13	Blocking Message Passing Safe Program Points Set . . . . .	47
3.14	Queues With Type Holes . . . . .	57
3.15	Semantics Of Non Blocking Message Queues . . . . .	57
3.16	Global Session Types Projection For Blocking Message Passing . . . . .	57
3.17	Non Blocking Message Passing Safe Program Points Set . . . . .	59
4.1	Maths Server . . . . .	63
4.2	Updated Maths Server . . . . .	63
4.3	Update Error Due To Update At Mismatched Code Points . . . . .	64
4.4	Producer/Consumer System With Acknowledgement . . . . .	65
4.5	Updated Producer/Consumer System With Acknowledgement . . . . .	65
4.6	Update Deadlock Due To Update At Mismatched Code Points . . . . .	65
4.7	Update Error Due To Inconsistent Update . . . . .	66
4.8	Update Success In Maths Server . . . . .	67
4.9	Update Calculus . . . . .	68

4.10	Update Calculus Operational Semantics	68
4.11	Update Function	69
4.12	updatable region Reduction Example	69
4.13	Example Update Reduction	70
4.14	Update Calculus	72
4.15	Extensions To Typing Rules	73
4.16	Update Extensions To Effect Semantics	73
4.17	Update Function On Effects	75
4.18	Global Update Compatibility	76
5.1	Global Session Types With Update Extensions	82
5.2	Extended Global Session Type Projection Function	82
5.3	Annotated Blocking Message Passing Safe Program Points Set	83
5.4	Update Success In Maths Server Using Global Typability	84
5.5	Buffer System	89
5.6	Infinite Delay Using Round Robin Fair Scheduling	89
5.7	Two Producer/Consumer System Without Acknowledgement	90
5.8	Introducing Deadlock Using Local Update	90
5.9	Update Success Using Local Update	91
5.10	Relaxed Valid Points Sets for Local Update	92
5.11	Definition of Participant Partial Order	93

## Declaration of Authorship

I, **Gabrielle Anderson**, declare that the thesis entitled *Behavioural Properties and Dynamic Software Update for Concurrent Programs* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as: ([Anderson and Rathke, 2009, 2011](#); [Anderson, 2011](#); [Anderson and Rathke, 2012](#))

Signed:.....

Date:.....



## Acknowledgements

The insight that the synchronisation behaviour inherent in blocking message passing programs could be used for co-ordinating updates came fairly early in the course of my doctoral work. One of the most challenging aspects of the work, although being preliminary to the problem of DSU for message passing programs, was how to formalise the synchronisation behaviour in static programs. Existing work on Session Typing provided a general description of the synchronisation behaviour. The way that the available formalisations relied on the synchronisation behaviour, however, was unclear. In order to use the behaviour in the more complex DSU setting, I had to unravel the details of the existing work, and create a new formalisation. One thing that I have learned is that, when building on the work of others, one must either work with the original authors, or budget a lot of time for the preliminary comprehension stage.

For most of the duration of my doctoral work the only people working on problems to do with type systems and static analyses in my university were my supervisor and myself. This meant that I had few opportunities, besides those with Julian, to bounce ideas around. On the occasions that I had the opportunity to attend conferences, workshops, and summer schools, I found the academic discussions that I had to be hugely fulfilling. I very much enjoyed hearing about how other people used such techniques, about the problems they were solving, and about the new ideas they had that I could include in my work. I now am very much of the opinion that having a thriving community in order to cross-pollinate ideas to be an essential part of the research process.

As this period of my life comes to a close, I am grateful for the lessons I've learned, the mistakes I've made, the people who've helped me through them, and the new opportunities on the horizon.



# Chapter 1

## Introduction

Software maintenance is a major part of the development cycle; at one point it was estimated that “60 percent of all business expenditures on computing are for maintenance of software written in Cobol” (Freeman, 1986; Bunker et al., 1993). Another study estimated that 67% of life-cycle costs are in the operations and maintenance phase (Arthur, 1988). The traditional methodology for rolling out an update to existing programs is to shut down the system, modify the binary, and restart the program. Halting a program to apply updates is, however, problematic. Downtime can have significant financial costs, often running into the millions of dollars (Scott, 1998). In addition, in some real-time and critical systems, downtime would break key safety properties, and would prevent updates from occurring at all. As a result, dynamic update of running code is the third most requested enhancement for Java (Sun-Microsystems, 2006). In response to such concerns, researchers and practitioners have investigated performing update on programs while they are still running.

Performing arbitrary updates to programs, at unspecified points during their execution, can lead to arbitrary program behaviour; this negates our attempt to maintain the program’s service and availability while performing an update. In order to maintain well formed programs it is therefore essential to place some restrictions on what updates can occur, and when they can occur. Dynamic Software Update (DSU) comprises a series of approaches that update running programs, while providing a degree of safety, flexibility, and efficiency (Subramanian et al., 2009).

Ensuring good behaviour of concurrent (multi-threaded) programs is notoriously difficult (Lea, 1999). When multiple threads can interleave in each other’s execution, a variety of interesting errors can occur. Data races consist of several processes accessing some shared state in such a way that each can see, and interfere with, intermediary values in each others’ state. Communication errors occur when values are sent over message queues and are subsequently used under incorrect typing assumptions, for example using an integer as an object reference. In addition, interleaving introduces a new potential

problem for concurrent programs: liveness. Deadlock occurs when a program can never proceed, and livelock occurs when part of a program continually executes, performing no useful computation. Concurrent programs often include programming constructs which can either reduce or block, depending on some accompanying state, such as those for locking mutexes or receiving data. In many cases this blocking behaviour is used to eliminate data races or communication errors. There is often a tension between providing safety guarantees and providing guarantees of deadlock absence (Lea, 1999), especially for mutual exclusion; in a program with more blocking it is easier to prove safety, but harder to prove liveness, and vice versa.

The majority of existing work on Dynamic Software Update concerns single-threaded programs (Gupta et al., 1996; Bierman et al., 2003; Boyapati et al., 2003; Stoyle et al., 2005; Neamtiu et al., 2008; Bierman et al., 2008; Kaashoek and Arnold, 2009). These have focused primarily on preserving type safety, both with respect to changing function signatures (Bierman et al., 2003; Stoyle et al., 2005) and changing class signatures (Boyapati et al., 2003). Whilst type safety is important in updates for concurrent programs, and indeed has been thoroughly considered (Neamtiu and Hicks, 2009; Subramanian et al., 2009), the concern of many concurrent program designers is focused elsewhere on behavioural properties such as liveness. In addition, the techniques developed for updating single-threaded programs rely on finding “safe” points during the execution, and applying the update at this point. The extension of this methodology to multi-threaded programs requires waiting until all threads happen to be at a safe point at the same time, or forcing threads to block when they reach a safe point, until all threads are globally synchronised. The result of this approach is delay during synchronisation, at best, and deadlock at worst, and can delay the application of updates indefinitely (Subramanian et al., 2009).

In this thesis we investigate how to prove that safety and liveness are preserved when updating a program. To this end we first consider how to prove these properties for programs in the absence of update (Chapter 3). We focus on programs that use message passing for shared state and synchronisation. We then consider how to update such programs while preserving safety and liveness (Chapter 4). We consider two specific approaches for update. The first is more straightforward and simpler to reason about, but requires global synchronisation. The second performs update using a less restrictive form of synchronisation, with which we can update threads separately; unfortunately this approach is significantly more complex to reason about, and is applicable to fewer programs than the first approach. Finally, we discuss our intuition on how to show that an update will not be delayed indefinitely.

## 1.1 Safety and Liveness of Concurrent Programs

Safety and liveness are two key issues for concurrent programs (Lea, 1999). Safety denotes that nothing bad ever happens. Liveness denotes that something good eventually happens. Safety is important for both single-threaded and multi-threaded programs, as many of the same errors can occur. The addition of concurrency, however, introduces a wide variety of issues for safety, and the issue of liveness.

Message passing is an idiom for concurrent programs, where processes use messages to communicate and synchronise. This obviates the need for direct shared state synchronisation. Session Typing is a formalism used to structure interaction between message passing processes, ranging from low level message passing programs to web orchestration and high level distributed web systems. By abstracting the details of such systems away to the core communication behaviour it is possible to prove a wide variety of properties. These properties include deadlock freedom and absence of communication errors (Honda, 1993; Honda et al., 1998; Bettini et al., 2008), safe re-orderings of communication actions (Gay and Hole, 2005; Mostrous et al., 2009; Deni  lou and Yoshida, 2010), and bounds on channel buffer sizes (Deni  lou and Yoshida, 2010; Gay and Vasconcelos, 2010).

The proofs of properties for message passing programs rely heavily on the synchronisation semantics of the communication actions; a program that does not block when it attempts to receive from an empty message queue has fewer places where it can deadlock, but the increased flexibility and decreased synchronisation makes communication safety more difficult to prove. Existing session typing work focuses on both synchronous and asynchronous communication, with blocking receives and non-blocking sends. The dependencies in the proofs on the communication semantics is often unclear and left implicit. This makes it difficult to re-factor results to message passing languages with even minor differences in communication semantics.

We consider how to prove communication safety and liveness in a general functional language with accompanying state. The general methodology amounts to model checking, and as such is computationally expensive. We then proceed to consider stronger properties that imply general safety, as these stronger properties are often less expensive to prove. We re-prove existing results of safety and liveness for message passing programs that use blocking receive semantics, using a novel and intuitive induction. We then prove safety and liveness for message passing programs that use a non-blocking receive semantics, using a similar induction. We argue that the methodology of a general property, and proving stronger inductive properties that imply it, aids the comprehension of the proofs, and aids the process of writing new proofs when small changes occur to the message passing semantics.

## 1.2 Dynamic Software Update

Safety and liveness are two key issues for concurrent programs. Hence when we update concurrent programs we should consider whether an update violates safety or liveness properties. We present a framework for updating programs written in a functional language with accompanying state based on that presented in Chapter 3. We show how to extend the general safety and liveness properties to accommodate updates. We then show how to prove safety for particular approaches of updates to message passing programs.

**Update Framework:** Updates can be introduced into a running program at any time. This reflects the fact that updates are written and rolled out at some arbitrary point after a program has begun execution. Updates are not necessarily applied as soon as they are introduced - in order to prove safety and liveness we must restrict when updates can occur. An update is only applied after some property of the code and the state becomes true, for example reaching a particular point in the code. An update replaces annotated regions of the code with new bodies of code.

**Global Typability Approach:** This approach consists of a form of global synchronisation. We assume that a program is following an iterative protocol, performing the same communication actions in a loop. We show that if all the threads are in the same iteration of the protocol, it is safe to perform the update. This approach, while requiring global synchronisation, is straightforward to understand and prove safe. We can consider the protocol being used and show how, in certain cases, the window when all threads are updatable occurs infinitely often, and hence an update will not be delayed indefinitely.

**Local Update Approach:** This approach is significantly more complex than the Global approach. We again assume that a program is following an iterative protocol, performing the same communication actions in a loop. The intuition of this approach is as follows. The thread that is at the largest iteration number is updated first. The remaining threads are updated when they reach the iteration number when the lead thread(s) updated. This approach ensures that all messages used under the old protocol are used up. Interestingly, we can perform this approach without actually keeping track of the iteration numbers of the individual threads; this information can be inferred from the code and the state. Using this approach we can update threads separately, without need for global synchronisation or reference execution point of other threads.

# Chapter 2

# Background Material and Related Work

In this chapter we present an overview of background material relevant to this thesis, and its context in the literature. In Section 2.1 we explore Type and Effect Systems: these are static analyses that can be used to abstract the side effecting (*impure*) behaviour of a program such as network or file accesses. In Section 2.2 we describe Session Typing analyses: these place constraints on the impure behaviour of message passing programs that are sufficient to prove various useful properties of concurrent programs such as liveness. In Section 2.3 we explore Dynamic Software Update. We describe common design decisions, with a particular emphasis on how these decisions influence interesting properties such as type safety.

## 2.1 Type and Effect Systems

An example type and effect analysis is one that determines the reference variables that are used in a program, and how they are used (Figure 2.3), as in (Nielson and Nielson, 1999). This information could be useful for static deallocation or for security to guarantee that a program only performs read actions on certain variables. We consider a simple call by value lambda calculus with reference variables, along the lines of ML (Figure 2.1). The language consists of values, functions, function applications, and reference variable constructs. The construct  $\mathbf{new}_\pi x := t_1 \mathbf{in} t_2$  creates a new reference variable  $x$  with an annotation  $\pi$ ; it evaluates  $t_1$  to a value and binds it as  $x$  within  $t_2$ . The

$$t ::= v \mid \lambda x.t \mid t t \mid \mathbf{new}_\pi x := t \mathbf{in} t \mid !x \mid x := t \quad v ::= n \mid b$$

Figure 2.1: Lambda Calculus

$$T ::= \text{INT} \mid \text{BOOL} \mid T \mathbf{ref} \pi \mid T \xrightarrow{\varphi} T \quad \varphi ::= \{\!\pi\!\} \mid \{\pi :=\} \mid \{\mathbf{new} \pi\} \mid \varphi \cup \varphi \mid \emptyset$$

Figure 2.2: Types and Effects

$$\begin{array}{c}
 \frac{}{\emptyset \wr \Gamma \vdash n: \text{INT}} \quad \frac{}{\emptyset \wr \Gamma \vdash b: \text{BOOL}} \quad \frac{}{\emptyset \wr \Gamma \vdash x: \Gamma(x)} \\
 \frac{\varphi \wr \Gamma, x: T_1 \vdash t: T_2}{\emptyset \wr \Gamma \vdash \lambda x. t: T_1 \xrightarrow{\varphi} T_2} \quad \frac{\varphi_1 \wr \Gamma \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \wr \Gamma \vdash t_2: T_2}{\varphi_1 \cup \varphi_2 \cup \varphi_3 \wr \Gamma \vdash t_1 t_2: T_1} \\
 \frac{\Gamma(x) = T \mathbf{ref} \pi \quad \varphi \wr \Gamma \vdash t: T \quad \Gamma(x) = T \mathbf{ref} \pi}{\{\!\pi\!\} \wr \Gamma \vdash !x: T \quad \varphi \cup \{\pi :=\} \wr \Gamma \vdash x := t: T} \\
 \frac{\varphi_1 \wr \Gamma \vdash t_1: T_1 \quad \varphi_2 \wr \Gamma, x: T_1 \mathbf{ref} \pi \vdash t_2: T_2}{\varphi_1 \cup \varphi_2 \cup \{\mathbf{new} \pi\} \wr \Gamma \vdash \mathbf{new}_\pi x := t_1 \mathbf{in} t_2: T_2}
 \end{array}$$

Figure 2.3: Side Effect Analysis

annotation  $\pi$  is unique, and is used to differentiate accesses to this reference variable from accesses to other reference variables. The construct  $!x$  dereferences the variable  $x$ , and  $x := t$  assigns to  $x$  the value obtained by reducing  $t$  to a value. This language has a conventional call by value semantics, and makes use of standard syntactic sugararing such as  $t_1; t_2$  for  $(\lambda x. t_2) t_1$  where  $x$  is not free in  $t_2$

In order to describe the side effects of programs we make use of types and effects (Figure 2.2). The effect  $\{\!\pi\!\}$  denotes the dereference of the variable that has annotation  $\pi$ . The effect  $\{\pi :=\}$  denotes assignment to the variable that has annotation  $\pi$ . The effect  $\{\mathbf{new} \pi\}$  denotes the creation of a new variable that has a new unique annotation  $\pi$ . The primitive types are straightforward. The reference variable type  $T \mathbf{ref} \pi$  denotes a reference variable that contains a value of type  $T$  and has annotation  $\pi$ . The function type includes an effect annotation  $\varphi$  on its arrow, referred to as the *latent effect*; this denotes the side effect that will occur by when the function is applied to an argument. The effect of the function does not appear in the effect of a program until the function is applied, because the body of the function (from which the effect  $\varphi$  comes) is not evaluated until that point.

We make use of the judgement  $\varphi \wr \Gamma \vdash t: T$  which denotes that, under type variable assumptions  $\Gamma$ , the program  $t$  will reduce to a value of type  $T$ , and its side effect is described by the effect  $\varphi$ . The effect  $\varphi$  is generally only a conservative approximation of the effect of a program due to undecidability of determining the actual effect (for example, we do not know which path will be taken at a conditional without running the program, which itself is undecidable). We specify the side effect analysis in Figure 2.3. Values and variables have no side effects. Functions, taken alone, also have no effects:

the effect of the function body is annotated on the type of the function. When a function is applied, the effect of the application construct is the effect  $\varphi_1$  of reducing the left hand side to a function, the effect  $\varphi_2$  of reducing the right hand side to a value, and the effect  $\varphi_3$  of reducing the function body, where  $\varphi_3$  is the latent effect annotated on the type of the function. The dereference of a variable  $x$  has the effect  $\{\!\pi\}$  where  $\pi$  is the annotation associated with  $x$ . Assignment of  $t$  to a variable  $x$  has the effect  $\varphi \cup \{\pi :=\}$ , where  $\pi$  is the annotation associated with the variable, and  $\varphi$  is the side effect of reducing  $t$  to the value which will be assigned to the variable. The creation a new reference variable has the effect  $\varphi_1 \cup \varphi_2 \cup \{\text{new } \pi\}$ , where  $\varphi_1$  is the effect of reducing  $t_1$  to the value assigned to  $x$ ,  $t_2$  is the effect of the body in which  $x$  is bound, and  $\{\text{new } \pi\}$  is the effect of creating a reference variable with annotation  $\pi$ .

An example program that we can analyse is the following:

$$(\lambda y. !y) (\text{new}_\pi x := 1 \text{ in } x) \quad (2.1)$$

This program creates a new reference variable annotated by  $\pi$ , assigns the value 1 to it, and passes it to a function which dereferences it. We can use the typing rules in Figure 2.3 to generate the possible side effect behaviour of the program:

$$\frac{\begin{array}{c} (y: \text{INT ref } \pi)(y) = \text{INT ref } \pi \\ \{\!\pi\} \wr y: \text{INT ref } \pi \vdash !y: \text{INT} \end{array}}{\emptyset \wr \emptyset \vdash \lambda y. !y: \text{INT ref } \pi \xrightarrow{\{\!\pi\}} \text{INT}} \quad \frac{\begin{array}{c} (x: \text{INT ref } \pi)(x) = \text{INT ref } \pi \\ \emptyset \wr x: \text{INT ref } \pi \vdash x: \text{INT ref } \pi \end{array}}{\emptyset \wr x: \text{INT ref } \pi \vdash x: \text{INT ref } \pi}$$

$$\frac{\emptyset \wr \emptyset \vdash \lambda y. !y: \text{INT ref } \pi \xrightarrow{\{\!\pi\}} \text{INT} \quad \{\text{new } \pi\} \wr \Gamma \vdash \text{new}_\pi x := 1 \text{ in } x: \text{INT ref } \pi}{\{\text{new } \pi, !\pi\} \wr \emptyset \vdash (\lambda y. !y) (\text{new}_\pi x := 1 \text{ in } x): T} \quad (2.2)$$

The above program has the side effects of creating a reference variable with the annotation  $\pi$ , and dereferencing of a reference variable with the same annotation, leading to an overall effect of  $\{\text{new } \pi, !\pi\}$ .

This analysis exemplifies the key aspects of type and effect systems. Certain primitives and constructs within the language are annotated with labels, and the remaining primitives have no effect. The effects of non-annotated constructs are constructed by combining the effects of the sub-components. Functions have their effect annotated on their type, and this effect is ignored until the function is applied to an argument.

In this thesis we use type and effect systems to abstract from a source language to an effect representation which is similar to a process calculus. The effect representation denotes both the communication actions that occur, and their temporal ordering. We do not make any contributions to the type and effect literature; we use it as a tool.

$t ::=$	<i>Terms</i>	$v ::=$	<i>Values</i>
$e$	Simple expressions	$n$	Integers
$\text{snd}(c, e).t$	Send	$b$	Booleans
$\text{rcv}(c)(x : T).t$	Receive	$()$	Unit
$\text{if } e \text{ then } t \text{ else } t$	Conditional		
$\text{sel}(c, l).t$	Selection	$\sigma ::=$	<i>Named channel queues</i>
$\text{case } c \text{ of } \{\widetilde{l \mapsto t}\}$	Branching	$c \mapsto q$	Map from channel name to queue
$\mu \underline{X}.t$	Recursion		
$\underline{X}$	Recursion variable	$\sigma, \sigma$	Set of named channel queues
$\text{error}$	Error		
$0$	End	$\emptyset$	Empty set of queues
$e ::= e + e, e \wedge e, \dots$	<i>Simple expressions</i>	$q ::=$	<i>Queues</i>
$P ::=$	<i>Processes</i>	$v$	Value
$t$	Single process	$l$	Label
$P \parallel P$	Parallel processes	$q, q$	Queue composition
		$\emptyset$	Empty queue

Figure 2.4: Message Passing Language and Resource Syntaxes

Type and effect systems are first proposed in (Gifford and Lucassen, 1986) in order to statically track the allocation, access, and deallocation of dynamically allocated memory. This work is expanded and simplified in (Nielson and Nielson, 1999) which shows how type and effect systems can also be used for information control in security, procedure call site analysis, side effect analysis, and simple message passing analysis. Type and effect systems are used many settings, including enforcing locking disciplines in Java (Flanagan and Freund, 2000), in policy-based access control (Bartoletti et al., 2007), and Dynamic Software Update (Neamtiu et al., 2008). A more detailed overview can be seen in (Marino and Millstein, 2008).

## 2.2 Session Typing

Session Typing analyses consist of a type and effect analysis, which obtains the communications behaviour of a program, and various constraints placed on what effects are accepted by the analysis. A key notion in Session Typing is duality, where certain effects fit together if the types of the values and labels which are sent and received are consistent. Systems with certain duality or complementarity properties can be shown to have certain safety and liveness properties, including communication safety, session fidelity, and deadlock freedom.

$$\begin{array}{c}
\frac{[\![e]\!] = v \quad \sigma' = \sigma[c \mapsto q, v]}{[\sigma[c \mapsto q]] \text{ snd}(c, e).t \rightarrow [\sigma'] t} \quad \frac{v: T \quad \sigma' = \sigma[c \mapsto q]}{[\sigma[c \mapsto v, q]] \text{ rcv}(c)(x : T).t \rightarrow [\sigma'] t[v/x]} \\
\frac{v: T' \neq T \quad \sigma' = \sigma[c \mapsto q]}{[\sigma[c \mapsto v, q]] \text{ rcv}(c)(x : T).t \rightarrow [\sigma'] \text{ error}} \quad \frac{e : T \quad T \neq \text{BOOL}}{[\sigma] \text{ if } e \text{ then } t_1 \text{ else } t_2 \rightarrow [\sigma] \text{ error}} \\
\frac{e \downarrow \text{true}}{[\sigma] \text{ if } e \text{ then } t_1 \text{ else } t_2 \rightarrow [\sigma] t_2} \quad \frac{e \downarrow \text{true}}{[\sigma] \text{ if } e \text{ then } t_1 \text{ else } t_2 \rightarrow [\sigma] t_3} \\
\frac{[\sigma[c \mapsto q]] \text{ sel}(c, l).t \rightarrow [\sigma[c \mapsto q, l]] t}{l_i \in \widetilde{l} \quad \sigma' = \sigma[c \mapsto q]} \quad \frac{[\sigma[c \mapsto l_i, q]] \text{ case } c \text{ of } \{\widetilde{l \mapsto t}\} \rightarrow [\sigma'] t_i}{l \notin \widetilde{l}} \\
\frac{[\sigma[c \mapsto l, q]] \text{ case } c \text{ of } \{\widetilde{l \mapsto t}\} \rightarrow [\sigma[c \mapsto l, q]] \text{ error}}{[\sigma] \mu \underline{X}.t \rightarrow [\sigma] t[\mu \underline{X}.t / \underline{X}]} \quad \frac{[\sigma] f(t) \rightarrow [\sigma] t}{[\sigma] P_1 \rightarrow [\sigma'] P'_1} \quad \frac{[\sigma] P_2 \rightarrow [\sigma'] P'_2}{[\sigma] P_1 \parallel P_2 \rightarrow [\sigma'] P'_1 \parallel P_2} \\
\frac{[\sigma] P_1 \parallel P_2 \rightarrow [\sigma'] P'_1 \parallel P_2}{[\sigma] P_1 \parallel P_2 \rightarrow [\sigma'] P_1 \parallel P'_2}
\end{array}$$

Figure 2.5: Message Passing Semantics

$$\begin{array}{l}
\varphi ::= c!(T); \varphi \mid c?(T); \varphi \mid \varphi \oplus \varphi \mid c\&\{\widetilde{l \mapsto \varphi}\} \mid \mu \underline{X}.\varphi \mid \underline{X} \mid 0 \\
\Phi ::= \varphi \mid \Phi \parallel \Phi \quad T ::= \text{INT} \mid \text{BOOL} \mid \text{UNIT}
\end{array}$$

Figure 2.6: Types and Effects for Message Passing Language

In order to discuss Session Typing we present a simple message passing system. We define the language in Figure 2.4. The transmission of values is performed by  $\text{snd}(c, e).t$ , which reduces  $e$  to a value, adds that value onto the end of the queue for channel  $c$ , and continues with  $t$ . The reception of values is performed by  $\text{rcv}(c)(x : T).t$ , which takes the value off the front of the message queue for  $c$ , and substitutes it for  $x$  in  $t$ . Conditional statements  $\text{if } e \text{ then } t \text{ else } t$  reduce the pure (non side-effecting) guard, and then branches depending on the value. Selection statements  $\text{sel}(c, l).t$  is used to signal a choice in protocol direction, in particular, which option a branch (that exists in another term with which the selecting term is interacting) should take. These put the relevant label on the end of the message queue for  $c$ . Branching statements  $\text{case } c \text{ of } \{\widetilde{l \mapsto t}\}$  provides a series of options between which another term can choose. These take the value off the front of the message queue for  $c$ , and branch over the value received. Recursion, errors, and blocked terms are standard. We specify the semantics formally in Figure 2.5.

Choice primitives can be used to specify more complex communication patterns, such as the following ATM example (Carbone et al., 2007):

$$\begin{aligned}
 & c_1!(\text{deposit}).c_1!(\text{amount}).c_1!(300).c_2!(\text{balance} \mapsto c_2?(x).()). \parallel \\
 & c_1\&\{\text{deposit} \mapsto (c_1\&\{\text{amount} \mapsto c_1?(x).c_2!(\text{balance}).c_2!(50+x).().\}), \\
 & \quad \text{withdraw} \mapsto c_1\&\{\text{amount} \mapsto c_1?(y).\text{if } (x \leq 50) \text{ then } c_2!(\text{dispense}).(). \\
 & \quad \quad \quad \text{else } c_2!(\text{overdraft}).(). \\
 & \quad \} \\
 & \}
 \end{aligned} \tag{2.3}$$

This example consists of two participants, a user (the first thread) and the ATM (the second thread). The ATM presents two services to the user: depositing money into the account, which returns the balance after the deposit, and withdrawing money using the ATM, both assume there is £50 in the account before any action. In this example the user chooses to deposit £300 into the account. The offering of different services, and the use of services to annotate the semantic meaning of values being sent back and forth, demonstrates the complexity of communications behaviour which can be represented using Session Types. Many Session Typing systems also include the delegation communication primitive; this permits channels (and the responsibilities entailed by the Session Typing analysis, below) to be sent over channels. For simplicity we do not consider delegation.

We represent the communication behaviour of a program by effects (which are referred to as Local Session Types) (Honda, 1993; Mostrous et al., 2009; Bettini et al., 2008; Yoshida and Vasconcelos, 2006; Honda et al., 2008; Gay et al., 2003; Gay and Vasconcelos, 2007). These are defined formally in Figure 2.6. The structure follows that of the message passing terms. Send and receive actions describe the transmission and reception of a value of type  $T$  on channel  $c$ , followed by the continuation  $\varphi$ . Non-deterministic choice between possible effects  $\varphi_1$  and  $\varphi_2$  is written as  $\varphi_1 \oplus \varphi_2$ . External choice, based on another participant's choice, is written as  $c\&\{\widetilde{l \mapsto t}\}$ . Recursion, errors, and blocked terms are standard. We define a simple type and effect analysis for message passing systems in Figure 2.7 to abstract a program's communication behaviour.

Communication Errors occur when a thread receives and uses a value which is of a type that it doesn't expect. An example reduction sequence with an error is:

$$\begin{aligned}
 1 & [c \mapsto \emptyset] \text{ snd}(c, \text{true}).0 \parallel \text{rcv}(c)(x : \text{INT}).x + 4 \\
 2 & \rightarrow [c \mapsto \text{true}] 0 \parallel \text{rcv}(c)(x : \text{INT}).x + 4 \\
 3 & \rightarrow [c \mapsto \emptyset] 0 \parallel \text{true} + 4 \\
 4 & \rightarrow [c \mapsto \emptyset] 0 \parallel \text{error}
 \end{aligned} \tag{2.4}$$

$$\begin{array}{c}
\frac{}{0 \wr \Gamma \vdash n: \text{INT}} \quad \frac{}{0 \wr \Gamma \vdash b: \text{BOOL}} \quad \frac{}{0 \wr \Gamma \vdash (): \text{UNIT}} \quad \frac{}{0 \wr \Gamma \vdash x: \Gamma(x)} \\
\frac{0 \wr \Gamma \vdash e: T_1 \quad \varphi \wr \Gamma \vdash t: T}{c!(T_1); \varphi \wr \Gamma \vdash \text{snd}(c, e).t: T} \quad \frac{}{c?(T_1); \varphi \wr \Gamma \vdash \text{rcv}(c)(x: T).t: T} \\
\frac{\varphi_i \wr \Gamma \vdash t_i: T}{\varphi_1 \oplus \varphi_2 \wr \Gamma \vdash \text{if } e \text{ then } t_1 \text{ else } t_2: T} \quad \frac{\varphi_i \wr \Gamma \vdash t_i: T}{c!(l); \varphi \vdash \text{sel}(c, l).t: T} \\
\frac{\varphi_i \wr \Gamma \vdash t_i: T}{c\&\{\widetilde{l \mapsto \varphi}\} \wr \Gamma \vdash c\&\{\widetilde{l \mapsto t}\}: T} \quad \frac{\varphi \wr \Gamma \vdash t: T}{\mu \underline{X}.\varphi \wr \Gamma \vdash \mu \underline{X}.t: T} \\
\frac{}{X \wr \Gamma \vdash X: T} \quad \frac{\varphi \wr \emptyset \vdash t: T}{\vdash t: \varphi} \quad \frac{\vdash P_1: \Phi_1 \quad \vdash P_2: \Phi_2}{\vdash P_1 \parallel P_2: \Phi_1 \parallel \Phi_2}
\end{array}$$

Figure 2.7: Type and Effect Analysis for Message Passing Language

Session Typing analyses can be used to rule out this type of error. These analyses can be simplified if they only need consider the effect of a program, rather than the entire program itself. Hence Session Typing analyses require that the behaviour of a program is completely encompassed in that program’s effect. In other words, whenever a program can perform a specific communication action, the effect of that program, as generated by the type and effect system, can perform an abstract version of the action. This property known as *fidelity* (Honda et al., 2008; Dezani-Ciancaglini and Liguoro, 2010). For example, an action  $c!(e)$ . is abstracted by  $c!(T)$ . where  $e: T$ .

The program in Example 2.4 can be represented using the effect:

$$c!(\text{BOOL}); 0 \parallel c?(\text{BOOL}); 0 \quad (2.5)$$

This denotes that the program consists of two threads, one which will send a boolean value on a channel, and one which will receive a value from the same channel that it expects to be an integer.

The two threads’ behaviour will obviously interact to cause an error, as the type being sent and the type expected to be received do not match up. We can formalise our intuition as to when a set of threads’ behaviour fits together; when they do so their behaviour is called *complementary*. We formally define complementarity in Figure 2.8, where channels are assumed to be used in a uni-directional manner by two participants.

The Communication Safety property denotes that a thread will never receive a value which is of a different type to that which it is expecting. When a message passing system conforms to the above assumptions, and when all pairs of effects (from the two participants) on channels are complementary, it is possible to prove communication safety (Honda, 1993; Honda et al., 1998; Gay et al., 2003; Gay and Vasconcelos, 2007). A

$$\begin{aligned}
\text{compl}(c!(T_1); \varphi_1, c?(T_1); \varphi_2) &\stackrel{\text{def}}{=} T_1 = T_2 \wedge \text{compl}(\varphi_1, \varphi_2) \\
\text{compl}(\bigoplus_I c!(T_1); \varphi_i, c\&\{l_j \mapsto \phi'_j\}_J) &\stackrel{\text{def}}{=} \forall i \in I \exists j \in J. l_i = l_j \wedge \text{compl}(\varphi_i, \varphi'_j) \\
\text{compl}(0, 0) &= \text{true}
\end{aligned}$$

Figure 2.8: Complementary Relation

basic Session Typing analysis can hence be seen as a type and effect analysis along with some restrictions on the permitted communication effects of message passing programs.

The concept of complementarity only fulfils our intuition for interactions which consist of two participants. Global Session Types generalise the notion of complementarity to describe how the interactions of an arbitrary number participants relate to each other (Honda et al., 2008). We define a simplified version of Global Session Types in Figure 2.9. The type  $d_1 \rightarrow d_2: c\langle T \rangle; G$  denotes that participant  $d_1$  sends a value of type  $T$  over channel  $c$  to participant  $d_2$ , and then continues with the behaviour specified in  $G$ . This describes the duality of sends and receives expressed in complementarity. The type  $d_1 \rightarrow d_2: \widetilde{c\langle T \mapsto G \rangle}; G$  denotes that participant  $d_1$  can make a choice (invoke a service) that is offered by participant  $d_2$ . This describes the duality of offering a service and service selection. The continuation for each construct may contain behaviour of participants other than  $d_1$  and  $d_2$ , and hence we can represent temporal information about the global communications behaviour. Consider the following Global Session Type:

$$\begin{aligned}
d_1 \rightarrow d_2: c\langle \text{INT} \rangle; \\
d_2 \rightarrow d_3: c\langle \text{BOOL} \rangle; \\
0
\end{aligned} \tag{2.6}$$

This denotes that participant  $d_2$  will receive an integer from  $d_1$ , and then proceed to send a boolean to  $d_3$ . Using the global description of the communications behaviour it is possible to obtain a thread local description of the communications behaviour (Honda et al., 2008). We define the projection function  $G \upharpoonright d$ , which generates the thread local behaviour described by a global protocol, in Figure 2.10. Global Session Types can be used to prove communication safety in systems with an arbitrary number of participants.

The earliest work on Session Typing focusses on two-party, synchronous communication (Honda, 1993). Honda considers how to rule out communication errors, how to rule out deadlock, and how to define bisimulation for Session Typing systems. Honda, Kubo, and Vasconcelos extend Honda's earlier work and describe how to represent method invocation, unbounded interaction patterns, and delegation using Session Typing (Honda et al., 1998). Some errors in (Honda, 1993; Honda et al., 1998) have been discovered and corrected (Gay and Hole, 2005; Yoshida and Vasconcelos, 2006).

$$G ::= d_1 \rightarrow d_2 : r\langle T \rangle; G \mid d_1 \rightarrow d_2 : r\langle \widetilde{l \mapsto G} \rangle \mid \mu \underline{X}.G \mid \underline{X} \mid 0$$

Figure 2.9: Global Session Types

$$\begin{aligned}
d_1 \rightarrow d_2 : c\langle T \rangle; G' \upharpoonright d_1 &\stackrel{\text{def}}{=} c!\langle T \rangle.G' \upharpoonright d_1 \\
d_1 \rightarrow d_2 : c\langle T \rangle; G' \upharpoonright d_2 &\stackrel{\text{def}}{=} c?(T).G' \upharpoonright d_2 \\
d_1 \rightarrow d_2 : c\langle \widetilde{l \mapsto G} \rangle_I \upharpoonright d_1 &\stackrel{\text{def}}{=} \bigoplus_I c!\langle T_1 \rangle; \varphi_i \\
d_1 \rightarrow d_2 : c\langle \widetilde{l \mapsto G} \rangle_I \upharpoonright d_2 &\stackrel{\text{def}}{=} c\&\{l \mapsto (G \upharpoonright d_2)\}_I \\
(\mu \underline{X}.G) \upharpoonright d &\stackrel{\text{def}}{=} \mu \underline{X}.(G \upharpoonright d) \\
\underline{X} \upharpoonright d &\stackrel{\text{def}}{=} \underline{X} \\
0 \upharpoonright d &\stackrel{\text{def}}{=} 0
\end{aligned}$$

Figure 2.10: Global Session Type Projection Function

Whilst most work on Session Typing has made use of  $\pi$  calculus style calculi (Honda, 1993; Mostrous et al., 2009; Bettini et al., 2008; Yoshida and Vasconcelos, 2006; Honda et al., 2008), in some a  $\lambda$  calculus formulation is used (Gay et al., 2003; Gay and Vasconcelos, 2007, 2010). In (Gay et al., 2003) Gay, Vasconcelos, and Ravara provide subtyping for Session Types. In (Gay and Vasconcelos, 2007) Gay and Vasconcelos include a linear type system. Additionally, they crucially introduce Session Typing for asynchronous systems.

Honda, Yoshida and Carbone generalise Session Typing to sessions that can include more than two parties (multi-party Session Typing) in a  $\pi$  calculus setting (Honda et al., 2008). They introduce the concept of a Global Session Type, which represents communications protocols between more than two participants. The global view can be projected down to the local view of each role and how it should communicate with the other roles. The communications mechanisms are also expanded to include multicast sending. They extend work on delegation, where an entity can pass responsibility for performing certain actions within a session on to another entity. The formulation of delegation, however, is such that the delegation is not completely transparent; it exposes implementation details of the delegator to the delegate which could be either technically or commercially undesirable. Bettini, Cocco, Luca, and Dezani-Ciancaglini solve the above limitations and provide the first robust, multi-party, Session Typing system (Bettini et al., 2008). They also guarantee global progress, the property that well typed programs will not deadlock, and specifically that different global sessions will not interfere with and deadlock each other.

Subtyping for Session Types is introduced in (Gay and Hole, 2005) which permits Session Type specifications to be refined into a richer behaviour. The subtyping is extended in (Mostrous et al., 2009) to thread local asynchronous subtyping, which permits the send

and receive actions performed by a thread to be re-ordered in certain circumstances. A simple example of this reordering is to permit moving sends in front of receives, which increases the asynchrony of the interleaved behaviour whilst maintaining the safety and liveness properties. In (Deniéou and Yoshida, 2010; Gay and Vasconcelos, 2010), it is shown how buffers with unbounded size can be identified, and how buffer bounds can be extracted for bounded buffers, from the Global Session Type representation. It is also shown how to perform thread local asynchronous subtyping without changing buffer bounds.

A more detailed review of Session Typing literature can be found in (Dezani-Ciancaglini and Liguoro, 2010).

## 2.3 Dynamic Software Update

The principle aim of Dynamic Software Update is to reduce the disruption that normally occurs when programs need to be updated. When a program needs to be updated, the traditional approach is to schedule some down time for the service, to take down the program, update it, and restart it. Downtime can have financial costs (Scott, 1998) and also presents concerns in safety critical systems.

If a program or service is stateless, such as an http server, a service on redundant hardware can be used to handle requests whilst the primary service is updated (Stoyle, 2006). If state or interaction behaviour (such as a communications protocol) needs to be preserved over the update, however, using redundant hardware to maintain a service preserves the problem of new programs interacting with old state or protocols (Stoyle, 2006). Hence the primary concern of Dynamic Software Update is how to co-ordinate the old and new code interacting with the old and new data.

In this section we describe the key design decisions surrounding Dynamic Software Update, and consider the implications of said choices, with a particular emphasis on how they influence safety properties. We begin by detailing the design decisions concerning the semantics of a system that can be dynamically updated (Section 2.3.1). We then consider safety for DSU, and how the prior design decision influence what safety properties can be demonstrated (Section 2.3.2). We conclude with other important considerations, including ease of use, efficiency, and case studies (Section 2.3.3).

### 2.3.1 Update Design Decisions

The semantics of Dynamic Software Update can be described in several key areas. Updating code and data are intimately interlinked. Update linking is important to implement code and data updating. Finally, update timing considers when an update is

applied, and how the structure of code can influence when it is practical to perform an update.

### 2.3.1.1 Updating Code

The most straightforward approach to updating code is to unload the entire binary of the old program and to load the entire binary of the new program. In order to structure and reason about updates code update is, for the most part, performed at a finer level of granularity, that of the function or method.

There are two principal approaches that can be taken when replacing a function (respectively method). The most common approach is to permit an update to replace the body of the function, but preserve (or refine/sub-type) the signature of the function (the types of the parameters, and the type of the return value) ([Gilmore et al., 1998](#); [Hjálmtýsson and Gray, 1998](#); [Orso et al., 2002](#); [Bierman et al., 2008](#); [Kaashoek and Arnold, 2009](#)). Whilst this approach might seem restrictive, it is sufficient to cover the majority of modifications ([Neamtiu et al., 2005](#); [Subramanian, 2010](#)). In particular it covers many security related updates ([Kaashoek and Arnold, 2009](#)); as delays to security updates until scheduled downtime leaves a large window of vulnerability, the ability to perform security updates without downtime is significant. The alternative is that an update can replace a function, and change the signature of the function, or even delete the function in its entirety. This approach is obviously more flexible, and is required for a significant minority of updates ([Neamtiu et al., 2005](#); [Subramanian, 2010](#)).

In order to update functions, we need a way to access and manipulate them. This is primarily achieved using indirection ([Subramanian et al., 2009](#); [Hayden et al., 2011d](#), 2009; [Neamtiu and Hicks, 2009](#); [Boyapati et al., 2003](#); [Sewell, 2001](#); [Cook and Lee, 1983](#); [Appel, 1994](#); [Kaashoek and Arnold, 2009](#); [Neamtiu et al., 2006](#); [Baumann et al., 2007](#)). Managed languages perform indirection as a matter of course, performing a table lookup on a function call to find the most relevant, or most recently compiled, version of a function. This functionality can be leveraged by rewriting the lookup table to point to the updated function body ([Boyapati et al., 2003](#); [Bierman et al., 2008](#); [Subramanian et al., 2009](#)). In non-managed languages such as C/C++, this indirection is often implemented on top of the program ([Neamtiu et al., 2006](#)), adding overhead during normal execution. An alternate approach is, when updating a function, to overwrite the first few assembly instructions of the function with a jump to the new definition; this approach has been implemented for the Linux kernel, and operates with almost no overhead ([Kaashoek and Arnold, 2009](#)).

It is also possible to update existing programs by directly manipulating functions that are on the stack ([Makris and Bazzi, 2009](#)). This approach uses annotation and special

compilation to create threads that can be interrupted and updated at runtime. This approach relies on specialist knowledge from the programmer, who has to write a program that transforms the existing stack frame.

The code of a program need not be updated all at once. Indeed, most Dynamic Software Update approaches that use indirection modify the function definition, and next time that function is called the new version is used, but does not modify any instances of the function that are on the stack. This approach can result in old code and new code existing in the same program.

An alternative to replacing individual functions is to load an entirely new binary into memory, and to start at some predefined point in that program's execution (Hayden et al., 2011b,d). The ability to resume execution at points other than the start of the new program permits designers to, for example, skip initialising variables, which would erase existing data which we want to keep. Hayden et al. argue that the approach of updating individual functions requires complex tool support, is challenging to reason about, and can add overhead due to the indirection discussed above. Whether the complexity and overhead required in more complex analyses is justified is an open research topic.

### 2.3.1.2 Updating Data

Unless application data or interactions with other programs needs to be preserved over the update application, it is easier to update a service using redundant hardware. Hence programs where the old and new code interacts with the old and new data, and the co-ordination thereof, are the key concerns and research areas for Dynamic Software Update.

An updated program may use data with a different format to that of the previous version, for example a record could be changed from `{FirstName : String, LastName : String}` to `{Name : String, Age : Integer}` (Gurtner, 2011). In order for the new code to interact with the old data, it must be transformed into the new format before the new code accesses it. In addition we do not want the transformation to throw away the old data; a transformation that initialises all data to a default value is comparable to stopping the program and restarting, without performing DSU (Subramanian et al., 2009). One possible transformation function would take a record in the old format, create a record in the new format, concatenate the `FirstName` and `LastName` fields, put the in the `Name` field, place a default value in the `Age` field, and return the new record. In systems where the old and new code co-exist, and the data can be accessed by both, we similarly need a transformation function from the new to the old format, to be applied before the old code uses a new data value (Gurtner, 2011; Wernli et al., 2011). The state transformers can be automatically generated using initialisation of default values or copying (Subramanian et al., 2009), or written by the programmer when more complex transformations are

required (Bierman et al., 2003; Stoyle et al., 2005; Soules et al., 2003; Hayden et al., 2011d).

### 2.3.1.3 Update Linking

Linking code to data at runtime is known as dynamic linking (Cardelli, 1997; Drossopoulou and Eisenbach, 2002). While the technology for linking identifiers to code and data exists, managing those identifiers, and co-ordinating when they can be updated, is not straightforward. Consider the following examples (Stoyle, 2006), where  $\pi_1(x, x')$  denotes choosing the first item in a pair and  $\pi_2(x, x')$  denotes choosing the second.

$$\text{let } x = (4, 5) \text{ in } \pi_1(\pi_1(x, x))x \quad (2.7)$$

Using a call by value reduction semantics this term could reduce as follows:

$$\text{let } x = (4, 5) \text{ in } \pi_1(\pi_1(x, x)) \rightarrow \pi_1(\pi_1(4, 5), (4, 5)) \quad (2.8)$$

After this reduction it is no longer meaningful to discuss updating  $x$ . In order to maintain updatable identifiers as long as possible Stoyle shows how to delay substitution but maintain contextual equivalence (Stoyle, 2006). These techniques can be used to administrate re-linking at runtime and to increase the scope of where updates can be applied.

### 2.3.1.4 Update Timing

Updatable programs typically make use of specialised API calls to apply updates, possibly performing some dynamic checks beforehand to ensure the update is safe (Hayden et al., 2011d, 2009; Neamtiu and Hicks, 2009; Boyapati et al., 2003; Sewell, 2001; Kaashoek and Arnold, 2009; Neamtiu et al., 2006; Appel, 1994). Explicit language calls to an update API requires the programmer to have a working knowledge of update and requires them to reason about when would be a good time to perform updates, a significant addition in complexity when writing a program. In some cases the API calls can be automatically inserted (Stoyle et al., 2005).

One other option is to provide some external update manager that introduces the update without any knowledge of the original program (Baumann et al., 2007; Subramanian et al., 2009). The manager monitors the program to be updated, waiting for a suitable update points. One way to do this is to insert breaks (return barriers) at potential update points that return control to the monitor. Potential points include function entry and exit points, and loop back-edges (Subramanian et al., 2009).

When code is replaced using indirection, old code in existing stack frames is usually left unchanged. In such cases, irrespective of the point when the function bodies are replaced, the update doesn't really occur until the old code has returned.

In order for old code to continue to execute safely it must either not use any new data, or any new data must be transformed into an old format before use. The reverse is true for new code and old data. In order to simplify updates many systems employ *Representation Consistency* (Stoyle et al., 2005). This denotes that, after an update data is only in the new format. The conversion from old to new data can either be performed eagerly (Stoyle et al., 2005; Sewell, 2001; Subramanian et al., 2009) or lazily (Appel, 1994; Boyapati et al., 2003; Neamtiu et al., 2006; Gurtner, 2011). When updates are performed lazily, the cost of performing the update may be amortised over program execution. The disadvantage of laziness is that it requires monitoring the format of data, which adds steady-state overhead. Eager updates can be implemented without steady state overhead. The disadvantage eagerness is that in order to update all the existing data a stop-the-world must be employed.

Programs that benefit most from dynamic updating are typically structured as long-running event processing loops (Subramanian, 2010). In order to make use of the function body replacement paradigm, long running loops can be extracted into recursive functions, so that a main loop can run to completion using the old code, and that the next time the loop starts (the next call of the loop function) the new code will be used. This can be done automatically (Neamtiu et al., 2006).

Multi-threaded programs provide a particular challenge for existing DSU systems. If updates are performed when the program reaches update point (either explicit or annotated by some monitor) as each thread can reduce independently there is no guarantee that all the threads will reach safe points at the same time. It is possible to block one thread when it reaches a safe point in the hope that the other threads will reach safe points. Other threads could, however, be reliant on the blocked thread to make progress. In some approaches the blocked thread is released after a certain amount of time (Neamtiu and Hicks, 2009). This technique can introduce overhead and delay of an update in the best case, and deadlock to an otherwise live program in the worst case. It is possible to require that no locks are held at update points, in order to try to avoid deadlock (Hayden et al., 2011b). In initial experiments this has not introduced deadlock, but this property has not been proved formally. In addition, stack transformation, if properly designed, can be immediately applied to multi-threaded programs without risk of introducing update-related deadlock (Makris and Bazzi, 2009).

### 2.3.2 Safety

While it is possible to naively modify any running program in arbitrary ways, the challenge is to do so in a way that leaves a valid updated program, for some definition of valid. The primary definition considered in existing research is type safety (Stoye, 2006; Bierman et al., 2008; Hjálmtýsson and Gray, 1998; Makris and Bazzi, 2009; Gartner, 2011; Subramanian, 2010; Neamtiu and Hicks, 2009; Boyapati et al., 2003; Appel, 1994; Neamtiu et al., 2006; Stoye et al., 2005). Other safety issues include transaction safety (Neamtiu et al., 2008; Neamtiu and Hicks, 2009), update validity (Gupta et al., 1996), assurance safety by testing (Hayden et al., 2009), and verification of specifications (Hayden et al., 2011a).

#### 2.3.2.1 Type Safety

Type safety is concerned with the prevention of applications from performing illegal operations (those that are disallowed in normal execution). Simple examples of these include performing integer arithmetic on boolean values, or attempting to call a function that doesn't exist. In order to prevent type errors we must co-ordinate how new and old code interact with each other, and with the data. This problem can be simplified by making use of representation consistency, the requirement that data or objects only ever be in one format.

One method of simply showing type safety is to require that any function that is to be updated is not on the call stack when the update is applied. This approach, known as *activeness safety*, neatly partitions old from new code, ensuring that at any one time code and data from only one version are being executed. Activeness safety is one of the most common approaches in DSU (Baumann et al., 2007; Soules et al., 2003; Baumann et al., 2005; Subramanian, 2010; Altekar et al., 2005; Kaashoek and Arnold, 2009) and can be implemented with a simple dynamic check. The requirement that no function that is to be updated is on the stack can lead to an update being delayed for an unbounded time (Subramanian, 2010).

If we permit active functions to be updated, and do not perform stack transformation on the code of these active functions, then the old code could access data that has been transformed into a new format. One solution (known as *con-freeness safety*) is to ensure that the remaining old code (that to be executed before a new function body is called) does not make use of the *format* of the data (referred to as using it *concretely*). Consider the following example:

$$\text{let } x = (4, 5) \text{ in } \pi_1(x, x) \quad (2.9)$$

If  $x$  were updated to an integer value (here 9) rather than a pair the above code would still be safe, as it does not directly use the fact that  $x$  is a pair:

$$\text{let } x = 9 \text{ in } \pi_1(x, x) \rightarrow \pi_1(9, 9) \rightarrow 9 \quad (2.10)$$

If the code made direct use of the structure, however, changing to an integer value would lead to an error:

$$\text{let } x = 9 \text{ in } \pi_1(\pi_1(x), x) \rightarrow \pi_1(\pi_1(9), 9) \rightarrow \pi_1(9) \rightarrow \text{error} \quad (2.11)$$

The original work states that “if code simply passes data around without relying on its representation, then updating that data poses no problem” (Stoyle et al., 2005). We can see that, in Equations 2.9 and 2.10, that the variable  $x$  is never used in a way that makes use of its tuple structure (it is only used to construct a new data structure), and hence the usage there is not concrete. In Equation 2.11, however, the tuple structure is used (by the splitting action  $\phi_1(x)$ ), and hence the usage there is concrete. When some code does not use a data type  $T$  concretely it is referred to as  $T$ -concreteness free, or  $T$ -con free. It is possible to automatically detect regions of code that are  $T$ -con free for all types that are modified by an update (Stoyle et al., 2005). Any update that occurs in one of these regions will be safe. It is possible to leverage these regions in a multi-threaded setting, waiting until all threads happen to be in a safe region, then pausing and updating (Neamtiu and Hicks, 2009). Preliminary experimental results show that deadlock or update delay does occur, but only rarely.

Some approaches do not require representation consistency, and permit code to interact with old and new versions of data (Hjálmtýsson and Gray, 1998; Duggan, 2005; Ajmani et al., 2006; Gurtner, 2011). One approach to ensure type-safe data usage is bi-directional transformation functions (Duggan, 2005). These functions transform the data from old to new formats and vice versa. A monitor can then be used to insert transformations when old code uses new data and vice-versa. A similar technique for object-oriented programs is proxy interfaces (Hjálmtýsson and Gray, 1998; Ajmani et al., 2006; Gurtner, 2011). When an update is introduced, proxies are created for each object, so that any calls to an old class interface are translated into calls to the new interface. The additional indirection required is the primary reason why most designs avoid the use of these techniques.

### 2.3.2.2 Transaction Safety

In addition to type safety there are other properties that we may wish to preserve, particularly over the update itself. Consider a program that adds a log entry once per

iteration of the main loop:

```

1  main
2      f()
3      ...
4      g()
5      h()
6      main()

```

(2.12)

where `f()` and `g()` are function calls. In the old code the log entry is performed by function `g`. Consider an update that changes the bodies of `f` and `g` but not the main function itself, and where in the updated program the log entry is performed by `f'` and not in `g'`. Consider also a reduction where the update is applied on line 4. The remaining code in this iteration is:

```

4  g'()
5  h()
6  main'()

```

(2.13)

So in this iteration of the `main` loop the log entry will not be performed by `g`, and `f` has already executed under the old definition. Hence the log entry will be skipped in this iteration. Transaction aids programmer reasoning about update by ensuring that annotated transactional regions appear to behave entirely with old code or entirely with new code (Neamtiu et al., 2008; Neamtiu and Hicks, 2009). They cannot, as is demonstrated here, behave with an old version of one function, and a new version of a different function. Note that when a given update does not update some function, such as `h`, then this function can be disregarded for the purposes of considering when the update should occur, as it is the same before and after such an update.

### 2.3.2.3 Update Validity

Gupta et al. introduce the notion of update validity (Gupta et al., 1996). Consider a program point with code  $P$  and its state  $\sigma$ . An update consists of some new code  $P'$ , and a state transformation function  $s$ . After an update the program point will be:  $[s(\sigma)] P'$ . An update is said to be valid if the update program can reduce to some other point  $[\sigma''] P''$  that is also reachable by reducing the new code with initial state, in effect that an update  $[s(\sigma)] P'$  is valid if  $[s(\sigma)] P' \rightarrow [\sigma''] P''$  and if  $[\sigma_{\text{init}}] P' \rightarrow [\sigma''] P''$ . Gupta et al. show that proving this is, in general, undecidable. They present restrictions to which programs and updates are admissible, that enables proof of update validity. The restrictions are rather severe, and Gupta et al. do not consider type safety issues.

#### 2.3.2.4 Assurance Safety By Testing and Verification

Many developers, rather than building formal models and ensuring that their program adheres to that model, use regression testing as the standard against which the program is measured. This consists of a suite of tests that are run any time a program is changed. The semantics of update itself can introduce unexpected bugs. Performing regression testing on all possible update points, particularly when programs contain concurrency, is prohibitively computationally expensive. In (Hayden et al., 2009), Hayden et. al. show how to reduce this state space into equivalence classes, so that only a representative member need be tested.

This approach is extendible to verification (Hayden et al., 2011a). Hayden et al. define a transformation that merges the old and new versions of a program to create a single program for verification. They describe properties that should be preserved across updates, such as the presence of key-value pairs in a table, even if the internal table format is changed. They prove their merging transformation correct, and show how the DSU specific properties can be checked using off-the-shelf tools that have no understanding of DSU. The is approach incurs an average slow down factor of only four, but can only handle single-threaded programs.

#### 2.3.2.5 No Safety

Some systems do not make any provisions for guaranteeing safety. Edit and continue development for integrated development environments provide some basic functionality to permit the programmer to change code and state during debugging, but makes no guarantees that the changes will be safe (Dmitriev, 2001; Eaddy and Feiner, 2005; Thomas et al., 2010). JRebel is another debugging tool, for Java web applications under development; it intercepts all method invocations and calls the most recent version, but does no state transformation, and hence is not type safe (ZeroTurnaround, 2011). In (Stewart and Chakravarty, 2005), which implements DSU for Haskell, any data migration is done by serialisation, transformation, and re-injection; these are outside the type system and hence not type safe.

One approach where the lack of safety guarantees is intentional is in Erlang (Wikstrom and Williams, 1993). The programming paradigm in Erlang is to assume that various parts of a program (threads, databases, etc.) will fail at some point, and to program defensively. The methodology for update in Erlang is to change the function body and perform any state transformations immediately, and if the remaining old code happens to touch the new data to simply handle the runtime exception that will then be thrown (Erlang, 2010).

### 2.3.3 Other Considerations

As the aim of DSU is to accommodate unforeseen changes, at unspecified points in execution, it can be difficult to conceptualise how and when updates will occur, and what they will do. Hence ease of use, for programmers, is a key issue in any production DSU approach. Tool support for automatically generating the updates themselves, using the difference between two versions of a program, is included in most DSU approaches aimed at developers (Neamtiu et al., 2006; Gregersen and Jorgensen, 2009; Subramanian, 2010; Thomas et al., 2010; Hayden et al., 2011d). In many systems, however, some programmer intervention is required, particularly in writing non-trivial state transformers (Neamtiu et al., 2006; Hayden et al., 2011d). Some authors argue that reasoning about the behaviour of an update program, particularly when modules or functions are replaced rather than the entire program, is the most difficult cognitive burden for users of DSU (Neamtiu et al., 2008; Hayden et al., 2011d). Others argue that transparency, the ability to use DSU without understanding the semantics or changing coding practices, for example by learning to write state transformation functions, is key to enabling access to DSU in production systems (Gregersen and Jorgensen, 2009). Further research into the practices of software engineers are clearly required.

There are several case studies relevant to DSU in general. The type of changes that occur in practice is the focus of the majority of studies. The incidence of updates that change function or method signatures is surprisingly high, and such changes are included in 10-60% of updates (Neamtiu et al., 2005; Baumann et al., 2007). In addition, the size of updates does not appear to strongly affect the proportion of updates that change signatures (Tempero et al., 2008). The majority of complexity of analyses for DSU, and inefficiencies in implementations, comes from supporting such changes. We consider that, given the proportion of updates that perform changes to signatures, some overhead and cognitive complexity is justified in order to provide update support.

The other case studies focus on dynamic behaviours that can occur when performing DSU. The presence of errors caused by update, and the capability of activeness and con-freeness properties to detect and prevent such errors, is explored in (Hayden et al., 2011c). The majority of errors that were not detected by these properties occurred in updates performed during the initialisation phase of a program. As the majority of updates are assumed to occur during long running main loops, this is not a serious problem. In addition, these properties permitted most updates to safely occur during the main loops, showing that the algorithm was not conservatively preventing updates in the key sections. In (Gregersen and Jorgensen, 2011), interesting, but non-erroneous (at least in terms of type or signature errors) behaviour that occurs when performing DSU in an object oriented system is documented. Transient inconsistency occurs when an updated application is temporally in an unreachable runtime state. An oblivious update is the absence of an expected runtime effect that would have occurred if the system was

started from scratch. Phantom objects appear when classes are removed by a dynamic update, but live objects that belong to that class remain. Gregersen et al. suggest various design patterns that can be used to make programs more amenable to updates and to avoid the described phenomena. They particularly emphasise loosely coupled designs, application logic being based on a dynamic interpretation of a program and its current state and not on static assumptions, and the use of declarative registrations of listener objects.

Systems that support DSU introduce overhead above static programs. The majority of the overhead concerns permitting changes to function or object signatures. In non-managed languages, such as those based on C and C++, the indirection necessary to update functions or modules introduces an overhead during steady state (non-updating) execution. Depending on the implementation details this overhead is often between 2-10% in most cases ([Neamtiu and Hicks, 2009](#); [Boyapati et al., 2003](#)). In situations where the overhead was higher (10-50%) the overhead appeared only in certain circumstances, and was reduced when a different compiler was used or the program was used in a distributed rather than single-computer setting ([Makris and Bazzi, 2009](#); [Neamtiu et al., 2006](#)). The interplay between network semantics, compiler semantics, and DSU would be an interesting path for future research. In managed languages such as Java, much of the indirection for function and method definitions is included by default in the program. It is possible to leverage these facilities to provide DSU with no steady-state overhead versus the static managed language ([Subramanian, 2010](#)). When whole program transformation is employed, rather than replacing functions or classes, the function indirection is not required, and DSU can be provide for non-managed languages without any steady-state overhead ([Hayden et al., 2011d,b](#))

## Chapter 3

# Safety and Liveness Of Concurrent Programs

In this chapter we consider how to represent safety and liveness properties for systems with a variety of side effecting semantics. We explore the communication safety and deadlock freedom properties for asynchronous message passing programs with blocking receives, and provide examples of erroneous programs and their reductions (Section 3.1). We provide a multi-threaded lambda calculus with operators used to access shared state, where the shared state can have an arbitrary semantics (Section 3.2). We show how to use a type and effect system to abstract the impure behaviour of a program, in the vein of Session Types (Section 3.3). We provide a general, inductive description of the communication safety and Deadlock Freedom properties, under the arbitrary shared state semantics (Section 3.4). We show how to re-formulate existing work on Global Session Types, and prove safety and liveness for blocking message passing programs, in an intuitive inductive manner (Section 3.5.1). We also show how to formulate Global Session Types, and prove safety and liveness, for non-blocking message passing, using a similar methodology to that used for blocking message passing (Section 3.5.2). We present our conclusions and describe on our novel contributions in Section 3.6.

```
rcv(c1)(x1 : INT).snd(c2, true).snd(c3, x1 + 1).0 || snd(c1, 2).rcv(c3)(x2 : INT).
  snd(c2, x2).0 || rcp(c2)(x3 : BOOL).rcp(c2)(x3 : INT).0
```

Figure 3.1: Message Passing Example

## 3.1 Motivation: Message Passing Programs

In this section we explore communication safety and liveness for asynchronous message passing programs with blocking receives, and provide examples of erroneous programs and their reductions.

### 3.1.1 Communication Errors

We introduce the concept of Communication Errors in Section 2.2. Intuitively, these occur when a thread receives a value which is not of the type that it expects. The simplest example of this is when the type of a value being sent does not match up with the expected receive type:

$$\begin{aligned}
 1 & [c \mapsto \emptyset] \text{ snd}(c, \text{true}).0 \parallel \text{rcv}(c)(x : \text{INT}).x + 4 \\
 2 & \rightarrow [c \mapsto \text{true}] 0 \parallel \text{rcv}(c)(x : \text{INT}).x + 4 \\
 3 & \rightarrow [c \mapsto \emptyset] 0 \parallel \text{true} + 4
 \end{aligned} \tag{3.1}$$

The program on line 3 is not typable. The effect of the program on line 1 can be expressed as  $c!(\text{BOOL}) \parallel c?(\text{INT})$ . When written in this form the inconsistency is clear.

The complementarity of the effect of a program is not, however, a sufficient condition to show an absence of Communication Errors; it only is so when there are no messages in the message queues. Existing messages can cause errors even for complementary programs:

$$\begin{aligned}
 1 & [c \mapsto \text{true}] \text{ snd}(c, 3).0 \parallel \text{rcv}(c)(x : \text{INT}).x + 4 \\
 2 & \rightarrow [c \mapsto \text{true}, 3] 0 \parallel \text{rcv}(c)(x : \text{INT}).x + 4 \\
 3 & \rightarrow [c \mapsto 3] 0 \parallel \text{true} + 4
 \end{aligned} \tag{3.2}$$

The program on line 3 is also not typable. The effect of the program on line 1 is  $c!(\text{BOOL}) \parallel c?(\text{BOOL})$ ; this effect is complementary, according to the standard definition (Honda, 1993). However, since there are messages in the queue, in order to rule out communication errors and deadlock we cannot consider solely the effect of the code. We must also consider the state of the message queues. We do this by ensuring that the messages in the queue are complementary to the effect of the code, and that after that point the effects are complementary (Bettini et al., 2008). We can show a similar program is safe:

$$\begin{aligned}
 1 & [c \mapsto 3] 0 \parallel \text{rcv}(c)(x : \text{INT}).x + 4 \\
 2 & [c \mapsto \emptyset] 0 \parallel 3 + 4 \\
 3 & [c \mapsto \emptyset] 0 \parallel 7
 \end{aligned} \tag{3.3}$$

```

1   $[\sigma_\emptyset]$   $\text{rcv}(c_1)(x_1 : \text{INT}).\text{snd}(c_2, \text{true}).\text{snd}(c_3, x_1 + 1).0 \parallel$ 
    $\text{snd}(c_1, 2).\text{rcv}(c_3)(x_2 : \text{INT}).\text{snd}(c_2, x_2).0 \parallel$ 
    $\text{rcv}(c_2)(x_3 : \text{BOOL}).\text{rcv}(c_2)(x_3 : \text{INT}).0$ 
2   $[c_1 \mapsto 2]$   $\text{rcv}(c_1)(x_1 : \text{INT}).\text{snd}(c_2, \text{true}).\text{snd}(c_3, x_1 + 1).0 \parallel$ 
    $\text{rcv}(c_3)(x_2 : \text{INT}).\text{snd}(c_2, x_2).0 \parallel \text{rcv}(c_2)(x_3 : \text{BOOL}).\text{rcv}(c_2)(x_3 : \text{INT}).0$ 
3   $[\sigma_\emptyset]$   $\text{snd}(c_2, \text{true}).\text{snd}(c_3, 2 + 1).0 \parallel \text{rcv}(c_3)(x_2 : \text{INT}).\text{snd}(c_2, x_2).0 \parallel$ 
    $\text{rcv}(c_2)(x_3 : \text{BOOL}).\text{rcv}(c_2)(x_3 : \text{INT}).0$ 
4   $[c_2 \mapsto \text{true}]$   $\text{snd}(c_3, 2 + 1).0 \parallel \text{rcv}(c_3)(x_2 : \text{INT}).\text{snd}(c_2, x_2).0 \parallel$ 
    $\text{rcv}(c_2)(x_3 : \text{BOOL}).\text{rcv}(c_2)(x_3 : \text{INT}).0$ 
5   $[\sigma_\emptyset]$   $\text{snd}(c_3, 2 + 1).0 \parallel \text{rcv}(c_3)(x_2 : \text{INT}).\text{snd}(c_2, x_2).0 \parallel \text{rcv}(c_2)(x_3 : \text{INT}).0$ 
6   $[c_3 \mapsto 3]$   $0 \parallel \text{rcv}(c_3)(x_2 : \text{INT}).\text{snd}(c_2, x_2).0 \parallel \text{rcv}(c_2)(x_3 : \text{INT}).0$ 
7   $[\sigma_\emptyset]$   $0 \parallel \text{snd}(c_2, 3).0 \parallel \text{rcv}(c_2)(x_3 : \text{INT}).0$ 
8   $[c_2 \mapsto 3]$   $0 \parallel 0 \parallel \text{rcv}(c_2)(x_3 : \text{INT}).0$ 
9   $[\sigma_\emptyset]$   $0 \parallel 0 \parallel 0$ 

```

Figure 3.2: Safe Blocking Semantics Reduction

The effect of this program is  $0 \parallel c?(\text{INT})$ . Here, the  $c?(\text{INT})$  is complementary to the  $\text{INT}$  value in the message queue.

The influence of the semantics of the shared state on Communication Safety can be illustrated using the program in Figure 3.1. Under a semantics with blocking receives there is no reduction sequence that contains Communication Errors; one such safe reduction is shown in Figure 3.2.

One alternate semantics for message passing has non-blocking receives; when a receive action is performed on an empty queue the action does not block. Normally, when performing a receive action on an empty queue, some specific error value would be returned, and the control flow would then handle the fact that a value of the expected type has not been returned. For the sake of the next example, however, we assume such a receive returns a default value of the expected type. Under such a semantics we can obtain a Communication Error, as in Figure 3.3. As the second thread does not synchronise by receiving the value from the first thread, it can bypass its receive and perform its send to the third thread ( $\text{snd}(c_2, x_2)....$ ) before the first thread performs its send to the third thread ( $\text{snd}(c_2, \text{true})....$ ). This leads to the third thread receiving an  $\text{INT}$  when it is expecting a  $\text{BOOL}$ . This exemplifies how synchronisation behaviour is key to proving safety in many circumstances.

```

1   $[\sigma_\emptyset]$   $\text{rcv}(c_1)(x_1 : \text{INT}).\text{snd}(c_2, \text{true}).\text{snd}(c_3, x_1 + 1).0 \parallel$ 
    $\text{snd}(c_1, 2).\text{rcv}(c_3)(x_2 : \text{INT}).\text{snd}(c_2, x_2).0 \parallel$ 
    $\text{rcv}(c_2)(x_3 : \text{BOOL}).\text{rcv}(c_2)(x_3 : \text{INT}).0$ 
2   $[c_1 \mapsto 2]$   $\text{rcv}(c_1)(x_1 : \text{INT}).\text{snd}(c_2, \text{true}).\text{snd}(c_3, x_1 + 1).0 \parallel$ 
    $\text{rcv}(c_3)(x_2 : \text{INT}).\text{snd}(c_2, x_2).0 \parallel \text{rcv}(c_2)(x_3 : \text{BOOL}).\text{rcv}(c_2)(x_3 : \text{INT}).0$ 
3   $[c_1 \mapsto 2]$   $\text{rcv}(c_1)(x_1 : \text{INT}).\text{snd}(c_2, \text{true}).\text{snd}(c_3, x_1 + 1).0 \parallel \text{snd}(c_2, 0).0 \parallel$ 
    $\text{rcv}(c_2)(x_3 : \text{BOOL}).\text{rcv}(c_2)(x_3 : \text{INT}).0$ 
4   $[c_1 \mapsto 2, c_2 \mapsto 0]$   $\text{rcv}(c_1)(x_1 : \text{INT}).\text{snd}(c_2, \text{true}).\text{snd}(c_3, x_1 + 1).0 \parallel 0 \parallel$ 
    $\text{rcv}(c_2)(x_3 : \text{BOOL}).\text{rcv}(c_2)(x_3 : \text{INT}).0$ 
5   $[c_1 \mapsto 2]$   $\text{rcv}(c_1)(x_1 : \text{INT}).\text{snd}(c_2, \text{true}).\text{snd}(c_3, x_1 + 1).0 \parallel 0 \parallel \text{error}$ 

```

Figure 3.3: Non Blocking Semantics Gives Rise to Communication Error

---

### 3.1.2 Deadlock

Deadlock, due to message passing, occurs when a communication action is blocked indefinitely. If no communication actions can block then we can trivially show an absence of deadlock. As we show in Example 3.3, however, intuition about which programs are safe can be incorrect when programs perform less synchronisation than we expect; in the above examples, synchronisation occurs when individual threads block temporarily at receives until messages become available. We therefore consider it important to work in systems which have some form of synchronisation, and which hence can deadlock.

The simplest example of deadlock is a program where one process expects a value to be sent and the others do not:

```

1   $[c \mapsto \emptyset]$   $0 \parallel \text{rcv}(c)(x : \text{INT}).x + 4$  (3.4)
2   $\not\rightarrow$ 

```

Here the two processes do not have complementary behaviour, and will obviously deadlock. Similarly to Example 3.2, complementary behaviour is not a sufficient condition to rule out deadlock when there are existing messages in the queue:

```

1   $[c_1 \mapsto l_2]$   $\text{case } c_1 \text{ of } \{l_1 \mapsto \text{snd}(c_2, 3).0, l_2 \mapsto 0\} \parallel \text{sel}(c_1, l_1).\text{rcv}(c_2)(x : \text{INT}).0$ 
2   $[c_1 \mapsto l_2, l_1]$   $\text{case } c_1 \text{ of } \{l_1 \mapsto \text{snd}(c_2, 3).0, l_2 \mapsto 0\} \parallel \text{rcv}(c_2)(x : \text{INT}).0$ 
3   $[c_1 \mapsto l_1]$   $0 \parallel \text{rcv}(c_2)(x : \text{INT}).0$ 
4   $\not\rightarrow$  (3.5)

```

In this case existing messages influence the control flow of the program, leading to the second process expecting the  $l_1$  service whilst the first process provides the  $l_2$  service.

$t ::=$	<i>Terms</i>	
$v$	Values	
$\underline{X}$	Recursion Variables	
$tt$	Application	
$(\alpha(\tilde{v}), T)$	State Access	
$\text{if } t \text{ then } t \text{ else } t$	Conditional	
$\text{case } \alpha(\tilde{v}) \text{ in } \{T \mapsto t\}$	Case Split	
$\text{error}$	Error	
$v ::=$	<i>Values</i>	
$n$	Integers	
$b$	Booleans	
$()$	Unit	
$r$	Resources	
$x$	Variables	
$l$	Service Labels	
$\text{rec } \underline{X}(x : T).t$	Recursive Functions	
$P ::=$	<i>Process Threads</i>	
$t$	Single Term Thread	
$P \parallel P$	Parallel Composition	

Figure 3.4: Language

We conclude that the safety and liveness of a program depends on: 1) the communication actions of the program, 2) the messages in the message queues, and 3) the semantics of the communication actions.

## 3.2 Language Definitions

We define our language in Figure 3.4 and our operational semantics in Figure 3.5. Values consist of integers, booleans, the unit value, variables, resources, and recursive functions. Resources are handles into the runtime state, such as a channel name  $c$  in previous examples. Labels are used to provide unique service labels. Recursive functions  $\text{rec } \underline{X}(x : T).t$  consist of a recursion variable  $\underline{X}$ , a formal parameter  $x$ , the expected type  $T$  of the variable (we define types in Section 3.3), and the body of the function  $t$ . Function application is denoted as  $t_1 t_2$ . The term  $(\text{rec } \underline{X}(x : T).t) v$  reduces to  $t[v/x][\text{rec } \underline{X}(x : T).t_2/\underline{X}]$ , assuming  $v : T$ , and  $\text{error}$  otherwise. We make use of an eager reduction strategy, and hence use  $t_1; t_2$  as syntactic sugar for  $(\text{rec } \underline{X}(x : T).t_2) t_1$ , where  $x$  and  $\underline{X}$  are not free in  $t_2$ . We also make use of the syntactic sugar  $\lambda x : T.t$  for  $\text{rec } \underline{X}(x : T).t$  when  $\underline{X}$  is not free in  $t$ .

Accesses to the shared state are denoted using  $(\alpha(\tilde{v}), T)$ . This denotes that action  $\alpha$  is performed using parameters  $\tilde{v}$ , and this access should return a value of type  $T$ .

$$\begin{array}{c}
\frac{[\sigma] \ t_1 \xrightarrow{\gamma} [\sigma'] t'_1}{[\sigma] \ t_1 t_2 \xrightarrow{\gamma} [\sigma'] t'_1 t_2} \quad \frac{[\sigma] \ t_2 \xrightarrow{\gamma} [\sigma'] t'_2}{[\sigma] \ v t'_2 \xrightarrow{\gamma} [\sigma'] v t'_2} \\
\\
\frac{}{[\sigma] \ \text{rec } \underline{X}(x : T).t \ v \xrightarrow{\tau} [\sigma] \ t[\text{rec } \underline{X}(x : T).t / \underline{X}][v/x]} \\
\\
\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T}{[\sigma] \ (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \ \sigma(\alpha(\tilde{v}))} \quad \frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T' \neq T}{[\sigma] \ (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \ \text{error}} \\
\\
\frac{}{[\sigma] \ \text{if true then } t_2 \ \text{else } t_3 \xrightarrow{\tau} [\sigma] \ t_2} \quad \frac{}{[\sigma] \ \text{if false then } t_2 \ \text{else } t_3 \xrightarrow{\tau} [\sigma] \ t_3} \\
\\
\frac{[\sigma] \ t_1 \xrightarrow{\gamma} [\sigma'] t'_1}{[\sigma] \ \text{if } t_1 \text{ then } t_2 \ \text{else } t_3 \xrightarrow{\gamma} [\sigma'] \ \text{if } t'_1 \text{ then } t_2 \ \text{else } t_3} \\
\\
\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T_i \quad T_i \in \tilde{T}}{[\sigma] \ \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} \xrightarrow{\alpha(\tilde{v})} [\sigma'] \ t_i} \quad \frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T' \quad T' \notin \tilde{T}}{[\sigma] \ \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} \xrightarrow{\alpha(\tilde{v})} [\sigma'] \ \text{error}} \\
\\
\frac{[\sigma] \ P_1 \xrightarrow{\gamma} [\sigma'] P'_1}{[\sigma] \ P_1 \parallel P_2 \xrightarrow{\gamma} [\sigma'] P'_1 \parallel P_2} \quad \frac{[\sigma] \ P_2 \xrightarrow{\gamma} [\sigma'] P'_2}{[\sigma] \ P_1 \parallel P_2 \xrightarrow{\gamma} [\sigma'] P_1 \parallel P'_2}
\end{array}$$

Figure 3.5: Operational Semantics

State accesses are not permitted to return values containing resources. Were we to permit resources to be passed around, then channels could be passed around, and we would need to consider session typing techniques for delegation. In a type and effect formulation, this would require dependent typing within the effect system, which whilst possible would significantly complicate our current analysis. Examples of state accesses include  $(\text{snd}(c, 3), \text{UNIT})$  and  $(\text{rcv}(c), \text{INT})$  which denote the communication actions of  $\text{snd}(c, 3)$  and  $\text{rcv}(c)(x : \text{INT})$  respectively, but not their continuations or the variable binding for  $\text{rcv}$ .

The conditional operator is standard. Service branching is provided using the case split operator  $\text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\}$ . This operator performs the communication action  $\alpha(\tilde{v})$ , as defined above, and then does a case split on the type of the result. If the value returned is not of an expected type then the case split reduces to **error**. We denote errors using the **error** term; errors cannot reduce, in and of themselves.

We parameterise our work on a *resource model*, which is a collection of *states* ranged over by  $\sigma$ , and a function that maps a state and a label to a return value and another state: we write  $\sigma(\alpha(\tilde{v})) = v$  and  $\sigma \xrightarrow{\alpha(\tilde{v})} \sigma'$  to denote this map. The resource model is a generalised model, and hence can be instantiated with either finite or infinite models. In particular, it is not restricted to queues, though in this thesis we will only consider examples that use queues.

$\llbracket \text{snd}(c, e).t \rrbracket$	$\stackrel{\text{def}}{=} (\lambda x : T. (\text{snd}(c, x), \text{UNIT}) \llbracket e \rrbracket); \llbracket t \rrbracket$	Sends
$\llbracket \text{rcv}(c)(x : T).t \rrbracket$	$\stackrel{\text{def}}{=} (\lambda x : T. \llbracket t \rrbracket) (\text{rcv}(c), T)$	Receives
$\llbracket \text{if } e \text{ then } t \text{ else } t \rrbracket$	$\stackrel{\text{def}}{=} \text{if } \llbracket e \rrbracket \text{ then } \llbracket t \rrbracket \text{ else } \llbracket t \rrbracket$	Conditionals
$\llbracket \text{sel}(c, l).t \rrbracket$	$\stackrel{\text{def}}{=} (\text{snd}(c, l), \text{UNIT}); \llbracket t \rrbracket$	Service Selection
$\llbracket \text{case } c \text{ of } \{l \mapsto t\} \rrbracket$	$\stackrel{\text{def}}{=} \text{case } \text{rcv}(c) \text{ in } \{ \text{Lab } l \mapsto \llbracket t \rrbracket \}$	Service Branching
$\llbracket \mu \underline{X}.t \rrbracket$	$\stackrel{\text{def}}{=} \text{rec } \underline{X}(x : \text{UNIT}). \llbracket t \rrbracket$	Recursion, where $x$ is not free in $t$
$\llbracket \underline{X} \rrbracket$	$\stackrel{\text{def}}{=} \underline{X}()$	Recursion Variable
$\llbracket \text{error} \rrbracket$	$\stackrel{\text{def}}{=} \text{error}$	Errors
$\llbracket e \rrbracket$	$\stackrel{\text{def}}{=} \dots$	Typical $\lambda$ -calculus embedding of numbers, etc.
$\llbracket P \parallel P \rrbracket$	$\stackrel{\text{def}}{=} \llbracket P \rrbracket \parallel \llbracket P \rrbracket$	Parallel Threads

$$\frac{}{\sigma[c \mapsto q] \xrightarrow{\text{snd}(c, v)} \sigma[c \mapsto (q, v)]} \quad \frac{}{\sigma[c \mapsto (v, q)] \xrightarrow{\text{rcv}(c)} \sigma[c \mapsto q]}$$

$$\sigma(\text{snd}(c, v)) \stackrel{\text{def}}{=} () \quad \sigma[c \mapsto v, q](\text{rcv}(c)) \stackrel{\text{def}}{=} v$$

Figure 3.6: Blocking Message Passing Calculus Embedding

We can embed the message passing calculus presented in Section 2.2 in our language. We describe this embedding in Figure 3.6; it is relatively straightforward, using functions to represent binding of received values. Recursion variables are mapped to an application of unit to the recursion variable.

### 3.3 Behavioural Abstraction

We use effects as an abstraction of the actions on the shared state. We describe this behaviour in both thread-local and global (parallel) settings. We also present an effect system that determines the effect of a program. In order to provide our safety and liveness guarantees, however, we need to do more than determine the effect of a program. We need to place some restrictions on which programs are valid. We describe these restrictions in Sections 3.4.

#### 3.3.1 Parallel and Local Session Types

In Figure 3.7 we present the syntax of our types and effects.

Types consist of base types, function types, and two nominal types (those where the type is defined by the singleton value which has said type). A function type  $T_1 \xrightarrow{\mathcal{L}} T_2$

$\varphi ::=$		$T ::=$	<i>Types</i>
	<i>Local Effects</i>		
$(\alpha(\tilde{T}), T)$	State Access	BOOL	Booleans
$\varphi; \varphi$	Sequence	INT	Integers
$\varphi \oplus \varphi$	Internal Choice	UNIT	Units
$\varphi \& \varphi$	External Choice	Res $r$	Resources
$\mu \underline{X}. \varphi$	Recursion	Lab $l$	Labels
$\underline{X}$	Recursion Variable	$T \xrightarrow{\varphi} T$	Functions
$\epsilon$	Empty Effect		
$\text{error}$	Error		
$\Phi ::=$		<i>Parallel Effects</i>	
		$\varphi$	Local effect
		$\Phi \parallel \Phi$	Parallel Composition

Figure 3.7: Types, and Local and Parallel Effects

is annotated with the latent effect  $\varphi$  of the function's body (as in Section 2.1). This denotes that when an argument is applied to the function, the body of the function will reduce with an effect described by  $\varphi$ . To aid our later analyses we wish to distinguish all resources and labels at the type level. We hence use nominal types for resources and labels; a resource  $r$  has type  $\mathbf{Res}\,r$ , and a label  $l$  has type  $\mathbf{Lab}\,l$ , respectively.

Local Effects describe the structure of actions performed by a single thread on the shared state. An abstract action  $(\alpha(\tilde{T}), T)$  denotes performing an action  $\alpha$ , where the types of the parameters to the action are  $\tilde{T}$ , and the expected type of the return value is  $T$ . For example, the action  $(\text{snd}(c, 3), \text{UNIT})$  is abstracted using  $(\text{snd}(\text{Res } c, \text{INT}), \text{UNIT})$ . Sequencing, recursion, empty effects, and errors are standard. Internal choice  $\varphi_1 \oplus \varphi_2$  denotes that a program can proceed with either the behaviour described in  $\varphi_1$  or the behaviour described in  $\varphi_2$ , and the choice will be made within the given process in an opaque manner. We will use internal choice to describe the behaviour of **if** statements. External choice  $\varphi_1 \& \varphi_2$  denotes that a program can proceed with either the behaviour described in  $\varphi_1$  or that described in  $\varphi_2$ , but that in order to choose a specific path the first action of that path must be reducible, and not blocked, at the time the choice is made. Parallel Effects are simply Local Effects in parallel with each other.

We make use of an equivalence relation over effects, which is the least relation defined structurally over Parallel Effects using:

$$\epsilon; \varphi \equiv \varphi \equiv \varphi; \epsilon \quad \mu X. \varphi \equiv \varphi[\mu X. \varphi / X] \quad (3.6)$$

### 3.3.2 Typing Rules

In Figure 3.8 we present our type and effect system. A typing judgement  $\varphi \wr \Gamma \vdash t : T$  denotes that, under typing assumptions  $\Gamma$  the communication behaviour of some code  $t$  is represented by the effect  $\varphi$ , and the code will return a value of type  $T$ . A typing judgement for parallel processes is similarly defined, without the typing environment or the return type. We make use of a type system for simple expressions, whose judgement  $t : T$  denotes that  $t$  is a term, that performs no actions on the shared state, whose return type is  $T$ .

Values are abstracted to empty effects. Resources and labels are nominally typed. Variables are typed using the typing assumptions. Actions are abstracted using the types of the parameters, and we require that an action does not return a resource type. Recursive functions are typed under the assumption that the variable  $x$  has type  $T_1$ , and that the function itself has type  $T_1 \xrightarrow{X} T_2$ . An action  $(\alpha(\tilde{v}), T)$ , that reduces the shared state using label  $\alpha(\tilde{v})$  and expects a return value of type  $T$ , assumes that the value returned will be of the expected type, and that it will not be a resource. Note that at this stage we have no guarantee that the return value actually will be of the expected type. In order to ensure that communication errors do not occur, such as that in Example 3.2, we need to verify that no such errors exist. We describe this verification in Section 3.4.

Recursive functions are abstracted using a function type, with a latent recursive effect. When typing the body of the function we assume that the recursion variable has the type of a function from  $T_1$  to  $T_2$ , with the latent effect  $\underline{X}$  to denote the recursion. We require that  $x$  is not free in  $\Gamma$ ,  $T_1$ ,  $T_2$ , or  $\varphi$ . This prevents the function from being passed a resource variable as a parameter, and effectively prohibits functions from having polymorphic effect dependent on the function parameter. Application is abstracted as a sequence of effects, firstly that of reducing the left hand side to a value, then that of reducing the right hand side to a value, then that of reducing the function body.

Conditional terms are abstracted as the effect of the term being split over, followed by the internal choice between the effects of the consequent and the alternative. We require that the predicate not access the shared state. A case split  $\mathbf{case} \alpha(\tilde{v}) \mathbf{in} \{\tilde{T} \mapsto t\}_M$  is abstracted as the external choice between, for each  $m \in M$ , performing action  $\alpha(\tilde{T})$  and receiving a value of type  $T_m$ , followed by the effect of the continuation  $t_m$ . As such  $\&\varphi_i$  is syntactic sugar for  $\varphi_1 \& \dots \& \varphi_n$  where  $I = \{1, \dots, n\}$ . We also assume that none of the values received are resources. Threads are typed using an empty type environment. Parallel threads are abstracted as the parallel composition of their constituent effects.

### 3.3.3 Semantics of Effects

As the type and effect system rules out simple type errors such as  $4 + \mathbf{true}$ , type errors can only be introduced by the reception of a value which is not of the type that a

$$\begin{array}{c}
\frac{}{\epsilon \wr \Gamma \vdash n: \text{INT}} \quad \frac{}{\epsilon \wr \Gamma \vdash b: \text{BOOL}} \quad \frac{}{\epsilon \wr \Gamma \vdash (): \text{UNIT}} \quad \frac{}{\epsilon \wr \Gamma \vdash r: \text{Res } r} \\
\frac{}{\epsilon \wr \Gamma \vdash x: \Gamma(x)} \quad \frac{}{\epsilon \wr \Gamma \vdash l: \text{Lab } l} \quad \frac{\epsilon \wr \Gamma \vdash v_i: T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \Gamma \vdash (\alpha(\tilde{v}), T): T} \\
\frac{\varphi \wr \Gamma, x: T_1, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash t: T_2 \quad x \notin \text{fv}(\Gamma), \text{fv}(T_1), \text{fv}(T_2), \text{fv}(\varphi) \quad \underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)}{\epsilon \wr \Gamma \vdash \text{rec } \underline{X}(x: T_1).t: T_1 \xrightarrow{\mu X. \varphi} T_2} \\
\frac{\varphi_1 \wr \Gamma \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \wr \Gamma \vdash t_2: T_2}{(\varphi_1; \varphi_2; \varphi_3) \wr \Gamma \vdash t_1 t_2: T_1} \\
\frac{\varphi_1 \wr \Gamma \vdash t_1: \text{BOOL} \quad \varphi_2 \wr \Gamma \vdash t_2: T \quad \varphi_3 \wr \Gamma \vdash t_3: T}{\varphi_1; (\varphi_2 \oplus \varphi_3) \wr \Gamma \vdash \text{if } t_1 \text{ then } t_2 \text{ else } t_3: T} \\
\frac{\varphi_i \wr \Gamma \vdash t_i: T \quad \epsilon \wr \Gamma \vdash v_i: T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \Gamma \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\}: T} \\
\frac{\varphi \wr \Gamma \vdash t: T \quad \varphi \equiv \varphi'}{\varphi \wr \Gamma \vdash t: \varphi'} \quad \frac{\varphi \wr \emptyset \vdash t: T \quad \vdash P_1: \Phi_1 \quad \vdash P_2: \Phi_2}{\vdash t: \varphi \quad \vdash P_1 \parallel P_2: \Phi_1 \parallel \Phi_2}
\end{array}$$

Figure 3.8: Typing Rules

program expects. In other words, the only source of type errors is actions performed on the shared state; we can therefore prove an absence of type errors by considering the actions on the shared state alone. In order to reason about how the effect of a program reduces, in the presence of some shared state, we also make use of an *abstract resource model* which abstracts away from particular values associated with resource states and provides a representation of the resource types. An abstract resource model is a collection of abstract states ranged over by  $\Sigma$  and a function that maps a state and an abstract label to a type and another abstract state: we write  $\Sigma(L) = T$  and  $\Sigma \xrightarrow{L} \Sigma'$  as above. Abstract actions are of the form  $\alpha(\tilde{T})$ , so that a label  $\alpha(\tilde{v})$  corresponds to a unique abstract label  $\alpha(\tilde{T})$  via simple typing of values. We assume the fact that we can relate the resource model and abstract resource model via an abstraction map  $\Lambda$  with the property that  $\sigma \xrightarrow{\alpha(\tilde{v})} \sigma'$  implies  $\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma')$  and  $\sigma(\alpha(\tilde{v})): \Lambda(\sigma)(\alpha(\tilde{T}))$ . We include a semantics for effects and abstract models in Figure 3.9, with a definition of the label abstraction rule in Figure 3.10.

If an action returns a value of the expected type, then that action reduces to the empty effect. If that action does not return a value of the expected type that action reduces to an error. Sequential composition and internal choice are standard. External choice can only reduce to one of the branches when that branch can also reduce (i.e. it is not blocked). When both options are enabled, if one can reduce without causing an error, that reduction will be chosen. If both can reduce without error then the path

$$\begin{array}{c}
\frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) = T}{[\Sigma] \ (\alpha(\tilde{T}), T) \rightarrow [\Sigma'] \ \epsilon} \quad \frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) \neq T}{[\Sigma] \ (\alpha(\tilde{T}), T) \rightarrow [\Sigma'] \ \text{error}} \\
\\
\frac{[\Sigma] \ \varphi_1 \xrightarrow{\gamma} [\Sigma'] \ \varphi'_1}{[\Sigma] \ \varphi_1 ; \varphi_2 \xrightarrow{\gamma} [\Sigma'] \ \varphi'_1 ; \varphi_2} \quad \frac{i \in 1, 2}{[\Sigma] \ \varphi_1 \oplus \varphi_2 \rightarrow [\Sigma] \ \varphi_i} \\
\\
\frac{\exists \varphi'_i, \alpha_i. \ [\Sigma] \ \varphi_i \xrightarrow{\alpha_i(\tilde{T})} [\Sigma'_i] \ \varphi'_i \wedge \text{error} \in \varphi'_i}{[\Sigma] \ \varphi_1 \& \varphi_2 \xrightarrow{\alpha_i(\tilde{T})} [\Sigma'_i] \ \varphi'_i} \\
\\
\frac{\forall \varphi'_i, \alpha_i. \ [\Sigma] \ \varphi_i \xrightarrow{\alpha_i(\tilde{T})} [\Sigma'_i] \ \varphi'_i \Rightarrow \text{error} \in \varphi'_i}{[\Sigma] \ \varphi_1 \& \varphi_2 \xrightarrow{\alpha_j(\tilde{T})} [\Sigma'_j] \ \varphi'_j} \\
\\
\frac{[\Sigma] \ \mu \underline{X}. \varphi \rightarrow [\Sigma] \ \varphi[\mu \underline{X}. \varphi / \underline{X}]}{\Phi \equiv \Phi_1 \quad [\Sigma] \ \Phi_1 \rightarrow [\Sigma'] \ \Phi_2 \quad \Phi_2 \equiv \Phi'} \quad \frac{}{[\Sigma] \ \Phi \rightarrow [\Sigma'] \ \Phi'} \\
\\
\frac{[\Sigma] \ \Phi_1 \xrightarrow{\gamma} [\Sigma'] \ \Phi'_1}{[\Sigma] \ \Phi_1 \parallel \Phi_2 \xrightarrow{\gamma} [\Sigma'] \ \Phi'_1 \parallel \Phi_2} \quad \frac{[\Sigma] \ \Phi_2 \xrightarrow{\gamma} [\Sigma'] \ \Phi'_2}{[\Sigma] \ \Phi_1 \parallel \Phi_2 \xrightarrow{\gamma} [\Sigma'] \ \Phi_1 \parallel \Phi'_2}
\end{array}$$

Figure 3.9: Effect Semantics

$$\frac{\epsilon \wr \emptyset \vdash v_i : T_i}{\alpha(\tilde{T}) \vdash \alpha(\tilde{v})}$$

Figure 3.10: Label Abstraction

is chosen non-deterministically. If both can reduce with an error then the path is also chosen non-deterministically. Recursion unfolding, equivalence reduction, and parallel reduction are standard.

One contribution in this area is the explicit definition of what type should be returned for a given shared state access. Whilst this is currently done in Session Typing analyses for receive actions, we extend it to all actions. One use for being able to return, and handle, different types for the same action permits us to provide information about the shared state and how it has responded to the action. This is particularly used in our handling of the control flow of non-blocking message passing (Section 3.5.2).

## 3.4 General Formulation

In order to provide safety and liveness guarantees we need to do more than determine the effect of a program. We must place some restrictions on which programs are valid, in

$$\begin{aligned}
\mathcal{G}^0(\Sigma, \Phi) &\stackrel{\text{def}}{=} \text{true} \\
\mathcal{G}^{k+1}(\Sigma, \Phi) &\stackrel{\text{def}}{=} \left[ \begin{array}{l} [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow \mathcal{G}^k(\Sigma', \Phi') \\ \wedge [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \vee \Phi = \Pi \epsilon \\ \wedge \text{error} \notin \Phi \end{array} \right] \\
&\text{where } k > 0
\end{aligned}$$

$$\mathcal{G}(\Sigma, \Phi) = \bigwedge_{k \in \mathcal{N}} \mathcal{G}^k(\Sigma, \Phi)$$

Figure 3.11: Global Compatibility

order to rule out errors and deadlock, such as those discussed in Section 3.1. We define a program point to consist of an effect  $\Phi$  and a shared state abstraction  $\Sigma$  (or a program and some shared state, respectively). A program point is safe and deadlock free if it does not contain an error, if all its reductions result in safe and deadlock free program points, and if it can perform at least one reduction (or consists of empty effects). We formalise these using the function in Figure 3.11. We refer to this property as *Global Compatibility*, and denote a globally compatible program point as  $\mathcal{G}(\Sigma, \Phi)$ . We describe this formalisation in more detail as follows:

$$\text{error} \notin \Phi \tag{3.7}$$

This requires that no errors exist in a globally compatible program point.

$$[\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow \mathcal{G}(\Sigma', \Phi') \tag{3.8}$$

This requires that for any reduction which can occur at this point that the reduced program must also be globally compatible.

$$[\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \vee \Phi \equiv \Pi \epsilon \tag{3.9}$$

This requires that either the effect can reduce, or it is equivalent to empty effects.

The predicate can alternatively be characterised as the greatest fixed point of the following function, defined over the state space  $S$  of all possible pairs  $(\Sigma, \Phi)$  of abstract states and parallel effects:

$$g(S) = \left\{ (\Sigma, \Phi) \mid \begin{array}{l} \forall \Sigma', \Phi'. [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow (\Sigma', \Phi') \in S \text{ and} \\ (\exists \Sigma', \Phi'. [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \text{ or } \Phi \equiv \Pi \epsilon) \text{ and } \text{error} \notin \Phi \end{array} \right\} \tag{3.10}$$

This is a monotone function on the powerset lattice of pairs of abstract states and effects. It has a greatest fixed point  $\nu g$  (which is just a set of tuples). Using the Knaster-Tarski fixed point theorem, we know that the predicate  $\mathcal{G}(\Sigma, \Phi)$  is defined if and only if  $(\Sigma, \Phi) \in \nu g$ .

### 3.4.1 Key Lemmas

Using the definition of Global Compatibility we can proceed to prove some key lemmas.

#### 3.4.1.1 Subject Reduction

If a term and its accompanying shared state's abstraction is globally compatible then we know that:

$$[\Sigma] \varphi \rightarrow [\Sigma'] \varphi' \Rightarrow \text{error} \notin \varphi' \quad (3.11)$$

This is as, if  $\mathcal{G}(\Sigma, \varphi)$ , then we know that  $\mathcal{G}(\Sigma', \varphi')$  and that if  $\mathcal{G}(\Sigma', \varphi')$  then  $\text{error} \notin \varphi'$ . Given that there are no type rules for `error` we know that all well typed programs are error free. We can use Equation 3.11 to show that a single reduction results in a well typed reduced term, and that the effect of the original program point can reduce to the effect of the new program point.

**Lemma 3.1.** *Expression Subject Reduction*

For all  $\varphi, \varphi'', t, t', T, \sigma, \sigma''$  if  $\varphi \wr \emptyset \vdash t: T$ , and  $[\sigma] t \xrightarrow{\gamma} [\sigma'] t'$  and if whenever  $\gamma' \vdash \gamma$  and  $[\Lambda(\sigma)] \varphi \xrightarrow{\gamma'} [\Lambda(\sigma'')] \varphi''$  then  $\text{error} \notin \varphi''$ , then there exists a  $\varphi'$  such that typing judgement  $\varphi' \wr \emptyset \vdash t': T$  holds.

*Proof.* Suppose the hypotheses. We then prove the conclusions simultaneously by induction over the derivation of  $[\sigma] t \xrightarrow{\gamma} [\sigma'] t'$ . The key cases are as follows.

Consider the case where the last reduction rule used in the derivation is:

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})): T}{[\sigma] (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \sigma(\alpha(\tilde{v}))}$$

As  $t = (\alpha(\tilde{v}), T)$  and  $\varphi \wr \emptyset \vdash (\alpha(\tilde{v}), T): T$ , then by the type rules we know that the following holds:

$$\frac{\epsilon \wr \emptyset \vdash v_i: T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \emptyset \vdash (\alpha(\tilde{v}), T): T}$$

where  $\varphi \equiv (\alpha(\tilde{T}), T)$ . Given that  $\gamma = \alpha(\tilde{v})$ , by the definition of  $\gamma' \vdash \gamma$  we know that then any relevant  $\gamma'$  must be  $\alpha(\tilde{T})$  where  $\tilde{v}: \tilde{T}$ . If  $[\Lambda(\sigma)] (\alpha(\tilde{T}), T) \xrightarrow{F} [\Lambda(\sigma'')] \varphi''$  then it will be by one of the following two rules:

$$\frac{\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \quad \Lambda(\sigma)(\alpha(\tilde{T})) = T}{[\Lambda(\sigma)] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')] \epsilon} \quad \frac{\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \quad \Lambda(\sigma)(\alpha(\tilde{T})) \neq T}{[\Lambda(\sigma)] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')] \text{error}}$$

As we assume that  $\mathbf{error} \notin \varphi''$  then we must use the first rule, and hence have that  $\varphi'' = \epsilon$  and that  $\Lambda(\sigma)(\alpha(\tilde{T})) = T$ . By our assumptions about the definition of the state projection function we know that  $\sigma(\alpha(\tilde{v})) : T$ . Then by Lemma A.9 we know that  $\epsilon \wr \emptyset \vdash \sigma(\alpha(\tilde{v})) : T$  holds, as required.

Consider the case where the last reduction rule used in the derivation is:

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T' \neq T}{[\sigma] \ (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \ \mathbf{error}}$$

As  $t = (\alpha(\tilde{v}), T)$  and  $\varphi \wr \emptyset \vdash (\alpha(\tilde{v}), T) : T$ , then by the type rules we know that the following holds:

$$\frac{\epsilon \wr \emptyset \vdash v_i : T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \emptyset \vdash (\alpha(\tilde{v}), T) : T}$$

where  $\varphi \equiv (\alpha(\tilde{T}), T)$ . Given that  $\gamma = \alpha(\tilde{v})$ , by the definition of  $\gamma' \vdash \gamma$  we know that then any relevant  $\gamma'$  must be  $\alpha(\tilde{T})$  where  $\tilde{v} : \tilde{T}$ . If  $[\Lambda(\sigma)] \ (\alpha(\tilde{T}), T) \xrightarrow{\gamma'}_F [\Lambda(\sigma'')] \ \varphi''$  then it will be by one of the following two rules:

$$\frac{\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \quad \Lambda(\sigma)(\alpha(\tilde{T})) = T}{[\Lambda(\sigma)] \ (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')] \ \epsilon} \quad \frac{\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \quad \Lambda(\sigma)(\alpha(\tilde{T})) \neq T}{[\Lambda(\sigma)] \ (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')] \ \mathbf{error}}$$

Consider when the first rule is used. Then we must have that  $\Lambda(\sigma)(\alpha(\tilde{T})) = T$ . However, by the reduction rule being considered we have that  $\Lambda(\sigma)(\alpha(\tilde{T})) \neq T$ , as  $\sigma(\alpha(\tilde{v})) : T'' \neq T$ , and by our assumptions we must have  $T''$  be the same as  $\Lambda(\sigma)(\alpha(\tilde{v}))$ . Hence we have a contradiction.

Consider when the second rule is used. Then we have that  $[\Lambda(\sigma)] \ (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')] \ \mathbf{error}$ . By our assumptions, however, we do not consider any such reductions. Hence we have a contradiction.

As our hypotheses are not fulfilled in this example, indeed they provide a contradiction, then we can disregard the  $[\sigma] \ (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \ \mathbf{error}$  reduction.

Consider the case where the last reduction rule used in the derivation is:

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T_i \quad T_i \in \tilde{T}}{[\sigma] \ \mathbf{case} \alpha(\tilde{v}) \ \mathbf{in} \ \{\tilde{T} \mapsto t\} \xrightarrow{\alpha(\tilde{v})} [\sigma'] \ t_i}$$

As  $t = \mathbf{case} \alpha(\tilde{v}) \ \mathbf{in} \ \{\tilde{T} \mapsto t\}$  and  $\varphi \wr \emptyset \vdash \mathbf{case} \alpha(\tilde{v}) \ \mathbf{in} \ \{\tilde{T} \mapsto t\} : T$ , then by the type rules we know that the following holds:

$$\frac{\varphi_i \wr \emptyset \vdash t_i : T \quad \epsilon \wr \emptyset \vdash v_i : T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \emptyset \vdash \mathbf{case} \alpha(\tilde{v}) \ \mathbf{in} \ \{\tilde{T} \mapsto t\} : T}$$

By the reduction rule we know that  $\exists T_i \in \tilde{T}$  such that  $\sigma(\alpha(\tilde{v})): T_i$ . By the same reduction rule we know that  $\text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\}$  reduces to the associated  $t_i$ . By the type rule we know that  $\varphi_i \setminus \emptyset \vdash t_i: T$ , as required.

Consider the case where the last reduction rule used in the derivation is:

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})): T' \quad T' \notin \tilde{T}}{[\sigma] \text{ case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} \xrightarrow{\alpha(\tilde{v})} [\sigma'] \text{ error}}$$

As  $t = \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\}$  and  $\varphi \setminus \emptyset \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\}: T$ , then by the type rules we know that the following holds:

$$\frac{\varphi_i \setminus \emptyset \vdash t_i: T \quad \epsilon \setminus \emptyset \vdash v_i: T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \setminus \emptyset \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\}: T}$$

Given that  $\gamma = \alpha(\tilde{v})$ , by the definition of  $\gamma' \vdash \gamma$  we know that then any relevant  $\gamma'$  must be  $\alpha(\tilde{T})$  where  $\tilde{v}: \tilde{T}$ . There are several possible derivations for the reduction  $[\Lambda(\sigma)] \&(\alpha(\tilde{T}), T_i); \varphi_i \xrightarrow{\gamma'} [\Lambda(\sigma'')] \varphi''$ . It starts by one of the following rules.

$$\frac{\begin{array}{c} \exists i. ([\Sigma] (\alpha(\tilde{T}), T_i); \varphi_i \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'_i \wedge \text{error} \notin \varphi'_i) \\ \hline [\Sigma] \&_I (\alpha(\tilde{T}), T_i); \varphi_i \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'_i \end{array}}{\begin{array}{c} \forall i. ([\Sigma] (\alpha(\tilde{T}), T_i); \varphi_i \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'_i \wedge \text{error} \in \varphi'_i) \\ \hline [\Sigma] \&_I (\alpha(\tilde{T}), T_i); \varphi_i \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'_j \end{array}}$$

As in the hypotheses we assume that  $\text{error} \notin \varphi''$  we can disregard the second rule. The derivation then proceeds by:

$$\frac{\begin{array}{c} [\Sigma] (\alpha(\tilde{T}), T_i) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \epsilon \\ \hline [\Sigma] (\alpha(\tilde{T}), T_i); \varphi_i \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi_i \end{array}}{}$$

and finally by one of the following two rules:

$$\frac{\begin{array}{c} \Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \quad \Lambda(\sigma)(\alpha(\tilde{T})) = T_i \\ \hline [\Lambda(\sigma)] (\alpha(\tilde{T}), T_i) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')] \epsilon \end{array}}{\begin{array}{c} \Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \quad \Lambda(\sigma)(\alpha(\tilde{T})) \neq T_i \\ \hline [\Lambda(\sigma)] (\alpha(\tilde{T}), T_i) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')] \text{ error} \end{array}}$$

Consider when the first rule is used. Then we must have that  $\Lambda(\sigma)(\alpha(\tilde{T})) = T_i$ . However, by the reduction rule being considered we have that  $\nexists i. \Lambda(\sigma)(\alpha(\tilde{T})) = T_i$ ,

as  $\nexists i . \sigma(\alpha(\tilde{v})) : T'' = T_i$ , and by our assumptions we must have  $T''$  be the same as  $\Lambda(\sigma)(\alpha(\tilde{v}))$ . Hence we have a contradiction.

Consider when the second rule is used. Then we have that  $[\Lambda(\sigma)] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')]$  **error**. By our assumptions, however, we do not consider any such reductions. Hence we have a contradiction.

As our hypotheses are not fulfilled in this example, indeed they provide a contradiction, then we can disregard the  $[\sigma] \text{ case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} \xrightarrow{\alpha(\tilde{v})} [\sigma'] \text{ error}$  reduction.

□

We can straightforwardly use Lemma 3.1 to prove a similar property for processes.

**Lemma 3.2.** *Thread Subject Reduction*

For all  $\Phi, \Phi'', P, P', \sigma, \sigma'', \gamma, \gamma'$  if  $\vdash P : \Phi$ , and  $[\sigma] P \xrightarrow{\gamma} [\sigma'] P'$ , and if whenever  $\gamma' \vdash \gamma$  and  $[\Lambda(\sigma)] \Phi \xrightarrow{\gamma'} [\Lambda(\sigma'')] \Phi''$  then **error**  $\notin \varphi''$ , then there exists a  $\Phi'$  such that typing judgement  $\vdash P' : \Phi'$  holds.

*Proof.* Suppose the hypotheses. We then prove the conclusions simultaneously by induction over the derivation of  $[\sigma] P \xrightarrow{\gamma} [\sigma'] P'$ . □

### 3.4.1.2 Liveness

By Equation 3.9 we know that the effect of a Globally Compatible program point can either perform a reduction, or it consists of empty effects and recursion variables. Using this fact we can prove that such program points have various liveness properties. Firstly we can show that if an expression has an empty effect that it can either perform a reduction, or it is a value.

**Lemma 3.3.** *Empty Effect Expression Liveness*

For all  $t, T$ , if  $\epsilon \wr \emptyset \vdash t : T$  then we know that either  $t = v$ , or there exists a  $t'$  such that  $[\sigma] t \xrightarrow{\tau} [\sigma] t'$ .

*Proof.* Suppose the hypotheses. We then prove the conclusions by induction over the derivation of  $\epsilon \wr \emptyset \vdash t : T$ . The key cases are as follows.

Consider the case where the last typing rule used in the derivation is:

$$\frac{\epsilon \wr \emptyset \vdash t_1 : T_2 \xrightarrow{\epsilon} T_1 \quad \epsilon \wr \emptyset \vdash t_2 : T_2}{(\epsilon; \epsilon; \epsilon) \wr \emptyset \vdash t_1 t_2 : T_1}$$

As  $\epsilon \wr \emptyset \vdash t_1: T_2 \xrightarrow{\epsilon} T_1$  and  $\epsilon \wr \emptyset \vdash t_2: T_2$  we can apply the inductive step to each of them and show that  $[\sigma] t_1 \xrightarrow{\tau} [\sigma] t'_1 \vee t_1 = v_1$ , and  $[\sigma] t_2 \xrightarrow{\tau} [\sigma] t'_2 \vee t_2 = v_2$ . If  $[\sigma] t_1 \xrightarrow{\tau} [\sigma] t'_1$  then we can perform the following reduction:

$$\frac{[\sigma] t_1 \xrightarrow{\tau} [\sigma'] t'_1}{[\sigma] t_1 t_2 \xrightarrow{\tau} [\sigma'] t'_1 t_2}$$

as required. If  $t_1 = v_1$  then we consider the properties of  $t_2$ . If  $[\sigma] t_2 \xrightarrow{\tau} [\sigma] t'_2$  then, given that  $t_1 = v_1$ , we can perform the following reduction:

$$\frac{[\sigma] t_2 \xrightarrow{\tau} [\sigma'] t'_2}{[\sigma] v t'_2 \xrightarrow{\tau} [\sigma'] v t'_2}$$

as required. The final case is when  $t_1 = v_1$  and  $t_2 = v_2$ . By the typing rule, we know that  $\epsilon \wr \emptyset \vdash t_1: T_2 \xrightarrow{\epsilon} T_1$ , and hence that  $t_1 = \mathbf{rec} \underline{X}(x: T).t$ . Hence we can perform the following reduction:

$$\overline{[\sigma] \mathbf{rec} \underline{X}(x: T).t v_2 \xrightarrow{\tau} [\sigma] t[\mathbf{rec} \underline{X}(x: T).t / \underline{X}][v_2/x]}$$

as required.

Consider the case where the last typing rule used in the derivation is:

$$\frac{\epsilon \wr \emptyset \vdash v_i: T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \emptyset \vdash (\alpha(\tilde{v}), T): T}$$

By the hypothesis we know that

$$(\alpha(\tilde{v}), T) \equiv \epsilon$$

This is a contradiction, and hence we can disregard this case.

□

We then consider non-empty effects. We show that if a program has an effect that can reduce then that program can always reduce.

**Lemma 3.4.** *Active Effect Expression Liveness*

For all  $\varphi, \varphi', t, T, \gamma', \sigma', \sigma$  if  $\varphi \wr \emptyset \vdash t: T$  and  $[\Lambda(\sigma)] \varphi \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'$  then there exists  $t', \gamma$  that  $[\sigma] t \xrightarrow{\gamma} [\sigma'] t'$ .

*Proof.* Suppose the hypotheses. We then prove the conclusions by induction over the derivation of  $\epsilon \wr \emptyset \vdash t: T$ . The key cases are as follows.

Consider the case where the last typing rule used in the derivation is:

$$\frac{\epsilon \vdash \emptyset \vdash v_i : T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \vdash \emptyset \vdash (\alpha(\tilde{v}), T) : T}$$

By the hypothesis we know that  $[\Lambda(\sigma)] (\alpha(\tilde{T}), T) \rightarrow [\Lambda(\sigma')] \varphi'$ . There are two relevant reduction rules for that derive this reduction:

$$\frac{\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \quad \Lambda(\sigma)(\alpha(\tilde{T})) = T}{[\Lambda(\sigma)] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')] \epsilon} \quad \frac{\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \quad \Lambda(\sigma)(\alpha(\tilde{T})) \neq T}{[\Lambda(\sigma)] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Lambda(\sigma')] \text{ error}}$$

Consider if the reduction is performed using the former. By our assumptions about the abstract and concrete state transition functions we know that the abstract state can reduce if and only if the concrete state can reduce, on a given action. By the premise of the reduction we know that  $\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma')$ , and hence that  $\sigma \xrightarrow{\alpha(\tilde{v})} \sigma'$ . We also know that  $\Lambda(\sigma)(\alpha(\tilde{T})) = T$ . By our assumptions about the abstract and concrete state projection function we know that  $\Lambda(\sigma)(\alpha(\tilde{T})) = T$  if and only if  $\sigma(\alpha(\tilde{v})) : T$ . We then have all the premises required in order to use the following rule:

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T}{[\sigma] (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \sigma(\alpha(\tilde{v}))}$$

to reduce  $(\alpha(\tilde{v}), T)$  as required.

Consider if the reduction is performed using the latter. By our assumptions about the abstract and concrete state transition functions we know that the abstract state can reduce if and only if the concrete state can reduce, on a given action. By the premise of the reduction we know that  $\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma')$ , and hence that  $\sigma \xrightarrow{\alpha(\tilde{v})} \sigma'$ . We also know that  $\Lambda(\sigma)(\alpha(\tilde{T})) \neq T$ . By our assumptions about the abstract and concrete state projection function we know that  $\Lambda(\sigma)(\alpha(\tilde{T})) = T$  if and only if  $\sigma(\alpha(\tilde{v})) : T'$  and  $T' \neq T$ . We then have all the premises required in order to use the following rule:

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T' \neq T}{[\sigma] (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \text{ error}}$$

to reduce  $(\alpha(\tilde{v}), T)$  as required.

Consider the case where the last typing rule used in the derivation is:

$$\frac{\varphi_1 \vdash \emptyset \vdash t_1 : T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \vdash \emptyset \vdash t_2 : T_2}{(\varphi_1; \varphi_2; \varphi_3) \vdash \emptyset \vdash t_1 t_2 : T_1}$$

By the hypotheses we know that  $[\Lambda(\sigma)] \varphi_1; \varphi_2; \varphi_3 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'$ , and by using a combination of the rule for sequence and equivalent effect reduction:

$$\frac{[\Sigma] \varphi_1 \xrightarrow{\gamma} [\Sigma'] \varphi'_1 \quad \Phi \equiv \Phi_1 \quad [\Sigma] \Phi_1 \xrightarrow{\gamma} [\Sigma'] \Phi_2 \quad \Phi_2 \equiv \Phi'}{[\Sigma] \varphi_1; \varphi_2 \xrightarrow{\gamma} [\Sigma'] \varphi'_1; \varphi_2} \quad \frac{[\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi'}{[\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi'}$$

we know that either  $[\Lambda(\sigma)] \varphi_1; \varphi_2; \varphi_3 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_1; \varphi_2; \varphi_3$  or  $[\Lambda(\sigma)] \varphi_1; \varphi_2; \varphi_3 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_2; \varphi_3$  where  $\varphi_1 \equiv \epsilon$ , or  $[\Lambda(\sigma)] \varphi_1; \varphi_2; \varphi_3 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_3$  where  $\varphi_1 \equiv \varphi_2 \equiv \epsilon$ . In each case we also know that  $[\Lambda(\sigma)] \varphi_i \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_i$ .

Consider the case where  $[\Lambda(\sigma)] \varphi_1; \varphi_2; \varphi_3 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_1; \varphi_2; \varphi_3$ . As  $\varphi_1 \setminus \emptyset \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1$  and  $[\Lambda(\sigma)] \varphi_1 \xrightarrow{\gamma} [\Lambda(\sigma')] \varphi'_1$  we can apply the inductive step and hence know that  $[\sigma] t_1 \xrightarrow{\gamma} [\sigma'] t'_1$ . Hence we can reduce  $t_1 t_2$  using the following rule:

$$\frac{[\sigma] t_1 \xrightarrow{\gamma} [\sigma'] t'_1}{[\sigma] t_1 t_2 \xrightarrow{\gamma} [\sigma'] t'_1 t_2}$$

as required.

Consider the case where  $[\Lambda(\sigma)] \varphi_1; \varphi_2; \varphi_3 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_2; \varphi_3$ . As  $\varphi_2 \setminus \emptyset \vdash t_2: T_2$  and  $[\Lambda(\sigma)] \varphi_2 \xrightarrow{\gamma} [\Lambda(\sigma')] \varphi'_2$  we can apply the inductive step and hence know that  $[\sigma] t_2 \xrightarrow{\gamma} [\sigma'] t'_2$ . By Lemma 3.3 we know that, as  $\varphi_1 \equiv \epsilon$  that either  $[\sigma] t_1 \xrightarrow{\gamma} [\sigma'] t'_1$  or  $t_1 = v_1$ . If the former then we can reduce  $t_1 t_2$  using the following rule:

$$\frac{[\sigma] t_1 \xrightarrow{\gamma} [\sigma'] t'_1}{[\sigma] t_1 t_2 \xrightarrow{\gamma} [\sigma'] t'_1 t_2}$$

as required. If the latter then we can reduce  $t_1 t_2$  using the following rule:

$$\frac{[\sigma] t_2 \xrightarrow{\gamma} [\sigma'] t'_2}{[\sigma] v t'_2 \xrightarrow{\gamma} [\sigma'] v t'_2}$$

as required. Consider the case when  $t_1 = v_1$ . By Lemma 3.3 we know that, as  $\varphi_2 \equiv \epsilon$  that either  $[\sigma] t_2 \xrightarrow{\gamma} [\sigma'] t'_2$  or  $t_2 = v_2$ . If the former then we can reduce  $t_1 t_2$  using the following rule:

$$\frac{[\sigma] t_2 \xrightarrow{\gamma} [\sigma'] t'_2}{[\sigma] v t'_2 \xrightarrow{\gamma} [\sigma'] v t'_2}$$

as required. If  $t_1 = v_1$  and  $t_2 = v_2$  then, as  $\varphi_1 \setminus \emptyset \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1$ , we know that  $t_1 = \mathbf{rec} \underline{X}(x: T).t$ , and hence can reduce  $t_1 t_2$  using the following rule:

$$\frac{[\sigma] \mathbf{rec} \underline{X}(x: T).t v \xrightarrow{\tau} [\sigma] t[\mathbf{rec} \underline{X}(x: T).t / \underline{X}][v/x]}{}$$

as required.

The case where  $[\Lambda(\sigma)] \varphi_1; \varphi_2; \varphi_3 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_3$  is similar to the case where  $[\Lambda(\sigma)] \varphi_1; \varphi_2; \varphi_3 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_2; \varphi_3$ .

□

We can extend these results straightforwardly to threads.

**Lemma 3.5.** *Empty Effect Thread Liveness Lemma*

For all  $P$ , if  $\vdash P: \Pi_I \epsilon$  then we know that either  $P = \Pi_I v$ , or there exists a  $P'$  such that  $[\sigma] P \xrightarrow{\tau} [\sigma] P'$ .

*Proof.* Suppose the hypotheses. We then simultaneously prove the conclusions by induction over the derivation of  $\vdash P: \Pi_I \epsilon$ . □

**Lemma 3.6.** *Active Effect Thread Liveness Lemma*

For all  $\Phi, \Phi', P, \gamma', \sigma', \sigma$  if  $\vdash P: \Phi$  and  $[\Lambda(\sigma)] \Phi \xrightarrow{\gamma'} [\Lambda(\sigma')] \Phi'$  then there exists  $P', \gamma$  that  $[\sigma] P \xrightarrow{\gamma} [\sigma'] P'$ .

*Proof.* Suppose the hypotheses. We then prove the conclusions by induction over the derivation of  $\vdash P: \Phi$ . □

### 3.4.2 Key Theorems

Checking whether Global Compatibility holds for a given system amounts to model checking, as the predicate involves considering all possible interleavings of effect reductions. As such a check is computationally expensive we permit an analysis to use stronger properties that imply Global Compatibility, which are often easier to prove. We express this using the following type rule:

$$\frac{\vdash P: \Phi \quad \exists \mathcal{C}. \mathcal{C}(\Lambda(\sigma), \Phi) \text{ and } (\mathcal{C} \implies \mathcal{G})}{\vdash [\sigma] P} \quad (3.12)$$

This denotes that, if we can find a property which implies Global Compatibility, then we consider programs and states that fulfil that property to be Globally Compatible. We refer to program points which can be validated using the rule in Equation 3.12 to be *valid*. We explore the stronger properties in Section 3.5.

We can prove key theorems using the lemmas in Section 3.4.1. Firstly we prove Subject Reduction. This states that if a valid program point performs a reduction, that the new program point is also valid.

**Theorem 3.7.** *Valid Program Points Subject Reduction*

$$\vdash [\sigma] P \wedge [\sigma] P \rightarrow [\sigma'] P' \Rightarrow \vdash [\sigma'] P'$$

*Proof.*

We know that:

$$\frac{\vdash P : \Phi \quad \exists \mathcal{C}. \mathcal{C}(\Lambda(\sigma), \Phi) \text{ and } (\mathcal{C} \implies \mathcal{G})}{\vdash [\sigma] P}$$

By the definition of global compatibility we know that Equation 3.8 holds. We can then show that:

$$[\Sigma] \varphi \rightarrow [\Sigma'] \varphi' \Rightarrow \#\text{error} \in \varphi' \quad (3.13)$$

Hence we can use Lemma 3.2 to show, by choosing  $\mathcal{C} = \mathcal{G}$  that:

$$\vdash P' : \Phi' \quad [\Lambda(\sigma)] \Phi \xrightarrow{\gamma'} [\Lambda(\sigma')] \Phi' \quad (3.14)$$

By Equation 3.8 we know that  $\mathcal{G}(\Sigma', \Phi')$  holds. Hence we can show that:

$$\frac{\vdash P' : \Phi' \quad \exists \mathcal{C}. \mathcal{C}(\Lambda(\sigma'), \Phi') \text{ and } (\mathcal{C} \implies \mathcal{G})}{\vdash [\sigma'] P'}$$

□

We also prove that all valid program points are live; this denotes that a set of threads can either reduce, or each thread consists of a value. In particular this denotes, if there are several possible actions in reducible positions, that at least one must not be blocked.

**Theorem 3.8.** *Liveness of Valid Program Points*

$$\vdash [\sigma] P \Rightarrow [\sigma] P \rightarrow [\sigma'] P' \vee P = \Pi v$$

*Proof.* By the definition of global compatibility we know that Equation 3.9 holds. We can perform a case analysis on the two disjuncts.

$$[\Sigma] \Phi \rightarrow [\Sigma'] \Phi \vee \Phi \equiv \Pi \epsilon \quad (3.15)$$

**Case**  $(\Phi = \Pi \epsilon)$

By Lemma 3.5 we have that:

$$[\sigma] P \xrightarrow{\gamma} [\sigma] P' \vee P = \Pi_I v \quad (3.16)$$

**Case**  $[\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$

$$\begin{aligned}
d_a \rightarrow d_b : c\langle \widetilde{T \mapsto G} \rangle_M \upharpoonright d_a &\stackrel{\text{def}}{=} \oplus_N(c!\langle T_n \rangle; G_n \upharpoonright d_a) \quad \emptyset \neq N \subseteq M \\
d_a \rightarrow d_b : c\langle \widetilde{T \mapsto G} \rangle_M \upharpoonright d_b &\stackrel{\text{def}}{=} \&_M(c?(T_m); G_m \upharpoonright d_b) \\
d_a \rightarrow d_b : c\langle \widetilde{T \mapsto G} \rangle_M \upharpoonright d &\stackrel{\text{def}}{=} G_1 \upharpoonright d \\
(\mu \underline{X}. G') \upharpoonright d &\stackrel{\text{def}}{=} \mu \underline{X}.(G' \upharpoonright d) \\
\underline{X} \upharpoonright d &\stackrel{\text{def}}{=} \underline{X} \\
0 \upharpoonright d &\stackrel{\text{def}}{=} 0
\end{aligned}$$

Figure 3.12: Global Session Types Projection

By Lemma 3.6.

$$[\sigma] \ P \xrightarrow{\gamma} [\sigma'] \ P' \quad (3.17)$$

□

Theorems 3.7 and 3.8 can also be proved as corollaries of Theorems 4.1 and 4.2 respectively

## 3.5 Specific Formulations

Instead of investigating all possible inter leavings of a system, we attempt to prove stronger properties that imply Global Compatibility, as these properties are often easier to prove. The challenge is then to define a  $\mathcal{C}$  such that  $\mathcal{C} \implies \mathcal{G}$ . In this section we describe two such compatibility predicates, one for message passing systems that use blocking receives (Section 3.5.1), and one for message passing systems that use non-blocking receives (Section 3.5.2).

### 3.5.1 Blocking Message Passing

Global Session Types (Figure 2.9) describe communication protocols between multiple role at a high level of abstraction (Bettini et al., 2008). Intuitively each participant role is played out by one thread in a program. We make use of the projection of Global Session Types to effects (Figure 3.12). Note that, in the case where the participant being projected to is neither the sender nor the receiver, we require that the projection of all the continuations be the same, so we (arbitrarily) project for the first one. We use the notation  $c!(T)$  as shorthand for  $(\text{snd}(\text{Res } c, T), \text{UNIT})$  and  $c?(T)$  as shorthand for  $(\text{rcv}(\text{Res } c), T)$ .

In order for a program to avoid communication errors and deadlock it is not sufficient for the effect of the program to be a projection of a Global Session Type. Consider the

$$\begin{aligned}
\Delta &\stackrel{\text{def}}{=} \bigcup_{n < \omega} \Delta_n & \Delta_0 &\stackrel{\text{def}}{=} \{ [\Sigma_\emptyset] \Phi^{\uparrow G} \} \\
\Delta_{n+1} &\stackrel{\text{def}}{=} \{ \delta([c \mapsto T; \Sigma^k], \&_J(c?(T_j); \varphi_b^j), \varphi_a^k) \mid \forall j \in J. \delta(\Sigma^j, \varphi_b^j, \varphi_a^j) \in \Delta_n \} \\
&\cup \{ \delta([\Sigma], \oplus_K(c!(T_k); \varphi_a^k), \&_J(c?(T_j); \varphi_b^j) \mid \\
&\quad \forall k \in K. \delta(\Sigma, \varphi_a^k, \varphi_b^k) \in \Delta_n \wedge \Sigma(c) = \emptyset) \\
&\cup \Delta_n
\end{aligned}$$

where  $c = d_a d_b$  and  $\emptyset \subset K \subseteq J$ .

Figure 3.13: Blocking Message Passing Safe Program Points Set

following program:

$$\begin{aligned}
1 & [c \mapsto 4] \text{ snd}(c, \text{true}).0 \parallel \text{rcv}(c)(x : \text{BOOL}).0 \\
2 & [\rightarrow] [c \mapsto 4, \text{true}] 0 \parallel \text{rcv}(c)(x : \text{BOOL}).0 \\
3 & [\rightarrow] [c \mapsto \text{true}] 0 \parallel \text{error}
\end{aligned} \tag{3.18}$$

The effect of the code on line 1 is  $c!(\text{BOOL}) \parallel c?(\text{BOOL})$  which is complementary, according to the standard definition (Honda, 1993). However, since there are messages in the queue, in order to rule out communication errors and deadlock we cannot consider solely the effect of the code. We must also consider the state of the message queues. We do this by ensuring that the messages in the queue are complementary to the effect of the code, and that after that point the effects are complementary (Bettini et al., 2008).

### 3.5.1.1 Safe Program Points

We use the notation  $\delta(\Sigma, \varphi_i, \varphi_j)$  to represent some parallel configuration with two particular participants  $p_i, p_j$  identified:

$$[\Sigma] \varphi_1 \parallel \dots \parallel \varphi_i \parallel \dots \parallel \varphi_j \parallel \dots \parallel \varphi_n$$

By writing  $\delta(\Sigma', \varphi'_i, \varphi'_j)$  we refer to this same configuration but with possible changes in the abstract state and identified participants' effects. We make use of  $\delta(\Sigma, \varphi_i)$  to similarly single out one participant from a parallel configuration.

We inductively define a set of safe program points, which can be reduced without causing error or deadlock (Figure 3.13). The base case  $\Delta_0$  consists of the projection of global session types, alongside empty message queues. The inductive step defines new valid points, using existing valid points, in one of two ways. Intuitively, the first adds a message to the front of the queue and a receive action that can expects a value of that type. The second adds a send and a receive that match up in terms of the type of values sent and received. The above intuition is complicated by the provision of choice, as is explained below. We explain these cases by example.

**Adding a message.** Consider the following Global Session Type, and a program point of its projection:

$$\begin{aligned} G = & d_1 \rightarrow d_2 : c \langle l_1 \mapsto \\ & d_1 \rightarrow d_2 : c \langle T_1 \mapsto 0 \rangle \\ & \rangle \end{aligned} \quad (3.19)$$

$$[\Sigma_\emptyset] c!(l_1); c!(T_1) \parallel c?(l_1); c?(T_1) \in \Delta_0 \quad (3.20)$$

This program point can perform a reduction that sends the label on the shared channel:

$$[\Sigma_\emptyset] c!(l_1); c!(T_1) \parallel c?(l_1); c?(T_1) \rightarrow [c \mapsto l_1] c!(T_1) \parallel c?(l_1); c?(T_1) \quad (3.21)$$

We can show that this new program point, derived from the first, is in the set of safe program points, using the extensions defined in Figure 3.13. Consider the global session type  $d_1 \rightarrow d_2 : c \langle T_1 \mapsto 0 \rangle$ . We can define a program point that is a projection of this global session type, and hence in  $\Delta_0$ :

$$[\Sigma_\emptyset] c!(T_1) \parallel c?(T_1) \quad (3.22)$$

This point can be extended by adding a message of the correct type to the queue, and adding a receive action of the same type, to the receiving process:

$$[c \mapsto \text{Lab } l_1] c!(T_1) \parallel c?( \text{Lab } l_1); c?(T_1) \quad (3.23)$$

This approach of adding a message onto a channel queue, and a corresponding receive, is relatively straightforward, and forms the basis of the first inductive definition in Figure 3.13. Were this the limit of what we wished to permit, we could define this first inductive step as:

$$\{ \delta([c \mapsto T; \Sigma], c?(T_j); \varphi_b, \varphi_a) \mid \delta(\Sigma, \varphi_b, \varphi_a) \in \Delta_n \} \quad (3.24)$$

When a value is received, however, it is possible for a case split to be performed on its type. In order to account for this case we must construct the set in a slightly more complex manner. Consider the following global session type:

$$\begin{aligned} G = & d_1 \rightarrow d_2 : c \langle \\ & l_1 \mapsto (d_1 \rightarrow d_2 : c \langle T_1 \mapsto 0 \rangle) \\ & l_2 \mapsto (d_1 \rightarrow d_2 : c \langle T_2 \mapsto 0 \rangle) \\ & \rangle \end{aligned} \quad (3.25)$$

and a program point with one of its projections:

$$[\Sigma_\emptyset] c!(\text{Lab } l_1); c!(T_1) \parallel (c?( \text{Lab } l_1); c?(T_1)) \& (c?( \text{Lab } l_2); c?(T_2)) \in \Delta_0 \quad (3.26)$$

This program point can perform a reduction that sends the label on the shared channel:

$$[\Sigma_\emptyset] \quad c! \langle \text{Lab } l_1 \rangle; c! \langle T_1 \rangle \parallel (c?(\text{Lab } l_1); c?(T_1)) \& (c?(\text{Lab } l_2); c?(T_2)) \rightarrow \quad (3.27)$$

$$[c \mapsto \text{Lab } l_1] \quad c! \langle T_1 \rangle \parallel (c?(\text{Lab } l_1); c?(T_1)) \& (c?(\text{Lab } l_2); c?(T_2))$$

Using the former approach we can straightforwardly show that the following program points are valid:

$$[c \mapsto \text{Lab } l_1] \quad c! \langle T_1 \rangle \parallel c?(\text{Lab } l_1); c?(T_1) \quad [c \mapsto \text{Lab } l_2] \quad c! \langle T_2 \rangle \parallel c?(\text{Lab } l_2); c?(T_2) \quad (3.28)$$

We can combine these two points, using the state and the effect of the sender from one program point, and performing an external choice between the effects of the receiver in each program point, to generate the program point:

$$[c \mapsto \text{Lab } l_1] \quad c! \langle T_1 \rangle \parallel (c?(\text{Lab } l_1); c?(T_1)) \& (c?(\text{Lab } l_2); c?(T_2)) \quad (3.29)$$

Note that the two points being combined have different accompanying states, and this is reflected in the precondition for this construction  $\delta(\Sigma^j, \varphi_b^j, \varphi_a^j) \in \Delta_n$ .

In the full version of the ‘adding a message’ case, we take a set of valid program points  $\{\varphi_{j_1}, \dots, \varphi_{j_n}\}$ , such as those in Equation 3.28, where the effects of all threads, apart from those of the proposed sender and receiver, are the same. We sequence a receive action of the associated type onto each continuation, for the receive effects, so that we obtain  $\{c?(T_{j_1}); \varphi_b^{j_1}, \dots, c?(T_{j_n}); \varphi_b^{j_n}\}$ . We define an external choice between these effects, and obtain

$$c?(T_{j_1}); \varphi_b^{j_1} \& \dots \& c?(T_{j_n}); \varphi_b^{j_n} \quad (3.30)$$

We then choose some  $k$  in  $j_1 \dots j_n$ , and add a message  $T_k$  to the front of the queue of the designated channel between the sender and the receiver. As the  $T_{j_1}, \dots, T_{j_n}$  are required to be distinct, we know that  $c?(T_{j_1}); \varphi_b^{j_1} \& \dots \& c?(T_{j_n}); \varphi_b^{j_n}$ , with message  $T_k$  at the front of its queue, can only ever evolve to  $\varphi_b^k$ . Therefore, in order to maintain duality, we pair this external choice with the effect  $\varphi_a^k$ , obtaining a final definition:

$$\{ \delta([c \mapsto T; \Sigma^k], \&_J (c?(T_j); \varphi_b^j), \varphi_a^k) \mid \forall j \in J. \delta(\Sigma^j, \varphi_b^j, \varphi_a^j) \in \Delta_n \} \quad (3.31)$$

**Adding a send-receive pair.** Consider the following Global Session Type:

$$d_1 \rightarrow d_2 : c_1(T_1); d_3 \rightarrow d_2 : c_2(T_3) \quad (3.32)$$

whose projection is:

$$[\Sigma_\emptyset] \quad c_1! \langle T_1 \rangle \parallel c_1? (T_1); c_2? (T_3) \parallel c_2! \langle T_3 \rangle \in \Delta_0 \quad (3.33)$$

We can show that we can derive this program point from the projection of the session:

$$d_3 \rightarrow d_2 : c_2(T_3) \quad (3.34)$$

whose projection is:

$$[\Sigma_\emptyset] \epsilon \parallel c_2?(T_3) \parallel c_2!(T_3) \in \Delta_0 \quad (3.35)$$

This point can be extended by adding the dual pair of  $c_1!(T_1)$  and  $c_1?(T_1)$  between the first and second processes

$$[\Sigma_\emptyset] c_1!(T_1) \parallel c_1?(T_1); c_2?(T_3) \parallel c_2!(T_3) \quad (3.36)$$

which returns us to the original projection. Given that the original was also a projection of a global session type, this may seem like a redundant construction, but it is essential when we consider how the  $c_2?(T_3) \parallel c_2!(T_3)$  interaction can occur with respect to the  $c_1!(T_1) \parallel c_1?(T_1)$  interaction.

We can show that the full system can reduce as:

$$[\Sigma_\emptyset] c_1!(T_1) \parallel c_1?(T_1); c_2?(T_3) \parallel c_2!(T_3) \xrightarrow{c_2!(T_3)} [c_2 \mapsto T_3] c_1!(T_1) \parallel c_1?(T_1); c_2?(T_3) \parallel \epsilon \quad (3.37)$$

This particular program point cannot be constructed by projecting a global session type and then using the emphadding a message technique, as that adds message/receive pairs to the *front* of the effect of the receiving process. Instead, we start with an empty session, and directly add a message/receive pair between the third and second processes:

$$[c_2 \mapsto T_3] \epsilon \parallel c_2?(T_3) \parallel \epsilon \quad (3.38)$$

We then make use of the *adding a send/receive pair* to add the not-yet-commenced dual interaction between the first and the second processes

$$[c_2 \mapsto T_3] c_1!(T_1) \parallel c_1?(T_1); c_2?(T_3) \parallel \epsilon \quad (3.39)$$

Note the side condition for this case, that the channel queue being added to must be empty. If we didn't have this side condition, then we could construct a program point such as

$$[c_2 \mapsto T_3] \epsilon \parallel c_2?(T_1); c_2?(T_3) \parallel c_2!(T_1); \epsilon \quad (3.40)$$

by adding the send/receive pair  $c_2!(T_1)$  and  $c_2?(T_1)$  to the third and second processes respectively. As the message queue for  $c_2$  is not empty, and we can place a receive expecting an arbitrary type at the front of the receiving process, the existing message in the queue will probably not be compatible with the first receive action of the receiving thread.

This construction makes use of a key design decision. We use a separate channel for each

direction of each pair of communicating roles (principle channel allocation (Deniéou and Yoshida, 2010)). This enables us to determine that it is safe to add a send/receive pair in front of existing message/receive pairs. The requirement that the channel being used is empty, and that the receiving process can only receive from the sending process on the channel being used, guarantees that the send/receive pair being added will not interact in error with any existing messages. For all message passing instances, for the remainder of the thesis, we assume that principle channel allocation is used.

### 3.5.1.2 Formal Properties

We can prove that if a program point is in the valid points set for a given  $G$ , and if the point reduces, then it stays in a valid points set, for a possibly different  $G'$ .

**Lemma 3.9.** *Valid Points Set Subject Reduction*

$$\forall \Sigma, \Phi, n, \Sigma', \Phi'. [\Sigma] \Phi \in \Delta_n \wedge [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow [\Sigma'] \Phi' \in \Delta_{n+1}$$

*Proof.* Suppose the hypotheses. We then prove the conclusion by induction over the  $n$  of  $\Delta_n$ .

Consider the case where  $n = 0$ . We then continue by a case analysis on the last reduction rule used in the derivation of  $[\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$ .

Consider the case when the last reduction performed is an internal choice. By the definition of  $\Delta_n$  in Figure 3.13, we know that  $[\Sigma] \Phi$  is of the form:

$$[\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel \oplus_N c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_m$$

Hence, by the rule for internal choice:

$$\frac{i \in 1, 2}{[\Sigma] \varphi_1 \oplus \varphi_2 \xrightarrow{\tau} [\Sigma] \varphi_i}$$

we know that:

$$[\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel \oplus_N c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_m \rightarrow [\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_m$$

By Lemma B.2 we know that:  $\exists G' = d_a \rightarrow d_b: c\langle \widetilde{T \mapsto G} \rangle_M . N \subseteq M \wedge \Pi G' \upharpoonright d_i \equiv \Pi \varphi_i$ . By the definition of the projection relation  $\upharpoonright$  we know that it is possible to project  $G'$  such that  $G' \upharpoonright d_a \equiv c!(T_n); \varphi_a^n$ . Hence we can show that  $\varphi_1 \parallel \dots \parallel c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_m \equiv \Pi G' \upharpoonright d_i$  and  $[\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_m \in \Delta_0$ , as required.

Consider the case when the last reduction performed is a send action, in sequence with some continuation. By the definition of  $\Delta_n$  in Figure 3.13, we know that  $[\Sigma] \Phi$  is of

the form  $[\Sigma] \varphi_1 \parallel \dots \parallel c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_m$ . By the action reduction and sequencing rules:

$$\frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) = T}{[\Sigma] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \epsilon} \quad \frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) \neq T}{[\Sigma] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \text{error}}$$

$$\frac{[\Sigma] \varphi_1 \xrightarrow{\gamma_F} [\Sigma'] \varphi'_1}{[\Sigma] \varphi_1; \varphi_2 \xrightarrow{\gamma_F} [\Sigma'] \varphi'_1; \varphi_2}$$

we know that, if  $[\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_o \rightarrow [c \mapsto T_n] \varphi_1 \parallel \dots \parallel \varphi_a^n \parallel \dots \parallel \varphi_o$  then  $[\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_o \rightarrow [c \mapsto T_n] \varphi_1 \parallel \dots \parallel \varphi_a^n \parallel \dots \parallel \varphi_o$ .

By the first part of the definition of  $\Delta_{n+1}$  in Figure 3.13, in order to show that  $\delta(c \mapsto T, \varphi_a^k, \&_J(c?(T_j); \varphi_b^j)) \in \Delta_1$ , we need to show that for each  $j \in J$  that  $\delta(\Sigma_\emptyset, \varphi_a^k, \varphi_b^j) \in \Delta_0$ .

By Lemma B.2 we know that  $\exists G' = d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle_M . \Pi G' \upharpoonright d_i \equiv \Pi \varphi_i$ . Hence we have that, for each of the threads that are not the sender or receiver threads, that  $\varphi_i \equiv G' \upharpoonright d_i \equiv G_d \upharpoonright d_i$ , for any  $d \in M$ . In such cases let  $\varphi'_i \equiv \varphi_i$ . By the definition of projection we know that  $G' \upharpoonright d_b \equiv \&_M c?(T_m); (G_m \upharpoonright d_b) \equiv \&_M c?(T_m); \varphi_b^m$  and  $G' \upharpoonright d_a \equiv c!(T_n); (G_n \upharpoonright d_a) \equiv c!(T_n); \varphi_a^n$ . Hence we have that  $[\Sigma_\emptyset] \Pi \varphi_i^m \equiv \Pi G_m \upharpoonright d_i \in \Delta_0$ , as required.

Consider the case where  $n > 0$ . By the structure of parallel effects we can identify one effect which is the one that reduces:

$$[\Sigma] \varphi_1 \parallel \dots \parallel \varphi_i \parallel \dots \parallel \varphi_m \rightarrow [\Sigma'] \varphi_1 \parallel \dots \parallel \varphi'_i \parallel \dots \parallel \varphi_m$$

As  $[\Sigma] \Phi \in \Delta_n$  we know that either in one of the constructs included in the definition of  $\Delta_n$ , or it is in  $\Delta_{n-1}$ .

Consider the case when  $[\Sigma] \Phi \in \{ \delta([c \mapsto T; \Sigma^k], \&_J(c?(T_j); \varphi_b^j)) \}$  where  $\forall j \in J. \delta(\Sigma^j, \varphi_b^j) \in \Delta_{n-1}$  (as in Figure 3.13). We can perform a case split over the index  $i$ . If  $i = b$  then, by the reduction rules, we know that  $\varphi'_i \equiv \varphi_b^k$  and that  $\Sigma' = \Sigma^k$ . Hence we have that  $[\Sigma'] \Phi' \equiv \delta(\Sigma^k, \varphi_a^k, \varphi_b^k)$ . Hence by our hypotheses we know that  $[\Sigma'] \Phi' \in \Delta_n$ . Consider the case when  $i \neq b$ . By Lemma B.1 we know that:

$$[c \mapsto T; \Sigma^k] \varphi_1 \parallel \dots \parallel \varphi_i \parallel \dots \parallel \varphi_m \rightarrow [c \mapsto T; \Sigma'^k] \varphi_1 \parallel \dots \parallel \varphi'_i \parallel \dots \parallel \varphi_m$$

if and only if:

$$[\Sigma^k] \varphi_1 \parallel \dots \parallel \varphi_i \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \parallel \varphi_m \rightarrow$$

$$[\Sigma'^k] \varphi_1 \parallel \dots \parallel \varphi'_i \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \parallel \varphi_m$$

Then, by the parallel reduction rules, we know that:

$$[\Sigma^k] \varphi_1 \parallel \dots \parallel \varphi_i \parallel \dots \parallel \varphi_b^k \parallel \dots \parallel \varphi_m \rightarrow [\Sigma'^k] \varphi_1 \parallel \dots \parallel \varphi'_i \parallel \dots \parallel \varphi_b^k \parallel \dots \parallel \varphi_m$$

By the hypotheses we know that  $[\Sigma^k] \varphi_1 \parallel \dots \parallel \varphi_i \parallel \dots \parallel \varphi_b^k \parallel \dots \parallel \varphi_m \in \Delta_{n-1}$ . Hence, by the inductive step, we know that  $[\Sigma'^k] \varphi_1 \parallel \dots \parallel \varphi'_i \parallel \dots \parallel \varphi_b^k \parallel \dots \parallel \varphi_m \in \Delta_{n-1}$ . By the definition of  $\Delta_n^{\uparrow G}$  in Figure 3.13 we hence know that:

$$[\Sigma'^k] \varphi_1 \parallel \dots \parallel \varphi'_i \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \parallel \varphi_m \in \Delta_n$$

as required.

Consider the case when  $[\Sigma] \Phi \in \{ \delta(\Sigma, \oplus_K(c!(T_k); \varphi_a^k), \&_J(c?(T_j); \varphi_b^j)) \}$  where  $\forall j \in J. \delta(\Sigma^j, \varphi_a^j, \varphi_b^j) \in \Delta_{n-1}^{\uparrow G}$  and  $\Sigma(c) = \emptyset$  (as in Figure 3.13). We can perform a case split over the index  $i$ . If  $i = b$  then, by the reduction rules we know that  $\Sigma(c) = T, q$ , which contradicts the hypotheses, and hence can disregard this case. Consider the case when  $i = a$ . If  $\oplus_K(c!(T_k); \varphi_a^k)$  is a choice then we can use the internal choice reduction rules to show that:

$$[\Sigma] \varphi_1 \parallel \dots \parallel \oplus_K(c!(T_k); \varphi_a^k) \parallel \dots \parallel \varphi_m \rightarrow [\Sigma] \varphi_1 \parallel \dots \parallel \oplus_L(c!(T_l); \varphi_a^l) \parallel \dots \parallel \varphi_m$$

for some  $L$  such that  $\emptyset \subset L \subseteq J$ . We straightforwardly have that  $\varphi_1 \parallel \dots \parallel \oplus_L(c!(T_l); \varphi_a^l) \parallel \dots \parallel \varphi_m \in \Delta_n$ . If  $\oplus_K(c!(T_k); \varphi_a^k)$  is a vacuous choice then it is simply of the form  $c!(T_k); \varphi_a^k$ . By the reduction rules then we know that  $\Sigma' = \Sigma_k; c \mapsto T_k$  and  $\varphi'_a = \varphi_a^k$ . By the hypotheses we know that  $\delta(\Sigma^k, \varphi_a^k, \varphi_b^k) \in \Delta_{n-1}$ . Then, by the first part of the definition of  $\Delta_n^{\uparrow G}$  in Figure 3.13 we know that  $[\Sigma'] \Phi' = \delta([c \mapsto T; \Sigma^k], \varphi_a^k, \&_J(c?(T_j); \varphi_b^j)) \in \Delta_n$  as required. Consider the case when  $i \neq b$  and  $i \neq a$ . This follows similarly to the case where  $[\Sigma] \Phi \in \{ \delta([c \mapsto T; \Sigma^k], \varphi_a^k, \&_J(c?(T_j); \varphi_b^j)) \}$  and  $i \neq b$ , above.

Consider the case when  $[\Sigma] \Phi \in \Delta_{n-1}$ . By the inductive hypothesis we have that  $[\Sigma'] \Phi' \in \Delta_{n-1}$ , and hence that  $[\Sigma'] \Phi' \in \Delta_n$  as required.  $\square$

We can also prove relatively straightforwardly that if program point is in one of the  $\Delta_n$  of  $\Delta$  then it contains no errors, and that it is live or empty.

**Lemma 3.10.** *Valid States Safety Lemma*

$$\forall \Sigma, \Phi, n. [\Sigma] \Phi \in \Delta_n \Rightarrow \text{error} \notin \Phi$$

*Proof.* Suppose the hypothesis. We then prove the conclusion by induction over  $n$ .  $\square$

**Lemma 3.11.** *Valid States Liveness*

$$\forall \Sigma, \Phi, n. [\Sigma] \Phi \in \Delta_n \Rightarrow \exists \Sigma', \Phi'. [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \vee \Phi \equiv \Pi \epsilon$$

*Proof.* Suppose the hypothesis. We then prove that one of the two conclusions hold, by induction over the  $n$  of  $\Delta_n$ .

Consider the case where  $n = 0$ . We then proceed by case analysis of  $G$ .

If  $[\Sigma_\emptyset] \Phi \in \Delta_0$  and  $\Phi \equiv \Phi'$  then  $[\Sigma_\emptyset] \Phi' \in \Delta_0$  and hence it is sufficient to consider direct reduction.

Consider the case where  $G = d_a \rightarrow d_b : c\langle \widetilde{T \mapsto G} \rangle$ . The global session type can be projected in different ways for the  $d_a$  component of the effect. If  $d_a \rightarrow d_b : c\langle \widetilde{T \mapsto G} \rangle \upharpoonright d_a \equiv \oplus_N(c!(T_n); G_n \upharpoonright d_a)$  then by the reduction rules we know that we can perform the following reduction:

$$[\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel \oplus_N(c!(T_n); G_n \upharpoonright d_a) \parallel \dots \parallel \varphi_o \rightarrow [\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel (c!(T_n); G_n \upharpoonright d_a) \parallel \dots \parallel \varphi_o$$

If, however, the global session type is projected to a send followed by a continuation, rather than a choice between such effects, such as  $d_a \rightarrow d_b : c\langle \widetilde{T \mapsto G} \rangle \upharpoonright d_a \equiv \varphi_1 \parallel \dots \parallel (c!(T_n); G_n \upharpoonright d_a) \parallel \dots \parallel \varphi_o$ , then we know that the effect can perform the following reduction:

$$[\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel (c!(T_n); G_n \upharpoonright d_a) \parallel \dots \parallel \varphi_o \rightarrow [c \mapsto T_n] \varphi_1 \parallel \dots \parallel G_n \upharpoonright d_a \parallel \dots \parallel \varphi_o$$

Hence we have that  $[\Sigma_\emptyset] \Phi$  is live, as required.

Consider the case where  $G = 0$ . Here we have that  $0 \upharpoonright d_i \equiv \epsilon$  for each  $d_i$ . Hence  $\Phi \equiv \Pi \epsilon$ , as required.

Consider the case where  $n > 0$ . As  $[\Sigma] \Phi \in \Delta_n$  then by Figure 3.13 we know that  $[\Sigma] \Phi$  is of one of the following forms.

If it is of the form:

$$\{ \delta([c \mapsto T; \Sigma^k], \&_J(c?(T_j); \varphi_b^j), \varphi_a^k) \mid \forall j \in J. \delta(\Sigma^j, \varphi_b^j, \varphi_a^j) \in \Delta_n \}$$

then we can perform the reduction:

$$[c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \parallel \varphi_m \rightarrow [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^k \parallel \dots \parallel \varphi_m$$

If it is of the form  $\delta(\Sigma, \oplus_K(c!(T_k); \varphi_a^k), \&_J(c?(T_j); \varphi_b^j))$  then we can perform the reduction:

$$[\Sigma] \varphi_1 \parallel \dots \parallel \oplus_N c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_m \rightarrow [\Sigma] \varphi_1 \parallel \dots \parallel c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_m$$

If it is of the form  $\delta(\Sigma, (c!(T_k); \varphi_a^k), \&_J(c?(T_j); \varphi_b^j))$  then we can perform the reduction:

$$[\Sigma] \varphi_1 \parallel \dots \parallel c!(T_n); \varphi_a^n \parallel \dots \parallel \varphi_m \rightarrow [\Sigma; c \mapsto T_n] \varphi_1 \parallel \dots \parallel \varphi_a^n \parallel \dots \parallel \varphi_m$$

Hence in each of the cases it is possible to reduce, as required.  $\square$

Using these lemmas we can prove that if a program point is in  $\Delta$  then it is globally compatible:

**Theorem 3.12.** *Sufficiency of Global Session Types*

$$\forall \Sigma, \Phi. \quad [\Sigma] \Phi \in \Delta \Rightarrow \exists k. \mathcal{G}^k(\Sigma, \Phi)$$

*Proof.* Suppose the hypothesis. We then prove the conclusion by induction over  $k$  in the definition of  $\mathcal{G}^k(\Sigma, \Phi)$ .

Consider the case where  $k = 0$ . By the definition of  $\Delta$  we know that  $\mathcal{G}^k(\Sigma', \Phi') \stackrel{\text{def}}{=} \text{true}$ , and hence the implication  $[\Sigma] \Phi \in \Delta \Rightarrow \mathcal{G}^0(\Sigma, \Phi)$  is trivially valid,

Consider the case where  $k > 0$ . By Figure 3.11 we know that

$$\mathcal{G}^k(\Sigma, \Phi) \stackrel{\text{def}}{=} \left[ \begin{array}{l} [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow \mathcal{G}^{k-1}(\Sigma', \Phi') \\ \wedge [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \vee \Phi = \Pi \epsilon \\ \wedge \# \text{error} \in \Phi \end{array} \right]$$

In order to show that  $\mathcal{G}^k(\Sigma, \Phi)$  holds we show that each of the constituents hold.

By Lemma 3.9 we know that, as  $[\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$ , then we have that  $[\Sigma'] \Phi' \in \Delta$ . By the inductive step we know that  $[\Sigma] \Phi \in \Delta \Rightarrow \mathcal{G}^l(\Sigma, \Phi)$ , for all  $\Sigma$  and  $\Phi$  and  $l < k$ , and hence that  $\mathcal{G}^{k-1}(\Sigma', \Phi')$ . We therefore have that the implication  $[\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow \mathcal{G}^{k-1}(\Sigma', \Phi')$  holds.

By Lemma 3.11 we know that:

$$[\Sigma] \Phi \in \Delta \Rightarrow [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \vee \Phi \equiv \Pi \epsilon$$

By Lemma 3.10 we know that  $\text{error} \notin \Phi$ . By Lemma 3.9 we know that:

$$[\Sigma] \Phi \in \Delta \wedge [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow [\Sigma'] \Phi' \in \Delta$$

which proves all the constituent conjuncts of  $\mathcal{G}^k(\Sigma, \Phi)$ , as required.  $\square$

We reiterate the key safety and liveness theorems from (Bettini et al., 2008) as instantiations of Theorems 3.7 and 3.8, defining  $\mathcal{C} \stackrel{\text{def}}{=} [\Sigma] \Phi$  using the fact that  $[\Sigma] \Phi \in \Delta \Rightarrow \mathcal{G}(\Sigma, \Phi)$  (by Theorem 3.12).

*Observation 3.13.* Safety of Static Message Passing Programs

If  $\vdash [\sigma] P$  using  $\mathcal{C} \stackrel{\text{def}}{=} [\Sigma] \Phi \in \Delta$  and  $[\sigma] P \rightarrow [\sigma'] P'$  then  $\#\text{error} \in P'$  and  $\vdash [\sigma'] P'$

*Observation 3.14.* Liveness of Static Message Passing Programs

If  $\vdash [\sigma] P$  using  $\mathcal{C} \stackrel{\text{def}}{=} [\Sigma] \Phi \in \Delta$  then  $[\sigma] P \rightarrow [\sigma'] P' \vee P = \Pi v$

### 3.5.2 Non-Blocking Message Passing

In order to increase asynchrony we may wish to permit receive actions to continue when they attempt to receive from an empty queue, rather than blocking and waiting for a value to become available. Whilst this asynchrony leads to less idling, the synchronisation behaviour required in order to avoid message passing errors is more complex. Consider the following program abstraction:

$$[\Sigma_\emptyset] c!(l_1); c!(T_1) \oplus c!(l_2); c!(T_2) \parallel c?(l_1); c?(T_1) \& c?(l_2); c?(T_2) \& c?(l_\emptyset); c?(T_3) \quad (3.41)$$

In this program the receiver presents two services to the sender. It also contains a third service, mapped to by  $l_\emptyset$ ; this is the value returned when a receive action is performed on an empty queue, to denote that case. We make use of a semantics that, upon performing a receive on an empty channel, puts a place-holder  $\emptyset$  into the queue that denotes that the next value sent on the queue should be immediately discarded. We suggest this as a possible semantics for non-blocking receives, which has not been considered in the session typing world, and as such is novel. Consider a reduction sequence, where the receiver acts before the sender:

$$\begin{aligned} & [\Sigma_\emptyset] c!(l_1); c!(T_1) \oplus c!(l_2); c!(T_2) \parallel c?(l_1); c?(T_1) \& c?(l_2); c?(T_2) \& c?(l_\emptyset); c?(T_3) \\ & \rightarrow [c \mapsto \emptyset] c!(l_1); c!(T_1) \oplus c!(l_2); c!(T_2) \parallel c?(T_3) \\ & \rightarrow [c \mapsto \emptyset] c!(l_1); c!(T_1) \parallel c?(T_3) \\ & \rightarrow [\Sigma_\emptyset] c!(T_1) \parallel c?(T_3) \\ & \rightarrow [c \mapsto T_1] \epsilon \parallel c?(T_3) \\ & \rightarrow [\Sigma_\emptyset] \epsilon \parallel \text{error} \end{aligned} \quad (3.42)$$

Here, as the receiver chooses which service is to be performed, and the sender is not informed of this choice, it goes on to perform a different service, and a communication error can occur. In order to avoid such errors we must provide feedback to the sender, so that its control flow choices will mirror those of the receiver. In the non-blocking semantics, whenever a send action is performed we automatically return a unit value; this denotes that a send action always succeeds at adding a value on to the end of the relevant queue. We modify this semantics (Figure 3.15). The modified semantics makes use of queues that consist entirely of place-holders or entirely of typed messages (Figure 3.14). When a send is performed on a queue without any place-holders, a unit

$$q_{[]} ::= \begin{array}{l} [] \\ | \quad \emptyset \\ | \quad q_{[]}, q_{[]} \end{array}$$

Figure 3.14: Queues With Type Holes

$$\begin{array}{c} \Sigma[c \mapsto q_{[]} \xrightarrow{c?(T)} \Sigma[c \mapsto [], q_{[]}]] \\ \Sigma[c \mapsto [], q_{[]} \xrightarrow{c!(T)} \Sigma[c \mapsto q_{[]}]] \\ \Sigma[c \mapsto q_{[]}](c?(T)) = l_{[]} \quad \Sigma[c \mapsto T', q](c?(T)) = T \\ \Sigma[c \mapsto [], q_{[]}](c!(T)) = l_{[]} \quad \Sigma[c \mapsto q](c?(T)) = \text{UNIT} \end{array}$$

Figure 3.15: Semantics Of Non Blocking Message Queues

$$\begin{array}{lcl} d_1 \rightarrow d_2: r\langle \widetilde{T \mapsto G} \rangle_M \upharpoonright d_1 & \stackrel{\text{def}}{=} & \oplus_N((c!(T_n), \text{UNIT}); G_n \upharpoonright d_1 \& (c!(T_n), l_{[]}); G_{l_{[]}} \upharpoonright d_1) \\ & & N \neq \emptyset \wedge (N \cup l_{[]} \mapsto G_{l_{[]}}) \subseteq M \\ d_1 \rightarrow d_2: r\langle \widetilde{T \mapsto G} \rangle \upharpoonright d_2 & \stackrel{\text{def}}{=} & \&_M(c?(T_m); G_m \upharpoonright d_1) \\ d_1 \rightarrow d_2: r\langle \widetilde{T \mapsto G} \rangle \upharpoonright d & \stackrel{\text{def}}{=} & G_1 \upharpoonright d \\ (\mu \underline{X}. G') \upharpoonright d & \stackrel{\text{def}}{=} & \mu \underline{X}. (G' \upharpoonright d) \\ \underline{X} \upharpoonright d & \stackrel{\text{def}}{=} & \underline{X} \\ 0 \upharpoonright d & \stackrel{\text{def}}{=} & 0 \end{array}$$

Figure 3.16: Global Session Types Projection For Blocking Message Passing

value is returned to indicate success, and that the receiver has not already passed this point. When a send is performed on a queue containing a place-holder  $[]$  the action returns  $l_{[]}$  to indicate that this has occurred. Then, for each possible service choice, we create an *external* choice between an action that transmits the service choice where a unit value returned, followed by the sender's side of that service, and an action that transmits the service choice but returns a place-holder  $l_{[]}$ , followed by the service for  $l_{[]}$ . We demonstrate this with an example:

$$\begin{aligned}
& [\Sigma_\emptyset] ((\mathbf{snd}(c, l_1), \text{UNIT}); c!(T_1)) \& ((\mathbf{snd}(c, l_1), l_\square); c!(T_3)) \parallel c?(l_1); c?(T_1) \& c?(l_\square); c?(T_3) \\
\rightarrow & [c \mapsto []] ((\mathbf{snd}(c, l_1), \text{UNIT}); c!(T_1)) \& ((\mathbf{snd}(c, l_1), l_\square); c!(T_3)) \parallel c?(T_3) \\
\rightarrow & [\Sigma_\emptyset] c!(T_3) \parallel c?(T_3) \\
\rightarrow & [c \mapsto T_1] \epsilon \parallel c?(T_3) \\
\rightarrow & [\Sigma_\emptyset] \epsilon \parallel \epsilon
\end{aligned} \tag{3.43}$$

Note that, in order to avoid similar errors for the second send we would have to also replace the send  $c!(T_3)$  with  $(\mathbf{snd}(c, l_1), \text{UNIT}) \& (\mathbf{snd}(c, l_1), l_\square)$ , the receive  $c?(T_1)$  with  $c?(T_1) \& c?(l_\square)$ , and the receive  $c?(T_3)$  with  $c?(T_3) \& c?(l_\square)$ ; we omit these in order to simplify this example.

We define the updated projection function of Global Session Types for non-blocking message passing in Figure 3.16. We add an additional requirement to the well formedness of Global Session Types, namely that each send/receive pair must include the case of when  $l_\square$  is received, as it can occur in any send/receive pair. Note that, in the projection  $d_1 \rightarrow d_2 : r\langle \widetilde{T \mapsto G} \rangle \upharpoonright d_2 \stackrel{\text{def}}{=} \&_M(c?(T_m); G_m \upharpoonright d_1)$ , the  $\widetilde{T \mapsto G}$  includes  $\mathbf{Lab} l_\square \mapsto G_{l_\square}$ .

We construct the set of safe program points (Figure 3.17) in a similar way to how we constructed the set for blocking message passing. The base case consists of effects that are projections of Global Session Types, alongside empty channel queues. The first half of the inductive case is similar to that in Figure 3.13, adding a message to the front of the queue and performing an external choice at the receiver. Note that this construction is only permitted if there are no place-holders in the queue that is being extended. The second half of the inductive case is formalises the approach described above of providing external choice at the receiver. A place-holder can be placed at the front of some message queues assuming no actual messages are in the queue being extended (but place-holders may be in the queue). The extended effect provides an external choice between a send action that receives a place-holder, indicating that the receiver has already skipped over this send, and a send action that receives the unit value, indicating that the message has been successfully added to the relevant queue.

We can prove to show that if a program point is in the valid points set, and if the point reduces, then it stays in the valid points set.

**Lemma 3.15.** *Valid Points Set Reduction*

$$\forall \Sigma, \Phi, n, \Sigma', \Phi'. [\Sigma] \Phi \in \Delta_n \wedge [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow [\Sigma'] \Phi' \in \Delta_{n+1}$$

*Proof.* The proof proceeds similarly to that in Lemma 3.9. Suppose the hypotheses. We then prove the conclusion by induction over the  $n$  of  $\Delta_n$ . The key differences are in the case where  $n > 0$ . By the structure of parallel effects we can identify one effect which is

$$\begin{aligned}
\Delta &\stackrel{\text{def}}{=} \bigcup_{n < \omega} \Delta_n & \Delta_0 &\stackrel{\text{def}}{=} \{ [\Sigma_\emptyset] \Phi^{\uparrow G} \} \\
\Delta_{n+1} &\stackrel{\text{def}}{=} \{ \delta([c \mapsto T; \Sigma^k], \&_J(c?(T_j); \varphi_b^j), \varphi_a^k) \mid \forall j \in J. \delta(\Sigma^j, \varphi_b^j, \varphi_a^j) \in \Delta_n \} \\
&\cup \{ \delta([\Sigma], \oplus_K(c!(T_k); \varphi_a^k), \&_J(c?(T_j); \varphi_b^j) \mid \\
&\quad \forall k \in K. \delta(\Sigma, \varphi_a^k, \varphi_b^k) \in \Delta_n \wedge \text{Lab } l_\square \in \tilde{T}_K \wedge \Sigma(c) = \emptyset) \\
&\cup \{ \delta([c \mapsto []; \Sigma^k], (\oplus_J(c!(T_j), \text{UNIT}); \varphi_a^j) \& ((c?(T_k), l_\square); \varphi_a^k)) \mid \\
&\quad \forall j \in J. \delta(\Sigma^j, \varphi_b^j, \varphi_a^j) \in \Delta_n \} \\
&\cup \Delta_n
\end{aligned}$$

where  $c = d_a d_b$  and  $\emptyset \subset K \subseteq J$ .

Figure 3.17: Non Blocking Message Passing Safe Program Points Set

the one that reduces:

$$[\Sigma] \varphi_1 \parallel \dots \parallel \varphi_i \parallel \dots \parallel \varphi_m \rightarrow [\Sigma'] \varphi_1 \parallel \dots \parallel \varphi'_i \parallel \dots \parallel \varphi_m$$

As  $[\Sigma] \Phi \in \Delta_n$  we know that either in one of the constructions in the definition of  $\Delta_n$ , or it is in  $\Delta_{n-1}$ . The case when  $[\Sigma] \Phi \in \{ \delta([c \mapsto T; \Sigma^k], \&_J(c?(T_j); \varphi_b^j)) \}$  is the same as in Lemma 3.9.

Consider the case when  $[\Sigma] \Phi \in \{ \delta(\Sigma, \oplus_K(c!(T_k); \varphi_a^k), \&_J(c?(T_j); \varphi_b^j)) \}$  where  $\forall j \in J. \delta(\Sigma^j, \varphi_a^j, \varphi_b^j) \in \Delta_{n-1}$  and  $l_\square \in \tilde{T}_J$  and  $\Sigma(c) = \emptyset$  (as in Figure 3.17). We can perform a case split over the index  $i$ . If  $i = b$  then, by the reduction rules we know that  $\Sigma' = \Sigma; c \mapsto l_\square$  and  $\varphi'_b = \varphi_b^{l_\square}$ . By the third part of the definition of  $\Delta_{n+1}$  we can show that  $\delta([c \mapsto []; \Sigma^k], \oplus_J(c!(T_j), \text{UNIT}); \varphi_a^j) \& ((c?(T_{l_\square}), l_\square); \varphi_a^{l_\square}, \varphi_b^{l_\square}) \in \Delta_n$ , as by the hypothesis we know that  $\forall j \in J. \delta(\Sigma^j, \varphi_a^j, \varphi_b^j) \in \Delta_n$ . Consider the case when  $i = a$ . If  $\oplus_K(c!(T_k); \varphi_a^k)$  is a choice then we can use the internal choice reduction rules to show that:

$$[\Sigma] \varphi_1 \parallel \dots \parallel \oplus_K(c!(T_k); \varphi_a^k) \parallel \dots \parallel \varphi_m \rightarrow [\Sigma] \varphi_1 \parallel \dots \parallel \oplus_L(c!(T_l); \varphi_a^l) \parallel \dots \parallel \varphi_m$$

for some  $L$  such that  $\emptyset \subset L \subseteq J$ . We straightforwardly have that  $\varphi_1 \parallel \dots \parallel \oplus_L(c!(T_l); \varphi_a^l) \parallel \dots \parallel \varphi_m \in \Delta_n$ . If  $\oplus_K(c!(T_k); \varphi_a^k)$  is a vacuous choice then it is simply of the form  $c!(T_k); \varphi_a^k$ . By the reduction rules then we know that  $\Sigma' = \Sigma_k; c \mapsto T_k$  and  $\varphi'_a = \varphi_a^k$ . By the hypotheses we know that  $\delta(\Sigma^k, \varphi_a^k, \varphi_b^k) \in \Delta_{n-1}$ . Then, by the first part of the definition of  $\Delta_n^{\uparrow G}$  in Figure 3.17 we know that  $[\Sigma'] \Phi' = \delta([c \mapsto T; \Sigma^k], \varphi_a^k, \&_J(c?(T_j); \varphi_b^j)) \in \Delta_n^{\uparrow G}$  as required. Consider the case when  $i \neq b$  and  $i \neq a$ . This follows similarly to the case where  $[\Sigma] \Phi \in \{ \delta([c \mapsto T; \Sigma^k], \varphi_a^k, \&_J(c?(T_j); \varphi_b^j)) \}$  and  $i \neq b$ , above.

Consider the case where, given  $\forall j \in J. \delta(\Sigma^j, \varphi_a^j, \varphi_b^j) \in \Delta_{n-1}$ , that  $[\Sigma] \Phi \in \{ \delta([c \mapsto []; \Sigma^k], \oplus_J(c!(T_j), \text{UNIT}); \varphi_a^j) \& ((c?(T_k), l_\square); \varphi_a^k) \}$ . We can perform a case split over the index  $i$ . Consider the case where  $i = a$ . By the reduction rules we know that  $\Sigma' = \Sigma_k$

and  $\varphi'_a = \varphi_a^k$ . By the hypotheses we know that  $\delta(\Sigma^k, \varphi_a^k, \varphi_b^k) \in \Delta_{n-1}$ , and hence that  $\delta(\Sigma^k, \varphi_a^k, \varphi_b^k) \in \Delta_n$  as required. Consider the case when  $i \neq a$ . This case is similar to the case where  $[\Sigma] \Phi \in \{ \delta([c \mapsto T; \Sigma^k], \&_J(c?(T_j); \varphi_b^j)) \}$  and  $i \neq b$ , above.

Consider the case when  $[\Sigma] \Phi \in \Delta_{n-1}$ . By the inductive hypothesis we have that  $[\Sigma'] \Phi' \in \Delta_{n-1}$ , and hence that  $[\Sigma'] \Phi' \in \Delta_n$  as required.  $\square$

We can also prove relatively straightforwardly that if program point is in  $\Delta$  then it contains no errors, and that it is live.

**Lemma 3.16.** *Non Blocking Valid States Safety*

$$\forall \Sigma, \Phi, n. [\Sigma] \Phi \in \Delta_n \Rightarrow \text{error} \notin \Phi$$

*Proof.* Suppose the hypothesis. We then prove the conclusion by induction over  $n$ .  $\square$

**Lemma 3.17.** *Non Blocking Valid States Liveness*

$$\forall \Sigma, \Phi, n. [\Sigma] \Phi \in \Delta_n \Rightarrow \exists \Sigma', \Phi'. [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \vee \Phi \equiv \Pi \epsilon$$

*Proof.* Suppose the hypothesis. We then prove that one or other of the conclusions are true, by induction over the  $n$  of  $\Delta_n$ . This proof essentially follows that in Lemma 3.11. The key difference is when  $n > 0$  and

$$[\Sigma] \Phi \in \{ \delta([c \mapsto []; \Sigma^k], \oplus_J((c!(T_j), \text{UNIT}); \varphi_a^j) \& ((c?(T_k), l_[]); \varphi_a^k)) \}$$

where  $\forall j \in J. \delta(\Sigma^j, \varphi_a^j) \in \Delta_n$ . By the reduction rules we know that:

$$[c \mapsto []; \Sigma^k] c?(T_k) \rightarrow [\Sigma^k] \epsilon$$

and hence that

$$\delta([c \mapsto []; \Sigma^k], \oplus_J((c!(T_j), \text{UNIT}); \varphi_a^j) \& ((c?(T_k), l_[]); \varphi_a^k)) \rightarrow \delta(\Sigma^k, \varphi_a^k)$$

as required.  $\square$

Using these lemmas we can prove that if a program point is in  $\Delta$  then it is globally compatible:

**Theorem 3.18.** *Sufficiency of Global Session Types*

$$\forall \Sigma, \Phi. [\Sigma] \Phi \in \Delta \Rightarrow \exists k. \mathcal{G}^k(\Sigma, \Phi)$$

*Proof.* This proof is similar to that of Theorem 3.12, except for it uses Lemmas 3.15, 3.16 and 3.17 rather than Lemmas 3.9, 3.10 and 3.11.  $\square$

We can hence guarantee an absence of communication errors in a message passing system with very little synchronisation. The synchronisation behaviour for blocking message passing relies on the fact that the participant that provides a service must wait for the participant that invokes the service to choose which service will be performed. When the participant that provides a service can choose which service will proceed, the participant that invokes the service must similarly be informed of the decision, to ensure that both participants proceed down the same service path. The proofs of safety and liveness for blocking message passing make use of an intuitive, inductive definition that determines which states are acceptable (those that can be reached by reduction from the projection of a Global Session Type). When we use this definition, the portions of the proofs that use the synchronisation behaviour are immediately apparent. Hence we can reuse the structure and large portions of the proof for blocking message passing.

### 3.6 Conclusions

In this chapter we consider communication safety and liveness for functional programs which access accompanying shared state with an arbitrary semantics. Our contributions are as follows. We extend previous work where only receive actions could return values of different types to a system where all accesses to the shared state can return typed values. We show how to prove safety and liveness in a system where the shared state has arbitrary semantics. We reprove existing results of safety and liveness for blocking message passing programs, using a novel, inductive definition, which we believe is significantly more comprehensible. We extend this technique to cover non-blocking message passing systems, and show how our inductive definition can be straightforwardly modified to cover non-blocking message passing.



## Chapter 4

# Dynamic Software Update

In this chapter we consider how to update message passing programs, whilst maintaining safety and liveness properties. We explore communications safety and deadlock freedom for updatable asynchronous message passing programs with blocking receives, and provide examples of erroneous updates to programs, and reductions of such systems (Section 4.1). We provide a calculus for updates to the language presented in Chapter 3, an operational semantics for update, and extend our behavioural abstraction (Section 4.2). We provide a general, inductive description of the communication safety and deadlock freedom properties in the presence of arbitrary updates (Section 4.3). We present our conclusions in Section 4.4

$$\begin{aligned} t_c &= \mu \underline{X}. \text{sel}(c_1, \text{prime?}). \text{snd}(c_1, \dots). \text{rcv}(c_2)(x : \text{INT}). \underline{X} \\ t_s &= \mu \underline{X}' . \text{case } c_1 \text{ of } \{ \begin{array}{l} \text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}). \text{snd}(c_2, \dots). \underline{X}', \\ \dots \end{array} \end{aligned}$$

Figure 4.1: Maths Server

---


$$\begin{aligned} t'_c &= \mu \underline{X}. \text{sel}(c_1, \text{prime?}). \text{snd}(c_1, \dots). \text{case } c_2 \text{ of } \{ \begin{array}{l} \text{result} \mapsto \text{rcv}(c_2)(x : \text{INT}). \underline{X}, \\ \text{err} \mapsto \underline{X} \end{array} \} \\ t'_s &= \mu \underline{X}' . \text{case } c_1 \text{ of } \{ \begin{array}{l} \text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}). \\ \quad \text{if } x < 3000 \text{ then} \\ \quad \quad \text{sel}(c_2, \text{result}). \text{snd}(c_2, \dots). \underline{X}' \\ \quad \text{else} \\ \quad \quad \text{sel}(c_2, \text{err}). \underline{X}' \\ \dots \end{array} \} \end{aligned}$$

Figure 4.2: Updated Maths Server

```

1   $[\sigma_\emptyset] (t_c \parallel t_s)$ 
2   $\rightarrow \dots$ 
3   $\rightarrow [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{BOOL}).t_c \parallel t_s)$ 
4   $\xrightarrow{\psi} [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{BOOL}).t'_c \parallel t'_s)$ 
5   $\equiv [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{BOOL}).t'_c \parallel$ 
    $\quad \text{if } x < 3000 \text{ then sel}(c_2, \text{result}).\text{snd}(c_2, \dots).t'_s \text{ else sel}(c_2, \text{err}).t'_s)$ 
6   $\rightarrow [c_1 \mapsto 3457] (\text{rcv}(c_2)(x : \text{BOOL}).t'_c \parallel \text{rcv}(c_1)(x : \text{INT}).$ 
    $\quad \text{if } x < 3000 \text{ then sel}(c_2, \text{result}).\text{snd}(c_2, \dots).t'_s \text{ else sel}(c_2, \text{err}).t'_s)$ 
7   $\rightarrow \dots$ 
8   $\rightarrow [\sigma_\emptyset] (\text{rcv}(c_2)(x : \text{BOOL}).t'_c \parallel \text{sel}(c_2, \text{result}).\text{snd}(c_2, \dots).t'_s)$ 
9   $\rightarrow [c_2 \mapsto \text{result}] (\text{rcv}(c_2)(x : \text{BOOL}).t'_c \parallel \text{snd}(c_2, \dots).t'_s)$ 
10  $\rightarrow [\sigma_\emptyset] (\text{error} \parallel \text{snd}(c_2, \dots).t'_s)$ 

```

Figure 4.3: Update Error Due To Update At Mismatched Code Points

## 4.1 Motivation: Message Passing Programs

We specify a maths server program in Figure 4.1. In this program the client  $t_c$  tells the server  $t_s$  that it would like the server to tell it if a given number is a prime. The client then sends the number to the server, and the server will send an integer reply. When the program is deployed it may be discovered that the server is often encountering errors due to overflow or timeouts. We may then wish to modify the program so that the server can choose to abort the calculation, and inform the client whether the computation produced a result or an error. We specify such a program in Figure 4.2. In order to migrate from one program to another we need some mechanism to modify the running code. We delay the full specification of this mechanism until Section 4.2 and, for the purposes of the motivation, describe updates informally.

The introduction of an update at an inopportune point can introduce a communication error into an otherwise safe program. Consider the reduction shown in Figure 4.3. The original maths server reduces until line 3. Then an update reduction, which is annotated with  $\psi$ , occurs on line 4; this modifies the continuations of both the client and the server to be the bodies of the updated maths server. Note that at line 3 the client has entered the body of its recursive loop, whilst the server has not. As the update occurs at this point we end up with a partial reduction of the old client code attempting to interact with the new version of the server; hence the server sends the `result` label to indicate that it has completed the computation successfully. Unfortunately the client is only expecting a boolean result, rather than a label and a result, and hence a communication error occurs. This illustrates how messages in the queue at the point an update is applied can affect whether the updated program is safe.

We specify a Producer/Consumer program, with an acknowledgement, in Figure 4.4. In this program the producer  $t_p$  sends one data item to the consumer  $t_{co}$ , and waits for an acknowledgement from the consumer that indicates that it is ready to continue. We may then wish to modify the program to permit the producer to send two data items before waiting for the acknowledgement, in order to increase asynchrony. We specify such a program in Figure 4.5.

The introduction of an update at an inopportune point can also introduce deadlock into an otherwise live program. Consider a reduction of the Producer/Consumer example, shown in Figure 4.6. The original Producer/Consumer program runs until line 3. Then an update reduction occurs on line 4; this modifies the continuations of both the client and the server to be the bodies of the updated Producer/Consumer program. Note that at line 3 the producer has entered the body of its recursive loop, while the consumer has not. As the update occurs at this point we end up with a partial reduction of the old producer attempting to interact with the new version of the server. The producer hence sends one data value and waits for an acknowledgement whilst the consumer receives the first data value and waits to receive the second before it sends the acknowledgement. This leads to a deadlocked system, where the producer and the consumer are both waiting for a message from the other.

$$\begin{aligned} t_p &= \mu \underline{X}.\text{snd}(c_1, v_1).\text{rcv}(c_2)(x : \text{UNIT}).\underline{X} \\ t_{co} &= \mu \underline{X}'.\text{rcv}(c_1)(x : T_1).\text{snd}(c_2, ()).\underline{X}' \end{aligned}$$

Figure 4.4: Producer/Consumer System With Acknowledgement

$$\begin{aligned} t'_p &= \mu \underline{X}.\text{snd}(c_1, v_1).\text{snd}(c_1, v_2).\text{rcv}(c_2)(x : \text{UNIT}).\underline{X} \\ t'_{co} &= \mu \underline{X}'.\text{rcv}(c_1)(x : T_1).\text{rcv}(c_1)(x : T_2).\text{snd}(c_2, ()).\underline{X}' \end{aligned}$$

Figure 4.5: Updated Producer/Consumer System With Acknowledgement

1	$[\sigma_\emptyset]$	$(t_p \parallel t_{co})$
2	$\rightarrow \dots$	
3	$\rightarrow [\sigma_\emptyset]$	$(\text{snd}(c_1, v_1).\text{rcv}(c_2)(x : \text{UNIT}).t_p \parallel t_{co})$
4	$\xrightarrow{\psi} [\sigma_\emptyset]$	$(\text{snd}(c_1, v_1).\text{rcv}(c_2)(x : \text{UNIT}).t'_p \parallel t'_{co})$
5	$\rightarrow [c_1 \mapsto v_1]$	$(\text{rcv}(c_2)(x : \text{UNIT}).t'_p \parallel t'_{co})$
6	$\rightarrow [c_1 \mapsto v_1]$	$(\text{rcv}(c_2)(x : \text{UNIT}).t'_p \parallel \text{rcv}(c_1)(x : T_1).\text{rcv}(c_1)(x : T_2).\text{snd}(c_2, ()).t'_{co})$
7	$\rightarrow [\sigma_\emptyset]$	$(\text{rcv}(c_2)(x : \text{UNIT}).t'_p \parallel \text{rcv}(c_1)(x : T_2).\text{snd}(c_2, ()).t'_{co})$
8	$\not\rightarrow$	

Figure 4.6: Update Deadlock Due To Update At Mismatched Code Points

```

1   $[\sigma_\emptyset] (t_p \parallel t_{co})$ 
2   $\rightarrow \dots$ 
3   $\rightarrow [\sigma_\emptyset] (\mathbf{snd}(c_1, v_1).\mathbf{rcv}(c_2)(x : \text{UNIT}).t_p \parallel \mathbf{rcv}(c_1)(x : T_1).\mathbf{snd}(c_2, ()).t_{co})$ 
4   $\xrightarrow{\psi} [\sigma_\emptyset] (\mathbf{snd}(c_1, v_1).\mathbf{rcv}(c_2)(x : \text{UNIT}).t'_p \parallel \mathbf{rcv}(c_1)(x : T_1).\mathbf{snd}(c_2, ()).t_{co})$ 
5   $\rightarrow \dots$ 
6   $\rightarrow [\sigma_\emptyset] (t'_p \parallel t_{co})$ 
7   $\rightarrow \dots$ 
8   $\rightarrow [\sigma_\emptyset] (\mathbf{snd}(c_1, v_2).\mathbf{rcv}(c_2)(x : \text{UNIT}).t'_p \parallel \mathbf{snd}(c_2, ()).t_{co})$ 
9   $\rightarrow \dots$ 
10  $\rightarrow [c_1 \mapsto v_2, c_2 \mapsto ()] (\mathbf{rcv}(c_2)(x : \text{UNIT}).t'_p \parallel \mathbf{rcv}(c_1)(x : T_1).\mathbf{snd}(c_2, ()).t_{co})$ 
11  $\rightarrow [c_2 \mapsto ()] (\mathbf{rcv}(c_2)(x : \text{UNIT}).t'_p \parallel \mathbf{error})$ 

```

Figure 4.7: Update Error Due To Inconsistent Update

---

We also wish to rule out trivial update errors. Consider an update that modifies the code of the producer to  $t'_p$  but leaves the code of the consumer as  $t_{co}$ . We present one reduction of this system in Figure 4.7. We apply the update on line 4. The system then reduces to complete one loop of each thread’s main loop. When the producer starts a new loop it uses the new code (line 6). The sender sends both its values, and the consumer performs two loops. On its second loop the consumer receives  $v_2$ , which is not of the type it is expecting.

It is, however, possible to safely perform the updates to the maths server and the Producer/Consumer system. Consider again the maths server and its update (Figures 4.1 & 4.2). We present a possible reduction which leaves the program both safe and live after the update occurs (Figure 4.8). The original maths server starts reducing; on line 6, part way through each thread’s recursive loop, the update is applied. Since the server and the client are each part way through their old code’s loop body they both continue communicating using the old code until they reach the point the point where they would recurse (line 11), at which point each continues using the new code.

To summarise, the effect of the code, and the messages in queue, at the point an update is applied will determine the safety and liveness of the updated program. In particular, all threads must update when they are on the same ‘run’ of an protocol. Possible techniques to achieve this include ensuring that all participants are always on the same run of a protocol, or to determine the run of the participant that is the furthest ahead, and to update each participant when it reaches that point.

```

1   $[\sigma_\emptyset] (t_c \parallel t_s)$ 
2   $\rightarrow \dots$ 
3   $\rightarrow [c_1 \mapsto \text{prime?}] (\text{snd}(c, 3457).\text{rcv}(c_2)(x : \text{INT}).t_c \parallel$ 
    $\mu \underline{X}' . \text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots). \underline{X}', \dots\})$ 
4   $\rightarrow [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{INT}).t_c \parallel$ 
    $\mu \underline{X}' . \text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots). \underline{X}', \dots\})$ 
5   $\rightarrow [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{INT}).t_c \parallel$ 
    $\text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots).t_s, \dots\})$ 
6   $\xrightarrow{\psi} [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{INT}).t'_c \parallel$ 
    $\text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots).t'_s, \dots\})$ 
7   $\rightarrow [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{INT}).t'_c \parallel$ 
    $\text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots).t'_s, \dots\})$ 
8   $\rightarrow [c_1 \mapsto 3457] (\text{rcv}(c_2)(x : \text{INT}).t'_c \parallel \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots).t'_s)$ 
9   $\rightarrow [\sigma_\emptyset] (\text{rcv}(c_2)(x : \text{INT}).t'_c \parallel \text{snd}(c_2, \dots).t'_s)$ 
10  $\rightarrow [c_2 \mapsto \text{true}] (\text{rcv}(c_2)(x : \text{INT}).t'_c \parallel t'_s)$ 
11  $\rightarrow [\sigma_\emptyset] (t'_c \parallel t'_s)$ 
12  $\rightarrow \dots$ 

```

Figure 4.8: Update Success In Maths Server

## 4.2 Update Calculus

In order to reason about how and when updates occur we introduce two additions to our calculus (Figure 4.9). We augment our language from Figure 3.4 with terms used to describe where updates can occur. We define our notion of updates, which describe the changes to a running program, and the program points at which said changes can be made. We also extend our behavioural abstraction to encompass updates.

### 4.2.1 Language Additions

A region annotation  $f(t)$  denotes a region or body of code, annotated with the name  $f$ , which can be replaced in its entirety. Such regions normally consist of a procedure or a function, but could be used to consider conceptual units of work, for example a series of processing calls, or a series of logging calls. When reducing this region, in the absence of updates, the annotation is stripped away as the region proper is reduced. Once the annotation has been stripped away the body cannot be replaced (though it may be possible to replace further annotated regions within the body). Updates replace the body of code inside an annotation with another body of code. This corresponds with our intuition that the body of the annotated region can only be replaced in its entirety. Annotations are a variant of the standard Dynamic Software Update approach which

$t$	$::=$	<i>Terms</i>
	$\dots$	
	$f(t)$	updatable regions
	$\text{dmod}_{\tilde{f}}(t)$	transactional regions
$\psi, \omega$	$::=$	<i>Predicated Updates</i>
	$p \mapsto u$	Map from predicate to update
	$\psi, \psi$	List of predicated updates
	$\emptyset$	Empty predicated update
$u$	$::=$	<i>Region updates</i>
	$f \mapsto t$	Map from region to new code
	$u, u$	Set of updates
	$\emptyset$	Empty update
$p$	$::=$	<i>Update predicates on code and message queues</i>
	$p(\sigma, P)$	

Figure 4.9: Update Calculus

$\frac{[\sigma] \quad t \xrightarrow{\tau}_F [\sigma'] \quad t'}{[\sigma'] \quad f(t) \xrightarrow{\tau}_F [\sigma'] \quad f(t')}$	$\frac{[\sigma] \quad t \xrightarrow{\alpha(\tilde{v})}_F [\sigma'] \quad t'}{[\sigma] \quad f(t) \xrightarrow{\alpha(\tilde{v})}_{F \cup \{f\}} [\sigma'] \quad t'}$	$\frac{}{[\sigma] \quad f(v) \xrightarrow{\tau} [\sigma] \quad v}$
$\frac{[\sigma] \quad t \xrightarrow{\gamma}_{F'} \quad [\sigma'] \quad t'}{[\sigma] \quad \text{dmod}_F(t) \xrightarrow{\gamma}_{F'} \quad [\sigma'] \quad \text{dmod}_{F \cup F'}(t')}$	$\frac{\text{regions}(t) \not\subseteq F}{[\sigma] \quad \text{dmod}_F(t) \xrightarrow{\gamma}_{F'} \quad [\sigma'] \quad \text{dmod}_{F \cup F'}(t')}$	$\frac{\text{regions}(t) \subseteq F}{[\sigma] \quad \text{dmod}_F(t) \xrightarrow{\tau} \quad [\sigma] \quad t}$
$\frac{[\sigma] \quad P \xrightarrow{\gamma}_F \quad [\sigma'] \quad P' \quad \nexists p \in \text{dom}(\psi). p(\sigma, P)}{[\sigma] \quad P, \psi \xrightarrow{\gamma} \quad [\sigma'] \quad P', \psi}$	$\frac{}{[\sigma] \quad P, \psi \xrightarrow{\tau} \quad [\sigma] \quad P, (\psi, \omega)}$	

Figure 4.10: Update Calculus Operational Semantics

permits the bodies of functions to be changed in their entirety (Hicks, 2005; Stoyle et al., 2005). The standard approach holds the body of a function in the heap and substitutes it for the function whenever a call is made. Instead, we simply annotate the code itself.

A transactional annotation  $\text{dmod}_{\tilde{f}}(t)$  denotes a region of code where, if an update modifies more than one annotated region, that it must be able to modify all annotated regions within the transactional annotation, or none of them. The subscript  $\tilde{f}$  records the annotations of named regions which have already been entered, by reducing the body

$\mathsf{upd}(v, u, b)$	$\stackrel{\mathsf{def}}{=} v$	$v \neq \mathsf{rec} \underline{X}(x : T).t$
$\mathsf{upd}((\alpha(\widetilde{v}), T), u, b)$	$\stackrel{\mathsf{def}}{=} (\alpha(\widetilde{v}), T)$	
$\mathsf{upd}(\mathsf{rec} \underline{X}(x : T).t, u, b)$	$\stackrel{\mathsf{def}}{=} \mathsf{rec} \underline{X}(x : T).\mathsf{upd}(t, u, b)$	
$\mathsf{upd}(t_1 t_2, u, b)$	$\stackrel{\mathsf{def}}{=} \mathsf{upd}(t_1, u, b) \mathsf{upd}(t_2, u, b)$	
$\mathsf{upd}(\mathsf{if } t_1 \mathsf{ then } t_2 \mathsf{ else } t_3, u, b)$	$\stackrel{\mathsf{def}}{=} \mathsf{if } t'_1 \mathsf{ then } t'_2 \mathsf{ else } t'_3$	$t'_i = \mathsf{upd}(t_i, u, b)$
$\mathsf{upd}(\mathsf{case } \alpha(\widetilde{v}) \mathsf{ in } \{\widetilde{T} \mapsto \widetilde{t}\}, u, b)$	$\stackrel{\mathsf{def}}{=} \mathsf{case } \alpha(\widetilde{v}) \mathsf{ in } \{\widetilde{T} \mapsto \widetilde{\mathsf{upd}(t_1, u, b)}\}$	
$\mathsf{upd}(\mathsf{dmod}_F(t_1), u, b)$	$\stackrel{\mathsf{def}}{=} \mathsf{dmod}_F(\mathsf{upd}(t_1, u, \mathsf{true}))$	$F \cap \mathsf{dom}(u) = \emptyset$
$\mathsf{upd}(\mathsf{dmod}_F(t_1), u, b)$	$\stackrel{\mathsf{def}}{=} \mathsf{dmod}_F(\mathsf{upd}(t_1, u, \mathsf{false}))$	$F \cap \mathsf{dom}(u) \neq \emptyset$
$\mathsf{upd}(f(t_1), u, b)$	$\stackrel{\mathsf{def}}{=} f(\mathsf{upd}(t_1, u, b))$	$\neg b \vee \exists t. (f \mapsto t) \in u$
$\mathsf{upd}(f(t_1), u, \mathsf{true})$	$\stackrel{\mathsf{def}}{=} f(t')$	$(f \mapsto t') \in u$

Figure 4.11: Update Function

1	$[\_] \mathsf{dmod}_\emptyset(\dots; f(t_1); \dots; g(t_2))$
2	$\rightarrow \dots$
3	$\rightarrow [\_] \mathsf{dmod}_\emptyset(f(t_1); \dots; g(t_2))$
4	$\rightarrow [\_] \mathsf{dmod}_f(t_1; \dots; g(t_2))$
5	$\rightarrow \dots$
6	$\rightarrow [\_] \mathsf{dmod}_f(g(t_2))$
7	$\rightarrow [\_] \mathsf{dmod}_{f,g}(t_2)$
9	$\rightarrow [\_] t_2$
8	$\rightarrow \dots$

Figure 4.12: updatable region Reduction Example

of the *transactional region*. We explain using the reduction in Figure 4.12. When an *updatable region* (named region) is reduced inside a transactional region its region name is added to the set of updatable regions which cannot be updated within the transactional region (lines 4 and 7). An update to  $f$  can occur at any point up to and including line 3, an update to  $g$  can occur at any point up to and including line 6, and an update to both  $f$  and  $g$  can occur at any point up to and including line 3. If a transactional region is nested within another transactional region it is handled separately to the body of the encapsulating region. In order to inform encapsulating transactional regions that an updatable region has been entered we annotate our reduction rules; compound operators pass through reductions of sub-terms, and all destructive operator reductions are annotated with  $\tau$  except the reduction for updatable regions.

In order to be able to update the bodies of the recursive loops of the Producer/Consumer example, as we did on line 4 of Figure 4.6, in our full calculus we write the original

1	$[\sigma_\emptyset] (t_p \parallel t_{co}), \psi_\emptyset$
2	$\rightarrow \dots$
3	$\rightarrow [c_1 \mapsto v_1] (\mathbf{rcv}(c_2)(x : \mathbf{UNIT}).t_p \parallel \mathbf{rcv}(c_1)(x : T_1).\mathbf{snd}(c_2, ()).t_{co}), \psi_\emptyset$
4	$\rightarrow [c_1 \mapsto v_1] (\mathbf{rcv}(c_2)(x : \mathbf{UNIT}).t_p \parallel \mathbf{rcv}(c_1)(x : T_1).\mathbf{snd}(c_2, ()).t_{co}), \psi$
5	$\rightarrow \dots$
6	$\rightarrow [\sigma_\emptyset] (\mu \underline{X}.\mathbf{dmod}_\emptyset(f(\mathbf{snd}(c_1, \dots).\mathbf{rcv}(c_2)(x : \mathbf{UNIT}).\underline{X})) \parallel \mu \underline{X}'.\mathbf{dmod}_\emptyset(g(\mathbf{rcv}(c_1)(x : T_1).\mathbf{snd}(c_2, ()).\underline{X}'))), \psi$
7	$\rightarrow [\sigma_\emptyset] (\mu \underline{X}.\mathbf{dmod}_\emptyset(f(\mu \underline{X}''.\mathbf{snd}(c_1, \dots).\mathbf{snd}(c_1, \dots).\mathbf{rcv}(c_2)(x : \mathbf{UNIT}).\underline{X}'')) \parallel \mu \underline{X}'.\mathbf{dmod}_\emptyset(g(\mu \underline{X}'''.\mathbf{rcv}(c_1)(x : T_1).\mathbf{rcv}(c_1)(x : T_2).\mathbf{snd}(c_2, ()).\underline{X}''')), \psi_\emptyset$
8	$\rightarrow \dots$
	$\psi \stackrel{\text{def}}{=} p \mapsto u \quad p(\sigma, P) \stackrel{\text{def}}{=} (P = t_p \parallel t_{co}) \quad u \stackrel{\text{def}}{=} f \mapsto t'_p, f \mapsto t'_{co}$
	$t_p = \mu \underline{X}.\mathbf{dmod}_\emptyset(f(\mathbf{snd}(c_1, \dots).\mathbf{rcv}(c_2)(x : \mathbf{UNIT}).\underline{X}))$
	$t_{co} = \mu \underline{X}'.\mathbf{dmod}_\emptyset(g(\mathbf{rcv}(c_1)(x : T_1).\mathbf{snd}(c_2, ()).\underline{X}'))$
	$t'_p = \mu \underline{X}''.\mathbf{snd}(c_1, \dots).\mathbf{snd}(c_1, \dots).\mathbf{rcv}(c_2)(x : \mathbf{UNIT}).\underline{X}''$
	$t'_{co} = \mu \underline{X}'''.\mathbf{rcv}(c_1)(x : T_1).\mathbf{rcv}(c_1)(x : T_2).\mathbf{snd}(c_2, ()).\underline{X}'''$

Figure 4.13: Example Update Reduction

program as:

$$\begin{aligned} t_p &= \mu \underline{X}.\mathbf{dmod}_\emptyset(f(\mathbf{snd}(c_1, \dots).\mathbf{rcv}(c_2)(x : \mathbf{UNIT}).\underline{X})) \\ t_{co} &= \mu \underline{X}'.\mathbf{dmod}_\emptyset(g(\mathbf{rcv}(c_1)(x : T_1).\mathbf{snd}(c_2, ()).\underline{X}')) \end{aligned} \quad (4.1)$$

#### 4.2.2 Updates

In Figure 4.9 we present our definition of updates. *Region updates*  $u$  are maps from region names  $f$  to new region bodies  $t$ . *Predicated updates*  $\psi$  are a list of mappings from *update predicates* to region updates. An update predicate  $p(\sigma, P)$  is a first order logic predicate, over the structure of the code, the shared state, and the structure of the predicated updates themselves. Update predicates denote when it is appropriate to apply a given region update to the existing program. Example predicates include:

$$p(\sigma, P) \stackrel{\text{def}}{=} (\sigma = \sigma_\emptyset) \quad p(\sigma, P) \stackrel{\text{def}}{=} (P = t_p \parallel t_{co}) \quad (4.2)$$

A region update is applied to the code as soon as its associated update predicate becomes true:

$$\frac{\psi = p_1 \mapsto u_1, \dots, p_n \mapsto u_n \quad i \text{ smallest in } 1 \dots n. p_i(\sigma, P)}{[\sigma] P, \psi \xrightarrow{\tau} [\sigma] \mathbf{upd}(P, u_i, \mathbf{false}), \psi \setminus p_i \mapsto u_i} \quad (4.3)$$

The order of the maps from predicates to region updates is important, as updates are applied from left to right (oldest to newest). We define an update function  $\text{upd}(t, u, b)$  (Figure 4.11) where  $b$  denotes whether  $t$  occurs in a context which is appropriate to update. The definition of update for transactional regions describes in which situations it is appropriate to update:

$$\begin{aligned}\text{upd}(\text{dmod}_{\tilde{f}}(t_1), u, b) &\stackrel{\text{def}}{=} \text{dmod}_{\tilde{f}}(\text{upd}(t_1, u, \text{true})) & \tilde{f} \cap \text{dom}(u) = \emptyset \\ \text{upd}(\text{dmod}_{\tilde{f}}(t_1), u, b) &\stackrel{\text{def}}{=} \text{dmod}_{\tilde{f}}(\text{upd}(t_1, u, \text{false})) & \tilde{f} \cap \text{dom}(u) \neq \emptyset\end{aligned}\quad (4.4)$$

When no updatable regions (that are being updated by  $u$ ) have been entered, within a given transactional region, then it is appropriate to update all possible named regions within the transactional region, otherwise it is not. When an update is applied to a named region, if we are in an appropriate context, and the update contains a replacement body for the named region, then we replace the body of the region in its entirety:

$$\text{upd}(f(t_1), u, \text{true}) \stackrel{\text{def}}{=} f(u(f)) \quad f \in \text{dom}(u) \quad (4.5)$$

Otherwise we continue recursively:

$$\text{upd}(f(t_1), u, b) \stackrel{\text{def}}{=} f(\text{upd}(t_1, u, b)) \quad (\neg b \vee f \notin \text{dom}(u)) \quad (4.6)$$

Updates can be introduced at any time, and are always appended onto the end of the update list:

$$\boxed{[\sigma] \ P, \psi} \xrightarrow{\tau} \boxed{[\sigma] \ P, (\psi, \omega)} \quad (4.7)$$

Here  $\omega$  is introduced externally. An external, or out-of-band, update (Soules et al., 2003; Ajmani, 2004; Stoyle et al., 2005; Nicoara et al., 2008) is one where the update is introduced using some mechanism that isn't considered as part of the system. An example of this would be the change of a physical piece of hardware, or making use of some system level code that isn't available to normal message passing programs. We use the notation  $\psi_\emptyset$  to indicate an empty list of Predicated Updates. In order to make our analysis tractable we assume that a program starts with no pending updates, and that only one update can be introduced into the system. That update, however, can be introduced at any point in the execution of the original program. This approach could be extended to multiple updates by considering an updated program as the initial program for the second update, and applying the analyses presented in this thesis for single updates, on the modified code with a new update.

Normal reductions (Figure 3.5) can only occur when no updates are applicable:

$$\frac{[\sigma] \ P \xrightarrow{\gamma} [\sigma'] \ P' \quad \nexists p \in \text{dom}(\psi). p(\sigma, P)}{[\sigma] \ P, \psi \xrightarrow{\gamma} [\sigma'] \ P', \psi} \quad (4.8)$$

$\psi, \omega ::=$	<i>Predicated Updates</i>
$p \mapsto u$	Map from predicate to update
$\psi, \psi$	List of predicate updates
$\emptyset$	Empty predicate update
$u ::=$	<i>Region updates</i>
$f \mapsto t$	Map from region to new code
$u, u$	Set of updates
$\emptyset$	Empty update
$p ::=$	<i>Update predicate</i>
$p(z_\Sigma, z_\Phi, z_\psi)$	on code and message queues and current updates

Figure 4.14: Update Calculus

---

We define all the additions to the operational semantics that handle updates in Figure 4.10. We make use of the auxiliary function `regions`( $t$ ) that is inductively defined over the structure of terms, and returns a set of the region annotations used in  $t$ .

We provide a full reduction example for update in Figure 4.13. The region update maps regions  $f$  and  $g$  to the new region bodies  $t'_p$  and  $t'_{co}$  respectively. The update predicate states that the only valid point at which the update can be applied is when both the producer and the consumer are at the top of their recursive loops. The update is introduced on line 4, but the update predicate doesn't become true until line 6, at which point the update is applied on line 7.

We assume that the new code included in a region update is all typable as `UNIT` under an empty type environment:

$$\forall f \mapsto t \in u. \varphi \wr \emptyset \vdash t: \text{UNIT} \quad (4.9)$$

There is a wide variety of work that considers how to update functions in a way that changes their type signature (Baumann et al., 2007; Soules et al., 2003; Baumann et al., 2005; Subramanian, 2010; Altekar et al., 2005; Kaashoek and Arnold, 2009). The techniques therewith are largely orthogonal to ours. We work under the simplifying assumption that the types of bodies of code will not change (in particular, that they have the unit assumption). This is as, to incorporate the techniques used to prove safety of changing type signatures as well, would have significantly complicated the presentation.

We refer to such a region update as *well formed*; we only consider well formed updates.

$$\begin{array}{c}
\varphi \wr \Gamma \vdash t : T \\
\hline
f(\varphi) \wr \Gamma \vdash f(t) : T
\end{array}
\qquad
\begin{array}{c}
\varphi \wr \Gamma \vdash t : T \\
\hline
\mathsf{dmod}_{\tilde{f}}(\varphi) \wr \Gamma \vdash \mathsf{dmod}_{\tilde{f}}(t) : T
\end{array}$$

Figure 4.15: Extensions To Typing Rules

$$\begin{array}{c}
\frac{}{[\Sigma] f(\epsilon) \xrightarrow{\tau} [\Sigma] \epsilon} \quad \frac{[\Sigma] \varphi \xrightarrow{\tau_F} [\Sigma'] \varphi'}{[\Sigma] f(\varphi) \xrightarrow{\tau_F} [\Sigma'] f(\varphi')} \quad \frac{[\Sigma] \varphi \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'}{[\Sigma] f(\varphi) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'} \\
\frac{[\Sigma] \varphi_1 \xrightarrow{\gamma_{F'}} [\Sigma'] \varphi'_1 \quad \mathsf{regions}(\varphi_1) \not\subseteq F}{[\Sigma] \mathsf{dmod}_F(\varphi_1) \xrightarrow{\gamma_{F'}} [\Sigma'] \mathsf{dmod}_{F \cup F'}(\varphi'_1)} \quad \frac{\mathsf{regions}(\varphi_1) \subseteq F}{[\Sigma] \mathsf{dmod}_F(\varphi_1) \xrightarrow{\tau} [\Sigma] \varphi_1} \\
\frac{[\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \quad \nexists p \in \mathsf{dom}(\psi). p(\Sigma, \Phi)}{[\Sigma] \Phi, \psi \xrightarrow{\gamma} [\Sigma'] \Phi', \psi} \quad \frac{\psi' \neq \emptyset \text{ introduced externally}}{[\Sigma] \Phi, \psi \xrightarrow{\tau} [\Sigma] \Phi, (\psi, \psi')} \\
\frac{\psi = p_1 \mapsto u_1, \dots, p_n \mapsto u_n \quad i \text{ smallest in } 1 \dots n. p_i(\Sigma, \Phi) \quad \Phi = \varphi_1 \parallel \dots \parallel \varphi_n \quad \varphi'_j = \mathsf{upd}(\varphi_j, u_i, \mathsf{false})}{[\Sigma] \Phi, \psi \xrightarrow{\tau} [\Sigma] \varphi'_1 \parallel \dots \parallel \varphi'_n, \psi \setminus p_i \mapsto u_i}
\end{array}$$

Figure 4.16: Update Extensions To Effect Semantics

### 4.2.3 Behavioural Abstraction

We extend the behavioural abstraction we defined in Section 3.3 to accommodate updates. We include typing rules for updatable regions and transactional regions (Figure 4.15). The effect reduction rules are extended to cover update (Figure 4.16); these mirror the concrete semantics. We abuse notation for the purposes of update predicates etc., assuming that:

$$\exists \Sigma, \Phi. p(\Sigma, \Phi) \Leftrightarrow \exists P, \sigma. \vdash P : \Phi \wedge \Sigma = \Lambda(\sigma) \wedge p(\sigma, P) \quad (4.10)$$

We define the update function over effects similarly to over terms (Figure 4.17). The most notable case is:

$$\mathsf{upd}(f(\varphi), u, b) \stackrel{\text{def}}{=} f(\varphi') \quad b \wedge f \in \mathsf{dom}(u) \wedge \varphi' \wr \emptyset \vdash u(f) : \_ \quad (4.11)$$

Here, instead of substituting the new body we instead determine its effect, under an empty type environment (which is possible for well formed updates), and substitute the effect.

## 4.3 General Formulation

In order to provide safety and liveness guarantees we must place restrictions on which programs are valid, in order to rule out errors and deadlock. We define a program point to consist of an effect  $\Phi$ , a shared state abstraction  $\Sigma$ , and possibly an update. A program point is safe and deadlock free if it does not contain an error, if all its reductions result in safe and deadlock free program points, and if it can perform at least one reduction (or consists of empty effects). We formalise these in the following section. We then proceed to prove key properties of our update system.

### 4.3.1 Global Update Compatibility

In order to rule out deadlock and errors we must prove three things: 1) that whenever the underlying language reductions are enabled that no errors are present and that the system is not deadlocked, 2) that update reductions leave a program and messages in a compatible state, and 3) that introducing the update at any point in the execution is safe. We refer the property that comprises these conditions as *Global Update Compatibility* and formalise them function that is inductively defined in Figure 4.18; this is an extension of our earlier work (Anderson and Rathke, 2011).

The predicate  $\mathcal{G}(\Sigma, \Phi, \psi, \omega)$  denotes that, given a program abstracted by  $\Phi$ , shared state  $\Sigma$ , a list of pending updates  $\psi$ , and an update  $\omega$  that can be introduced into the list of pending updates at any time, that no reductions will reduce in type errors, that the system is not deadlocked, and that any updates will leave the system in a safe state. We describe the restrictions in more detail below.

The predicate can alternatively be characterised as the greatest fixed point of the following function, defined over the state space  $S$  of tuples  $(\Sigma, \Phi, \psi, \omega)$  of abstract states, parallel effects, and pairs of updates.

$$g(S) = \left\{ (\Sigma, \Phi, \psi, \omega) \mid \left[ \begin{array}{l} \forall p \in \text{dom}(\psi). \neg p(\Sigma, \Phi) \text{ and } \forall \Sigma', \Phi'. [\Sigma]\Phi \rightarrow [\Sigma']\Phi' \Rightarrow \\ (\Sigma', \Phi', \psi, \omega) \in S \text{ and } (\exists \Sigma', \Phi'. [\Sigma]\Phi \rightarrow [\Sigma']\Phi' \\ \text{or } \Phi \equiv \Pi \epsilon) \\ \text{or, for the first } p \mapsto u \text{ in } \psi \text{ such that } p(\Sigma, \Phi) \text{ we have} \\ (\Sigma, \text{upd}(\Phi, u), (\psi \setminus p \mapsto u), \omega) \in S \\ \text{and } \exists \text{error} \in \Phi \\ \text{and } (\Sigma, \Phi, (\psi, \omega), \emptyset) \in S \end{array} \right] \right\} \quad (4.12)$$

This is a monotone function on the powerset lattice of tuples of abstract states, effects, updates, and updates. It has a greatest fixed point  $\nu g$  (which is just a set of tuples). Using the Knaster-Tarski fixed point theorem, we know that the predicate  $\mathcal{G}(\Sigma, \Phi, \psi, \omega)$  is defined if and only if  $(\Sigma, \Phi, \psi, \omega) \in \nu g$ .

$$\begin{aligned}
\mathsf{upd}(\epsilon, u, b) &\stackrel{\text{def}}{=} \epsilon \\
\mathsf{upd}(\underline{X}, u, b) &\stackrel{\text{def}}{=} \underline{X} \\
\mathsf{upd}((\alpha(\widetilde{T}), T), u, b) &\stackrel{\text{def}}{=} (\alpha(\widetilde{T}), T) \\
\mathsf{upd}(\varphi_1; \varphi_2, u, b) &\stackrel{\text{def}}{=} \mathsf{upd}(\varphi_1, u, b); \mathsf{upd}(\varphi_2, u, b) \\
\mathsf{upd}(\varphi_1 \oplus \varphi_2, u, b) &\stackrel{\text{def}}{=} \mathsf{upd}(\varphi_1, u, b) \oplus \mathsf{upd}(\varphi_2, u, b) \\
\mathsf{upd}(\varphi_1 \& \varphi_2, u, b) &\stackrel{\text{def}}{=} \mathsf{upd}(\varphi_1, u, b) \& \mathsf{upd}(\varphi_2, u, b) \\
\mathsf{upd}(\mu \underline{X}. \varphi, u, b) &\stackrel{\text{def}}{=} \mu \underline{X}. \mathsf{upd}(\varphi, u, b) \\
\mathsf{upd}(\mathsf{error}, u, b) &\stackrel{\text{def}}{=} \mathsf{error} \\
\mathsf{upd}(f(\varphi), u, b) &\stackrel{\text{def}}{=} f(\mathsf{upd}(\varphi, u, b)) \quad (\neg b \vee f \mapsto \notin u) \\
\mathsf{upd}(f(\varphi), u, \mathsf{true}) &\stackrel{\text{def}}{=} f(\varphi') \quad \exists t. f \mapsto t \in u \wedge \varphi' \wr \emptyset \vdash t' : \mathsf{UNIT} \\
\mathsf{upd}(\mathsf{dmod}_\gamma(\varphi), u, b) &\stackrel{\text{def}}{=} \mathsf{dmod}_\gamma(\mathsf{upd}(\varphi, u, \mathsf{true})) \quad \gamma \cap \mathsf{dom}(u) = \emptyset \\
\mathsf{upd}(\mathsf{dmod}_\gamma(\varphi), u, b) &\stackrel{\text{def}}{=} \mathsf{dmod}_\gamma(\mathsf{upd}(\varphi, u, \mathsf{false})) \quad \gamma \cap \mathsf{dom}(u) \neq \emptyset
\end{aligned}$$

Figure 4.17: Update Function On Effects

## Language Reductions

An error occurs when a thread receives a value it does not expect. In order to rule out such cases we need to ensure that, at every point which a program can reach, if the program can perform a reduction it reduces to another valid state:

$$[\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \Rightarrow \mathcal{G}(\Sigma', \Phi', \psi, \omega) \quad (4.13)$$

Deadlock occurs when a process wants to access the shared state but the action is blocked. In order to rule out deadlock we need to ensure, at every point  $[\Sigma] \Phi$  which a program can reach, that either it is able to perform a reduction, or it consists of empty effects:

$$[\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \vee \Phi = \Pi_I \epsilon \quad (4.14)$$

We can only perform reductions on the Message Passing Calculus if there are no updates which are applicable at that point:

$$\frac{[\sigma] P \xrightarrow{\gamma} [\sigma'] P' \quad \nexists p \in \mathsf{dom}(\psi). p(\sigma, P)}{[\sigma] P, \psi \xrightarrow{\gamma} [\sigma'] P', \psi} \quad (4.15)$$

Hence we only need to ensure that (4.13) and (4.14) hold in situations where  $\nexists p \in \mathsf{dom}(\psi). p(\Sigma, P)$ :

$$\begin{aligned}
&\nexists p \in \mathsf{dom}(\psi). p(\Sigma, \Phi) \wedge (\exists \Sigma', \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \vee \Phi \equiv \Pi_I \epsilon) \wedge \\
&(\forall \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \Rightarrow \mathcal{G}^k(\Sigma', \Phi', \psi, \omega))
\end{aligned} \quad (4.16)$$

$$\mathcal{G}^0(\Sigma, \Phi, \psi, \omega) \stackrel{\text{def}}{=} \text{true}$$

$$\mathcal{G}^k(\Sigma, \Phi, \psi, \omega) \stackrel{\text{def}}{=} \left[ \begin{array}{l} \#p \in \text{dom}(\psi). p(\Sigma, \Phi) \wedge \\ (\forall \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \Rightarrow \mathcal{G}^{k-1}(\Sigma', \Phi', \psi, \omega)) \\ \wedge (\exists \Sigma', \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \vee \Phi = \Pi_I \epsilon) \\ \wedge \forall \psi = p_1 \mapsto u_1, \dots, p_n \mapsto u_n \wedge i \text{ smallest in } 1 \dots n. \\ \quad p_i(\Sigma, \Phi) \wedge \mathcal{G}^{k-1}(\Sigma, \text{upd}(\Phi, u_i, \text{false}), \psi \setminus p_i \mapsto u_i, \omega) \\ \wedge \# \text{error} \in \Phi \wedge \mathcal{G}^{k-1}(\Sigma, \Phi, (\psi, \omega), \emptyset) \end{array} \right] \text{ where } k > 0$$

$$\mathcal{G}(\Sigma, \Phi, \psi, \omega) = \bigwedge_{k \in \mathcal{N}} \mathcal{G}^k(\Sigma, \Phi, \psi, \omega)$$

Figure 4.18: Global Update Compatibility

---

#### 4.3.1.1 Applying Updates

If a pending Predicated Update is applicable (if its update predicate is true) we apply it straight away. Hence we require, whenever a pending Predicate Update is applicable, that it leaves the system in a Globally Consistent state:

$$\begin{aligned} \psi &= p_1 \mapsto u_1, \dots, p_n \mapsto u_n \wedge i \text{ smallest in } 1 \dots n. \\ p_i(\Sigma, \Phi) &\wedge \mathcal{G}(\Sigma, \text{upd}(\Phi, u_i, \text{false}), (\psi \setminus p_i \mapsto u_i), \omega) \end{aligned} \tag{4.17}$$

#### 4.3.1.2 Introducing Updates

Updates can be introduced into the system at any time:

$$[\sigma] P, \psi \xrightarrow{\tau} [\sigma] P, (\psi, \omega) \tag{4.18}$$

In order to make our work tractable we assume that a program commences execution with no pending updates, and that only one update,  $\omega$ , is introduced in the entire course of the program. We require that if we move the update we are considering into the list of pending updates, that the system remains Globally Compatible:

$$\mathcal{G}(\Sigma, \Phi, (\psi, \omega), \emptyset) \tag{4.19}$$

#### 4.3.2 Key Properties

As discussed in Section 3.4, verifying that Global Update Compatibility holds for a given system amounts to model checking, and such a check is computationally expensive.

We again permit stronger properties which imply General Compatibility for updatable programs:

$$\frac{\vdash P : \Phi \quad \exists \mathcal{C}. \mathcal{C}(\Lambda(\sigma), \Phi, \psi, \omega) \text{ and } (\mathcal{C} \implies \mathcal{G})}{\vdash_{\omega} [\sigma] P, \psi} \quad (4.20)$$

The judgement  $\vdash_{\omega} [\sigma] P, \psi$  defines that a program  $P$ , with shared state  $\sigma$ , pending updates  $\psi$ , and an update which can become pending at some arbitrary point in the future  $\omega$ , has Global Update Compatibility. We refer to program points that can be validated using the rule in Equation 4.20 to be *update valid*.

We can prove that an update valid program points will retain safety in the presence of update.

**Theorem 4.1.** *Subject Reduction of Update Valid Program Points*

For all  $\omega, \sigma, P, \psi, \gamma, \sigma', P', \psi'$  if:

$$\vdash_{\omega} [\sigma] P, \psi \quad [\sigma] P, \psi \xrightarrow{\gamma} [\sigma'] P', \psi'$$

then there exists  $\omega'$  such that

$$\vdash_{\omega'} [\sigma'] P', \psi'$$

*Proof.* Suppose the hypotheses. We then prove the conclusion by induction over the derivation tree of the reduction  $[\sigma] P, \psi \xrightarrow{\gamma} [\sigma'] P', \psi'$ .

Consider the case where the last reduction rule used is RNoUPD:

$$\frac{[\sigma] P \xrightarrow{\gamma}_F [\sigma'] P' \quad \nexists p \in \text{dom}(\psi). p(\sigma, P)}{[\sigma] P, \psi \xrightarrow{\gamma} [\sigma'] P', \psi}$$

By the reduction rule we know that  $\nexists p \in \text{dom}(\psi). p(\Sigma, \Phi)$ . We can use Lemma 3.2 to show that as  $[\sigma] P \xrightarrow{\gamma} [\sigma'] P'$  then we know that  $\vdash P' : \Phi'$  and  $[\Lambda(\sigma)] \Phi, \psi \xrightarrow{\gamma'} [\Lambda(\sigma')] \Phi', \psi'$ . Then, by the definition of  $\mathcal{G}$ , we can know that:  $\nexists \text{error} \in \Phi$  and  $[\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow \mathcal{G}(\Sigma', \Phi', \psi, \omega)$  and hence that  $\nexists \text{error} \in \Phi'$ . Finally, given that  $\mathcal{G}(\Sigma', \Phi', \psi, \omega)$ , then letting  $\mathcal{C} = \mathcal{G}$ , we can show that:  $\vdash_{\omega} [\sigma'] P', \psi$ , using the rule in Equation 4.20.

Consider the case where the last reduction rule used is RINTRUPD. By the definition of  $\mathcal{G}$  we know that  $\mathcal{G}(\Sigma, \Phi, (\psi, \omega), \emptyset)$ . Hence we can straightforwardly show that  $\vdash_{\emptyset} [\sigma] P, \psi, \omega$ , using the rule in Equation 4.20.

Consider the case where the last reduction rule used is RUPD:

$$\frac{\psi = p_1 \mapsto u_1, \dots, p_n \mapsto u_n \quad i \text{ smallest in } 1 \dots n. p_i(\sigma, P)}{[\sigma] P, \psi \xrightarrow{\tau} [\sigma] \text{upd}(P, u_i, \text{false}), \psi \setminus p_i \mapsto u_i}$$

By Lemma A.5 we can show that as  $\vdash P: \Phi$  and  $\forall f \mapsto t. \varphi \wr \emptyset \vdash t: \text{UNIT}$  we know that  $\vdash \text{upd}(P, u, \text{false}): \text{upd}(\Phi, u, \text{false})$ . By the definition of  $\mathcal{G}$  we know that given that  $\psi = p_1 \mapsto u_1, \dots, p_n \mapsto u_n$  there exists a smallest  $i$  such that  $p_i(\Sigma, \Phi)$  holds, and furthermore that  $\mathcal{G}(\Sigma, \text{upd}(\Phi, u_i, \text{false}), \psi \setminus p_i \mapsto u_i, \omega)$ . Hence we can show that  $\vdash_{\omega} [\sigma'] P', \psi$ , using the rule in Equation 4.20.  $\square$

We can prove that an update valid program points will retain liveness in the presence of update.

**Theorem 4.2.** *Liveness of Update Valid Program Points*

For all  $\omega, \sigma, P, \psi$ , if:

$$\vdash_{\omega} [\sigma] P, \psi$$

then either  $P = \Pi v$ , or there exists  $\gamma, \sigma', P', \psi'$  such that

$$[\sigma] P, \psi \xrightarrow{\gamma} [\sigma'] P', \psi'$$

*Proof.* It is the case that either there exists an update predicate that is fulfilled by the current shared state and program point, or there is not. We perform a case split on these conditions.

Consider the case where there is no applicable update. As  $\vdash_{\omega} [\sigma] P, \psi$  holds then by the definition of  $\mathcal{G}$ , we know that either there exists  $\Sigma', \Phi'$  such that  $[\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi'$ , or that  $\Phi \equiv \Pi_I \epsilon$ . We perform a case split between these two possibilities.

Consider the case where  $[\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi'$ . By Lemma 3.6 we know that  $[\sigma] P \xrightarrow{\gamma} [\sigma'] P'$ , and hence, using the (RNoUPD) rule we can show that  $[\sigma] P, \psi \xrightarrow{\gamma} [\sigma'] P', \psi'$ , as required.

Consider the case where  $\Phi = \Pi_I \epsilon$ . By Lemma 3.5 we know that either  $[\sigma] P \xrightarrow{\gamma} [\sigma'] P'$  or  $P = \Pi_I v$ , as required.

Consider the case where there is an applicable update. We know that  $\mathcal{G}$  holds. As there exists an update predicate that is true for the current program point we can choose the smallest  $i$  such that  $p_i(\Sigma, \Phi)$  holds. Hence, using the (RNoUPD) rule, we can show that  $[\sigma] P, \psi \xrightarrow{\gamma} [\sigma'] \text{upd}(P, u, \text{false}), \psi \setminus p_i \mapsto u_i$ , as required  $\square$

## 4.4 Conclusions

In this chapter we consider communication safety and liveness for side effecting programs that can be updated at runtime. Our contributions are as follows. We demonstrate several examples of how update can introduce errors and deadlock into otherwise safe and live programs. In particular we show how if a consumer is updated before a producer, in a producer consumer system with acknowledgement, deadlock can be introduced into the system. We show how to extend general compatibility to incorporate dynamic programs. This approach is of even higher complexity than that for global compatibility presented in Chapter 3. This is as each at each reachable program point obtained by interleaving of the program, we must check to see if it is safe to introduce an update, and that all interleavings of the updated program are also safe. We again rely on the implication that we will develop specific systems that have reduced verification complexity. These are explored in the next chapter.



## Chapter 5

# Specific Formulations for DSU

The cost of verifying general compatibility is un-necessarily expensive, due to the high checking complexity. We explore stronger properties that imply Global Update Compatibility, as these properties are often easier to prove, and reduce the cost of verification. In this chapter we present two techniques to statically prove safety of updates to message passing programs, where we informally argue that the cost of verification is linear in the modified program. The first technique is a form of Global Synchronisation that is straightforward to reason about and prove safe (Section 5.1). The second technique uses the semantics of blocking message passing to provide synchronisation so that we can update groups of threads separately rather than all threads at a synchronised point (Section 5.2). We contrast their trade-off between simplicity and restrictions on the programs and updates allowed, present our conclusions, and describe our novel contributions in Section 5.3.

$$G ::= d_1 \rightarrow d_2 : \widetilde{\langle T \mapsto G \rangle} \mid \mu \underline{X}.G \mid \underline{X} \mid 0 \mid f(G) \mid \mathbf{dmod}(G)$$

Figure 5.1: Global Session Types With Update Extensions

$$f(G) \upharpoonright d \stackrel{\text{def}}{=} f_d(G \upharpoonright d) \quad \mathbf{dmod}_F(G) \upharpoonright d \stackrel{\text{def}}{=} \mathbf{dmod}_{\{f_d \mid f \in F\}}(G \upharpoonright d)$$

Figure 5.2: Extended Global Session Type Projection Function

## 5.1 Global Typability

We extend Global Session Types to include updatable and transactional regions (Figure 5.1). The programs that most benefit from DSU are those that are structured as long running processing loops (Subramanian, 2010). Hence we consider pre-update message passing programs which can be abstracted using Global Session Types of the form  $\mu \underline{X}.\mathbf{dmod}(f(G))$ , where there are no additional region annotations in  $G$ ; this denotes a long running loop that can be updated by replacing the all future iterations of a processing loop with another processing loop.

We define the projection of the extended Global Session Types in Figure 5.2 (extending the function in Figure 2.10). We handle projection of updatable and transactional regions as follows. The projection of annotation  $f(G)$  onto a participant  $d$  is defined to be  $f_d(G \upharpoonright d)$ . Note the subscript  $d$  on the region name for the local Session Type of  $d$ ; this simply permits us to ensure the region annotation in the Global Session Type is projected to syntactically separate region annotations in each of the participant’s local Session Types. We do this in order to be able to update each process independently. The projection of transactional regions is straightforward.

### 5.1.1 Formalisation

One important feature of the Global Session Type approach is that each communication action in the Global Type projects to a *pair* of send/receive actions in the local Session Types. As this pair of actions is decoupled into independent processes, in terms of protocol a point where one of the actions has been executed and the other hasn’t is not a suitable point for update. Checking for emptiness of queues would be sufficient to rule out this case, however this Global action may be a part of a larger transaction, as indicated by the regions. In order to maintain transactional sanity we do not want to permit one process to be updated when another is not, as in Figure 4.3. As processes can enter the transactional regions in the local types independently, we must also verify that either all processes involved in that particular transaction have entered the region, or all processes have not. A straightforward way of doing this is to check that the processes

$$\begin{aligned}
\Delta^{\uparrow G} &\stackrel{\text{def}}{=} \bigcup_{n < \omega} \Delta_n^{\uparrow G} & \Delta_0^{\uparrow G} &\stackrel{\text{def}}{=} \{ [\Sigma_\emptyset] \Phi^{\uparrow G} \} \\
\Delta_{n+1}^{\uparrow G} &\stackrel{\text{def}}{=} \{ \delta([c \mapsto T; \Sigma^k], \varphi'_b, \varphi_a^k) \mid \forall j \in J. \delta(\Sigma^j, \varphi_b^j, \varphi_a^j) \in \Delta_n^{\uparrow G} \} \\
&\cup \{ \delta([\Sigma], \varphi'_a, \varphi_b^j) \mid \forall k \in K. \delta(\Sigma, \varphi_a^k, \varphi_b^k) \in \Delta_n^{\uparrow G} \wedge \Sigma(c) = \emptyset \} \\
&\cup \{ \delta([\Sigma], \mathbf{dmod}_F(\varphi)) \mid \delta(\Sigma, \varphi) \in \Delta_n^{\uparrow G} \wedge F \subseteq \mathbf{regions}(\varphi) \} \\
&\cup \Delta_n^{\uparrow G}
\end{aligned}$$

where  $c = d_a db, \emptyset \subset K \subseteq J$ , all  $G$  are of the form  $\mathbf{dmod}(f(G'))$ , that the only place that the  $\mathbf{dmod}()$  and  $f()$  annotations appear is at top level, and

$$\begin{aligned}
\varphi'_a &= \begin{cases} \mathbf{dmod}_\emptyset(f(\oplus_K(c!(T_k); \varphi_a^k))) & \text{if } G \upharpoonright d_a \equiv \mathbf{dmod}_\emptyset(f_{d_i}(\oplus_K(c!(T_k); \varphi_a^k))) \\ \oplus_K(c!(T_k); \varphi_a^k) & \text{otherwise} \end{cases} \\
\varphi'_b &= \begin{cases} \mathbf{dmod}_\emptyset(f(\&_J(c?(T_j); \varphi_b^j))) & \text{if } G \upharpoonright d_b \equiv \mathbf{dmod}_\emptyset(f_{d_i}(\&_J(c?(T_j); \varphi_b^j))) \\ \&_J(c?(T_j); \varphi_b^j) & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 5.3: Annotated Blocking Message Passing Safe Program Points Set

are typable using *some* Global Session Type, and that the message queues are empty. This guarantees that communications actions are in parity, and also that processes are in consistent states with respect to transactions. We call this approach *Global Typability*, and formalise its constraints on the updates that are permissible as:

$$\mathcal{GT}(\psi) \text{ iff } (\psi = p_{gt}(z_\Sigma, z_\Phi, z_\psi) \mapsto \{ \widetilde{f_{d_i} \mapsto t_{d_i}} \} \text{ and } \exists \Phi. \vdash \prod t_{d_i} : \Phi^\uparrow) \quad (5.1)$$

where  $p_{gt}(z_\Sigma, z_\Phi, z_\psi)$  is the predicate  $(z_\Sigma = \Sigma_\emptyset) \wedge z_{\Phi^\uparrow}$ , and  $z_{\Phi^\uparrow}$  denotes that  $z_\Phi$  is the projection of some Global Session Type.

In order to ensure that updates transform programs from (provably) safe and live states to (provably) safe and live states, we must extend the consistency definition from Figure 3.13 to account for possible program states that can be reached by reducing a program with  $\mathbf{dmod}()$  and  $f()$  annotations. We modify the inductive definition of a safe program points sets to comprise such points (Figure 5.3). Note that, unlike in Figure 3.13, the safe points set is parameterised by a global session type  $G$ . This is used to maintain correct encapsulation with respect to updateable and transactional regions. We return to this in the description of the second construction, below. The base case consists of program points that are projections of Global Session Types, with accompanying empty message queues. Here we make explicit our assumption that programs are of the form  $\mu \underline{X}. \mathbf{dmod}(f(G_1))$ , and that the only place that the  $\mathbf{dmod}()$  and  $f()$  annotations appear is at top level. In addition, we require that the body of the protocol only appear at top level, under the recursive binding. We will return to this requirement momentarily.

```

1   $[\sigma_\emptyset] (t_c \parallel t_s), \psi_\emptyset$ 
2   $\rightarrow \dots$ 
3   $\rightarrow [c_1 \mapsto \text{prime?}] (\text{snd}(c, 3457).\text{rcv}(c_2)(x : \text{INT}).t_c \parallel$ 
    $\mu \underline{X}' . g(\text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots). \underline{X}', \dots\}), \psi_\emptyset$ 
4   $\rightarrow [c_1 \mapsto \text{prime?}] (\text{snd}(c, 3457).\text{rcv}(c_2)(x : \text{INT}).t_c \parallel$ 
    $\mu \underline{X}' . g(\text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots). \underline{X}', \dots\}), \psi$ 
5   $\rightarrow [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{INT}).t_c \parallel$ 
    $\mu \underline{X}' . g(\text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots). \underline{X}', \dots\}), \psi$ 
6   $\rightarrow [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{INT}).t_c \parallel$ 
    $g(\text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots).t_s, \dots\}), \psi$ 
7   $\rightarrow [c_1 \mapsto \text{prime?}, 3457] (\text{rcv}(c_2)(x : \text{INT}).t_c \parallel$ 
    $\text{case } c_1 \text{ of } \{\text{prime?} \mapsto \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots).t_s, \dots\}), \psi$ 
8   $\rightarrow [c_1 \mapsto 3457] (\text{rcv}(c_2)(x : \text{INT}).t_c \parallel \text{rcv}(c_1)(x : \text{INT}).\text{snd}(c_2, \dots).t_s), \psi$ 
9   $\rightarrow [\sigma_\emptyset] (\text{rcv}(c_2)(x : \text{INT}).t_c \parallel \text{snd}(c_2, \dots).t_s), \psi$ 
10  $\rightarrow [\sigma_\emptyset] (\text{rcv}(c_2)(x : \text{INT}).t'_c \parallel \text{snd}(c_2, \dots).t'_s), \psi_\emptyset$ 
11  $\rightarrow \dots$ 

```

where

$$\psi = p \mapsto (f \mapsto t'_c, g \mapsto t'_s) \quad p_s(\Sigma, \Phi) \stackrel{\text{def}}{=} (\Sigma = \Sigma_\emptyset \wedge \exists G. \Phi = G \upharpoonright d_1 \parallel \dots \parallel G \upharpoonright d_n)$$

Figure 5.4: Update Success In Maths Server Using Global Typability

The first and second constructions are similar to those for blocking message passing in Figure 3.13, in that they add a value to the front of the queue with a complementary external choice on receive, and add complementary internal choice between sends and external choice between receives, respectively. There is an additional aspect to the construction, however, that we explain by example. We want to guarantee that all threads are updated, and that an update does not result in old code interacting with new code. Specifically, we wish to rule out situations such as:

$$[\Sigma_\emptyset] \mu \underline{X} . \text{dmod}_\emptyset(f(c! \langle T \rangle; \underline{X})) \parallel \mu \underline{X} . (c? \langle T \rangle; \underline{X}) \quad (5.2)$$

In this case we could update the producer to be empty, thus introducing deadlock into the system. Hence we require that any time we use a construction rule, that if the extended process is the projection of the top level protocol  $\mu \underline{X} . \text{dmod}(f(G_1))$ , that the relevant `dmod()` and `f()` annotations must be added. The third construction permits us to add a `dmodF()` where the region set  $F$  comprises all regions within the enclosed process; this means that the next reduction for this process will simply strip away the extraneous annotation.

We can use Global Typability to perform the update to the maths server (from Figure 4.1 to Figure 4.2) using the update predicate for Global Typability. An example reduction

for such an update is shown in Figure 5.4. On line 4 we introduce the update when there are values in the message queues. The program then continues evaluating until the update predicate becomes true on line 9, after which the update is performed on line 10. The Global Update predicate will rule out example erroneous reductions such as in Figures 4.3.

### 5.1.2 Properties and Proofs

In order to prove that Global Typability implies Global Compatibility we make use of the following lemmas. First we show that any program point within the set of safe program points reduces, then the new program point is still within the set of safe program points.

**Lemma 5.1.** *Valid Simple Points Reduction*

$$\forall \Sigma, \Phi, G_1, n, \Sigma', \Phi'. [\Sigma] \Phi \in \Delta_n^{G_1} \wedge [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow [\Sigma'] \Phi' \in \Delta_n^{G_1}$$

*Proof.* The proof proceeds similarly to that in Lemma 3.9. Suppose the hypotheses. We then prove the conclusion by induction over the  $n$  of  $\Delta_n^{G_1}$ . The key differences are in the case where  $n > 0$ . By the structure of parallel effects we can identify one effect which is the one that reduces:

$$[\Sigma] \varphi_1 \parallel \dots \parallel \varphi_i \parallel \dots \parallel \varphi_m \rightarrow [\Sigma'] \varphi_1 \parallel \dots \parallel \varphi'_i \parallel \dots \parallel \varphi_m$$

As  $[\Sigma] \Phi \in \Delta_n^{G_1}$  we know that  $[\Sigma] \Phi$  is either in one of the sets included in the definition of  $\Delta_n^{G_1}$ , or it is in  $\Delta_{n-1}^{G_1}$ .

Consider the case when  $[\Sigma] \Phi \in \{\delta([c \mapsto T; \Sigma^k], \varphi'_b)\}$  where  $\forall j \in J. \delta(\Sigma^j, \varphi_b^j) \in \Delta_{n-1}^{G_1}$  (as in Figure 5.3). We can perform a case split over the index  $i$ . If  $i = b$  then  $\varphi_b$  is of one of two forms. If  $\varphi \equiv \text{dmod}_\emptyset(f(\&_J(c!(T_j); \varphi_b^j)))$  then by the reduction rules, we know that  $\varphi'_i \equiv \text{dmod}_{f_{d_b}}(\varphi_b^k)$  and that  $\Sigma' = \Sigma_k$ . As we have that  $\delta(\Sigma^k, \varphi_a^k, \varphi_b^k) \in \Delta_{n-1}^{G_1}$ , then by the third construction in Figure 5.3 we have that  $\delta(\Sigma^k, \text{dmod}_{f_{d_b}}(\varphi_b^k)) \in \Delta_{n-1}^{G_1}$ , as required. The cases where  $\varphi \equiv \&_J(c!(T_j); \varphi_b^j)$  and  $i \neq b$  continue as in Lemma 3.9.

Consider the case when  $[\Sigma] \Phi \in \{\delta(\Sigma, \varphi'_a, \varphi'_b)\}$  where  $\forall j \in J. \delta(\Sigma^j, \varphi_a^j, \varphi_b^j) \in \Delta_{n-1}^{G_1}$  and  $\Sigma(c) = \emptyset$  (as in Figure 5.3). This follows the first case, and has four permutations (rather than two), as  $\varphi'_a$  can be of form  $\text{dmod}_\emptyset(f(\oplus_K(c!(T_k); \varphi_a^k)))$  and of form  $\oplus_K(c!(T_k); \varphi_a^k)$ , and  $\oplus_K(c!(T_k); \varphi_a^k)$  can either be a true choice, or a degenerate choice  $c!(T_k); \varphi_a^k$ , again making use of the third construction in Figure 5.3 when  $\varphi'_a \equiv \text{dmod}_\emptyset(f(\oplus_K(c!(T_k); \varphi_a^k)))$ .

Consider the case when  $[\Sigma] \Phi \in \{\delta(\Sigma, \text{dmod}_F(\varphi_i))\}$  where  $\delta(\Sigma, \varphi_i) \in \Delta_{n-1}^{G_1}$  and  $F \subseteq \text{regions}(\varphi_a)$  (as in Figure 5.3). By the (DMod) rule we know that  $[\Sigma] \varphi_i \rightarrow [\Sigma'] \varphi'_i$ . As  $\delta(\Sigma, \varphi_i) \in \Delta_{n-1}^{G_1}$ , then by the inductive hypothesis we can show that  $\delta(\Sigma', \varphi'_i) \in$

$\Delta_{n-1}^{|G_1|}$ . Then by the third construction we can show that  $\delta(\Sigma', \text{dmod}_F(\varphi'_i)) \in \Delta_{n-1}^{|G_1|}$ , as required  $\square$

We can prove that valid program points are safe and live.

**Lemma 5.2.** *Valid States Safety Lemma*

$$\forall \Sigma, \Phi, G, n. \quad [\Sigma] \Phi \in \Delta_n^{|G|} \Rightarrow \text{error} \notin \Phi$$

*Proof.* Similar to that of Lemma 3.10.  $\square$

**Lemma 5.3.** *Valid States Liveness*

$$\forall \Sigma, \Phi, G, n. \quad [\Sigma] \Phi \in \Delta_n^{|G|} \Rightarrow \exists \Sigma', \Phi'. \quad [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \vee \Phi \equiv \Pi \epsilon$$

*Proof.* The proof proceeds similarly to that of Lemma 3.11, by induction over  $n$ . The key difference is where  $n > 0$  and where  $[\Sigma] \Phi \in \{\delta(\Sigma, \text{dmod}_F(\varphi_i))\}$ , given that  $\delta(\Sigma, \varphi_i) \in \Delta_{n-1}^{|G|}$  and  $F \subseteq \text{regions}(\varphi)$ . By the reduction rules, as  $F \subseteq \text{regions}(\varphi)$ , we know that  $[\Sigma] \text{dmod}_F(\phi) \rightarrow [\Sigma] \phi$ , and hence

$$[\Sigma] \phi_1 || \dots || \text{dmod}_F(\phi) || \dots || \phi_m \rightarrow [\Sigma] \phi_1 || \dots || \phi || \dots || \phi_m$$

as required.  $\square$

We prove that, if we apply a well formed update to a valid (abstract) process point, that the updated point is still valid.

**Lemma 5.4.** *Global Session Type Updatability*

$\forall \Phi, G, u, f_i, t_i, \phi_i, G', b. \text{ if}$

$$\Phi \equiv \Phi^{\uparrow G} \quad u = \widetilde{f_i \mapsto t_i}_I \quad \varphi_i \wr \emptyset \vdash t_i : \text{UNIT} \quad G' \upharpoonright d_i \equiv \varphi_i$$

*Then  $\exists G'', \Phi'$  such that*

$$\text{upd}(\Phi, u, b) \equiv \Phi'^{\uparrow G''}$$

*Proof.* Suppose the hypotheses. We then prove the conclusion by induction over the structure of  $G$ . The key cases are as follows.

*Consider the case where  $G = f(G_1)$ . The update function depends on the boolean value  $b$ . Hence we perform a case split on its possible values. Consider the case  $b = \text{true}$ . By*

the definition of  $\upharpoonright$  we know that  $\Pi f(G_1) \upharpoonright d \equiv \Pi f_d(G_1 \upharpoonright d)$ . By Lemma A.4 and the definition of  $\text{upd}$  we know that:

$$\begin{aligned}\Pi \text{upd}(f(G_1) \upharpoonright d, u, \text{true}) &\equiv \Pi \text{upd}(\varphi'_i, u, \text{true}) \\ &\equiv \Pi \text{upd}(G' \upharpoonright d, u, \text{true})\end{aligned}$$

as required. Consider the case where  $b = \text{false}$ . By Lemma A.4 and the definition of  $\text{upd}$  we know that  $\Pi \text{upd}(f(G_1) \upharpoonright d, u, \text{false}) \equiv \Pi f_d(\text{upd}(G_1 \upharpoonright d, u, \text{false}))$ . We can use the inductive hypothesis to show that  $\Pi \text{upd}(G_1 \upharpoonright d, u, \text{false}) \equiv \Pi G''' \upharpoonright d$ . Hence we have that  $\Pi \text{upd}(f(G_1) \upharpoonright d, u, \text{false}) \equiv \Pi f_d(\text{upd}(G''' \upharpoonright d, u, \text{false}))$  and  $G'' = f(G''')$ , as required.

Consider the case where  $G = \text{dmod}_F(G_1)$ . By Lemma A.4 and the definition of  $\text{upd}$  we know that  $\Pi \text{upd}(\text{dmod}_F(G_1) \upharpoonright d, u, b) \equiv \Pi \text{dmod}_{F_d}(\text{upd}(G_1 \upharpoonright d, u, \text{true}))$ , where  $\text{regions}(G_1) \cap F = \emptyset$  and  $\Pi \text{upd}(\text{dmod}_F(G_1) \upharpoonright d, u, b) \equiv \Pi \text{dmod}_{F_d}(\text{upd}(G_1 \upharpoonright d, u, \text{false}))$ , where  $\text{regions}(G_1) \cap F \neq \emptyset$ . By the inductive hypothesis we can show that  $\Pi \text{upd}(G_1 \upharpoonright d, u, b) \equiv \Pi G''' \upharpoonright d$ . Hence we have that  $\Pi \text{upd}(\text{dmod}_F(G_1) \upharpoonright d, u, b) \equiv \Pi (\text{dmod}_F(G''')) \upharpoonright d$ , as required.  $\square$

We can show that if we apply a well formed update to a valid program point then the updated point is still valid.

**Lemma 5.5.** *Global Typability Update Consistency*

If:

$$\begin{aligned}\vdash P: \Phi \quad \exists G. \Phi = G \upharpoonright d_1 \parallel \dots \parallel G \upharpoonright d_n \\ u = f_{d_1} \mapsto t_{d_1}, \dots, f_{d_n} \mapsto t_{d_n}. \exists G'. \forall i \in 1 \dots n. G \upharpoonright d_i \setminus \emptyset \vdash t_{d_i}: T\end{aligned}$$

then:

$$\vdash \text{upd}(P, u, \text{false}): \Phi' \quad \exists G'. \Phi' = G' \upharpoonright d_1 \parallel \dots \parallel G' \upharpoonright d_n$$

*Proof.* By induction over the structure of  $G$ .  $\square$

We can then prove that if such an update is introduced into a (valid) message passing program that has reduced from the projection of a Global Session Type, that whenever the update is applied it will result in another valid program. An informal explanation of this proof is as follows. Up until the point that an update is introduced, and its update predicate has become true, we can use the fact that the program point is a member of the valid points set to show that the program is live and safe, as in Observation 3.13 and Observation 3.14. At the point the update is applied we know that the message queues are empty, and the effect of the process is the projection of a Global Session Type. When the well formed update is applied to the program we can use Lemma 5.4 to show that the updated program's effect is still the projection of a Global Session Type. As the message queues are still empty then we can show that  $\vdash \text{upd}(P, u, \text{false}): \Phi'^\dagger$ .

and hence that  $[\Sigma] \mathsf{upd}(\Phi, u, \mathsf{false}) \in \Delta_0^{G_2}$ . After the update is applied we know that no updates will be introduced or applied again. hence we can show that the resulting program is safe and live, as in Section 3.5.1.

**Theorem 5.6.** *Global Typability Theorem*

For all  $k, \Sigma, P, \psi, \omega$ , if  $(\mathcal{GT}(\psi) \text{ or } \phi = \psi = \emptyset)$  and  $[\Sigma] \Phi \in \Delta^{G_1}$  then  $\mathcal{G}^k(\Sigma, \Phi, \psi, \omega)$ .

*Proof.* Suppose the hypotheses. We then prove the conclusion by induction over the  $k$  in  $\mathcal{G}^k(\Sigma, \Phi, \psi, \omega)$ .

Consider the case where  $k = 0$ . We trivially have that  $\mathcal{G}^0(\Sigma, \Phi, \psi, \omega) = \mathsf{true}$ .

Consider the case where  $k > 0$ . We prove each of the conjuncts of  $\mathcal{G}$  separately. In order to prove the second conjunct: by Lemma 5.2 we can show that  $\#\mathsf{error} \in \Phi'$ , as required. We can show the third conjunct directly using the inductive hypothesis, and hence have  $\mathcal{G}^{k-1}(\Sigma, \Phi, (\psi, \omega), \emptyset)$  as required, as  $k - 1 < k$ . In order to prove the third conjunct we must either show that no update is applicable, and for all possible reductions the resulting program points are globally compatible, or if an update is applicable then the updated program point is globally compatible.

Consider the case where there are no applicable updates. By Lemma 5.1 we know that if there exists  $\Sigma'$  and  $\Phi'$  such that there is a reduction  $[\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$  then the resulting program point  $[\Sigma'] \Phi'$  are still in the valid points set,  $[\Sigma'] \Phi' \in \Delta^{G_1}$ . By the inductive hypothesis, as  $\mathcal{GT}(\psi)$  and  $[\Sigma'] \Phi' \in \Delta^{G_1}$  then we know that  $\mathcal{G}^{k-1}(\Sigma', \Phi', \psi, \omega)$ . By Lemma 5.3, then as  $[\Sigma] \Phi \in \Delta_n^G$  then we know that either  $[\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$  or  $\Phi \equiv \Pi \epsilon$ , as required.

Consider the case where there is an applicable update. We can choose the first update such that its predicate is applicable, and let  $i$  be this update's index. By Lemma 5.4 we know that  $\mathsf{upd}(\Phi, u, b) \equiv \Phi^i$ , and hence that  $[\Sigma_\emptyset] \mathsf{upd}(\Phi, u, b) \in \Delta_0^{G_2}$ . We can use the inductive hypothesis to show that  $\mathcal{G}^{k-1}(\Sigma_\emptyset, \mathsf{upd}(\Phi, u, b), \emptyset, \emptyset)$ , as required.  $\square$

### 5.1.3 Summary

We can update programs whose effects are projections of Global Session Types; such situations correspond to a degree of global synchronisation. As such, at each point during the reduction we must consider whether the whole program can be so abstracted. At first glance this might seem a heavy requirement, however in practice we believe that we could use a check in/check out approach similar to that used in (Neamtiu and Hicks, 2009) to mitigate this cost. As an example we present an annotation of the program

$$\begin{aligned}
t_p &= \mu \underline{X}.\text{dmod}_\emptyset(f(\text{snd}(c_1, v_1).\underline{X})) \\
t_b &= \mu \underline{X}'.\text{dmod}_\emptyset(f'(\text{rcv}(c_1)(x : T_1).\text{snd}(c_2, v_2).\underline{X}')) \\
t_{co} &= \mu \underline{X}''.\text{dmod}_\emptyset(f''(\text{rcv}(c_2)(x : T_2).\underline{X}''))
\end{aligned}$$

Figure 5.5: Buffer System

---


$$\begin{aligned}
1 & [ \sigma_\emptyset ] (t_p \parallel t_b \parallel t_{co}) \\
2 & \rightarrow [ c_1 \mapsto v_1 ] (\text{dmod}_f(t_p) \parallel t_b \parallel t_{co}) \\
3 & \rightarrow [ \sigma_\emptyset ] (\text{dmod}_f(t_p) \parallel \text{dmod}_{f'}(\text{snd}(c_2, v_2).t_b) \parallel t_{co}) \\
4 & \rightarrow [ \sigma_\emptyset ] (t_p \parallel \text{dmod}_{f'}(\text{snd}(c_2, v_2).t_b) \parallel t_{co}) \\
5 & \rightarrow [ \sigma_\emptyset ] (t_p \parallel \text{snd}(c_2, v_2).t_b \parallel t_{co}) \\
6 & \rightarrow [ c_1 \mapsto v_1 ] (\text{dmod}_f(t_p) \parallel \text{snd}(c_2, v_2).t_b \parallel t_{co}) \\
7 & \rightarrow [ c_1 \mapsto v_1, c_1 \mapsto v_2 ] (\text{dmod}_f(t_p) \parallel t_b \parallel t_{co}) \\
8 & \rightarrow [ c_1 \mapsto v_1 ] (\text{dmod}_f(t_p) \parallel t_b \parallel \text{dmod}_{f''}(t_{co})) \\
9 & \rightarrow \dots
\end{aligned}$$

Figure 5.6: Infinite Delay Using Round Robin Fair Scheduling

---

defined in Figure 4.4.

$$\begin{aligned}
t_p &= {}^1\mu \underline{X}. {}^2\text{snd}(c_1, v_1). {}^3\text{rcv}(c_2)(x : \text{UNIT}). \underline{X} \\
t_{co} &= {}^1\mu \underline{X}' . {}^2\text{rcv}(c_1)(x : T_1). {}^3\text{snd}(c_2, ()). \underline{X}' 
\end{aligned} \tag{5.3}$$

We can then represent the code point  $t_p \parallel t_{co}$  as 11, the code point:

$${}^2\text{snd}(c_1, v_1). {}^3\text{rcv}(c_2)(x : \text{UNIT}). t_p \parallel t_{co} \tag{5.4}$$

as 21, etc. We can then easily show that the effect of a program is the projection of a Global Session Type when a program is at point 11, 22, etc.

## 5.2 Local Update

While the Global Typability approach is applicable to all long running event processing loops, it suffers from the disadvantage that in some settings the update can be delayed indefinitely. One example where we may reasonably expect the indefinite delay of an update's application, even in the presence of fair scheduling, is a producer/consumer example with a buffer (Figure 5.5). If we take a round robin scheduling of this process, we obtain the reduction sequence shown in Figure 5.6. Here, after the initial point, we never reach a state where the system is the projection of a global session type; either a value is in the channel queues, or the threads are out of sync with respect to their  $\text{dmod}()$  annotations.

$$\begin{aligned}
t_{p1} &= \mu \underline{X}.\text{dmod}_\emptyset(f(\text{snd}(c_1, v_1).\underline{X})) \\
t_{p2} &= \mu \underline{X}'.\text{dmod}_\emptyset(f'(\text{snd}(c_2, v_2).\underline{X}')) \\
t_{co} &= \mu \underline{X}''.\text{dmod}_\emptyset(f''(\text{rcv}(c_1)(x : T_1).\text{rcv}(c_2)(x : T_2).\underline{X}''))
\end{aligned}$$

Figure 5.7: Two Producer/Consumer System Without Acknowledgement

1	$[\sigma_\emptyset]$	$(t_{p1} \parallel t_{p2} \parallel t_{co}), \psi_\emptyset$
2	$\rightarrow$	$[c_1 \mapsto v_1]$ $(\text{dmod}_f(t_{p1}) \parallel t_{p2} \parallel t_{co}), \psi_\emptyset$
3	$\rightarrow$	$[c_1 \mapsto v_1]$ $(\text{dmod}_f(t_{p1}) \parallel t_{p2} \parallel t_{co}), \psi$
4	$\rightarrow$	$[c_1 \mapsto v_1]$ $(\text{dmod}_f(t_{p1}) \parallel t'_{p2} \parallel t_{co}), \psi'$
5	$\rightarrow$	$[c_1 \mapsto v_1]$ $(t_{p1} \parallel t'_{p2} \parallel t_{co}), \psi'$
6	$\rightarrow$	$[c_1 \mapsto v_1]$ $(t'_{p1} \parallel t'_{p2} \parallel t_{co}), \psi''$

Figure 5.8: Introducing Deadlock Using Local Update

We should be able to update the system in Figure 5.5; intuitively we can see that if we update the producer first, then the buffer, then finally the consumer, that we should avoid errors and deadlock. This, however, requires updating the threads separately, and risks the old protocol interacting with the new.

Communication errors occur when a value is received on a channel that is not of the expected type. If we assume that the new program uses different channels to those in the old program, then neither can receive a value sent by the other, and we rule out communication errors completely. The problem then consists of guaranteeing liveness.

In this example, the producer never receives any values, and hence cannot be blocked (and hence deadlocked) while it is performing the old protocol. If we are to update the threads separately, this therefore seems like a good candidate to be the first updated. Once updated it may block waiting for values sent by other threads under the new protocol; this is not problematic if we can show that all threads will be updated to the new protocol, and that once they all are that the system behaves as a static message passing programs as in Chapter 3. We then have to guarantee that each thread in turn migrates to the new protocol.

The key restriction that we introduce in order to update a program using Local Update is that only one process can start the protocol. Consider the system in Figure 5.7. This program is a producer/consumer protocol with two senders and one receiver. Hence either the first or the second producer can evaluate first, and start the protocol. Consider an update that consists of three separate predicated updates. The predicates for  $t_{pi}$  are true when the thread reaches the top of its loop. The predicate for  $t_{co}$  is true when it reaches the top of its loop, its channel queues are empty, and the producer threads have already performed their updates. Then consider one reduction of this system

1	$[\sigma_\emptyset] (t_p \parallel t_b \parallel t_{co}), \psi_\emptyset$
2	$\rightarrow \dots$
3	$\rightarrow [c_1 \mapsto v_1] (t_p \parallel t_b \parallel t_{co}), \psi_\emptyset$
4	$\rightarrow [c_1 \mapsto v_1] (t_p \parallel t_b \parallel t_{co}), \psi_1, \psi_2, \psi_3$
5	$\rightarrow [c_1 \mapsto v_1] (t'_p \parallel t_b \parallel t_{co}), \psi_2, \psi_3$
6	$\rightarrow [\sigma_\emptyset] (t'_p \parallel \mathbf{dmod}_{f_2}(c_2!(v_2); t_b) \parallel t_{co}), \psi_2, \psi_3$
7	$\rightarrow [\sigma_\emptyset] (t'_p \parallel c_2!(v_2); t_b \parallel t_{co}), \psi_2, \psi_3$
8	$\rightarrow [c_2 \mapsto v_2] (t'_p \parallel t_b \parallel t_{co}), \psi_2, \psi_3$
9	$\rightarrow [c_2 \mapsto v_2] (t'_p \parallel t'_b \parallel t_{co}), \psi_3$
10	$\rightarrow [\sigma_\emptyset] (t'_p \parallel t'_b \parallel \mathbf{dmod}_{f_3}(t_{co})), \psi_3$
11	$\rightarrow [\sigma_\emptyset] (t'_p \parallel t'_b \parallel t_{co}), \psi_3$
12	$\rightarrow [\sigma_\emptyset] (t'_p \parallel t'_b \parallel t'_{co}), \psi_\emptyset$

where:

$$\begin{aligned}\psi_i &= p_i \mapsto (f_i \mapsto t_i) & p_1([\Sigma] \Phi, \psi) &\stackrel{\text{def}}{=} (\varphi_1 \equiv \mu \underline{X}.\mathbf{dmod}_\emptyset(f_1(\varphi'_1))) \\ p_2([\Sigma] \Phi, \psi) &\stackrel{\text{def}}{=} (\varphi_2 \equiv \mu \underline{X}.\mathbf{dmod}_\emptyset(f_2(\varphi'_2)) \wedge \Sigma(c_1) = \emptyset \wedge \psi_1 \notin u) \\ p_3([\Sigma] \Phi, \psi) &\stackrel{\text{def}}{=} (\varphi_3 \equiv \mu \underline{X}.\mathbf{dmod}_\emptyset(f_3(\varphi'_3)) \wedge \Sigma(c_2) = \emptyset \wedge \psi_1 \notin u \wedge \psi_2 \notin u)\end{aligned}$$

Figure 5.9: Update Success Using Local Update

(Figure 5.8). After the first producer sends a value, the update is introduced. The second consumer updates straight away. The first consumer then drops its old  $\mathbf{dmod}()$  and updates (line 6). As the first producer is now working under the new protocol, it will never send a value on  $c_1$ . As the second producer has already sent a value under the old protocol, however, the consumer cannot update, but also will never reduce as it will never receive on channel  $c_1$ , hence the consumer is deadlocked.

We describe an example where Local Update preserves safety and liveness. We abuse notation and define the update predicates over effects rather than terms. Consider an update to the system defined in Figure 5.5, which updates the producer whenever it is at the top of its main loop, which updates the buffer when it is at the top of its main loop, it has no messages in its buffers, and the producer has already been updated, and updates the consumer when it is at the top of its main loop, it has no messages in its buffers, and both the producer and the buffer have already been updated. We present one possible reduction sequence of this system in Figure 5.9. Note how the producer is updated immediately (line 3), but the buffer must reduce and consumer its messages before it can update (line 9). Note also how the consumer waits until the buffer has updated before it performs its update, despite the fact it is at the top of its main loop and has no messages in its channel queues. In this example the consumer could update before the buffer without incident, however in a similar system, where the consumer

$$\begin{aligned}
\Delta_{\underline{D}}^{\uparrow G} &\stackrel{\text{def}}{=} \bigcup_{n < \omega} {}_n\Delta_{\underline{D}}^{\uparrow G} \quad {}_0\Delta_{\underline{D}}^{\uparrow G} \stackrel{\text{def}}{=} \{ [\Sigma_{\emptyset}] \Phi_{\underline{D}}^{\uparrow G} \} \\
{}_n\Delta_{\underline{D}}^{\uparrow G} &\stackrel{\text{def}}{=} \{ \delta([c \mapsto T; \Sigma^k], \varphi'_b) \mid \forall j \in J. \delta(\Sigma^j, \varphi_b^j) \in {}_n\Delta_{\underline{D}}^{\uparrow G} \wedge k \in J \} \\
&\cup \{ \delta(\Sigma, \varphi'_a, \varphi'_b) \mid \forall j \in J. \delta(\Sigma, \varphi_a^j, \varphi_b^j) \in {}_n\Delta_{\underline{D}}^{\uparrow G} \wedge \Sigma(c) = \emptyset \wedge d_a, d_b \in \underline{D} \} \\
&\cup \{ \delta(\Sigma, \mathbf{dmod}_{\{f_{d_b}\}}(\varphi_b^k)) \mid \delta(\Sigma, \varphi_b^k) \in {}_n\Delta_{\underline{D}}^{\uparrow G} \wedge \\
&\quad G \upharpoonright d_b \equiv \mathbf{dmod}_{\emptyset}(f_{d_b}(\&_J(c?(T_j); \varphi_b^j))) \} \\
&\cup \{ \delta(\Sigma, \mathbf{dmod}_{\{f_{d_a}\}}(\varphi_a^k)) \mid \delta(\Sigma, \varphi_a^k) \in {}_n\Delta_{\underline{D}}^{\uparrow G} \wedge \\
&\quad G \upharpoonright d_a \equiv \mathbf{dmod}_{\emptyset}(f_{d_a}(\oplus_K(c!(T_k); \varphi_a^k))) \} \\
&\cup {}_n\Delta_{\underline{D}}^{\uparrow G}
\end{aligned}$$

where  $c = d_a d_b$ ,  $\emptyset \subset K \subseteq J$ ,  $D = \overline{D} \cup \underline{D}$ , and  $\overline{D} \cap \underline{D} = \emptyset$ , and

$$\begin{aligned}
\varphi'_a &= \begin{cases} \mathbf{dmod}_{\emptyset}(f_{d_a}(\oplus_K(c!(T_k); \varphi_a^k))) & \text{if } G \upharpoonright d_a \equiv \mathbf{dmod}_{\emptyset}(f_{d_a}(\oplus_K(c!(T_k); \varphi_a^k))) \\ \oplus_K(c!(T_k); \varphi_a^k) & \text{otherwise} \end{cases} \\
\varphi'_b &= \begin{cases} \mathbf{dmod}_{\emptyset}(f_{d_b}(\&_J(c?(T_j); \varphi_b^j))) & \text{if } G \upharpoonright d_b \equiv \mathbf{dmod}_{\emptyset}(f_{d_b}(\&_J(c?(T_j); \varphi_b^j))) \\ \&_J(c?(T_j); \varphi_b^j) & \text{otherwise} \end{cases} \\
\Delta_{\overline{D}}^{\uparrow G} &\stackrel{\text{def}}{=} \{ [\Sigma] \prod_{\overline{D}} \phi_{d_i} \mid [\Sigma] \prod_{\overline{D}} \in \Delta^{\uparrow G} \wedge \\
&\quad \forall d \in \underline{D}. (\phi_d \equiv G \upharpoonright d \wedge \forall c \in \mathbf{S}(\phi_d). \Sigma(c) = \emptyset) \}
\end{aligned}$$

Figure 5.10: Relaxed Valid Points Sets for Local Update

sends an acknowledgement to the buffer after receiving a value, if the consumer updates before the buffer then the buffer could be left part way through its execution, waiting for the acknowledgement, introducing deadlock.

Before we define Local Update we need to make some simplifying assumptions. We again will consider looping global types  $G^o$  that are of the form  $\mu \underline{X}. f(G_o)$  where  $G_o$  does not contain any instances of  $\mu$  nor any region annotations, as in the previous section. We assume that the old program is one that can evolve to the projection of a global type  $G_1$ . We assume that the new program is a projection of the global type  $G_2$ , that no channels are shared between  $G_1$  and  $G_2$ , and that  $G_1$  and  $G_2$  share the same set of roles. We assume that  $G_1$  and  $G_2$  are both of form  $G_o$ . We define a partial order  $<_{G_1}$  over participants  $d$  of a protocol  $G_1$  denoting the order under which the participants can begin execution of a protocol  $G_1$  (Figure 5.11). For example, for the Producer/Buffer/Consumer example we would define  $d_p < d_b < d_{co}$ . We prove that the relation defined in Figure 5.11 is a strict partial order in Lemma D.1.

$$\begin{array}{c}
\frac{d_1 <_G d_2}{d_1 <_{\mu \underline{X}.G} d_3} \quad \frac{d_1 <_G d_2}{d_1 <_{f(G)} d_3} \quad \frac{d_1 <_G d_2}{d_1 <_{\mathbf{dmod}_F(G)} d_3} \quad \frac{d_1 <_G d_2 \quad d_2 <_G d_3}{d_1 <_G d_3} \\
\\
\frac{d_1 <_{G_k} d_2 \quad d_2 \neq d, d'}{d_1 <_{d \rightarrow d':c\langle \widetilde{T \mapsto G} \rangle} d_2} \quad \frac{}{d_1 <_{d_1 \rightarrow d_2:c\langle \widetilde{T \mapsto G_k} \rangle} d_2}
\end{array}$$

Figure 5.11: Definition of Participant Partial Order

We require that the order has a unique minimum element. In particular this will rule out examples such as:

$$\begin{aligned}
\mu \underline{X}.\mathbf{dmod}(f( & \quad d_1 \rightarrow d_3: c_1\langle T_1 \rangle; \\
& \quad d_2 \rightarrow d_3: c_2\langle T_2 \rangle; \\
& \quad 0 \\
& ))
\end{aligned} \tag{5.5}$$

which represents the system in Figure 5.6. We omit the subscript when the meaning is clear. We only define  $<_{G_1}$  for the Global Session Type  $G_1$  of the original program. This order does not change as a given  $[\Sigma] \Phi$  reduces. We assume that we can partition the roles in  $G_1$  (and hence  $G_2$ ) into two disjoint sets, where  $\underline{D}$  denotes the roles of processes that have not yet been updated, and  $\overline{D}$  denote the roles of processes that have been updated. We write  $\underline{D} < \overline{D}$  when  $\forall d \in \underline{D}, d' \in \overline{D}$  we have that  $d \not< d'$ .

We refer to the set of channels that are received upon in an effect  $\phi$  as  $\mathbf{R}(\phi)$ , and refer to the set of channels that are sent upon by an effect  $\phi$  as  $\mathbf{S}(\phi)$ . Define the Local Update predicate  $(LU)\phi$  on code updates as:

$$\mathcal{LU}(\psi) \text{ iff } (\psi = (p_i(z_\Sigma, z_\Phi, z_\psi) \mapsto (f_{d_i} \mapsto \widetilde{t_{d_i}})_D) \text{ and } \forall d_i \in D. \vdash t_{d_i} : G \upharpoonright d_i) \tag{5.6}$$

where  $p_i(z_\Sigma, z_\Phi, z_\psi)$  is  $\forall c \in \mathbf{R}(z_{\phi_i}).z_\Sigma(c) = \emptyset \wedge \forall d_j <_{G_1} d_i. z_{\psi_j}$  is empty,  $\psi_j$  is the update for process projected from  $d_j$ , and  $D$  is the set of all roles used in  $G_1$ .

We relax the definition of the set of abstract configurations reducing towards  $G$ . We define the relaxed projection  $\Phi_{\underline{D}}^{|G}$  to mean that  $\Phi = \prod_D G \upharpoonright G \upharpoonright d_i$  for all participants  $d_i$  in a given  $D$ . Note that  $D$  need not comprise all roles in the given  $G$ .

We define the relaxed set of configurations for the non-updated roles  $\underline{D}$  as  $\Delta_{\underline{D}}^{|G_1}$ , where  $G_1$  is the protocol of the old program (Figure 5.10). This definition permits the prefix of messages from processes whose role is not present in  $\underline{D}$  when prefixing a receive action to a process that is present, along with prefixing send and receive actions to processes whose roles are both in  $\underline{D}$ . We perform the choice between possible branches, and the correct annotation of regions as previously.

We define the relaxed set of configurations for the updated roles  $\overline{D}$  as  $\Delta_{\overline{D}}^{|G_2}$ , where  $G_2$  is the protocol of the new program (Figure 5.10). This set filters the non-relaxed definition

$\Delta^{\uparrow G_2}$  (Figure 5.3) for those configurations where the effects of all the roles that are not in  $\overline{D}$  are projections of the protocol  $G_2$ . That is, any roles that have not yet been updated should start at the beginning of their runs of  $G_2$ .

### 5.2.1 Properties and Proofs

In order to prove that Local Update implies Global Compatibility we make use of some technical lemmas.

**Lemma 5.7.** *Old Program Subject Reduction*

$$\forall \Sigma, \Phi, G, n, \underline{D}, \Sigma', \Phi'. \quad [\Sigma] \Phi \in {}_n \Delta_{\underline{D}}^{\uparrow G} \wedge [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow [\Sigma'] \Phi' \in \Delta_{\underline{D}}^{\uparrow G}$$

where the participants of  $G$  are  $D$ , and  $\underline{D} <_G \overline{D}$ .

*Proof.* This proof is similar that of Lemma 5.1, and proceeds by induction over  $n$ . The key differences are in the case when  $n > 0$ , in the third and fourth productions of  ${}_n \Delta_{\underline{D}}^{\uparrow G}$ , as in Figure 5.10. Let  $[\Sigma] \Phi \in \{ \delta(\Sigma, \text{dmod}_{\{f_{d_b}\}}(\varphi_b^k)), \text{ where } \delta(\Sigma, \varphi_b^k) \in {}_n \Delta_{\underline{D}}^{\uparrow G} \text{ and } G \upharpoonright d_b \equiv \text{dmod}_{\emptyset}(f_{d_b}(\&_J(c?(T_j); \varphi_b^j))) \}$ . If the process that is reduced is  $\text{dmod}_{\{f_{d_b}\}}(\varphi_b^k)$  then, by the reduction rules, we know that  $[\sigma] \text{dmod}_{\{f_{d_b}\}}(\varphi_b^k) \rightarrow [\Sigma] \varphi_b^k$ . By the hypotheses we know that  $\delta(\Sigma, \varphi_b^k) \in {}_n \Delta_{\underline{D}}^{\uparrow G}$ , as required. If another process is reduced than we proceed as in Lemma 5.1. The fourth case is similar to the third.

□

**Lemma 5.8.** *New Program Subject Reduction*

$$\forall \Sigma, \Phi, n, G, \overline{D}, \Sigma', \Phi'. \quad [\Sigma] \Phi \in {}_n \Delta_{\overline{D}}^{\uparrow G} \wedge [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow [\Sigma'] \Phi' \in \Delta_{\overline{D}}^{\uparrow G}$$

where the participants of  $G$  are  $D$ , and  $\underline{D} <_G \overline{D}$ .

*Proof.* This proof directly follows that of Lemma 5.1, and proceeds by induction over  $n$ . Note that the key differences between the definition of  ${}_n \Delta_{\overline{D}}^{\uparrow G}$  and  $\Delta^{\uparrow G}$  (as in Lemma 5.1) are the restriction on which processes are included, and the fact that channels for non-included roles must have empty queues. As reduction does not add or remove processes, and cannot add messages to channels of processes that are not included (by principle channel allocation), then we can use essentially the same proof technique. □

We prove the straightforward property that program points in the valid set do not contain errors.

**Lemma 5.9.** *Local Update Safety*

$$\forall \Sigma, \Phi, G, \underline{D}, \overline{D}. \quad [\Sigma] \Phi \in {}_n\Delta_{\underline{D}}^{|G} \Rightarrow \text{error} \notin \Phi$$

$$\forall \Sigma, \Phi, G, \underline{D}, \overline{D}. \quad [\Sigma] \Phi \in {}_n\Delta_{\overline{D}}^{|G} \Rightarrow \text{error} \notin \Phi$$

where the participants of  $G$  are  $\underline{D}$ , and  $\underline{D} <_G \overline{D}$ .

*Proof.* Assume the hypotheses. We then prove the conclusions by straightforward induction over  $n$ . This proof is similar to that of Lemma 3.10.  $\square$

We prove that each program point in the old program points set is either live, or has are relaxed projections of a global type:

**Lemma 5.10.** *Old Program Liveness*

$\forall \Sigma, \Phi, n, G, \underline{D}$  If  $[\Sigma] \Phi \in {}_n\Delta_{\underline{D}}^{|G}$  then either  $\exists \Sigma', \Phi'. [\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$  or  $[\Sigma] \Phi = [\Sigma_\emptyset] \Phi_{\underline{D}}^{|G}$ , where the participants of  $G$  are  $\underline{D}$ , and  $\underline{D} <_G \overline{D}$ .

*Proof.* Assume the hypothesis. We then prove that one of the two conclusions hold, by induction over  $n$ .

Consider the case where  $n = 0$ . We straightforwardly have, by the definition of  ${}_0\Delta_{\underline{D}}^{|G}$ , that  $[\Sigma] \Phi = [\Sigma_\emptyset] \Phi_{\underline{D}}^{|G_1}$ .

Consider the case where  $n > 0$ . We consider each of the cases of the definition of  ${}_n\Delta_{\underline{D}}^{|G}$  in Figure 5.10. Consider the case where  $[\Sigma] \Phi \in \{ \delta([c \mapsto T; \Sigma^k], \varphi'_b) \}$ . We know that  $[c \mapsto T_k] \&_J (c?(T_j); \varphi'_b) \rightarrow [\Sigma] \varphi'_b$ , and hence  $\exists \Sigma', \Phi'. [\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$ , as required. Consider the case where  $[\Sigma] \Phi \in \{ \delta(\Sigma, \varphi'_a, \varphi'_b) \}$ . By the reduction rules we know that  $[\Sigma] \oplus_K (c!(T_k); \varphi'_a) \rightarrow [\Sigma; c \rightarrow T_k] \varphi'_a$ , and hence  $\exists \Sigma', \Phi'. [\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$ , as required. Consider the case where  $[\Sigma] \Phi \in \{ \delta(\Sigma, \text{dmod}_{\{f_{d_b}\}}(\varphi'_b)) \}$ . By the definition of well formed global session types we know that the only region annotation that each process (with role  $d_i$ ) will have is  $f_{d_b}$ . Hence by the reduction rules for  $\text{dmod}$  we can show that  $\Sigma, \text{dmod}_{\{f_{d_b}\}}(\varphi'_b) \rightarrow \Sigma, \varphi'_b$ , and hence  $\exists \Sigma', \Phi'. [\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$ , as required. Consider the case where  $[\Sigma] \Phi \in \{ \delta(\Sigma, \text{dmod}_{\{f_{d_a}\}}(\varphi'_a)) \}$ . By the definition of well formed global session types we know that the only region annotation that each process (with role  $d_i$ ) will have is  $f_{d_a}$ . Hence by the reduction rules for  $\text{dmod}$  we can show that  $\Sigma, \text{dmod}_{\{f_{d_a}\}}(\varphi'_a) \rightarrow \Sigma, \varphi'_a$ , and hence  $\exists \Sigma', \Phi'. [\Sigma] \Phi \rightarrow [\Sigma'] \Phi'$ , as required. The case where  $[\Sigma] \Phi \in {}_{n-1}\Delta_{\underline{D}}^{|G}$  proceeds immediately from the inductive hypothesis.  $\square$

We need to make sure that updating a process of role  $d_i$  with an update  $\mathcal{LU}(\psi)$  results in a process that conforms to the new protocol  $G_2$ .

**Lemma 5.11.**  $\forall G_1, d_i, T, \psi. \text{ if } G_1 \upharpoonright d_i \setminus \emptyset \vdash t_{d_i} : T \text{ and } \mathcal{LU}(\psi) \text{ then } \exists G_2. G_2 \upharpoonright d_i \setminus \emptyset \vdash \text{upd}(t_i, f_{d_i} \mapsto t'_{d_i}, \text{false}) : T.$

Where  $\psi = (p_i(z_\Sigma, z_\Phi, z_\psi) \rightsquigarrow \widetilde{(f_{d_i} \mapsto t_{d_i})_D})$ , as  $\mathcal{LU}(\psi)$  holds.

*Proof.* This proof follows that of Lemma 5.4.  $\square$

We need to show that if we extend a relaxed configuration of updated processes with a newly updated process then the new configuration will be in relaxed set, extended by the new role:

**Lemma 5.12.**  $\forall \Sigma, \Phi, G_2, \overline{D}, d', \overline{D'}. \text{ if } [\Sigma] \Phi \in \Delta_{\overline{D}}^{|G_2} \text{ and } \overline{D'} = \overline{D} \cup \{d'\} \text{ then } [\Sigma] (\Phi \parallel G_2 \upharpoonright d') \in \Delta_{\overline{D'}}^{|G_2}.$

*Proof.* By the definition of  $\Delta_{\overline{D}}^{|G_2}$  we know that  $[\Sigma] \Phi = [\Sigma] \prod_{\overline{D}} \phi_{d_i}$ , where  $[\Sigma] \prod_{\overline{D}} \in \Delta^{|G}$  and  $\forall d \in \overline{D}. (\phi_d \equiv G \upharpoonright d \wedge \forall c \in S(\phi_d). \Sigma(c) = \emptyset)$ . Hence we can straightforwardly show that  $[\Sigma] (\Phi \parallel G_2 \upharpoonright d') \in \Delta_{\overline{D'}}^{|G_2}$ .  $\square$

We also need to show that if we redact a relaxed configuration of non-updated processes by a newly updated process then the new configuration will be in the relaxed set, redacted by the updated role:

**Lemma 5.13.**  $\forall \Sigma, \Phi, G_1, \underline{D}, n, d', \underline{D'}. \text{ if } [\Sigma] (\Phi \parallel G_1 \upharpoonright d') \in {}_n \Delta_{\underline{D'}}^{|G_1} \text{ and } \forall c \in R(G \upharpoonright d'). \Sigma(c) = \emptyset \text{ and } \underline{D'} = \underline{D} \cup \{d'\} \text{ and } \exists d \in \underline{D} \text{ such that } d < d' \text{ then } [\Sigma] \Phi \in {}_n \Delta_{\underline{D}}^{|G_1}.$

*Proof.* Assume the hypotheses. We then show the conclusion by induction over the  $n$  in  ${}_n \Delta_{\underline{D}}^{|G_1}$ . The case where  $n = 0$  is immediate by the definition of  ${}_0 \Delta_{\underline{D}}^{|G_1}$ . We consider the case when  $n > 0$ . We then proceed by case analysis over the definition of  ${}_n \Delta_{\underline{D'}}^{|G_1}$ , as in Figure 5.10.

Consider the case where  $[\Sigma] \Phi \in \{ \delta([c \mapsto T; \Sigma^k], \varphi'_b) \}$ . We can then consider the identity of  $d'$ . If  $d' = d_b$  then we have a contradiction, as the hypotheses state that  $(c \mapsto T_k; \Sigma)(c) = \emptyset$ , which is clearly not true. If  $d' = d_b$  then by the inductive hypothesis we know that as  $\delta(\Sigma, \phi_{d_b}) = [\Sigma] \dots ||\phi_b|| \dots ||\phi_{d'}|| \dots \in {}_{n-1} \Delta_{\underline{D'}}^{|G_1}$  then we know that  $[\Sigma] \dots ||\phi_b|| \dots \in {}_{n-1} \Delta_{\underline{D}}^{|G_1}$ , and hence by the first construction of Figure 5.10 that  $[\Sigma] \dots ||\phi'_b|| \dots \in {}_{n-1} \Delta_{\underline{D}}^{|G_1}$ , as required.

Consider the case where  $[\Sigma] \Phi \in \{ \delta(\Sigma, \varphi'_a, \varphi'_b) \}$ . We can then consider the identity of  $d'$ . If  $d' = d_b$  and  $\phi'_b = \&_J(c?(T_j); \varphi'_b)$  then we have a contradiction, as the hypothesis requires that  $\phi_{d'} \equiv G_1 \upharpoonright d' \equiv \text{dmod}_\emptyset(f_{d_b}(\&_J(c?(T_j); \varphi'_b)))$ . If  $d' = d_b$  and  $\phi'_b = \text{dmod}_\emptyset(f_{d_b}(\&_J(c?(T_j); \varphi'_b)))$  then we either have that  $d_a \in \underline{D'}$  or not. If  $d_a \in \underline{D'}$  then we have a contradiction, as  $d_a < d_b$ , and the hypothesis states that  $\exists d \in \underline{D'} \text{ such that } d < d'$ . Consider the case where  $d_a \notin \underline{D'}$ . This is a contradiction, as the hypothesis

states that  $d_a, d_b \in \underline{D}$ . If  $d' \neq d_b$  and  $d' \neq d_a$  then the proof proceeds as when  $d' \neq d_b$  in the first case.

The cases where  $[\Sigma] \Phi \in \{ \delta(\Sigma, \text{dmod}_{\{f_{d_b}\}}(\varphi_b^k)) \}$  and where  $[\Sigma] \Phi \in \{ \delta(\Sigma, \text{dmod}_{\{f_{d_a}\}}(\varphi_a^k)) \}$  proceed similarly to the first case. The case where  $[\Sigma] \Phi \in {}_n\Delta_{\underline{D}}^{\text{IG}}$  proceeds directly from the inductive hypothesis.  $\square$

Finally, we can then use these lemmas to prove that Local Update implies Global Compatibility:

**Theorem 5.14. Local Update Theorem**

For all  $k, \Sigma, P, \psi, \omega$ , if  $\mathcal{LU}(\psi, \omega)$  and we can form the partitions  $\Sigma = \Sigma_1, \Sigma_2$  and  $D = \underline{D} \uplus \overline{D}$  such that  $[\Sigma_1] \prod_{\underline{D}} \phi_d \in \Delta_{\underline{D}}^{\text{IG}}$  and  $[\Sigma_2] \prod_{\overline{D}} \phi_d \in \Delta_{\overline{D}}^{\text{IG}}$  and  $\overline{D} < \underline{D}$  then  $\mathcal{G}^k(\Sigma, \Phi, \psi, \omega)$ .

*Proof.* Assume the hypotheses. We then proceed by induction over  $k$  in  $\mathcal{G}^k(\Sigma, \Phi, \psi, \omega)$ . The base case is trivial. Consider the case where  $k < 0$ . We prove each of the conjuncts of  $\mathcal{G}^k(\Sigma, \Phi, \psi, \omega)$ , as defined in Figure 4.18, separately. In order to show  $\text{error} \notin \Phi$  we appeal to Lemma 5.9 that shows that there are no errors in either  $\prod_{\underline{D}} \phi_d$  or  $\prod_{\overline{D}} \phi_d$ . In order to show that  $\mathcal{G}^k(\Sigma, \Phi, \psi, \omega, \emptyset)$  we appeal to the fact that  $\emptyset$  is the unit in update composition, and we have that  $\mathcal{LU}(\psi, \omega, \emptyset)$  and hence can use the inductive hypothesis to show that  $\mathcal{G}^{k-1}(\Sigma, \Phi, \psi, \omega)$ . In order to show the first conjunct we perform a case analysis on whether an update is applicable or not.

Consider the case where an update is applicable. We must show that all reductions lead to safe program states, and the program is either live or empty. Consider a given reduction. This reduction must either be in the old program ( $d \in \underline{D}$ ) or in the new program ( $d \in \overline{D}$ ). If the former ( $[\Sigma_1] \prod_{\underline{D}} \phi_d \rightarrow [\Sigma'_1] \prod_{\underline{D}} \phi'_d$ ) then by Lemma 5.8 we know that  $[\Sigma'_1] \prod_{\underline{D}} \phi'_d \in {}_n\Delta_{\underline{D}}^{\text{IG}}$ . We then have a decomposition of  $[\Sigma'] \Phi'$  such that  $\Sigma' = \Sigma'_1, \Sigma_2$  and  $[\Sigma'_1] \prod_{\underline{D}} \phi'_d \in \Delta_{\underline{D}}^{\text{IG}}$  and  $[\Sigma_2] \prod_{\overline{D}} \phi_d \in \Delta_{\overline{D}}^{\text{IG}}$ , and hence can use the inductive hypothesis to show that  $\mathcal{G}^{k-1}(\Sigma', \Phi', \psi, \omega)$ . In order to show that the program is live or empty we appeal to Lemma 5.10, which denotes that the old program partition is either live or one of its processes are updatable. As there are no processes that are eligible for update, then the old program partition is live, as required.

Consider the case where an update is applicable. Let this process's role be  $d$ . By the update predicates we know that  $\phi_d \equiv G_1 \upharpoonright d$ . By Lemma 5.11 we know that  $G_2 \upharpoonright d_i \wr \emptyset \vdash \text{upd}(t_i, f_{d_i} \mapsto t'_{d_i}, \text{false}) : T$ . Let  $\underline{D} = \underline{D}' \uplus \{d\}$  and  $\overline{D}' = \overline{D} \uplus \{d\}$ . Then, by Lemmas 5.12 and 5.13 we can show that given the new partitions  $\underline{D}'$  and  $\overline{D}'$  that we have that  $[\Sigma_1] \prod_{\underline{D}'} \phi_d \in \Delta_{\underline{D}'}^{\text{IG}}$  and  $[\Sigma_2] \prod_{\overline{D}'} \phi'_d \in \Delta_{\overline{D}'}^{\text{IG}}$ , where for each  $d' \neq d$  we have that  $\phi'_{d'} = \phi_{d'}$ , and  $\phi'_d = \text{upd}(\phi_d, f_d \mapsto t'_{d_i}, \text{false})$ , and that  $\overline{D}' < \underline{D}'$ . Hence can use the inductive hypothesis to show that  $\mathcal{G}^{k-1}(\Sigma, \text{upd}(\Phi, f_{d_i} \mapsto t'_{d_i}, b), \psi \setminus f_{d_i} \mapsto t'_{d_i}, \emptyset)$ , as required.  $\square$

### 5.2.2 Summary

In order to be able to update threads separately we restrict the possible programs that be updated. In particular we consider only programs where each iteration of the event processing loop has a clear leader. The resulting system permits separate updates of individual threads, without any global synchronisation. In particular, this simplifies the checking of when the relevant update predicates are fulfilled: a separate monitor can be maintained for each thread, which only activates when a thread is at the top of its event processing loop, and in that case only needs to examine its channel queues and whether or not certain other threads have been updated.

### 5.3 Conclusions

We consider two possible approaches to providing update to blocking message passing programs: Global Typability and Local Update.

The main advantages of Global Typability are that it can handle most message passing programs that can be abstracted using Global Session Typing, that it can update to any program which can be abstracted using Global Session Typing, and that the update semantics is relatively straightforward to understand. The main disadvantages of Global Typability are that there is no guarantee that an update will ever be applied, and the cost of checking the update predicates.

The main advantages of Local Update are that threads can be updated separately, reducing the cost of checking update predicates (as no global information is required), and that assuming fair scheduling that updates will not be infinitely delayed. The disadvantage is that we restrict the form of updatable programs to those which have a unique thread that starts each iteration of the event processing loop. The degree to which this restriction will real out desired programs, in actual system development, is a topic for further research.



# Chapter 6

## Conclusions and Future Work

In this thesis we present techniques to prove communication safety and liveness in message passing programs. We re-prove key results from the session typing literature, using a novel inductive technique. We extend existing results of liveness and communication safety to non-blocking message passing systems.

We extend the above techniques to accommodate message passing programs that can be dynamically updated. Our key contribution are techniques which can be used to update multi-threaded programs without introducing deadlock. In particular we show how to, in certain situations, update threads separately, without requiring global synchronisation. We also present our intuition behind Update Liveness, the property that an update will not be delayed indefinitely.

The discovery of points at which update can be performed on more than one thread is widely accepted to be very difficult. We can utilise insights illuminated by our analysis to provide design patterns that are amenable to update. This has the potential to widely extend the ability to update multi-threaded programs.

In this chapter we provide a summary of how the above was achieved (Section 6.1) and the possibilities for future work (Section 6.2).

### 6.1 Summary

We began by considering examples that displayed type safety, type errors, liveness, and deadlock, for blocking message passing programs. From this we defined a general technique for ensuring safety and liveness, that amounts to model checking. We reproved existing results in Session Types work for blocking message passing programs using an intuitive inductive technique. In the course of this proof we identified the aspects that relied on the synchronisation behaviour of the blocking semantics. Then, when

we wanted to prove similar results for non-blocking message passing, we were able to straightforwardly define the synchronisation behaviour required.

We then considered type safety and liveness for Dynamic Software Update. We analysed examples that display type safety, type errors, liveness, and deadlock, for updatable message passing programs. We then extended our general technique for proving safety and liveness of static programs to encompass updatable programs. These updates could be introduced at any time and applied at an arbitrary point specified by a predicate over the code and the state. We then considered techniques for performing update to programs that can be abstracted using Global Session Types. We proved that updates using these techniques would not introduce communication errors or deadlock. We made use of the synchronisation behaviour inherent in blocking message passing to co-ordinate when updates occurred. In certain cases this permitted us to update mutually dependent threads separately.

## 6.2 Future Work

While Dynamic Software Update has been informally considered since the early 1980s, it has only begun to be strongly focused upon in the last 10 years. As such a young field of inquiry, there are a wide variety of paths currently open for future work.

At a high methodological level it is unclear how much complexity is acceptable for programmers using DSU frameworks. The semantics of when an update occurs, particularly in systems where individual functions or class definitions are transformed rather than the entire program, is complex. As such, it is difficult for a non-specialist to reason about how state transformation functions behave at runtime. When updates are made more transparent, however, to the degree that state transformation functions are no longer needed, updates that change function or class signatures must be ruled out ([Gregersen and Jorgensen, 2009](#)). Clearly, given studies that indicate that at least 10% of all updates require such flexibility, this is not an acceptable restriction. Further study is needed to determine the acceptable level of complexity for users of DSU. One possible line of research is to give DSU tools to real programmers, and to get them experiment with updates to real programs, and to write the relevant transformation functions. One could then compare the functions written by programmers to those written by DSU specialists, and gather feedback on what was straightforward, and what was complex.

One approach to decrease the complexity of programming using DSU is to perform whole program transformation, starting the update at specified update points in the code, and re-entering the new code at specified entry points. While this technique significantly reduces the conceptual burden, it relies on programs reaching specific update points. This restriction is more consequential in multi-threaded programs; in a naive implementation there is no guarantee that the threads will ever all be at an updatable point.

We could incorporate our analysis, which uses synchronisation behaviour, into work on whole program transformation in order to provide confidence that updates will not be delayed indefinitely.

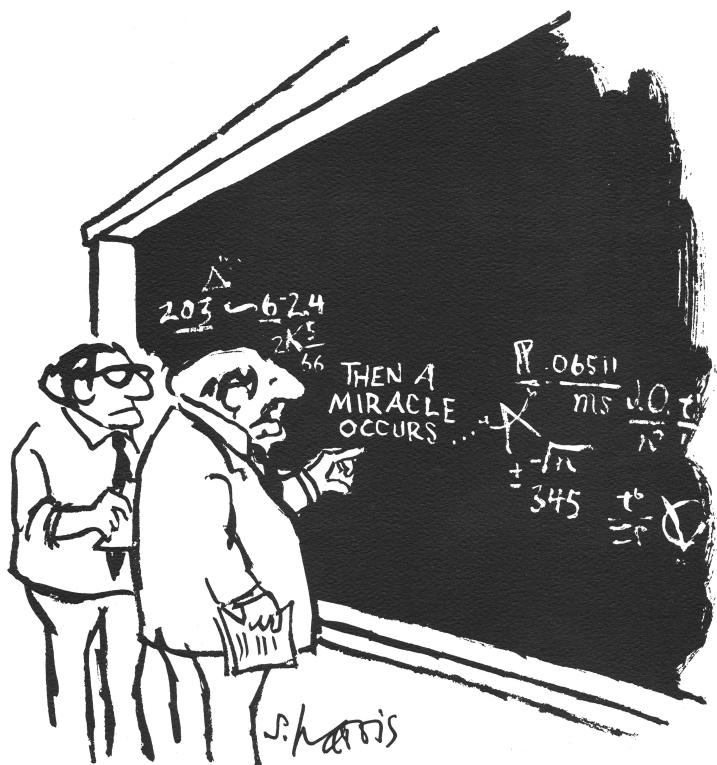
Mobile computation devices, such as smart phones, run many programs (known as *apps*) concurrently. Unfortunately, due to hardware constraints, it is currently not possible to keep all of these apps in memory at the same time. As a result the OS will often shut down an app when the user wishes to run another. In order to maintain the illusion that an app has been in the background, rather than shut down and restarted, the app designers can make use of design pattern that allows them to save existing state to backing store, and to automatically load in that state when the app is restarted. This is a form of DSU ([Hayden et al., 2011d](#)), in particular one that uses a form of whole program transformation. As the mobile apps must run on limited hardware they are typically smaller than many production systems. As such it may be possible to perform case studies on how amenable they are to DSU on a large number of such apps, in order to obtain more general results.

Various behaviours have been identified when performing Dynamic Software Update that, while they are not necessarily errors in traditional safety properties, cause confusing or intuitively incorrect behaviour ([Gregersen and Jorgensen, 2011](#)). Some of these behaviours we may wish to consider erroneous, such as when classes are removed by a dynamic update, but live objects that belong to that class remain. Some of these behaviours may be erroneous or not, depending on the situation. Transient inconsistency occurs when an updated application is temporally in an unreachable runtime state; in Local Update, in the period between when the leading thread updates and the following threads do, the system is in a transiently inconsistent state, but will eventually reach a consistent state. Further study is required to determine where and when these phenomena are actually erroneous, particularly by examining the behaviour of real updated programs. Gregersen and Jorgensen also describe several design patterns that they posit will ameliorate such phenomena. Study into the impact of using such patterns on programmers, and on the efficacy of these patterns, is also required.

Session Typing techniques can be used to analyse the call graph of object oriented programs. As such, we should be able to straightforwardly extend our techniques to updating object oriented programs. There exists a significant body of work on updating object oriented programs, in particular ([Subramanian, 2010](#)) which utilises interrupts to detect safe update points and the garbage collector to provide object transformation. These techniques provide update facilities with minimal overhead. One major drawback, however, is that, in the interest of maintaining type safety in a straightforward manner, Subramanian only permits update to methods that are not active. This leads to some updates being delayed indefinitely. We believe that we could combine our more complex safety analysis with their efficient implementation to provide an update system that permits update to active methods, but still has minimal overhead. In addition, we

would like to formalise our work on Update Liveness, and prove, for certain classes of Global Typability and Local Update, that updates are never delayed indefinitely.

# Appendices



"I THINK YOU SHOULD BE MORE EXPLICIT HERE IN STEP TWO."

Copyright © 2012 by Sidney Harris, ScienceCartoonsPlus.com



## Appendix A

# General Formulation

In this chapter we provide the proofs of communication safety and liveness for the general formulation of dynamic programs. The proofs for the general formulation of static programs are instantiations of these proofs.

### A.1 Environment Extension Lemma

**Lemma A.1.** *If  $\varphi \wr \Gamma \vdash t: T$  then  $\varphi \wr \Gamma, \Gamma' \vdash t: T$ .*

*Proof.* By simultaneous induction over  $t$  and  $\Gamma$ . Note that  $\text{dom}(\Gamma) \cap \text{dom}(\Gamma') = \emptyset$

NB: as we type  $v$  using an empty environment it cannot be another variable  $x'$ .

**Case:**  $t = n$

$$\overline{\epsilon \wr \Gamma \vdash n: \text{INT}}$$

We can trivially say that:

$$\overline{\epsilon \wr \Gamma, \Gamma' \vdash n: \text{INT}}$$

**Case:**  $t = b, t = (), t = r, t = l$

Similar to Case:  $t = n$ .

**Case:**  $t = x$

$$\overline{\epsilon \wr \Gamma \vdash x: \Gamma(x)}$$

As  $\Gamma(x)$  is defined we know that  $x \in \text{dom}(\Gamma)$ . Since composition of type environments assumes disjoint domains then we can say that:

$$\overline{\epsilon \wr (\Gamma, \Gamma') \vdash x: (\Gamma, \Gamma')(x)}$$

**Case:**  $t = X$

$$\frac{\Gamma(\underline{X}) = T_1 \xrightarrow{X} T_2}{\underline{X} \wr \Gamma \vdash \underline{X}: \Gamma(\underline{X})}$$

As  $\Gamma(\underline{X})$  is defined we know that  $\underline{X} \in \text{dom}(\Gamma)$ . Since composition of type environments assumes disjoint domains then we can say that:

$$\frac{(\Gamma, \Gamma')(\underline{X}) = T_1 \xrightarrow{X} T_2}{\underline{X} \wr (\Gamma, \Gamma') \vdash \underline{X}: (\Gamma, \Gamma')(\underline{X})}$$

**Case:**  $t = (\alpha(\tilde{v}), T)$

$$\frac{\epsilon \wr \Gamma \vdash v_i: T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \Gamma \vdash l: T}$$

As this type rule places no restrictions on  $\Gamma$  then we can trivially show that:

$$\frac{\epsilon \wr \Gamma, \Gamma' \vdash v_i: T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \Gamma, \Gamma' \vdash l: T}$$

**Case:**  $l$

$$\frac{\epsilon \wr \Gamma \vdash v_i: T_i}{\Gamma \vdash \alpha(\tilde{v}): \alpha(\tilde{T})}$$

By induction we know that  $\epsilon \wr \Gamma, \Gamma' \vdash v_i: T_i$ . Hence we can say that:

$$\frac{\epsilon \wr \Gamma, \Gamma' \vdash v_i: T_i}{\Gamma, \Gamma' \vdash_{\psi} \alpha(\tilde{v}): \alpha(\tilde{T})}$$

**Case:**  $t = \text{rec } X(x : T).t$

$$\frac{\varphi \wr \Gamma, x: T_1, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash t: T_2 \quad x \notin \text{fv}(\Gamma), \text{fv}(T_1), \text{fv}(T_2), \text{fv}(\varphi) \quad \underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)}{\epsilon \wr \Gamma \vdash \text{rec } \underline{X}(x: T_1).t: T_1 \xrightarrow{\mu X \cdot \varphi} T_2}$$

By induction we know that  $\varphi \vdash \Gamma, x: T_1, \underline{X}: T_1 \xrightarrow{\varphi} T_2, \Gamma' \vdash v_i: T_i$ , where we  $\alpha$ -convert  $x$  as necessary to ensure that they do not appear in  $\Gamma'$ . Hence we can say that:

$$\frac{\varphi \vdash \Gamma, \Gamma', x: T_1, \underline{X}: T_1 \xrightarrow{\varphi} T_2 \vdash t: T_2 \\ x \notin \text{fv}(\Gamma, \Gamma'), \text{fv}(T_1), \text{fv}(T_2), \text{fv}(\varphi) \quad \underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)}{\epsilon \vdash \Gamma, \Gamma' \vdash \text{rec } \underline{X}(x: T_1). t: T_1 \xrightarrow{\mu \underline{X}, \varphi} T_2}$$

**Case:**  $t = t_1 t_2$

$$\frac{\varphi_1 \vdash \Gamma \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \vdash \Gamma \vdash t_2: T_2}{(\varphi_1; \varphi_2; \varphi_3) \vdash \Gamma \vdash t_1 t_2: T_1}$$

By induction we know that:

$$\varphi_1 \vdash \Gamma, \Gamma' \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \vdash \Gamma, \Gamma' \vdash t_2: T_2$$

Hence we can say that:

$$\frac{\varphi_1 \vdash \Gamma, \Gamma' \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \vdash \Gamma, \Gamma' \vdash t_2: T_2}{(\varphi_1; \varphi_2; \varphi_3) \vdash \Gamma, \Gamma' \vdash t_1 t_2: T_1}$$

**Case:**  $t = \text{casel in } \{\widetilde{T \mapsto t}\}$

$$\frac{\varphi_i \vdash \Gamma \vdash t_i: T \quad \epsilon \vdash \Gamma \vdash v_i: T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\widetilde{T}), T_i); \varphi_i \vdash \Gamma \vdash \text{casel in } \{\widetilde{T \mapsto t}\}: T}$$

By induction we know that

$$\varphi_i \vdash \Gamma, \Gamma' \vdash t_i: T$$

Hence we can say that:

$$\frac{\varphi_i \vdash \Gamma, \Gamma' \vdash t_i: T \quad \epsilon \vdash \Gamma, \Gamma' \vdash v_i: T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\widetilde{T}), T_i); \varphi_i \vdash \Gamma, \Gamma' \vdash \text{casel in } \{\widetilde{T \mapsto t}\}: T}$$

**Case:**  $t = \text{error}$

This case can never be typed and hence can be disregarded as the hypothesis assumes that  $t$  is typed.

**Case:**  $t = f(t)$

$$\frac{\varphi \vdash \Gamma \vdash t: T}{f(\varphi) \vdash \Gamma \vdash f(t): T}$$

By induction we know that

$$\varphi \wr \Gamma, \Gamma' \vdash t : T$$

Hence we can say that:

$$\frac{\varphi \wr \Gamma, \Gamma' \vdash t : T}{f(\varphi) \wr \Gamma, \Gamma' \vdash f(t) : T}$$

**Case:**  $t = \text{dmod}_{\tilde{f}}(t)$

$$\frac{\varphi \wr \Gamma \vdash t : T}{\text{dmod}_{\tilde{f}}(\varphi) \wr \Gamma \vdash \text{dmod}_{\tilde{f}}(t) : T}$$

By induction we know that:

$$\varphi \wr \Gamma, \Gamma' \vdash t : T$$

Hence we can say that:

$$\frac{\varphi \wr \Gamma, \Gamma' \vdash t : T}{\text{dmod}_{\tilde{f}}(\varphi) \wr \Gamma, \Gamma' \vdash \text{dmod}_{\tilde{f}}(t) : T}$$

□

## A.2 Substitution Lemma

**Lemma A.2.** *If*

$$\varphi \vdash \Gamma, x: T' \vdash t: T \quad \epsilon \vdash \emptyset \vdash v: T'' \quad T'[v/x] = T''$$

*then*

$$\varphi[v/x] \vdash \Gamma[v/x] \vdash t[v/x]: T[v/x]$$

*Proof.* By induction over  $t$

**Case:**  $t = n$

$$\overline{\epsilon \vdash \Gamma, x: T' \vdash n: \text{INT}}$$

we know that:

$$\epsilon[v/x] = \epsilon \quad n[v/x] = n \quad \text{INT}[v/x] = \text{INT} = T''$$

hence we can say:

$$\overline{\epsilon \vdash \Gamma[v/x] \vdash n: \text{INT}}$$

**Case:**  $t = b, t = (), t = l, t = X$

Similar to Case:  $t = n$ .

**Case:**  $t = x'$

$$\overline{\epsilon \vdash \Gamma' \vdash x': \Gamma'(x')}$$

where  $\Gamma' = \Gamma, x: T'$ .

If  $x' = x$ :

$$\epsilon[v/x] = \epsilon \quad x'[v/x] = v$$

We want to prove that  $\epsilon \vdash \Gamma[v/x] \vdash v: T'[v/x]$ . We do this by a case split over the structure of  $T'$

Case  $T' = \text{INT}, \text{BOOL}, \text{UNIT}$

$$T'[v/x] = T' = T''$$

$\epsilon \vdash \Gamma[v/x] \vdash v: T'$  using the  $\text{TINT}$ ,  $\text{TBOOL}$ , or  $\text{TUNIT}$  rule respectively

Case  $T' = \text{Res } r$

By the hypothesis we know that  $\epsilon \wr \emptyset \vdash v : \text{Res } r$ .

As the type environment is empty of this type judgement is empty, the only way we could derive this fact is by using the TRES rule, and hence  $v = r = r$ .

The only  $T'$  such that  $\text{Res } r[v/x] = \text{Res } v = T''$  are  $\text{Res } v$  and  $\text{Res } x$ .

As here  $T = T'$  we then know that  $T$  must be either  $\text{Res } v$  or  $\text{Res } x$ .

We can easily show that:

$$\text{Res } v[v/x] = \text{Res } v \quad \text{Res } x[v/x] = \text{Res } v$$

We then can use the TRES rule to show that:

$$\overline{\epsilon \wr \Gamma[v/x] \vdash v : \text{Res } v}$$

Case  $T' = T_1 \xrightarrow{\varphi} T_2$

By the hypothesis we know that:

$$\epsilon \wr \emptyset \vdash v : (T_1 \xrightarrow{\varphi} T_2)[v/x]$$

By the Environment Extension Lemma we can say:

$$\epsilon \wr \Gamma[v/x] \vdash v : (T_1 \xrightarrow{\varphi} T_2)[v/x]$$

as required.

If  $x' \neq x$ :

$$\epsilon[v/x] = \epsilon \quad x'[v/x] = x' \quad \Gamma(x')[v/x] = \Gamma[v/x](x')$$

$$\overline{\epsilon \wr \Gamma[v/x] \vdash x' : \Gamma[v/x](x')}$$

**Case:**  $t = r$

$$\overline{\epsilon \wr \Gamma, x : T' \vdash r : \text{Res } r}$$

We know that  $r \neq x$  and hence that

$$\epsilon[v/x] = \epsilon \quad r[v/x] = r \quad \text{Res } r[v/x] = \text{Res } r$$

We can then simply use the  $\text{TR}_{\text{ES}}$  rule:

$$\frac{}{\epsilon \wr \Gamma[v/x] \vdash r : \text{Res } r}$$

**Case:**  $t = \text{rec } X(x' : T_1).t'$

$$\frac{\varphi \wr \Gamma, x : T', x' : T_1, \underline{X} : T_1 \xrightarrow{X} T_2 \vdash t' : T_2 \\ x' \notin \text{fv}(\Gamma, x : T'), \text{fv}(T_1), \text{fv}(T_2), \text{fv}(\varphi) \quad \underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)}{\epsilon \wr \Gamma, x : T' \vdash \text{rec } \underline{X}(x' : T_1).t' : T_1 \xrightarrow{\mu \underline{X} \cdot \varphi} T_2}$$

By induction, where  $T'[v/x] = T''$ :

$$\varphi[v/x] \wr (\Gamma, x' : T_1, \underline{X} : T_1 \xrightarrow{x'} T_2)[v/x] \vdash t'[v/x] : T_2[v/x]$$

As  $x'$  is not free in  $t, (\Gamma, x : T')$  we can  $\alpha$  convert it to ensure it does not clash with  $x$ , hence we have that:

$$\frac{\varphi[v/x] \wr \Gamma[v/x], x'' : T_1[v/x], \underline{X} : T_1[v/x] \xrightarrow{X} T_2[v/x] \vdash t'[v/x] : T_2[v/x] \\ x'' \notin \text{fv}(\Gamma[v/x]), \text{fv}(T_1[v/x]), \text{fv}(T_2[v/x]), \text{fv}(\varphi[v/x]) \quad \underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)}{\epsilon \wr \Gamma[v/x] \vdash \text{rec } \underline{X}(x'' : T_1[v/x]).t'[v/x] : T_1[v/x] \xrightarrow{\mu \underline{X} \cdot \varphi[v/x]} T_2[v/x]}$$

and hence:

$$\epsilon \wr \Gamma[v/x] \vdash (\text{rec } \underline{X}(x'' : T_1).t')[v/x] : (T_1 \xrightarrow{\mu x' \cdot \varphi} T_2)[v/x]$$

**Case:**  $t = t_1 t_2$

$$\frac{\varphi_1 \wr \Gamma, x : T' \vdash t_1 : T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \wr \Gamma, x : T' \vdash t_2 : T_2}{(\varphi_1; \varphi_2; \varphi_3) \wr \Gamma, x : T' \vdash t_1 t_2 : T_1}$$

By induction:

$$\begin{aligned} \varphi_1[v/x] \wr \Gamma[v/x] \vdash t_1[v/x] : T_1 &\xrightarrow{\varphi_3} T_1[v/x] \\ \varphi_2[v/x] \wr \Gamma[v/x] \vdash t_2[v/x] : T_2[v/x] \end{aligned}$$

Therefore we can say:

$$\frac{\varphi_1[v/x] \wr \Gamma[v/x] \vdash t_1[v/x] : T_2[v/x] \xrightarrow{\varphi_3[v/x]} T_1[v/x] \quad \varphi_2[v/x] \wr \Gamma[v/x] \vdash t_2[v/x] : T_2[v/x]}{(\varphi_1[v/x]; \varphi_2[v/x]; \varphi_3[v/x]) \wr \Gamma[v/x] \vdash t_1[v/x] t_2[v/x] : T_1[v/x]}$$

and hence:

$$(\varphi_1; \varphi_2; \varphi_3)[v/x] \wr \Gamma[v/x] \vdash (t_1 t_2)[v/x]: T_1[v/x]$$

**Case:**  $t = (\alpha(\tilde{v}), T)$

$$\frac{\epsilon \wr \Gamma, x: T' \vdash v_i: T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \Gamma, x: T' \vdash (\alpha(\tilde{v}), T): T}$$

By induction we have that  $\epsilon \wr \Gamma[v/x] \vdash v_i[v/x]: T_i[v/x]$ . Hence we can show, using the TAcc rule, that:

$$(\alpha(\tilde{T}), T)[v/x] \wr \Gamma[v/x] \vdash (\alpha(\tilde{T}), T)[v/x]: T[v/x]$$

**Case:**  $t = \text{if } t_1 \text{ then } t_2 \text{ else } t_3$

$$\frac{\varphi_1 \wr \Gamma, x: T' \vdash t_1: \text{BOOL} \quad \varphi_2 \wr \Gamma, x: T' \vdash t_2: T \quad \varphi_3 \wr \Gamma, x: T' \vdash t_3: T}{\varphi_1; (\varphi_2 \oplus \varphi_3) \wr \Gamma, x: T' \vdash \text{if } t_1 \text{ then } t_2 \text{ else } t_3: T}$$

By induction:

$$\begin{aligned} \varphi_1[v/x] \wr \Gamma[v/x] &\vdash t_1[v/x]: \text{BOOL}[v/x] \\ \varphi_2[v/x] \wr \Gamma[v/x] &\vdash t_2[v/x]: T[v/x] \quad \varphi_3[v/x] \wr \Gamma[v/x] \vdash t_3[v/x]: T[v/x] \end{aligned}$$

We can the use the TIF rule:

$$\frac{\varphi_1[v/x] \wr \Gamma[v/x] \vdash t_1[v/x]: \text{BOOL} \quad \varphi_2[v/x] \wr \Gamma[v/x] \vdash t_2[v/x]: T[v/x] \quad \varphi_3[v/x] \wr \Gamma[v/x] \vdash t_3[x/x]: T[v/x]}{\varphi_1[v/x]; (\varphi_2[v/x] \oplus \varphi_3[v/x]) \wr \Gamma[v/x] \vdash \text{if } t_1[v/x] \text{ then } t_2[v/x] \text{ else } t_3[x/x]: T[v/x]}$$

and hence:

$$\varphi_1; (\varphi_2 \oplus \varphi_3)[v/x] \wr \Gamma[v/x] \vdash (\text{if } t_1 \text{ then } t_2 \text{ else } t_3)[v/x]: T[v/x]$$

**Case:**  $t = \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T}_i \mapsto \tilde{t}_i\}$

$$\frac{\varphi_i \wr \Gamma, x: T' \vdash t_i: T \quad \epsilon \wr \Gamma, x: T' \vdash v_i: T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \Gamma, x: T' \vdash \text{case } l \text{ in } \{\tilde{T} \mapsto \tilde{t}\}: T}$$

By induction we know that:

$$\varphi_i[v/x] \wr \Gamma[v/x] \vdash t_i[v/x]: T[v/x]$$

We can then say, using the T<sub>CASE</sub> rule:

$$\frac{\varphi_i[v/x] \wr \Gamma[v/x] \vdash t_i[v/x]: T[v/x] \quad T[v/x] \mapsto \tilde{t}[v/x] = T_1[v/x] \mapsto t_1[v/x] \dots T_n[v/x] \mapsto t_n[v/x]}{\&(\alpha(\tilde{v})[v/x], T_i[v/x]); \varphi_i[v/x] \wr \Gamma[v/x] \vdash \text{case } \alpha(\tilde{v})[v/x] \text{ in } \{\tilde{T}[v/x] \mapsto \tilde{t}[v/x]\}: T[v/x]} \text{ (:T<sub>CASE</sub>)}$$

which implies:

$$(\&(\alpha(\tilde{v}), T_i); \varphi_i)[v/x] \wr \Gamma[v/x] \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto \tilde{t}\}[v/x]: T[v/x]$$

**Case:**  $t = \text{error}$

This case can never be typed and hence can be disregarded as the hypothesis assumes that  $t$  is typed.

**Case:**  $t = \text{dmod}_{F_1}(t')$

$$\frac{\varphi \wr \Gamma, x: T' \vdash t': \text{UNIT}}{\text{dmod}_{\gamma_1}(\varphi) \wr \Gamma, x: T' \vdash \text{dmod}_{\gamma_1}(t'): \text{UNIT}}$$

By induction we know that:

$$\varphi[v/x] \wr \Gamma[v/x] \vdash t'[v/x]: \text{UNIT}$$

Hence we can say that:

$$\frac{\varphi[v/x] \wr \Gamma[v/x] \vdash t'[v/x]: \text{UNIT}}{\text{dmod}_{\gamma_1}(\varphi[v/x]) \wr \Gamma[v/x] \vdash \text{dmod}_{\gamma_1}(t'[v/x]): \text{UNIT}}$$

which implies:

$$\text{dmod}_{\gamma_1}(\varphi)[v/x] \wr \Gamma[v/x] \vdash \text{dmod}_{F_1}(t')[v/x]: \text{UNIT}$$

**Case:**  $t = f(t)$

Similar to case  $t = \text{dmod}_{F_1}(t)$ . □

### A.3 Recursive Variable Substitution Lemma

**Lemma A.3.** *If*

$$\varphi \wr \Gamma, \underline{X} : T_1 \xrightarrow{X} T_2 \vdash t : T \quad \epsilon \wr \emptyset \vdash v : T_1 \xrightarrow{\varphi'} T_2 \quad \underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)$$

*then:*

$$\varphi[\varphi'/\underline{X}] \wr \Gamma \vdash t[v/\underline{X}] : T[\varphi'/\underline{X}]$$

*Proof.* By induction over  $t$

**Case:**  $t = n$

$$\overline{\epsilon \wr \Gamma, \underline{X} : T_1 \xrightarrow{X} T_2 \vdash n : \text{INT}}$$

we know that:

$$\epsilon[\varphi'/\underline{X}] = \epsilon \quad n[v/\underline{X}] = n \quad \text{INT}[\varphi'/\underline{X}] = \text{INT}$$

hence we can say:

$$\overline{\epsilon \wr \Gamma \vdash n : \text{INT}}$$

**Case:**  $t = b, t = (), t = l, t = x, t = r$

Similar to Case:  $t = n$ .

**Case:**  $t = X'$

$$\frac{\Gamma'(\underline{X}') = T_1 \xrightarrow{X'} T_2}{\underline{X}' \wr \Gamma' \vdash \underline{X}' : \Gamma'(\underline{X}')}}$$

where  $\Gamma' = \Gamma, \underline{X} : T_1 \xrightarrow{X} T_2$ .

If  $\underline{X}' = \underline{X}$ :

$$\underline{X}[\varphi'/\underline{X}] = \varphi' \quad \underline{X}'[v/\underline{X}] = v \quad T_1 \xrightarrow{X} T_2[\varphi'/\underline{X}] = T_1[\varphi'/\underline{X}] \xrightarrow{\varphi'} T_2[\varphi'/\underline{X}] = T_1 \xrightarrow{\varphi'} T_2$$

as  $\underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)$ .

By the hypothesis we know that:

$$\epsilon \wr \emptyset \vdash v : T_1 \xrightarrow{\varphi'} T_2$$

and hence, using the Environment Extension Lemma, that:

$$\underline{X}[\epsilon/\underline{X}] \wr \Gamma \vdash v : T_1 \xrightarrow{\underline{X}} T_2[\varphi'/\underline{X}]$$

If  $\underline{X}' \neq \underline{X}$ :

$$\underline{X}'[\varphi'/\underline{X}] = \underline{X}' \quad \underline{X}'[v/\underline{X}] = \underline{X}' \quad T_1 \xrightarrow{\underline{X}'} T_2[\varphi/\underline{X}] = T_1 \xrightarrow{\underline{X}'} T_2$$

and hence that:

$$\underline{X}'[\varphi'/\underline{X}] \wr \Gamma \vdash \underline{X}' : T_1 \xrightarrow{\underline{X}'} T_2[\varphi'/\underline{X}]$$

**Case:**  $t = \mathbf{rec} \underline{X}'(x' : T_1).t'$

$$\frac{\varphi \wr \Gamma, \underline{X} : T_1 \xrightarrow{\underline{X}} T_2, x' : T_3, \underline{X}' : T_3 \xrightarrow{\underline{X}'} T_4 \vdash t' : T_4 \\ x' \notin \mathbf{fv}(\Gamma, \underline{X} : T_1 \xrightarrow{\underline{X}} T_2), \mathbf{fv}(T_3), \mathbf{fv}(T_4), \mathbf{fv}(\varphi) \quad \underline{X}' \notin \mathbf{fv}(T_1), \mathbf{fv}(T_2)}{\epsilon \wr \Gamma, \underline{X} : T_1 \xrightarrow{\underline{X}} T_2 \vdash \mathbf{rec} \underline{X}'(x' : T_3).t' : T_3 \xrightarrow{\mu \underline{X}' \cdot \varphi} T_4}$$

By induction, where  $T'[v/x] = T''$ :

$$\varphi[\varphi'/\underline{X}] \wr (\Gamma, x' : T_3, \underline{X} : T_3 \xrightarrow{\underline{X}'} T_4)[v/x] \vdash t'[v/\underline{X}] : T_4[\varphi'/\underline{X}]$$

As  $\underline{X}'$  is not free in  $t, (\Gamma, x : T')$  we can  $\alpha$  convert it to ensure it does not clash with  $\underline{X}$ , hence we have that:

$$\frac{\varphi[\varphi'/\underline{X}] \wr \Gamma, x' : T_1[v/\underline{X}], \underline{X} : T_1[v/\underline{X}] \xrightarrow{\underline{X}} T_2[\varphi'/\underline{X}] \vdash t'[v/\underline{X}] : T_2[\varphi'/\underline{X}] \\ x' \notin \mathbf{fv}(\Gamma), \mathbf{fv}(T_1[v/\underline{X}]), \mathbf{fv}(T_2[\varphi'/\underline{X}]), \mathbf{fv}(\varphi[\varphi'/\underline{X}]) \quad \underline{X} \notin \mathbf{fv}(T_1), \mathbf{fv}(T_2)}{\epsilon \wr \Gamma \vdash \mathbf{rec} \underline{X}(x' : T_1[v/\underline{X}]).t'[v/\underline{X}] : T_1[v/\underline{X}] \xrightarrow{\mu \underline{X} \cdot \varphi[\varphi'/\underline{X}]} T_2[\varphi'/\underline{X}]}$$

and hence:

$$\epsilon[\varphi'/\underline{X}] \wr \Gamma \vdash (\mathbf{rec} \underline{X}(x'' : T_1).t')[v/\underline{X}] : (T_1 \xrightarrow{\mu x' \cdot \varphi} T_2)[\varphi'/\underline{X}]$$

**Case:**  $t = t_1 t_2$

$$\frac{\varphi_1 \wr \Gamma, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \wr \Gamma, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash t_2: T_2}{(\varphi_1; \varphi_2; \varphi_3) \wr \Gamma, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash t_1 t_2: T_1}$$

By induction:

$$\begin{aligned} \varphi_1[\varphi'/\underline{X}] \wr \Gamma \vdash t_1[v/\underline{X}]: T_1 \xrightarrow{\varphi_3} T_1[\varphi'/\underline{X}] \\ \varphi_2[\varphi'/\underline{X}] \wr \Gamma \vdash t_2[v/\underline{X}]: T_2[\varphi'/\underline{X}] \end{aligned}$$

Therefore we can say:

$$\frac{\varphi_1[\varphi'/\underline{X}] \wr \Gamma \vdash t_1[v/\underline{X}]: T_2[\varphi'/\underline{X}] \xrightarrow{\varphi_3[\varphi'/\underline{X}]} T_1[\varphi'/\underline{X}] \quad \varphi_2[\varphi'/\underline{X}] \wr \Gamma \vdash t_2[v/\underline{X}]: T_2[\varphi'/\underline{X}]}{(\varphi_1[\varphi'/\underline{X}]; \varphi_2[\varphi'/\underline{X}]; \varphi_3[\varphi'/\underline{X}]) \wr \Gamma \vdash t_1[v/\underline{X}] t_2[v/\underline{X}]: T_1[\varphi'/\underline{X}]}$$

and hence:

$$(\varphi_1; \varphi_2; \varphi_3)[\varphi'/\underline{X}] \wr \Gamma \vdash (t_1 t_2)[v/\underline{X}]: T_1[\varphi'/\underline{X}]$$

**Case:**  $t = (\alpha(\tilde{v}), T)$

$$\frac{\epsilon \wr \Gamma, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash v_i: T_i \quad \mathsf{Res} r \notin T}{(\alpha(\tilde{T}), T) \wr \Gamma, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash (\alpha(\tilde{v}), T): T}$$

By induction we have that  $\epsilon \wr \Gamma \vdash v_i[v/\underline{X}]: T_i[\varphi'/\underline{X}]$ . Hence we can show, using the TAcc rule,

$$(\alpha(\tilde{T}), T)[\varphi'/\underline{X}] \wr \Gamma \vdash (\alpha(\tilde{T}), T)[v/\underline{X}]: T[\varphi'/\underline{X}]$$

**Case:**  $t = \mathbf{if} \ t_1 \ \mathbf{then} \ t_2 \ \mathbf{else} \ t_3$

$$\frac{\varphi_1 \wr \Gamma, x: T' \vdash t_1: \mathbf{BOOL} \quad \varphi_2 \wr \Gamma, x: T' \vdash t_2: T \quad \varphi_3 \wr \Gamma, x: T' \vdash t_3: T}{\varphi_1; (\varphi_2 \oplus \varphi_3) \wr \Gamma, x: T' \vdash \mathbf{if} \ t_1 \ \mathbf{then} \ t_2 \ \mathbf{else} \ t_3: T}$$

By induction:

$$\begin{aligned} \varphi_1[\varphi'/\underline{X}] \wr \Gamma \vdash t_1[v/\underline{X}]: \mathbf{BOOL}[\varphi'/\underline{X}] \\ \varphi_2[\varphi'/\underline{X}] \wr \Gamma \vdash t_2[v/\underline{X}]: T[\varphi'/\underline{X}] \quad \varphi_3[\varphi'/\underline{X}] \wr \Gamma \vdash t_3[v/\underline{X}]: T[\varphi'/\underline{X}] \end{aligned}$$

We can use the  $\text{TI}_F$  rule:

$$\frac{\varphi_1[\varphi'/\underline{X}] \wr \Gamma \vdash t_1[v/\underline{X}]: \text{BOOL} \quad \varphi_2[v/x] \wr \Gamma \vdash t_2[v/\underline{X}]: T[\varphi'/\underline{X}] \quad \varphi_3[v/x] \wr \Gamma \vdash t_3[x/\underline{X}]: T[\varphi'/\underline{X}]}{\varphi_1[\varphi'/\underline{X}]; (\varphi_2[v/x] \oplus \varphi_3[v/x]) \wr \Gamma \vdash \text{if } t_1[v/\underline{X}] \text{ then } t_2[v/\underline{X}] \text{ else } t_3[x/\underline{X}]: T[\varphi'/\underline{X}]}$$

and hence:

$$\varphi_1; (\varphi_2 \oplus \varphi_3)[\varphi'/\underline{X}] \wr \Gamma \vdash (\text{if } t_1 \text{ then } t_2 \text{ else } t_3)[v/\underline{X}]: T[\varphi'/\underline{X}]$$

**Case:**  $t = \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T}_i \mapsto t_i\}$

$$\frac{\varphi_i \wr \Gamma, \underline{X}: T_1 \xrightarrow{\underline{X}} T_2 \vdash t_i: T \quad \epsilon \wr \Gamma, \underline{X}: T_1 \xrightarrow{\underline{X}} T_2 \vdash v_i: T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \Gamma, \underline{X}: T_1 \xrightarrow{\underline{X}} T_2 \vdash \text{case } \alpha(\tilde{T}) \text{ in } \{\tilde{T} \mapsto t\}: T}$$

By induction we know that:

$$\varphi_i[\varphi'/\underline{X}] \wr \Gamma \vdash t_i[v/\underline{X}]: T[\varphi'/\underline{X}]$$

We can then say, using the  $\text{TCASE}$  rule:

$$\frac{\varphi_i[\varphi'/\underline{X}] \wr \Gamma \vdash t_i[v/\underline{X}]: T[\varphi'/\underline{X}] \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{v})[\varphi'/\underline{X}], T_i[\varphi'/\underline{X}]); \varphi_i[\varphi'/\underline{X}] \wr \Gamma \vdash \text{case } \alpha(\tilde{v})[v/\underline{X}] \text{ in } \{\tilde{T} \mapsto t\}: T[\varphi'/\underline{X}]} \text{ (:TCASE)}$$

which implies that:

$$(\&(\alpha(\tilde{v}), T_i); \varphi_i)[\varphi'/\underline{X}] \wr \Gamma[v/x] \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\}[v/\underline{X}]: T[\varphi'/\underline{X}]$$

**Case:**  $t = \text{error}$

This case can never be typed and hence can be disregarded as the hypothesis assumes that  $t$  is typed.

**Case:**  $t = \text{dmod}_{F_1}(t')$

$$\frac{\varphi \wr \Gamma, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash t': \text{UNIT}}{\text{dmod}_{\gamma_1}(\varphi) \wr \Gamma, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash \text{dmod}_{\gamma_1}(t'): \text{UNIT}}$$

By induction we know that:

$$\varphi[\varphi'/\underline{X}] \wr \Gamma \vdash t'[v/\underline{X}]: \text{UNIT}[\varphi'/\underline{X}]$$

Hence we can say that:

$$\frac{\varphi[\varphi'/\underline{X}] \wr \Gamma \vdash t'[v/\underline{X}]: \text{UNIT}[\varphi'/\underline{X}]}{\text{dmod}_{\gamma_1}(\varphi[\varphi'/\underline{X}]) \wr \Gamma \vdash \text{dmod}_{\gamma_1}(t'[v/\underline{X}]): \text{UNIT}[\varphi'/\underline{X}]}$$

which implies:

$$\text{dmod}_{\gamma_1}(\varphi)[\varphi'/x] \wr \Gamma \vdash \text{dmod}_{F_1}(t')[v/\underline{X}]: \text{UNIT}$$

**Case:**  $t = f(t)$

Similar to case  $t = \text{dmod}_{F_1}(t)$ . □

## A.4 Update Equivalence Lemma

**Lemma A.4.**

$$\varphi \equiv \varphi' \Rightarrow \mathsf{upd}(\varphi, u, b) \equiv \mathsf{upd}(\varphi', u, b)$$

*Proof.* By induction over  $\equiv$

**Case**  $\varphi \equiv \epsilon; \varphi$

$$\begin{aligned} \mathsf{upd}(\epsilon; \varphi, u, b) &= \epsilon; \mathsf{upd}(\varphi, u, b) \\ &\equiv \mathsf{upd}(\varphi, u, b) \end{aligned}$$

**Case**  $\mu X. \varphi \equiv \varphi[\mu X. \varphi / X]$

$$\begin{aligned} \mathsf{upd}(\mu \underline{X}. \varphi, u, b) &= \mu \underline{X}. \mathsf{upd}(\varphi, u, b) \\ &\equiv \mathsf{upd}(\varphi, u, b)[\mu \underline{X}. \mathsf{upd}(\varphi, u, b) / \underline{X}] \end{aligned}$$

**Case**  $\varphi = \varphi_1; \varphi_2$

By induction we know that:

$$\varphi'_i \equiv \varphi_i \Rightarrow \mathsf{upd}(\varphi_i, u, b) \equiv \mathsf{upd}(\varphi'_i, u, b)$$

We know that:

$$\mathsf{upd}(\varphi_1; \varphi_2, u, b) = \mathsf{upd}(\varphi_1, u, b); \mathsf{upd}(\varphi_2, u, b)$$

Hence we have that:

$$\mathsf{upd}(\varphi_1; \varphi_2, u, b) \equiv \mathsf{upd}(\varphi'_1; \varphi'_2, u, b)$$

**Case**  $\varphi = \varphi_1 \oplus \varphi_2, \varphi_1 \& \varphi_2, \mu X. \varphi_1, f(\varphi_1), \mathsf{dmod}_F(\varphi_1)$

Similar to case  $\varphi = \varphi_1; \varphi_2$ .

**Case**  $\varphi = f(\varphi_1)$

**Case**  $b = \text{true}$

By the definition of `upd` we know that:

$$\text{upd}(f(\varphi_1), u, \text{false}) = f(\varphi')$$

where:

$$f \mapsto t' \notin u \quad \varphi' \wr \emptyset \vdash t':$$

Hence we have that:

$$\text{upd}(f(\varphi_1), u, \text{true}) = \text{upd}(f(\varphi'_1), u, \text{true}) = f(\varphi')$$

**Case**  $b = \text{false}$

By induction we know that:

$$\varphi'_1 \equiv \varphi_1 \Rightarrow \text{upd}(\varphi_1, u, b) \equiv \text{upd}(\varphi'_1, u, b)$$

By the definition of `upd` we know that:

$$\text{upd}(f(\varphi_1), u, \text{false}) = f(\text{upd}(\varphi_1, u, b))$$

Hence we have that:

$$\text{upd}(f(\varphi_1), u, \text{false}) \equiv \text{upd}(f(\varphi'_1), u, \text{false})$$

**Case**  $\varphi = \text{dmod}_F(\varphi_1)$

**Case**  $\gamma \cap \text{dom}(u) = \emptyset$

By the definition of `upd` we know that:

$$\text{upd}(\text{dmod}_\gamma(\varphi_1), u, b) = \text{dmod}_\gamma(\text{upd}(\varphi_1, u, \text{true}))$$

By induction we know that:

$$\varphi'_1 \equiv \varphi_1 \Rightarrow \text{upd}(\varphi_1, u, b) \equiv \text{upd}(\varphi'_1, u, b)$$

Hence we have that:

$$\text{upd}(\text{dmod}_\gamma(\varphi_1), u, \text{true}) \equiv \text{upd}(\text{dmod}_\gamma(\varphi'_1), u, \text{true})$$

**Case**  $\gamma \cap \text{dom}(u) \neq \emptyset$

By the definition of `upd` we know that:

$$\text{upd}(\text{dmod}_\gamma(\varphi_1), u, b) = \text{dmod}_\gamma(\text{upd}(\varphi_1, u, \text{false}))$$

By induction we know that:

$$\varphi'_1 \equiv \varphi_1 \Rightarrow \text{upd}(\varphi_1, u, b) \equiv \text{upd}(\varphi'_1, u, b)$$

Hence we have that:

$$\text{upd}(\text{dmod}_\gamma(\varphi_1), u, \text{false}) \equiv \text{upd}(\text{dmod}_\gamma(\varphi'_1), u, \text{false})$$

□

## A.5 Expression Update Consistency Lemma

**Lemma A.5.** *If:*

$$\varphi \wr \Gamma \vdash t : T$$

*then:*

$$\mathsf{upd}(\varphi, u, b) \wr \Gamma \vdash \mathsf{upd}(t, u, b) : T' \quad T \equiv_{\mathit{eff}} T'$$

$$\frac{T_1 = T_2 \quad T_1 \equiv_{\mathit{eff}} T'_1 \quad T_2 \equiv_{\mathit{eff}} T'_2}{T_1 \xrightarrow{\varphi} T_2 \equiv_{\mathit{eff}} T'_1 \xrightarrow{\varphi'} T'_2}$$

Proof: by induction over  $t$ .

**Case:**  $t = n$

$$\overline{\epsilon \wr \Gamma \vdash n : \mathbf{INT}}$$

By the definition of  $\mathsf{upd}$ :

$$\mathsf{upd}(n, u, b) \stackrel{\mathsf{def}}{=} n \quad \mathsf{upd}(\epsilon, u, b) \stackrel{\mathsf{def}}{=} \epsilon$$

Hence we can say that:

$$\mathsf{upd}(\epsilon, u, b) \wr \Gamma \vdash \mathsf{upd}(n, u, b) : \mathbf{INT}$$

**Case:**  $t = b$ ,  $t = ()$ ,  $t = r$ ,  $t = x$ ,  $t = (\alpha(\tilde{v}), T)$ ,  $t = X$

Similar to case  $t = n$ .

**Case:**  $t = \mathsf{rec} X(x : T).t_1$

$$\frac{\varphi \wr \Gamma, x : T_1, \underline{X} : T_1 \xrightarrow{X} T_2 \vdash t_1 : T_2 \quad x \notin \mathsf{fv}(\Gamma), \mathsf{fv}(T_1), \mathsf{fv}(T_2), \mathsf{fv}(\varphi) \quad \underline{X} \notin \mathsf{fv}(T_1), \mathsf{fv}(T_2)}{\epsilon \wr \Gamma \vdash \mathsf{rec} \underline{X}(x : T_1).t_1 : T_1 \xrightarrow{\mu X \cdot \varphi} T_2}$$

By the definition of  $\mathsf{upd}$ :

$$\mathsf{upd}(\mathsf{rec} \underline{X}(x : T).t, u, b) \stackrel{\mathsf{def}}{=} \mathsf{rec} \underline{X}(x : T).\mathsf{upd}(t, u, b) \quad \mathsf{upd}(\epsilon, u, b) \stackrel{\mathsf{def}}{=} \epsilon$$

By induction we know that:

$$\mathsf{upd}(\varphi, u, b) \wr \Gamma, x : T_1, \underline{X} : T_1 \xrightarrow{\varphi} T_2 \vdash \mathsf{upd}(t_1, u, b) : T'_2 \quad T_2 \equiv_{\mathit{eff}} T'_2$$

Hence we can say that:

$$\frac{\begin{array}{c} \mathsf{upd}(\varphi, u, b) \wr \Gamma, x: T_1, \underline{X}: T_1 \xrightarrow{X} T'_2 \vdash \mathsf{upd}(t_1, u, b): T'_2 \\ x \notin \mathsf{fv}(\Gamma), \mathsf{fv}(T_1), \mathsf{fv}(T'_2), \mathsf{fv}(\mathsf{upd}(\varphi, u, b)) \quad \underline{X} \notin \mathsf{fv}(T_1), \mathsf{fv}(T_2) \end{array}}{\epsilon \wr \Gamma \vdash \mathbf{rec} \underline{X}(x: T_1). \mathsf{upd}(t_1, u, b): T_1 \xrightarrow{\mu \underline{X}. \mathsf{upd}(\varphi, u, b)} T'_2}$$

which implies that:

$$\mathsf{upd}(\epsilon, u, b) \wr \Gamma \vdash \mathsf{upd}(\mathbf{rec} \underline{X}(x: T). t, u, b): T_1 \xrightarrow{\mu \underline{X}. \mathsf{upd}(\varphi, u, b)} T'_2$$

where  $T_1 \xrightarrow{\varphi} T_2 \equiv_{eff} T_1 \xrightarrow{\mu \underline{X}. \mathsf{upd}(\varphi, u, b)} T'_2$  as  $T_1 = T_1$  and  $T_2 \equiv_{eff} T'_2$  and

$$\frac{T_1 \equiv_{eff} T'_1 \quad T_2 \equiv_{eff} T'_2}{T_1 \xrightarrow{\varphi} T_2 \equiv_{eff} T'_1 \xrightarrow{\varphi'} T'_2}$$

**Case:**  $t = t_1 t_2$

$$\frac{\varphi_1 \wr \Gamma \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \wr \Gamma \vdash t_2: T_2}{(\varphi_1; \varphi_2; \varphi_3) \wr \Gamma \vdash t_1 t_2: T_1}$$

By the definition of  $\mathsf{upd}$ :

$$\begin{aligned} \mathsf{upd}(t_1 t_2, u, b) &\stackrel{\text{def}}{=} \mathsf{upd}(t_1, u, b) \mathsf{upd}(t_2, u, b) \\ \mathsf{upd}(\varphi_1; \varphi_2; \varphi_3, u, b) &\stackrel{\text{def}}{=} \mathsf{upd}(\varphi_1, u, b); \mathsf{upd}(\varphi_2, u, b); \mathsf{upd}(\varphi_3, u, b) \end{aligned}$$

By induction we know that:

$$\mathsf{upd}(\varphi_1, u, b) \wr \Gamma \vdash \mathsf{upd}(t_1, u, b): T'_2 \xrightarrow{\mathsf{upd}(\varphi_3, u, b)} T'_1$$

$$\mathsf{upd}(\varphi_2, u, b) \wr \Gamma \vdash \mathsf{upd}(t_2, u, b): T'_2$$

where  $T_i \equiv_{eff} T'_i$ . Hence we can say that:

$$\frac{\varphi'_1 \wr \Gamma \vdash t'_1: T'_2 \xrightarrow{\varphi'_3} T'_1 \quad \varphi'_2 \wr \Gamma \vdash t'_2: T'_2}{(\varphi'_1; \varphi'_2; \varphi'_3) \wr \Gamma \vdash t'_1 t'_2: T'_1}$$

where  $t'_i = \mathsf{upd}(t_i, u, b), \varphi'_i = \mathsf{upd}(\varphi_i, u, b), T_i \equiv_{eff} T'_i$ .

**Case:**  $t = \mathbf{if} \ t \ \mathbf{then} \ t \ \mathbf{else} \ t, \mathbf{dmod}_{F_1}(t)$

Similar to case  $t = t_1 t_2$

**Case:**  $t = \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t'\}$

$$\frac{\varphi_i \wr \Gamma \vdash t_i : T \quad \epsilon \wr \Gamma \vdash v_i : T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \Gamma \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t'\} : T}$$

By the definition of **upd**:

$$\text{upd}(\text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t'\}, u, b) \stackrel{\text{def}}{=} \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto \text{upd}(t_i, u, b)\}$$

$$\text{upd}(\&\varphi_i, u, b) \stackrel{\text{def}}{=} \&\text{upd}(\varphi_i, u, b)$$

By induction we have that:

$$\text{upd}(\varphi_i, u, b) \wr \Gamma \vdash \text{upd}(t_i, u, b) : T'$$

where  $T \equiv_{\text{eff}} T'$ . Hence we can say that:

$$\frac{\text{upd}(\varphi_i, u, b) \wr \Gamma \vdash t_i : T' \quad \epsilon \wr \Gamma \vdash v_i : T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \text{upd}(\varphi_i, u, b) \wr \Gamma \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto \text{upd}(t_i, u, b)\} : T'}$$

where  $T \equiv_{\text{eff}} T'$ , which implies

$$\text{upd}(\&(\alpha(\tilde{T}), T_n); \varphi_n, u, b) \wr \Gamma \vdash \text{upd}(\text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t'\}, u, b) : T'$$

**Case:**  $t = \text{error}$

This case can never be typed and hence can be disregarded as the hypothesis assumes that  $t$  is typed.

**Case:**  $t = f(t_1)$

$$\frac{\varphi \wr \Gamma \vdash t : \text{UNIT}}{f(\varphi) \wr \Gamma \vdash f(t) : \text{UNIT}}$$

We can perform a case split on the relevant conditions in the definition of **upd**:

**Case**  $b \wedge f \in \text{dom}(u)$

$$\begin{aligned} \text{upd}(f(t_1), u, \text{true}) &\stackrel{\text{def}}{=} f(t') \quad (f \mapsto t') \in u \\ \text{upd}(f(\varphi_1), u, \text{true}) &\stackrel{\text{def}}{=} f(\varphi') \quad (f \mapsto t') \in u \wedge \varphi' \wr \emptyset \vdash t' : \gamma \end{aligned}$$

As we assume updates are well formed then we know that we can make the judgement:

$$\varphi' \wr \emptyset \vdash t' : \text{UNIT}$$

We can then use the  $\text{TR}_{\text{REG}}$  typing rule to show that:

$$\frac{\varphi' \wr \Gamma \vdash t' : \text{UNIT}}{f(\varphi') \wr \Gamma \vdash f(t') : \text{UNIT}}$$

and hence that:

$$\text{upd}(f(\varphi_1), u, \text{true}) \wr \Gamma \vdash \text{upd}(f(t_1), u, \text{true}) : \text{UNIT}$$

**Case**  $\neg(b \wedge f \in \text{dom}(u))$

$$\begin{aligned} \text{upd}(f(t_1), u, b) &\stackrel{\text{def}}{=} f(\text{upd}(t_1, u, b)) \quad (\neg b \vee (f \mapsto \gamma) \notin u) \\ \text{upd}(f(\varphi_1), u, b) &\stackrel{\text{def}}{=} f(\text{upd}(\varphi_1, u, b)) \quad (\neg b \vee (f \mapsto \gamma) \notin u) \end{aligned}$$

By induction we know that

$$\text{upd}(\varphi_1, u, b) \wr \Gamma \vdash \text{upd}(t_1, u, b) : T'$$

where  $\text{UNIT} \equiv_{\text{eff}} T'$ . The only possible such  $T'$  is  $\text{UNIT}$ . Hence we can straightforwardly say that:

$$\frac{\text{upd}(\varphi, u, b) \wr \Gamma \vdash \text{upd}(t', u, b) : \text{UNIT}}{f(\text{upd}(\varphi, u, b)) \wr \Gamma \vdash f(\text{upd}(t', u, b)) : \text{UNIT}}$$

**Case:**  $t = \text{dmod}_F(t_1)$

$$\frac{\varphi_1 \wr \Gamma \vdash t_1 : \text{UNIT}}{\text{dmod}_\gamma(\varphi_1) \wr \Gamma \vdash \text{dmod}_\gamma(t_1) : \text{UNIT}} \quad \frac{\varphi \wr \Gamma \vdash t : T \quad \varphi \equiv \varphi'}{\varphi \wr \Gamma \vdash t : \varphi'}$$

We can perform a case split on the relevant conditions in the definition of  $\text{upd}$ :

**Case**  $\tilde{f} \cap \text{dom}(u) = \emptyset$

By the definition of  $\text{upd}$ :

$$\begin{aligned} \text{upd}(\text{dmod}_F(t_1), u, b) &\stackrel{\text{def}}{=} \text{dmod}_F(\text{upd}(t_1, u, \text{true})) & \tilde{f} \cap \text{dom}(u) = \emptyset \\ \text{upd}(\text{dmod}_\gamma(\varphi), u, b) &\stackrel{\text{def}}{=} \text{dmod}_\gamma(\text{upd}(\varphi, u, \text{true})) & \tilde{f} \cap \text{dom}(u) = \emptyset \end{aligned}$$

By induction we know that

$$\mathsf{upd}(\varphi_1, u, \mathsf{true}) \wr \Gamma \vdash \mathsf{upd}(t_1, u, \mathsf{true}) : T'$$

where  $\mathsf{UNIT} \equiv_{eff} T'$ . The only possible such  $T'$  is  $\mathsf{UNIT}$ . Hence we can straightforwardly say that:

$$\frac{\mathsf{upd}(\varphi_1, u, \mathsf{true}) \wr \Gamma \vdash \mathsf{upd}(t_1, u, \mathsf{true}) : \mathsf{UNIT}}{\mathsf{dmod}_\gamma(\mathsf{upd}(\varphi_1, u, \mathsf{true})) \wr \Gamma \vdash \mathsf{dmod}_\gamma(\mathsf{upd}(t_1, u, \mathsf{true})) : \mathsf{UNIT}}$$

**Case**  $\tilde{f} \cap \mathsf{dom}(u) \neq \emptyset$

By the definition of  $\mathsf{upd}$ :

$$\begin{aligned} \mathsf{upd}(\mathsf{dmod}_F(t_1), u, b) &\stackrel{\text{def}}{=} \mathsf{dmod}_F(\mathsf{upd}(t_1, u, \mathsf{false})) \quad \tilde{f} \cap \mathsf{dom}(u) \neq \emptyset \\ \mathsf{upd}(\mathsf{dmod}_\gamma(\varphi), u, b) &\stackrel{\text{def}}{=} \mathsf{dmod}_\gamma(\mathsf{upd}(\varphi, u, \mathsf{false})) \quad \tilde{f} \cap \mathsf{dom}(u) \neq \emptyset \end{aligned}$$

By induction we know that

$$\mathsf{upd}(\varphi_1, u, \mathsf{false}) \wr \Gamma \vdash \mathsf{upd}(t_1, u, \mathsf{false}) : T'$$

where  $\mathsf{UNIT} \equiv_{eff} T'$ . The only possible such  $T'$  is  $\mathsf{UNIT}$ . Hence we can straightforwardly say that:

$$\frac{\mathsf{upd}(\varphi_1, u, \mathsf{false}) \wr \Gamma \vdash \mathsf{upd}(t_1, u, \mathsf{false}) : \mathsf{UNIT}}{\mathsf{dmod}_\gamma(\mathsf{upd}(\varphi_1, u, \mathsf{false})) \wr \Gamma \vdash \mathsf{dmod}_\gamma(\mathsf{upd}(t_1, u, \mathsf{false})) : \mathsf{UNIT}}$$

## A.6 Thread Update Consistency Lemma

**Lemma A.6.** *If:*

$$\vdash P: \Phi \quad \forall f \mapsto t. \varphi \wr \emptyset \vdash t: \text{UNIT}$$

*then:*

$$\vdash \text{upd}(P, u, \text{false}) : \text{upd}(\Phi, u, \text{false})$$

Proof: by induction over  $P$

**Case:**  $P = t$

$$\frac{\varphi \wr \emptyset \vdash t: T}{\vdash t: \varphi}$$

By the Expression Update Consistency Lemma:

$$\text{upd}(\varphi, u, b) \wr \emptyset \vdash \text{upd}(t, u, b) : T' \quad T \equiv_{eff} T'$$

Hence:

$$\frac{\text{upd}(\varphi, u, b) \wr \emptyset \vdash \text{upd}(t, u, b) : T'}{\vdash \text{upd}(t, u, b) : \text{upd}(\varphi, u, b)}$$

**Case:**  $P = P_1 \parallel P_2$

$$\frac{\vdash P_1: \Phi_1 \quad \vdash P_2: \Phi_2}{\vdash P_1 \parallel P_2: \Phi_1 \parallel \Phi_2}$$

By induction we know that:

$$\vdash \text{upd}(P_1, u, \text{false}) : \text{upd}(\Phi_1, u_i, \text{false}) \quad \vdash \text{upd}(P_2, u, \text{false}) : \text{upd}(\Phi_2, u_i, \text{false})$$

Hence we can say that:

$$\vdash \text{upd}(P_1 \parallel P_2, u, \text{false}) : \text{upd}(\Phi_1 \parallel \Phi_2, u, \text{false})$$

## A.7 Expression Safety Lemma

**Lemma A.7.**

$$\varphi \wr \Gamma \vdash t : T \Rightarrow \#\text{error} \in t$$

*Proof.* By induction over  $t$

**Case**  $t = n, P = b, P = (), P = r, P = l, P = x, P = X, P = (\alpha(\tilde{v}), T)$

We trivially know that none of these include `error`.

**Case**  $t = \text{rec } X(x : T).t_1$

$$\frac{\varphi_1 \wr \Gamma, x : T_1, \underline{X} : T_1 \xrightarrow{X} T_2 \vdash t_1 : T_2 \quad x \notin \text{fv}(\Gamma), \text{fv}(T_1), \text{fv}(T_2), \text{fv}(\varphi_1) \quad \underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)}{\epsilon \wr \Gamma \vdash \text{rec } \underline{X}(x : T_1).t_1 : T_1 \xrightarrow{\mu X \cdot \varphi_1} T_2}$$

We know that  $\varphi_1 \wr \Gamma, x : T_1, \underline{X} : T_1 \xrightarrow{X} T_2 \vdash t_1 : T_2$ . We can then use the inductive hypothesis to show that  $\#\text{error} \in t_1$ . We know straightforwardly that  $\#\text{error} \in x, T_1, \underline{X}$ . Hence we can say that  $\#\text{error} \in \text{rec } \underline{X}(x : T).t_1$ .

**Case:**  $t = t_1 t_2$

$$\frac{\varphi_1 \wr \Gamma \vdash t_1 : T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \wr \Gamma \vdash t_2 : T_2}{(\varphi_1; \varphi_2; \varphi_3) \wr \Gamma \vdash t_1 t_2 : T_1}$$

By induction we know that  $\#\text{error} \in t_1, t_2$ . Hence we have that  $\#\text{error} \in t_1 t_2$

**Case:**  $t = (\alpha(\tilde{v}), T)$

$$\frac{\epsilon \wr \Gamma \vdash v_i : T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \Gamma \vdash (\alpha(\tilde{v}), T) : T}$$

We know straightforwardly, and by induction, that  $\#\text{error} \in v_i$ . Hence we have that  $\#\text{error} \in (\alpha(\tilde{v}), T)$

**Case:**  $t = \text{if } t_1 \text{ then } t_2 \text{ else } t_3$

$$\frac{\varphi_1 \wr \Gamma \vdash t_1: \text{BOOL} \quad \varphi_2 \wr \Gamma \vdash t_2: T \quad \varphi_3 \wr \Gamma \vdash t_3: T}{\varphi_1; (\varphi_2 \oplus \varphi_3) \wr \Gamma \vdash \text{if } t_1 \text{ then } t_2 \text{ else } t_3: T}$$

By induction we know that  $\# \text{error} \in t_1, t_2, t_3$ . Hence we have that  $\# \text{error} \in \text{if } t_1 \text{ then } t_2 \text{ else } t_3$ .

**Case:**  $t = \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T_i \mapsto t_i}\}$

$$\frac{\varphi_i \wr \Gamma \vdash t_i: T \quad \epsilon \wr \Gamma \vdash v_i: T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \Gamma \vdash \text{case } \alpha(\tilde{T}) \text{ in } \{\tilde{T \mapsto t}\}: T}$$

We know straightforwardly that  $\# \text{error} \in v_i, T_i$ . By induction we know that  $\# \text{error} \in t_i$ . Hence we have that  $\# \text{error} \in \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T_i \mapsto t_i}\}$ .

**Case:**  $t = \text{let } x = \text{new int}'$

Similar to case  $t = \text{let } x = \alpha(\tilde{v}) \text{ int}'$ .

**Case:**  $t = \text{error}$

This case can never be typed and hence can be disregarded as the hypothesis assumes that  $t$  is typed.

**Case:**  $t = \text{dmod}_{F_1}(t_1)$

$$\frac{\varphi_1 \wr \Gamma \vdash t_1: \text{UNIT}}{\text{dmod}_{\gamma_1}(\varphi_1) \wr \Gamma \vdash \text{dmod}_{\gamma_1}(t_1): \text{UNIT}}$$

We know straightforwardly that  $\# \text{error} \in \gamma$ . By induction we know that  $\# \text{error} \in t_1$ . Hence we have that  $\# \text{error} \in \text{dmod}_{F_1}(t_1)$ .

**Case:**  $t = f(t_1)$

We know straightforwardly that  $\# \text{error} \in f$ . By induction we know that  $\# \text{error} \in t_1$ . Hence we have that  $\# \text{error} \in f(t_1)$ .

□

## A.8 Safety Theorem

**Theorem A.8.**

$$\vdash P: \Phi \Rightarrow \#\text{error} \in P$$

*Proof.* By induction over  $P$

**Case  $P = t$**

$$\frac{\varphi \wr \emptyset \vdash t: T}{\vdash t: \varphi}$$

As  $\varphi \wr \emptyset \vdash t: T$ , by Lemma A.7 we know that  $\#\text{error} \in t$

**Case  $P = P_1 \parallel P_2$**

$$\frac{\vdash P_1: \Phi_1 \quad \vdash P_2: \Phi_2}{\vdash P_1 \parallel P_2: \Phi_1 \parallel \Phi_2}$$

By induction we know that  $\#\text{error} \in P_1, P_2$ . Hence we have that  $\#\text{error} \in P_1 \parallel P_2$

□

## A.9 Value Effect Lemma

**Lemma A.9.**

$$t = v \Rightarrow \epsilon \wr \Gamma \vdash t : T$$

*Proof.* Proof: by case split over  $v$

**Case**  $v = n$

$$\overline{\epsilon \wr \Gamma \vdash n : \text{INT}}$$

**Case**  $v = b$

$$\overline{\epsilon \wr \Gamma \vdash b : \text{BOOL}}$$

**Case**  $v = ()$

$$\overline{\epsilon \wr \Gamma \vdash () : \text{UNIT}}$$

**Case**  $v = r$

$$\overline{\epsilon \wr \Gamma \vdash r : \text{Res } r}$$

**Case**  $v = x$

$$\overline{\epsilon \wr \Gamma \vdash x : \Gamma(x)}$$

**Case**  $v = l$

$$\overline{\epsilon \wr \Gamma \vdash l : \text{Lab } l}$$

**Case**  $v = \text{rec } X(x : T).t_1$

$$\frac{\varphi \wr \Gamma, x : T_1, \underline{X} : T_1 \xrightarrow{X} T_2 \vdash t : T_2 \quad x \notin \text{fv}(\Gamma), \text{fv}(T_1), \text{fv}(T_2), \text{fv}(\varphi) \quad \underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)}{\epsilon \wr \Gamma \vdash \text{rec } \underline{X}(x : T_1).t : T_1 \xrightarrow{\mu X \cdot \varphi} T_2}$$

□

## A.10 General Compatibility Well Definedness Lemma

**Lemma A.10.**  $\mathcal{G}(\Sigma, \Phi, \psi, \omega)$  is well defined fixed point satisfying  $\bigwedge_{k \in \mathcal{N}} \mathcal{G}^k(\Sigma, \Phi, \psi, \omega)$ .

$\mathcal{G}(\Sigma, \Phi, \psi, \omega) = \bigwedge_{k \in \mathcal{N}} \mathcal{G}^k(\Sigma, \Phi, \psi, \omega)$ :

$$\mathcal{G}^0(\Sigma, \Phi, \psi, \omega) = \mathbf{true}$$

$$\mathcal{G}^k(\Sigma, \Phi, \psi, \omega) = \left[ \begin{array}{l} \nexists p \in \mathbf{dom}(\psi). p(\Sigma, \Phi) \wedge \\ (\forall \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \Rightarrow \mathcal{G}^{k-1}(\Sigma', \Phi', \psi, \omega)) \\ \wedge (\exists \Sigma', \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \vee \Phi = \Pi_I \epsilon) \\ \vee \psi = p_1 \mapsto u_1, \dots, p_n \mapsto u_n \wedge i \text{ smallest in 1...n.} \\ p_i(\Sigma, \Phi) \wedge \mathcal{G}^{k-1}(\Sigma, \mathbf{upd}(\Phi, u_i, \mathbf{false}), \psi \setminus p_i \mapsto u_i, \omega) \\ \wedge \nexists \mathbf{error} \in \Phi \wedge \mathcal{G}^{k-1}(\Sigma, \Phi, (\psi, \omega), \emptyset) \end{array} \right]$$

where  $k > 0$

Proof: We prove (anti) monotonicity by induction over  $k$ : if  $\mathcal{G}^{k+1}(\Sigma, \Phi, \psi, \omega)$  then  $\mathcal{G}^k(\Sigma, \Phi, \psi, \omega)$ .

**Case**  $k = 0$

By the definition  $\mathcal{G}^0(\Sigma, \Phi, \psi, \psi') = \mathbf{true}$ , and hence we trivially have that given  $\mathcal{G}^1(\Sigma, \Phi, \psi, \omega)$  then  $\mathcal{G}^0(\Sigma, \Phi, \psi, \omega)$ .

**Case**  $k = n > 0$

By the hypothesis we know that:

$$\left[ \begin{array}{l} \nexists p \in \mathbf{dom}(\psi). p(\Sigma, \Phi) \wedge \\ (\forall \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \Rightarrow \mathcal{G}^n(\Sigma', \Phi', \psi, \omega)) \\ \wedge (\exists \Sigma', \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \vee \Phi = \Pi_I \epsilon) \\ \vee \psi = p_1 \mapsto u_1, \dots, p_n \mapsto u_n \wedge i \text{ smallest in 1...n.} \\ p_i(\Sigma, \Phi) \wedge \mathcal{G}^n(\Sigma, \mathbf{upd}(\Phi, u_i, \mathbf{false}), \psi \setminus p_i \mapsto u_i, \omega) \\ \wedge \nexists \mathbf{error} \in \Phi \wedge \mathcal{G}^n(\Sigma, \Phi, (\psi, \omega), \emptyset) \end{array} \right]$$

holds.

By induction we know that if  $k = n - 1$  that  $\mathcal{G}^n(\Sigma, \Phi, \psi, \omega)$  then  $\mathcal{G}^{n-1}(\Sigma, \Phi, \psi, \omega)$ . Hence we can show that:

$$\left[ \begin{array}{l} \nexists p \in \text{dom}(\psi). p(\Sigma, \Phi) \wedge \\ (\forall \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \Rightarrow \mathcal{G}^{n-1}(\Sigma', \Phi', \psi, \omega)) \\ \wedge (\exists \Sigma', \Phi'. [\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi' \vee \Phi = \Pi_I \epsilon) \\ \vee \psi = p_1 \mapsto u_1, \dots, p_n \mapsto u_n \wedge i \text{ smallest in } 1\dots n. \\ p_i(\Sigma, \Phi) \wedge \mathcal{G}^{n-1}(\Sigma, \text{upd}(\Phi, u_i, \text{false}), \psi \setminus p_i \mapsto u_i, \omega) \\ \wedge \nexists \text{error} \in \Phi \wedge \mathcal{G}^{n-1}(\Sigma, \Phi, (\psi, \omega), \emptyset) \end{array} \right]$$

which, as  $n > 0$ , is the definition of  $\mathcal{G}^n(\Sigma, \Phi, \psi, \omega)$ .

## A.11 Expression Subject Reduction Lemma

Lemma 3.1

If:

$$\varphi \vdash \emptyset \vdash t: T \quad [\Sigma] \varphi \rightarrow [\Sigma'] \varphi' \Rightarrow \text{#error} \in \varphi' \quad [\sigma] t \xrightarrow{\gamma} F [\sigma'] t'$$

Then

$$\varphi' \vdash \emptyset \vdash t': T \quad [\Lambda(\sigma)] \varphi \xrightarrow{\gamma} F [\Lambda(\sigma')] \varphi' \quad \gamma' \vdash \gamma$$

*Proof.* By induction over  $[\sigma] t \rightarrow [\sigma'] t'$ .

**Case RAppOne**

$$\frac{\varphi_1 \vdash \emptyset \vdash t_1: T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \vdash \emptyset \vdash t_2: T_2}{(\varphi_1; \varphi_2; \varphi_3) \vdash \emptyset \vdash t_1 t_2: T_1}$$

By induction we know that if:

$$[\Lambda(\sigma)] \varphi \xrightarrow{\gamma'} F [\Lambda(\sigma')] \varphi' \quad \text{error} \notin \varphi'_1 \quad [\sigma] t_1 \xrightarrow{\gamma} F [\sigma'] t'_1$$

then:

$$\varphi'_1 \vdash \emptyset \vdash t'_1: T \quad \gamma' \vdash \gamma$$

Hence we can show that:

$$\frac{[\sigma] t_1 \xrightarrow{\gamma} F [\sigma'] t'_1}{[\sigma] t_1 t_2 \xrightarrow{\gamma} F [\sigma'] t'_1 t_2}$$

**Case RAppTwo**

Similar to case RAPPONE

**Case RAppThree**

$$\frac{\overline{[\sigma] \text{rec } \underline{X}(x: T).t v \xrightarrow{\tau} [\sigma] t[\text{rec } \underline{X}(x: T).t / \underline{X}][v/x]}}{\frac{\epsilon \vdash \emptyset \vdash \text{rec } \underline{X}(x: T).t: T_2 \xrightarrow{\mu \underline{X} \cdot \varphi_3} T_1 \quad \epsilon \vdash \emptyset \vdash v: T_2}{(\epsilon; \epsilon; \mu \underline{X} \cdot \varphi_3) \vdash \emptyset \vdash \text{rec } \underline{X}(x: T).t v: T_1}} \\ \frac{\varphi \vdash \emptyset, x: T_1, \underline{X}: T_1 \xrightarrow{X} T_2 \vdash t: T_2 \quad x \notin \text{fv}(\emptyset), \text{fv}(T_1), \text{fv}(T_2), \text{fv}(\varphi) \quad \underline{X} \notin \text{fv}(T_1), \text{fv}(T_2)}{\epsilon \vdash \emptyset \vdash \text{rec } \underline{X}(x: T_1).t: T_1 \xrightarrow{\mu \underline{X} \cdot \varphi} T_2}$$

Hence we know that:

$$\mu \underline{X}.\varphi_3 \wr \emptyset \vdash \mathbf{rec} \underline{X}(x : T).t v : T_1$$

By Lemma A.2 and Lemma A.3 we know that:

$$\varphi_3[\mu \underline{X}.\varphi_3 / \underline{X}][v/x] \wr \emptyset[v/x] \vdash t[\mathbf{rec} \underline{X}(x : T).t / \underline{X}][v/x] : T_2[v/x][\mu \underline{X}.\varphi_3 / \underline{X}]$$

as  $x \notin \text{fv}(\Gamma), \text{fv}(T_1), \text{fv}(T_2), \text{fv}(\varphi_3)$  and  $\underline{X} \notin \text{fv}(T_2)$  we can say that:

$$\varphi_3[\mu \underline{X}.\varphi_3 / \underline{X}] \wr \emptyset \vdash t[\mathbf{rec} \underline{X}(x : T).t / \underline{X}][v/x] : T_2$$

As

$$\varphi_3[\mu \underline{X}.\varphi_3 / \underline{X}] \equiv \varphi_3$$

we can use RTREC to show that:

$$\overline{[\Sigma] \mu \underline{X}.\varphi_3 \xrightarrow{\tau} [\Sigma] \varphi_3[\mu \underline{X}.\varphi_3 / \underline{X}]}$$

where  $\gamma = \tau$ .

### Case RAccOne

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T}{[\sigma](\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma']\sigma(\alpha(\tilde{v}))} \quad \frac{\epsilon \wr \emptyset \vdash v_i : T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \emptyset \vdash (\alpha(\tilde{v}), T) : T}$$

By our assumptions we know that  $\sigma(\alpha(\tilde{v})) = v$ . By Lemma A.9 we know that:

$$t = v \Rightarrow \epsilon \wr \Gamma \vdash t : T$$

We can show that  $[\Lambda(\sigma)]((\alpha(\tilde{v}), T), T) \rightarrow [\Lambda(\sigma')]\epsilon$  using the RTAccSUCC rule:

$$\frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) = T}{[\Sigma](\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Sigma']\epsilon}$$

According to TAcc we know that:

$$\epsilon \wr \Gamma \vdash v_i : T_i$$

and hence we can show that:

$$\frac{\epsilon \wr \emptyset \vdash v_i : T_i}{\alpha(\tilde{T}) \vdash \alpha(\tilde{v})}$$

### Case RAccTwo

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})): T' \neq T}{[\sigma] (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \text{error}} \quad \frac{\epsilon \wr \emptyset \vdash v_i: T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{T}), T) \wr \emptyset \vdash (\alpha(\tilde{v}), T): T}$$

Using the premises of RAccTwo we can show that:

$$\frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) \neq T}{[\Sigma] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \text{error}}$$

By the premises we know that:

$$[\Lambda(\sigma)] (\alpha(\tilde{T}), T) \rightarrow [\Lambda(\sigma')] \varphi'' \Rightarrow \text{error} \in \varphi''$$

This is a contradiction. Hence this reduction cannot occur, and we can disregard it.

### Case RIfOne

$$\frac{\overline{[\sigma] \text{if true then } t_2 \text{ else } t_3 \xrightarrow{\tau} [\sigma] t_2} \quad \varphi_1 \wr \emptyset \vdash \text{true: BOOL} \quad \varphi_2 \wr \emptyset \vdash t_2: T \quad \varphi_3 \wr \emptyset \vdash t_3: T}{\varphi_1; (\varphi_2 \oplus \varphi_3) \wr \emptyset \vdash \text{if true then } t_2 \text{ else } t_3: T}$$

By Lemma A.9 we know that  $\varphi_1 = \epsilon$ , and hence that:

$$\varphi_2 \oplus \varphi_3 \wr \Gamma \vdash \text{if true then } t_2 \text{ else } t_3: T$$

By TINT we know that:

$$\varphi_2 \wr \emptyset \vdash t_2: T$$

We can show that  $[\Lambda(\sigma)] \varphi_2 \oplus \varphi_3 \xrightarrow{\gamma'} [\Lambda(\sigma)] \varphi_2$  using RTINTCHOICE:

$$\frac{i \in 1, 2}{[\Sigma] \varphi_1 \oplus \varphi_2 \xrightarrow{\tau} [\Sigma] \varphi_i}$$

### Case RIfTwo

Similar to case RIfOne.

### Case RIfThree

$$\frac{\varphi_1 \wr \emptyset \vdash t_1: \text{BOOL} \quad \varphi_2 \wr \emptyset \vdash t_2: T \quad \varphi_3 \wr \emptyset \vdash t_3: T}{\varphi_1; (\varphi_2 \oplus \varphi_3) \wr \emptyset \vdash \text{if } t_1 \text{ then } t_2 \text{ else } t_3: T}$$

By induction we know that if:

$$[\Lambda(\sigma)] \varphi \xrightarrow{\gamma'}_F [\Lambda(\sigma')] \varphi' \quad \text{error} \notin \varphi' \quad [\sigma] t_1 \xrightarrow{\gamma}_F [\sigma'] t'_1$$

then:

$$\varphi'_1 \wr \emptyset \vdash t'_1: T \quad \gamma' \vdash \gamma$$

Hence we can show that:

$$\frac{[\sigma'] t_1 \xrightarrow{\gamma}_F [\sigma] t'_1}{[\sigma] \text{if } t_1 \text{ then } t_2 \text{ else } t_3 \xrightarrow{\gamma}_F [\sigma'] \text{if } t'_1 \text{ then } t_2 \text{ else } t_3}$$

### Case RCaseOne

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})): T_i \quad T_i \in \tilde{T}}{[\sigma] \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} \xrightarrow{\alpha(\tilde{v})} [\sigma'] t_i} \quad \frac{\varphi_i \wr \emptyset \vdash t_i: T \quad \epsilon \wr \emptyset \vdash v_i: T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \emptyset \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\}: T}$$

By TCASE we know that  $\varphi_i \wr \emptyset \vdash t_i: T$ . As  $\sigma(l): T_i$  we can use RTAccSUCC to show that:

$$\frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) = T}{[\Sigma] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \epsilon}$$

We can then use RTEXTCHOICEONE and RTEXTCHOICETWO to show that:

$$[\Lambda(\sigma)] \&(\alpha(\tilde{T}), T); \varphi_i \rightarrow [\Lambda(\sigma'')] \varphi_i$$

By our assumptions we know that if:

$$\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \alpha(\tilde{T}) \vdash \alpha(\tilde{v})$$

that

$$\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma')$$

By TCASE we know that  $\epsilon \wr \emptyset \vdash v_i: T_i$ , and hence  $\sigma'' = \sigma'$ .

### Case RCaseTwo

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})): T' \quad T' \notin \tilde{T}}{[\sigma] \text{ case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} \xrightarrow{\alpha(\tilde{v})} [\sigma'] \text{ error}} \quad \frac{\varphi_i \wr \emptyset \vdash t_i: T \quad \epsilon \wr \emptyset \vdash v_i: T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \emptyset \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\}: T}$$

As  $T' \notin \tilde{T}$ , by RTACCEERR we know that  $\forall T_i \in \tilde{T}$ :

$$\frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) \neq T}{[\Sigma] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \text{ error}}$$

Hence we can use

$$\frac{\forall i. ([\Sigma] \varphi_i \xrightarrow{\alpha(\tilde{T})}_F [\Sigma_1] \varphi'_i \wedge \text{error} \in \varphi'_i)}{[\Sigma] \varphi_1 \& \varphi_2 \xrightarrow{\alpha(\tilde{T})_j}_F [\Sigma_i] \varphi'_j}$$

to show that:

$$[\Lambda(\sigma)] \&(\alpha(\tilde{T}), T); \varphi_i \rightarrow [\Lambda(\sigma')] \text{ error}; \varphi_i$$

By the hypotheses we know that:

$$[\Lambda(\sigma)] (\alpha(\tilde{T}), T) \rightarrow [\Lambda(\sigma')] \varphi'' \Rightarrow \text{error} \in \varphi''$$

This is a contradiction. Hence this reduction cannot occur, and we can disregard it.

### Case RRegOne

$$\frac{\varphi_1 \wr \Gamma \vdash t_1: \text{UNIT}}{f(\varphi_1) \wr \Gamma \vdash f(t_1): \text{UNIT}}$$

By induction we know that if:

$$[\Lambda(\sigma)] \varphi \xrightarrow{\gamma'}_F [\Lambda(\sigma')] \varphi' \quad \text{error} \notin \varphi'_1 \quad [\sigma] t_1 \xrightarrow{\gamma}_F [\sigma'] t'_1$$

then:

$$\varphi'_1 \wr \emptyset \vdash t'_1: T \quad \gamma' \vdash \gamma$$

Hence we can show that:

$$\frac{[\sigma] t_1 \xrightarrow{\tau}_F [\sigma'] t'_1}{[\sigma'] f(t_1) \xrightarrow{\tau}_F [\sigma'] f(t'_1)}$$

### Case RRegTwo

$$\frac{\varphi_1 \wr \Gamma \vdash t_1: \text{UNIT}}{f(\varphi_1) \wr \Gamma \vdash f(t_1): \text{UNIT}}$$

By induction we know that if:

$$[\Lambda(\sigma)] \varphi \xrightarrow{\gamma'}_F [\Lambda(\sigma')] \varphi' \quad \text{error} \notin \varphi'_1 \quad [\sigma] t_1 \xrightarrow{\gamma}_F [\sigma'] t'_1$$

then:

$$\varphi'_1 \wr \emptyset \vdash t'_1 : T \quad \gamma' \vdash \gamma$$

Hence we can show that:

$$\frac{[\sigma] t \xrightarrow{\alpha(\tilde{v})}_F [\sigma'] t'}{[\sigma] f(t) \xrightarrow{\alpha(\tilde{v})}_{F \cup \{f\}} [\sigma'] t'}$$

### Case RRegThree

$$\frac{}{[\sigma] f(t) \xrightarrow{\tau} [\sigma] t} \quad \frac{\varphi_1 \wr \Gamma \vdash v : \text{UNIT}}{f(\varphi_1) \wr \Gamma \vdash f(v) : \text{UNIT}}$$

By Lemma A.9 we know that  $\varphi_1 = \epsilon$ . We can straightforwardly show that:

$$\overline{[\Sigma] f(\epsilon) \xrightarrow{\tau} [\Sigma] \epsilon}$$

### Case RDmodOne

$$\frac{\varphi \wr \Gamma \vdash t : \text{UNIT}}{\text{dmod}_{\gamma_1}(\varphi) \wr \Gamma \vdash \text{dmod}_{\gamma_1}(t) : \text{UNIT}}$$

By induction we know that if:

$$[\Lambda(\sigma)] \varphi \xrightarrow{\gamma'}_F [\Lambda(\sigma')] \varphi' \quad \text{error} \notin \varphi'_1 \quad [\sigma] t_1 \xrightarrow{\gamma}_F [\sigma'] t'_1$$

then:

$$\varphi'_1 \wr \emptyset \vdash t'_1 : T \quad \gamma' \vdash \gamma$$

Hence we can show that:

$$\frac{[\sigma] t \xrightarrow{\gamma}_{F'} [\sigma'] t' \quad \text{regions}(t) \not\subseteq F}{[\sigma] \text{dmod}_F(t) \xrightarrow{\gamma}_{F'} [\sigma'] \text{dmod}_{F \cup F'}(t')}$$

**Case RDmodTwo**

$$\frac{\mathbf{regions}(t) \subseteq F}{[\sigma] \mathbf{dmod}_F(t) \xrightarrow{\tau} [\sigma] t} \quad \frac{\varphi \wr \Gamma \vdash t: \mathbf{UNIT}}{\mathbf{dmod}_{\gamma_1}(\varphi) \wr \Gamma \vdash \mathbf{dmod}_{\gamma_1}(t): \mathbf{UNIT}}$$

We can straightforwardly use RTDMODTwo to show that:

$$\frac{\mathbf{regions}(\varphi_1) \subseteq F}{[\Lambda(\sigma)] \mathbf{dmod}_F(\varphi_1) \xrightarrow{\tau} [\Lambda(\sigma)] \varphi_1}$$

□

## A.12 Process Subject Reduction Lemma

Lemma 3.2

If:

$$\vdash P: \Phi \quad [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \Rightarrow \text{#error} \in \Phi' \quad [\sigma] P \xrightarrow{\gamma} [\sigma'] P'$$

then

$$\vdash P': \Phi' \quad [\Lambda(\sigma)] \Phi \xrightarrow{\gamma'} [\Lambda(\sigma')] \Phi'$$

where

$$\gamma' \vdash \gamma$$

*Proof.* By induction over  $P$

**Case:**  $P = t$

This case can be typed using:

$$\frac{\varphi \wr \emptyset \vdash t: T \quad \vdash t: \varphi \quad \text{res}(t) \cup \text{res}(\sigma) \subseteq R \quad \exists \mathcal{C}. \mathcal{C}(\Lambda(\sigma), \varphi, \psi, \omega) \text{ and } (\mathcal{C} \implies \mathcal{S})}{\vdash_{\omega}^{\mathcal{S}} [\sigma] t, \psi: \varphi}$$

Using Lemma 3.1 we know that, if:

$$\varphi \wr \emptyset \vdash t: T \quad \mathcal{S}(\Sigma, \varphi, \psi, \omega) \quad [\sigma] t \xrightarrow{\gamma} [\sigma'] t'$$

then

$$\varphi' \wr \emptyset \vdash t': T \quad [\Lambda(\sigma)] \varphi \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi' \quad \gamma' \vdash \gamma$$

We can then apply the `TTHREAD` typing rule:

$$\frac{\varphi' \wr \emptyset \vdash t': T}{\vdash t': \varphi'}$$

**Case:**  $P = P_1 \parallel P_2$

This case can be typed using:

$$\frac{\vdash P_1 : \Phi_1 \quad \vdash P_2 : \Phi_2}{\vdash P_1 \parallel P_2 : \Phi_1 \parallel \Phi_2} \quad \frac{\vdash t : \varphi \quad \text{res}(t) \cup \text{res}(\sigma) \subseteq R \quad \exists \mathcal{C}. \mathcal{C}(\Lambda(\sigma), \varphi, \psi, \omega) \text{ and } (\mathcal{C} \implies \mathcal{S})}{\vdash_{\omega}^{\mathcal{S}} [\sigma] t, \psi : \varphi}$$

By the reduction rules:

$$\frac{[\Sigma] \Phi_1 \xrightarrow{\gamma} [\Sigma'] \Phi'_1}{[\Sigma] \Phi_1 \parallel \Phi_2 \xrightarrow{\gamma} [\Sigma'] \Phi'_1 \parallel \Phi_2} \quad \frac{[\Sigma] \Phi_2 \xrightarrow{\gamma} [\Sigma'] \Phi'_2}{[\Sigma] \Phi_1 \parallel \Phi_2 \xrightarrow{\gamma} [\Sigma'] \Phi'_1 \parallel \Phi'_2}$$

we know that if:

$$\# \text{error} \in \Phi_1 \parallel \Phi_2 \quad [\Sigma] \Phi_1 \parallel \Phi_2 \rightarrow [\Sigma'] \Phi' \Rightarrow \# \text{error} \in \varphi'$$

then:

$$\# \text{error} \in \Phi_1, \Phi_2 \quad [\Sigma] \Phi_1 \rightarrow [\Sigma'] \Phi'_1 \Rightarrow \# \text{error} \in \varphi' \quad [\Sigma] \Phi_2 \rightarrow [\Sigma'] \Phi'_2 \Rightarrow \# \text{error} \in \varphi'$$

We can then use the inductive hypothesis to show that if:

$$[\sigma] P_1 \xrightarrow{\gamma} [\sigma'] P'_1$$

then

$$\vdash P'_1 : \Phi'_1 \quad [\Lambda(\sigma)] \Phi \xrightarrow{\gamma'} [\Lambda(\sigma')] \Phi'$$

We can then show that:

$$\frac{\vdash P_1 : \Phi_1 \quad \vdash P_2 : \Phi_2}{\vdash P_1 \parallel P_2 : \Phi_1 \parallel \Phi_2}$$

and:

$$\frac{[\Lambda(\sigma)] \Phi_1 \xrightarrow{\gamma} [\Lambda(\sigma')] \Phi'_1}{[\Lambda(\sigma)] \Phi_1 \parallel \Phi_2 \xrightarrow{\gamma} [\Lambda(\sigma')] \Phi'_1 \parallel \Phi_2}$$

The case for  $P_2$  is symmetrical.

□

## A.13 Empty Effect Expression Liveness Lemma

Lemma 3.3

If

$$\epsilon \wr \emptyset \vdash t : T$$

then

$$[\sigma] t \xrightarrow{\gamma'} [\sigma] t' \vee t = v$$

*Proof.* By induction over  $t$ .

**Case**  $t = v$

We trivially have that  $t = v$ .

**Case:**  $t = t_1 t_2$

$$\frac{\varphi_1 \wr \emptyset \vdash t_1 : T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \wr \emptyset \vdash t_2 : T_2}{(\varphi_1; \varphi_2; \varphi_3) \wr \emptyset \vdash t_1 t_2 : T_1}$$

We perform a case split on the structure of  $t_1 t_2$

**Case**  $t_1 \neq v$

By the hypothesis and the TAPP rule we know that:

$$\varphi_1 \wr \emptyset \vdash t_1 : T_2 \xrightarrow{\varphi_3} T_1 \quad [\Sigma] \varphi \rightarrow [\Sigma'] \varphi' \quad \varphi \equiv \varphi_1; \varphi_2; \varphi_3$$

By the inductive hypothesis we know that:

$$[\sigma] t_1 \xrightarrow{\gamma'} [\sigma] t'_1 \vee t_1 = v$$

We know the latter cannot be the case, as it contradicts the assumption that  $t_1 \neq v$ .  
Hence:

$$[\sigma] t_1 \xrightarrow{\gamma'} [\sigma] t'_1$$

Hence we can show that:

$$\frac{[\sigma] t_1 \xrightarrow{\gamma_F} [\sigma'] t'_1}{[\sigma] t_1 t_2 \xrightarrow{\gamma_F} [\sigma'] t'_1 t_2}$$

**Case**  $t_1 = v, t_2 \neq v'$

By the hypothesis and the TAPP rule we know that:

$$\varphi_2 \wr \emptyset \vdash t_2 : T_2$$

By the inductive hypothesis we know that:

$$[\sigma] t_2 \xrightarrow{\gamma'} [\sigma] t'_2 \vee t_2 = v$$

We know the latter cannot be the case, as it contradicts the assumption that  $t_2 \neq v$ .

Hence:

$$[\sigma] t_2 \xrightarrow{\gamma'} [\sigma] t'_2$$

Hence we can show that:

$$\frac{[\sigma] t_2 \xrightarrow{\gamma} [\sigma'] t'_2}{[\sigma] v t_2 \xrightarrow{\gamma} [\sigma'] v t'_2}$$

**Case**  $t_1 = v, t_2 = v'$

As  $t_1 = v$  then we know that  $t_1 = \mathbf{rec} \underline{X}(x : T).t_3$ . Hence we can straightforwardly apply the RAPPTHREE rule:

$$\overline{[\sigma] \mathbf{rec} \underline{X}(x : T).t v \xrightarrow{\tau} [\sigma] t[\mathbf{rec} \underline{X}(x : T).t / \underline{X}][v/x]}$$

**Case:**  $t = (\alpha(\tilde{v}), T)$

$$\frac{\epsilon \wr \emptyset \vdash v_i : T_i \quad \text{Res } r \notin T}{(\alpha(\tilde{v}), T) \wr \emptyset \vdash (\alpha(\tilde{v}), T) : T}$$

By the hypothesis we know that

$$(\alpha(\tilde{v}), T) \equiv \epsilon$$

This is a contradiction, and hence we can disregard this case.

**Case:**  $t = \mathbf{if} t_1 \mathbf{then} t_2 \mathbf{else} t_3$

$$\frac{\varphi_1 \wr \emptyset \vdash t_1 : \text{BOOL} \quad \varphi_2 \wr \emptyset \vdash t_2 : T \quad \varphi_3 \wr \emptyset \vdash t_3 : T}{\varphi_1; (\varphi_2 \oplus \varphi_3) \wr \emptyset \vdash \mathbf{if} t_1 \mathbf{then} t_2 \mathbf{else} t_3 : T}$$

By the inductive hypothesis we know that:

$$[\sigma] t_1 \xrightarrow{\gamma'} [\sigma] t'_1 \vee t_1 = v$$

If the former is the case then we can show that:

$$[\sigma] t_1 \xrightarrow{\gamma'} [\sigma] t'_1$$

and hence that:

$$\frac{[\sigma'] t_1 \xrightarrow{\gamma} [\sigma] t'_1}{[\sigma] \text{if } t_1 \text{ then } t_2 \text{ else } t_3 \xrightarrow{\gamma} [\sigma'] \text{if } t'_1 \text{ then } t_2 \text{ else } t_3}$$

If the latter is the case then we can show that:

$$\frac{}{[\sigma] \text{if true then } t_2 \text{ else } t_3 \xrightarrow{\tau} [\sigma] t_2} \quad \frac{}{[\sigma] \text{if false then } t_2 \text{ else } t_3 \xrightarrow{\tau} [\sigma] t_3}$$

**Case:**  $t = \text{case } \alpha(\tilde{T}) \text{ in } \{\tilde{T} \mapsto t\}$

$$\frac{\varphi_i \wr \Gamma \vdash t_i : T \quad \epsilon \wr \Gamma \vdash v_i : T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \Gamma \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} : T}$$

By the hypothesis we know that

$$\&(\alpha(\tilde{T}), T); \varphi_i \equiv \epsilon$$

This is a contradiction, and hence we can disregard this case.

**Case:**  $t = f(t_1)$

$$\frac{\varphi_1 \wr \Gamma \vdash t_1 : \text{UNIT}}{f(\varphi_1) \wr \Gamma \vdash f(t_1) : \text{UNIT}}$$

By the hypothesis we know that

$$f(\varphi_1) \equiv \epsilon$$

This is a contradiction, and hence we can disregard this case.

**Case:**  $t = \text{dmod}_{F_1}(t_1)$

$$\frac{\varphi_1 \wr \emptyset \vdash t_1 : \text{UNIT}}{\text{dmod}_{\gamma_1}(\varphi_1) \wr \emptyset \vdash \text{dmod}_{\gamma_1}(t_1) : \text{UNIT}}$$

By the hypothesis we know that

$$\mathbf{dmod}_F(\varphi_1) \equiv \epsilon$$

This is a contradiction, and hence we can disregard this case.  $\square$

## A.14 Active Effect Expression Liveness Lemma

Lemma 3.4

If

$$\varphi \wr \emptyset \vdash t : T \quad [\Lambda(\sigma)] \varphi \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'$$

then

$$[\sigma] t \xrightarrow{\gamma'} [\sigma'] t'$$

*Proof.* By induction over  $t$ .

**Case**  $t = v$

By Lemma A.9 we know that  $\varphi = \epsilon$ . There are no reduction rules for  $\epsilon$ . Therefore the premises of the lemma are not fulfilled, and we can disregard the case.

**Case:**  $t = t_1 t_2$

$$\frac{\varphi_1 \wr \emptyset \vdash t_1 : T_2 \xrightarrow{\varphi_3} T_1 \quad \varphi_2 \wr \emptyset \vdash t_2 : T_2}{(\varphi_1; \varphi_2; \varphi_3) \wr \emptyset \vdash t_1 t_2 : T_1}$$

We perform a case split on the structure of  $t_1 t_2$

**Case**  $t_1 \neq v$

By the hypothesis and the TAPP rule we know that:

$$\varphi_1 \wr \emptyset \vdash t_1 : T_2 \xrightarrow{\varphi_3} T_1 \quad [\Sigma] \varphi \rightarrow [\Sigma'] \varphi' \quad \varphi \equiv \varphi_1; \varphi_2; \varphi_3$$

Either  $\varphi_1 \equiv \epsilon$  or  $\varphi_1 \not\equiv \epsilon$ .

**Case:**  $\varphi_1 \equiv \epsilon$  By Lemma 3.3 we know that:

$$[\sigma] t_1 \xrightarrow{\gamma'} [\sigma'] t'_1 \vee t_1 = v$$

We know the latter cannot be the case, as it contradicts the assumption that  $t_1 \neq v$ .  
Hence:

$$[\sigma] t_1 \xrightarrow{\gamma'} [\sigma'] t'_1$$

**Case:**  $\varphi_1 \not\equiv \epsilon$  The reduction rule for sequencing is:

$$\frac{[\Sigma] \varphi_1 \xrightarrow{F} [\Sigma'] \varphi'_1}{[\Sigma] \varphi_1; \varphi_2; \varphi_3 \xrightarrow{F} [\Sigma'] \varphi'_1; \varphi_2; \varphi_3}$$

By this reduction rule we know that:

$$[\Lambda(\sigma)] \varphi_1 \xrightarrow{\gamma} [\Lambda(\sigma')] \varphi'_1$$

Hence we can use the inductive hypothesis to show that:

$$[\sigma] t_1 \xrightarrow{\gamma} [\sigma'] t'_1$$

**Hence we can show that**, irrespective of whether or not  $\varphi_1 \equiv \epsilon$  that:

$$\frac{[\sigma] t_1 \xrightarrow{F} [\sigma'] t'_1}{[\sigma] t_1 t_2 \xrightarrow{F} [\sigma'] t'_1 t_2}$$

**Case**  $t_1 = v, t_2 \neq v'$

By the hypothesis and the TAPP rule we know that:

$$\varphi_2 \wr \emptyset \vdash t_2 : T_2$$

By Lemma A.9, we know that  $\varphi_1 = \epsilon$ . Either  $\varphi_2 \equiv \epsilon$  or  $\varphi_2 \not\equiv \epsilon$ .

**Case:**  $\varphi_2 \equiv \epsilon$  By Lemma 3.3 we know that:

$$[\sigma] t_2 \xrightarrow{\gamma'} [\sigma'] t'_2 \vee t_2 = v$$

We know the latter cannot be the case, as it contradicts the assumption that  $t_2 \neq v$ .

Hence:

$$[\sigma] t_2 \xrightarrow{\gamma'} [\sigma'] t'_2$$

**Case:**  $\varphi_2 \not\equiv \epsilon$  The reduction rule for sequencing is:

$$\frac{[\Sigma] \varphi_2 \xrightarrow{F} [\Sigma'] \varphi'_2}{[\Sigma] \varphi_2; \varphi_3 \xrightarrow{F} [\Sigma'] \varphi'_2; \varphi_3}$$

By this reduction rule we know that:

$$[\Lambda(\sigma)] \varphi_1 \xrightarrow{\gamma} [\Lambda(\sigma')] \varphi'_1$$

Hence we can use the inductive hypothesis to show that:

$$[\sigma] t_1 \xrightarrow{\gamma} [\sigma'] t'_1$$

**Hence we can show that**, irrespective of whether or not  $\varphi_2 \equiv \epsilon$  that:

$$\frac{[\sigma] t_2 \xrightarrow{\gamma} [\sigma'] t'_2}{[\sigma] v t_2 \xrightarrow{\gamma} [\sigma'] v t'_2}$$

**Case**  $t_1 = v, t_2 = v'$

As  $t_1 = v$  then we know that  $t_1 = \mathbf{rec} X(x : T).t_3$ . Hence we can straightforwardly apply the RAPP THREE rule:

$$\overline{[\sigma] \mathbf{rec} \underline{X}(x : T).t v \xrightarrow{\tau} [\sigma] t[\mathbf{rec} \underline{X}(x : T).t / \underline{X}][v/x]}$$

**Case:**  $t = (\alpha(\tilde{v}), T)$

$$\frac{\epsilon \wr \emptyset \vdash v_i : T_i \quad \mathbf{Res} r \notin T}{(\alpha(\tilde{T}), T) \wr \emptyset \vdash (\alpha(\tilde{v}), T) : T}$$

By the hypothesis we know that

$$[\Sigma] (\alpha(\tilde{T}), T) \rightarrow [\Sigma'] \varphi'$$

The reduction rules for  $(\alpha(\tilde{T}), T)$  that fulfil these conditions are:

$$\frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) = T}{[\Sigma] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \epsilon} \quad \frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) \neq T}{[\Sigma] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \mathbf{error}}$$

Hence we can show that  $\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma')$ . By our assumptions we know that:

$$\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \Leftrightarrow \sigma \xrightarrow{\alpha(\tilde{v})} \sigma'$$

Hence we can reduce  $(\alpha(\tilde{v}), T)$  using either the RAccSUCC or RAccERR rule:

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T}{[\sigma] (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \sigma(\alpha(\tilde{v}))} \quad \frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})) : T' \neq T}{[\sigma] (\alpha(\tilde{v}), T) \xrightarrow{\alpha(\tilde{v})} [\sigma'] \mathbf{error}}$$

**Case:**  $t = \text{if } t_1 \text{ then } t_2 \text{ else } t_3$

$$\frac{\varphi_1 \wr \emptyset \vdash t_1: \text{BOOL} \quad \varphi_2 \wr \emptyset \vdash t_2: T \quad \varphi_3 \wr \emptyset \vdash t_3: T}{\varphi_1; (\varphi_2 \oplus \varphi_3) \wr \emptyset \vdash \text{if } t_1 \text{ then } t_2 \text{ else } t_3: T}$$

Either  $\varphi_1 \equiv \epsilon$  or  $\varphi_1 \not\equiv \epsilon$ .

**Case:**  $\varphi_1 \equiv \epsilon$  By Lemma 3.3 we know that:

$$[\sigma] t_1 \xrightarrow{\gamma'} [\sigma'] t'_1 \vee t_1 = v$$

If the former is the case then we can show that:

$$[\sigma] t_1 \xrightarrow{\gamma'} [\sigma'] t'_1$$

and hence that:

$$\frac{[\sigma'] t_1 \xrightarrow{\gamma} [\sigma] t'_1}{[\sigma] \text{if } t_1 \text{ then } t_2 \text{ else } t_3 \xrightarrow{\gamma} [\sigma'] \text{if } t'_1 \text{ then } t_2 \text{ else } t_3}$$

If the latter is the case then we can show that:

$$\frac{}{[\sigma] \text{if true then } t_2 \text{ else } t_3 \xrightarrow{\tau} [\sigma] t_2} \quad \frac{}{[\sigma] \text{if false then } t_2 \text{ else } t_3 \xrightarrow{\tau} [\sigma] t_3}$$

**Case:**  $\varphi_1 \not\equiv \epsilon$  The reduction rule for sequencing is:

$$\frac{[\Sigma] \varphi_1 \xrightarrow{\gamma} [\Sigma'] \varphi'_1}{[\Sigma] \varphi_1; \varphi_2 \oplus \varphi_3 \xrightarrow{\gamma} [\Sigma'] \varphi'_1; \varphi_2 \oplus \varphi_3}$$

By this reduction rule we know that:

$$[\Lambda(\sigma)] \varphi_1 \xrightarrow{\gamma} [\Lambda(\sigma')] \varphi'_1$$

Hence we can use the inductive hypothesis to show that:

$$[\sigma] t_1 \xrightarrow{\gamma} [\sigma'] t'_1$$

and hence that:

$$\frac{[\sigma'] t_1 \xrightarrow{\gamma} [\sigma] t'_1}{[\sigma] \text{if } t_1 \text{ then } t_2 \text{ else } t_3 \xrightarrow{\gamma} [\sigma'] \text{if } t'_1 \text{ then } t_2 \text{ else } t_3}$$

**Case:**  $t = \text{case } \alpha(\tilde{T}) \text{ in } \{\tilde{T} \mapsto t\}$

$$\frac{\varphi_i \wr \Gamma \vdash t_i : T \quad \epsilon \wr \Gamma \vdash v_i : T_i \quad T_i \neq \text{Res } r}{\&(\alpha(\tilde{T}), T_i); \varphi_i \wr \Gamma \vdash \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} : T}$$

By the hypothesis we know that

$$[\Sigma] \&(\alpha(\tilde{T}), T); \varphi_i \rightarrow [\Sigma'] \varphi'$$

The reduction rules for  $\&(\alpha(\tilde{T}), T); \varphi_i$  that fulfils these conditions are:

$$\begin{array}{c} \frac{\exists i. ([\Sigma] \varphi_i \xrightarrow{\alpha(\tilde{T})} F [\Sigma'] \varphi'_i \wedge \text{#error} \in \varphi'_i)}{[\Sigma] \varphi_1 \& \varphi_2 \xrightarrow{\alpha(\tilde{T})} F [\Sigma'] \varphi'_i} \quad \frac{[\Sigma] \varphi_2 \xrightarrow{\alpha(\tilde{T})} F [\Sigma'] \varphi'_2 \quad \text{#error} \in \varphi'_2}{[\Sigma] \varphi_1 \& \varphi_2 \xrightarrow{\alpha(\tilde{T})} F [\Sigma'] \varphi'_2} \\ \\ \frac{\forall i. ([\Sigma] \varphi_i \xrightarrow{\alpha(\tilde{T})} F [\Sigma_1] \varphi'_i \wedge \text{error} \in \varphi'_i)}{[\Sigma] \varphi_1 \& \varphi_2 \xrightarrow{\alpha(\tilde{T})_j} F [\Sigma_i] \varphi'_j} \end{array}$$

Each of these can only reduce if:

$$[\Sigma] (\alpha(\tilde{T}), T_i) \rightarrow$$

By our assumptions we know that  $\Lambda(\sigma) \xrightarrow{\alpha(\tilde{T})} \Lambda(\sigma') \Leftrightarrow \sigma \xrightarrow{\alpha(\tilde{v})} \sigma'$ .

Hence we can reduce  $(\alpha(\tilde{v}), T)$  using either the RCASESUCC or the RCASEERR rule:

$$\frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})): T_i \quad T_i \in \tilde{T}}{[\sigma] \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} \xrightarrow{\alpha(\tilde{v})} [\sigma'] t_i} \quad \frac{\sigma \xrightarrow{\alpha(\tilde{v})} \sigma' \quad \sigma(\alpha(\tilde{v})): T' \quad T' \notin \tilde{T}}{[\sigma] \text{case } \alpha(\tilde{v}) \text{ in } \{\tilde{T} \mapsto t\} \xrightarrow{\alpha(\tilde{v})} [\sigma'] \text{error}}$$

**Case:**  $t = f(t_1)$

$$\frac{\varphi_1 \wr \Gamma \vdash t_1 : \text{UNIT}}{f(\varphi_1) \wr \Gamma \vdash f(t_1) : \text{UNIT}}$$

By the hypothesis we know that:

$$[\Sigma] f(\varphi_1) \rightarrow [\Sigma'] \varphi'$$

There are three possible reduction rules for  $f(\varphi_1)$ .

**Case RTRegOne**

$$\frac{[\Sigma] \varphi_1 \xrightarrow{\tau} F [\Sigma'] \varphi'}{[\Sigma] f(\varphi_1) \xrightarrow{\tau} F [\Sigma'] f(\varphi')}$$

As

$$[\Lambda(\sigma)] \varphi_1 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_1$$

we can apply the inductive hypothesis and show that:

$$[\sigma] t_1 \xrightarrow{\gamma} [\sigma'] t'_1$$

We can then show that:

$$\frac{[\sigma] t_1 \xrightarrow{\tau} [\sigma'] t'_1}{[\sigma'] f(t_1) \xrightarrow{\tau} [\sigma'] f(t'_1)}$$

**Case RTRegTwo**

$$\frac{[\Sigma] \varphi_1 \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'}{[\Sigma] f(\varphi_1) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'_1}$$

As

$$[\Lambda(\sigma)] \varphi_1 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_1$$

we can apply the inductive hypothesis and show that:

$$[\sigma] t_1 \xrightarrow{\gamma} [\sigma'] t'_1$$

We can then show that:

$$\frac{[\sigma] t_1 \xrightarrow{\alpha(\tilde{v})} [\sigma'] t'_1}{[\sigma] f(t_1) \xrightarrow{\alpha(\tilde{v})} [\sigma'] t'_1}$$

**Case RTRegThree**

$$\overline{[\Sigma] f(\epsilon) \xrightarrow{\tau} [\Sigma] \epsilon}$$

When  $\exp = v$  We can straightforwardly use the following reduction rule:

$$\overline{[\sigma] f(v) \xrightarrow{\tau} [\sigma] v}$$

Otherwise we can use RTREGONE or RTREGTWO.

**Case:**  $t = \text{dmod}_{F_1}(t_1)$

$$\frac{\varphi_1 \wr \emptyset \vdash t_1 : \text{UNIT}}{\text{dmod}_{\gamma_1}(\varphi_1) \wr \emptyset \vdash \text{dmod}_{\gamma_1}(t_1) : \text{UNIT}}$$

There are two possible reductions for  $\text{dmod}_{F_1}(\varphi_1)$ . We can perform a case split on the rules' premises:

**Case  $\text{regions}(\varphi) \not\subseteq \gamma_1$**

$$\frac{[\Sigma] \varphi_1 \xrightarrow{\gamma} [\Sigma'] \varphi'_1 \quad \text{regions}(\varphi_1) \not\subseteq F}{[\Sigma] \text{dmod}_F(\varphi_1) \xrightarrow{\gamma} [\Sigma'] \text{dmod}_{F \cup F'}(\varphi'_1)}$$

As

$$[\Lambda(\sigma)] \varphi_1 \xrightarrow{\gamma'} [\Lambda(\sigma')] \varphi'_1$$

we can apply the inductive hypothesis and show that:

$$[\sigma] t \xrightarrow{\gamma} [\sigma'] t'$$

and hence can show that:

$$\frac{[\sigma] t \xrightarrow{\gamma} [\sigma'] t' \quad \text{regions}(t) \not\subseteq F}{[\sigma] \text{dmod}_F(t) \xrightarrow{\gamma} [\sigma'] \text{dmod}_{F \cup F'}(t')}$$

**Case  $\text{regions}(\varphi) \subseteq \gamma_1$**

$$\frac{\text{regions}(\varphi_1) \subseteq F}{[\Sigma] \text{dmod}_F(\varphi_1) \xrightarrow{\tau} [\Sigma] \varphi_1}$$

In this case we can straightforwardly apply the following reduction rule:

$$\frac{\text{regions}(t) \subseteq F}{[\sigma] \text{dmod}_F(t) \xrightarrow{\tau} [\sigma] t}$$

□

## A.15 Empty Effect Thread Liveness Lemma

Lemma 3.5 If:

$$\vdash P: \Phi \quad \Phi = \Pi_I \varphi_i \quad (\varphi_i \equiv \epsilon \vee \varphi_i \equiv \underline{X})$$

then

$$[\sigma] P \xrightarrow{\gamma} [\sigma] P' \vee P = \Pi_I v$$

*Proof.* **Case**  $\Phi = \varphi$

$$\frac{\varphi \wr \emptyset \vdash t: T}{\vdash t: \varphi}$$

By Lemma 3.3 we know that:

$$[\sigma] t \xrightarrow{\gamma} [\sigma] t' \vee t = v$$

**Case:**  $\Phi = \Phi_1 \parallel \Phi_2$

$$\frac{\vdash P_1: \Phi_1 \quad \vdash P_2: \Phi_2}{\vdash P_1 \parallel P_2: \Phi_1 \parallel \Phi_2}$$

We can then use the inductive hypothesis to show that:

$$([\sigma] P_1 \xrightarrow{\gamma} [\sigma] P'_1 \vee P_1 = \Pi v)$$

$$([\sigma] P_2 \xrightarrow{\gamma} [\sigma] P'_2 \vee P_2 = \Pi v)$$

and hence that:

$$[\sigma] P_1 \parallel P_2 \xrightarrow{\gamma} [\sigma] P'_1 \parallel P_2 \vee [\sigma] P_1 \parallel P_2 \xrightarrow{\gamma} [\sigma] P_1 \parallel P'_2 \vee P_1 \parallel P_2 = \Pi v$$

□

## A.16 Active Effect Thread Liveness Lemma

Lemma 3.6

If:

$$\vdash P: \Phi \quad [\Lambda(\sigma)] \Phi \xrightarrow{\gamma'} [\Lambda(\sigma')] \Phi'$$

then

$$[\sigma] P \xrightarrow{\gamma} [\sigma'] P'$$

*Proof.* By induction over  $[\Lambda(\sigma)] \Phi \xrightarrow{\gamma'} [\Lambda(\sigma')] \Phi'$

**Case RTAccSucc, RTSeq, RTIntChoice, RTEExtChoiceOne, RTEExtChoiceTwo, RTEExtChoiceThree, RTRegOne, RTRegTwo, RTRegThree, RTDmodOne, RTDmodTwo**

$$\frac{\varphi \wr \emptyset \vdash t: T}{\vdash t: \varphi}$$

By Lemma 3.4 we know that:

$$[\sigma] t \xrightarrow{\gamma} [\sigma'] t'$$

**Case: RTParOne**

$$\frac{\vdash P_1: \Phi_1 \quad \vdash P_2: \Phi_2}{\vdash P_1 \parallel P_2: \Phi_1 \parallel \Phi_2} \quad \frac{[\Sigma] \Phi_1 \xrightarrow{F} [\Sigma'] \Phi'_1}{[\Sigma] \Phi_1 \parallel \Phi_2 \xrightarrow{F} [\Sigma'] \Phi'_1 \parallel \Phi_2}$$

By the reduction rule we know that:

$$[\Lambda(\sigma)] \Phi_1 \xrightarrow{\gamma'} [\Lambda(\sigma')] \Phi'_1$$

We can then use the inductive hypothesis to show that:

$$[\sigma] P_1 \xrightarrow{\gamma} [\sigma'] P'_1$$

and hence that:

$$\frac{[\sigma] P_1 \xrightarrow{F} [\sigma'] P'_1}{[\sigma] P_1 \parallel P_2 \xrightarrow{F} [\sigma'] P'_1 \parallel P_2}$$

**Case: RTParTwo**

Similar to case RTPARONE.

**Case: RTEquiv**

$$\frac{\Phi \equiv \Phi_1 \quad [\Sigma] \Phi_1 \xrightarrow{\gamma} [\Sigma'] \Phi_2 \quad \Phi_2 \equiv \Phi'}{[\Sigma] \Phi \xrightarrow{\gamma} [\Sigma'] \Phi'}$$

By induction. □

## Appendix B

# Blocking Message Passing

In this chapter we provide the full proofs of how the blocking message passing approach in Section 3.5.1 implies General Compatibility of static programs.

### B.1 Separate Channels Lemma

**Lemma B.1.**

$$[\Sigma] \varphi_i \rightarrow [\Sigma'] \varphi'_i \Leftrightarrow [c \mapsto q; \Sigma] \varphi_i \rightarrow [c \mapsto q; \Sigma'] \varphi'_i$$

where

$$c \neq dd_i \wedge \varphi_i \text{ has principle channel allocation}$$

*Proof.* By induction over  $\rightarrow$

**Case: RTAccSucc**

$$\frac{\Sigma \xrightarrow{\alpha(\tilde{T})} \Sigma' \quad \Sigma(\alpha(\tilde{T})) = T}{[\Sigma] (\alpha(\tilde{T}), T) \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \epsilon}$$

The only actions we can perform are sending and receiving.

**Case  $c! \langle T \rangle$**

Here  $\Sigma' = \Sigma[c \mapsto q', T]$ . Hence:

$$[\Sigma[c \mapsto q']] c! \langle T \rangle \rightarrow [\Sigma[c \mapsto q', T]] \epsilon \Leftrightarrow [\Sigma[c \mapsto q, q']] c! \langle T \rangle \rightarrow [\Sigma[c \mapsto q, q', T]] \varphi'_1$$

### Case $c'?(T)$

As  $c \neq dd_i$  and  $\varphi_i$  has principle channel allocation we know that  $c' \neq c$ . Hence:

$$[\Sigma[c' \mapsto q']] c?(T) \rightarrow [\Sigma[c' \mapsto q'', T]] \epsilon \Leftrightarrow [c \mapsto q; \Sigma[c' \mapsto q']] c?(T) \rightarrow [c \mapsto q; \Sigma[c' \mapsto q'']] \varphi'_1$$

### Case: RTAccErr

Similar to case RTAccSucc

### Case: RTSeq

$$\frac{[\Sigma] \varphi_1 \xrightarrow{\gamma} [\Sigma'] \varphi'_1}{[\Sigma] \varphi_1; \varphi_2 \xrightarrow{\gamma} [\Sigma'] \varphi'_1; \varphi_2}$$

By induction we know that:

$$[\Sigma] \varphi_1 \rightarrow [\Sigma'] \varphi'_1 \Leftrightarrow [c \mapsto q; \Sigma] \varphi_1 \rightarrow [c \mapsto q; \Sigma'] \varphi'_1$$

Hence we have that:

$$[\Sigma] \varphi_1; \varphi_2 \rightarrow [\Sigma'] \varphi'_1; \varphi_2 \Leftrightarrow [c \mapsto q; \Sigma] \varphi_1; \varphi_2 \rightarrow [c \mapsto q; \Sigma'] \varphi'_1; \varphi_2$$

### Case: RTIntChoice

$$\frac{i \in 1, 2}{[\Sigma] \varphi_1 \oplus \varphi_2 \xrightarrow{\tau} [\Sigma] \varphi_i}$$

Trivial

### Case: RTExtChoiceOne

$$\frac{\exists i . ([\Sigma] \varphi_i \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'_i \wedge \text{#error} \in \varphi'_i)}{[\Sigma] \varphi_1 \& \varphi_2 \xrightarrow{\alpha(\tilde{T})} [\Sigma'] \varphi'_i}$$

By induction we know that:

$$[\Sigma] \varphi_i \rightarrow [\Sigma'] \varphi'_i \Leftrightarrow [c \mapsto q; \Sigma] \varphi_i \rightarrow [c \mapsto q; \Sigma'] \varphi'_i$$

Hence we have that:

$$[\Sigma] \varphi_1 \& \varphi_2 \rightarrow [\Sigma'] \varphi'_i \Leftrightarrow [c \mapsto q; \Sigma] \varphi_1 \& \varphi_2 \rightarrow [c \mapsto q; \Sigma'] \varphi'_i$$

### Case: RTExtChoiceTwo

Similar to case RTExtChoiceOne.

### Case: RTUnfold

$$\overline{[\Sigma] \mu \underline{X}. \varphi \xrightarrow{\tau} [\Sigma] \varphi[\mu \underline{X}. \varphi / \underline{X}]}$$

Trivial

### Case: RTEquiv

$$\frac{\varphi \equiv \varphi_1 \quad [\Sigma] \varphi_1 \xrightarrow{\gamma} [\Sigma'] \varphi_2 \quad \varphi_2 \equiv \varphi'}{[\Sigma] \varphi \xrightarrow{\gamma} [\Sigma'] \varphi'}$$

By induction we know that:

$$[\Sigma] \varphi_1 \rightarrow [\Sigma'] \varphi_2 \Leftrightarrow [c \mapsto q; \Sigma] \varphi_1 \rightarrow [c \mapsto q; \Sigma'] \varphi_2$$

Hence we have that:

$$[\Sigma] \varphi \rightarrow [\Sigma'] \varphi' \Leftrightarrow [c \mapsto q; \Sigma] \varphi \rightarrow [c \mapsto q; \Sigma'] \varphi'$$

□

## B.2 Send Action Equivalence Lemma

**Lemma B.2.** *If*

$$\Pi G \upharpoonright d_i \equiv \Pi \varphi_i \wedge \varphi_a \equiv \oplus_N(c! \langle T_n \rangle; \varphi_b^n) \wedge N \neq \emptyset$$

*Then*

$$\exists G' = d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle_M . N \subseteq M \wedge \Pi G' \upharpoonright d_i \equiv \Pi \varphi_i$$

*Proof.* By induction over  $G$

**Case**  $G = 0, X$

$$G \upharpoonright d_a \not\equiv \oplus_N(c! \langle T_n \rangle; \varphi_b^n)$$

Hence we can disregard these cases.

**Case**  $G = \mu X.G''$

By equivalence we know that:

$$\begin{aligned} \mu \underline{X}.G'' \upharpoonright d_a &\equiv \mu \underline{X}. \oplus_N (c! \langle T_n \rangle; \varphi_b^n) \\ &\equiv \oplus_N (c! \langle T_n \rangle; \varphi_b^n) [\mu \underline{X}. \oplus_N (c! \langle T_n \rangle; \varphi_b^n) / \underline{X}] \\ &\equiv (G''[\mu \underline{X}.G'' / \underline{X}]) \upharpoonright d_a \end{aligned}$$

Hence we can consider  $(G''[\mu \underline{X}.G'' / \underline{X}]) \upharpoonright d_a$  directly.

**Case**  $G = d_{a'} \rightarrow d_{b'} : c' \langle \widetilde{T \mapsto G} \rangle_M$

$$c' = c$$

We let  $G' = G$  directly.

$$c' \neq c$$

Hence  $a' \neq a$ ,  $b' \neq b$ . By the definition of  $\upharpoonright$  we know that:

$$G_m \upharpoonright d_a \equiv \oplus_N (c! \langle T_n \rangle; \varphi_b^n)$$

Let  $\Phi_m \equiv \Pi G_m \upharpoonright d_i$ . Then, by induction we can show that:

$$\exists G'_m . \Phi_m \equiv \Pi G'_m \upharpoonright d_i \quad G'_m = d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle_O \quad N \subseteq O$$

Then let:

$$G' = d_a \rightarrow d_b : c \langle T_o \mapsto d_{a'} \rightarrow d_{b'} : c' \langle \widetilde{T_m \mapsto G_o} \rangle_M \rangle_O$$

Here:

$$\Pi G' \upharpoonright d_i \equiv \Pi \varphi_i$$

□

### B.3 Receive Action Equivalence Lemma

**Lemma B.3.** *If*

$$\Pi G \upharpoonright d_i \equiv \Pi \varphi_i \wedge \varphi_b \equiv \&_M(c! \langle T_m \rangle; \varphi_b^m)$$

*Then*

$$\exists G' = d_a \rightarrow d_b : \widetilde{c \langle T \mapsto G \rangle_N} \cdot \Pi G' \upharpoonright d_i \equiv \Pi \varphi_i$$

*Proof.* Similar to Lemma B.2. □

## B.4 Valid States Safety Lemma

Lemma 3.10

$$[\Sigma] \Phi \in \Delta \Rightarrow \text{error} \notin \Phi$$

*Proof.* By induction over  $n$ .

**Case**  $n = 0$

$$\Delta_0 \stackrel{\text{def}}{=} \{[\Sigma_\emptyset] \Phi \mid \Phi \equiv \Pi G \upharpoonright d_i\}$$

By induction over  $G$

**Case**  $G = 0, X$

Trivial.

**Case**  $G = \mu X.G$

By induction we have that:

$$\Phi' = \Pi \varphi_i \equiv \Pi G \upharpoonright d_i \Rightarrow \text{error} \notin \varphi$$

Hence:

$$\Phi = \Pi \mu \underline{X}. \varphi_i \equiv \Pi \mu \underline{X}. G \upharpoonright d_i \Rightarrow \text{error} \notin \varphi'$$

**Case**  $G = d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle$

$$\begin{aligned} d_1 \rightarrow d_2 : r \langle \widetilde{T \mapsto G} \rangle_M \upharpoonright d_1 &\stackrel{\text{def}}{=} \oplus_N (c! \langle T_n \rangle; G_n \upharpoonright d_1) \quad \emptyset \neq N \subseteq M \\ d_1 \rightarrow d_2 : r \langle \widetilde{T \mapsto G} \rangle \upharpoonright d_2 &\stackrel{\text{def}}{=} \&_M (c? \langle T_m \rangle; G_m \upharpoonright d_1) \\ d_1 \rightarrow d_2 : r \langle \widetilde{T \mapsto G} \rangle \upharpoonright d &\stackrel{\text{def}}{=} G_1 \upharpoonright d \end{aligned}$$

By induction we know that:

$$\Phi' \Pi \varphi'_i \equiv \Pi G_m \upharpoonright d_i \Rightarrow \text{error} \notin \varphi'_i$$

Hence:

$$\varphi_i^a \equiv \oplus_N (c! \langle T_n \rangle; G_n \upharpoonright d_1) \Rightarrow \text{error} \notin \varphi_i^a$$

$$\varphi_i^n \equiv \&_M (c? \langle T_m \rangle; G_m \upharpoonright d_1) \Rightarrow \text{error} \notin \varphi_i^n$$

Therefore we have that:

$$\Phi \equiv \Pi d_a \rightarrow d_b : \widetilde{c \langle T \mapsto G \rangle \upharpoonright d_i} \Rightarrow \text{error} \notin \Phi$$

**Case**  $n > 0$

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_b^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\delta'_k = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \varphi_m \quad k \in J$$

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^j \parallel \dots \parallel \varphi_a^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\delta'_k = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \parallel \varphi_a^k \parallel \dots \varphi_m \quad k \in J$$

By induction we know that:

$$[\Sigma] \Phi' = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_b^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n \Rightarrow \text{error} \notin \Phi'$$

Hence we have that:

$$\text{error} \notin \&_J(c?(T_j); \varphi_b^j)$$

Therefore we can show that:

$$[\_] \Phi = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \varphi_m \Rightarrow \text{error} \notin \Phi$$

□

## B.5 Valid States Liveness Lemma

Lemma 3.11

$$[\Sigma] \Phi \in \Delta \Rightarrow [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \vee \Phi \equiv \Pi \epsilon \vee \Phi \equiv \Pi \underline{X}$$

*Proof.* By induction over  $n$ .

**Case**  $n = 0$

By case analysis of  $G$ .

If  $[\Sigma_\emptyset] \Phi \in \Delta_0$  and  $\Phi \equiv \Phi'$  then  $[\Sigma_\emptyset] \Phi' \in \Delta_0$  and hence it is sufficient to consider direct reduction.

**Case**  $G = d_a \rightarrow d_b : c\langle \widetilde{T \mapsto G} \rangle$

If

$$d_a \rightarrow d_b : c\langle \widetilde{T \mapsto G} \rangle \upharpoonright d_a \equiv \oplus_N (c! \langle T_n \rangle; G_n \upharpoonright d_a)$$

then:

$$[\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel \oplus_N (c! \langle T_n \rangle; G_n \upharpoonright d_a) \parallel \dots \parallel \varphi_o \rightarrow [\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel (c! \langle T_n \rangle; G_n \upharpoonright d_a) \parallel \dots \parallel \varphi_o$$

If

$$d_a \rightarrow d_b : c\langle \widetilde{T \mapsto G} \rangle \upharpoonright d_a \equiv \varphi_1 \parallel \dots \parallel (c! \langle T_n \rangle; G_n \upharpoonright d_a) \parallel \dots \parallel \varphi_o$$

then:

$$[\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel (c! \langle T_n \rangle; G_n \upharpoonright d_a) \parallel \dots \parallel \varphi_o \rightarrow [c \mapsto T_n] \varphi_1 \parallel \dots \parallel G_n \upharpoonright d_a \parallel \dots \parallel \varphi_o$$

**Case**  $G = 0$

$$\Phi \equiv \Pi \epsilon$$

**Case**  $G = X$

$$\Phi \equiv \Pi \underline{X}$$

**Case**  $n > 0$

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_b^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\delta'_k = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \varphi_m \quad k \in J$$

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^j \parallel \dots \parallel \varphi_a^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\delta'_k = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \parallel \varphi_a^k \parallel \dots \varphi_m \quad k \in J$$

We can show that:

$$[c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \varphi_m \rightarrow [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \varphi_b^j \parallel \dots \varphi_m$$

□

## Appendix C

# Non Blocking Message Passing

In this chapter we provide the full proofs of how the non-blocking message passing approach in Section 3.5.2 implies General Compatibility of static programs.

### C.1 Non Blocking Send Action Equivalence Lemma

**Lemma C.1.** *If*

$$\varphi_a \equiv \oplus_N((c!\langle T_n \rangle, \text{UNIT}); G_n \upharpoonright d_1 \& (c!\langle T_n \rangle, l_{[]}); G_{l_{[]}} \upharpoonright d_1)$$

$$\Pi G \upharpoonright d_i \equiv \Pi \varphi_i \quad N \neq \emptyset$$

*Then*

$$\exists G' = d_a \rightarrow d_b : \widetilde{c\langle T \mapsto G \rangle_M} \cdot N \subseteq M \quad \Pi G' \upharpoonright d_i \equiv \Pi \varphi_i$$

*Proof.* By induction over  $G$

**Case**  $G = 0, X$

$$G \upharpoonright d_a \not\equiv \oplus_N(c!\langle T_n \rangle; \varphi_b^n)$$

Hence we can disregard these cases.

**Case**  $G = \mu X.G''$

By equivalence we know that:

$$\begin{aligned}\mu \underline{X}.G'' \upharpoonright d_a &\equiv \mu \underline{X}.\oplus_N (c! \langle T_n \rangle; \varphi_b^n) \\ &\equiv \oplus_N (c! \langle T_n \rangle; \varphi_b^n) [\mu \underline{X}.\oplus_N (c! \langle T_n \rangle; \varphi_b^n) / \underline{X}] \\ &\equiv (G''[\mu \underline{X}.G'' / \underline{X}]) \upharpoonright d_a\end{aligned}$$

Hence we can consider  $(G'[\mu \underline{X}.G' / \underline{X}]) \upharpoonright d_a$  directly.

**Case**  $G = d_{a'} \rightarrow d_{b'} : c' \langle \widetilde{T \mapsto G} \rangle_M$

$$c' = c$$

We let  $G' = G$  directly.

$$c' \neq c$$

Hence  $a' \neq a$ ,  $b' \neq b$ . By the definition of  $\upharpoonright$  we know that:

$$G_m \upharpoonright d_a \equiv \oplus_N ((c! \langle T_n \rangle, \text{UNIT}); G_n \upharpoonright d_1 \& (c! \langle T_n \rangle, l_{[]})); G_{l_{[]}} \upharpoonright d_1$$

Let  $\Phi_m \equiv \Pi G_m \upharpoonright d_i$ . Then, by induction we can show that:

$$\exists G'_m . \Phi_m \equiv \Pi G'_m \upharpoonright d_i \quad G'_m = d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle_O \quad N \subseteq O$$

Then let:

$$G' = d_a \rightarrow d_b : c \langle T_o \mapsto d_{a'} \rightarrow d_{b'} : c' \langle \widetilde{T_m \mapsto G_o} \rangle_M \rangle_O$$

Here:

$$\Pi G' \upharpoonright d_i \equiv \Pi \varphi_i$$

□

## C.2 Non Blocking Receive Action Equivalence Lemma

**Lemma C.2.** *If*

$$\Pi G \upharpoonright d_i \equiv \Pi \varphi_i \wedge \varphi_b \equiv \&_M(c! \langle T_m \rangle; \varphi_b^m)$$

*Then*

$$\exists G' = d_a \rightarrow d_b : \widetilde{c \langle T \mapsto G \rangle_N} . \Pi G' \upharpoonright d_i \equiv \Pi \varphi_i$$

*Proof.* Similar to Lemma C.1. □

### C.3 Valid States Safety Lemma

Lemma 3.16

$$[\Sigma] \Phi \in \Delta \Rightarrow \text{error} \notin \Phi$$

*Proof.* By induction over  $n$ .

**Case**  $n = 0$

$$\Delta_0 \stackrel{\text{def}}{=} \{[\Sigma_\emptyset] \Phi \mid \Phi \equiv \Pi G \upharpoonright d_i\}$$

By induction over  $G$

**Case**  $G = 0, X$

Trivial.

**Case**  $G = \mu X.G$

By induction we have that:

$$\Phi' = \Pi \varphi_i \equiv \Pi G \upharpoonright d_i \Rightarrow \text{error} \notin \varphi$$

Hence:

$$\Phi = \Pi \mu \underline{X}. \varphi_i \equiv \Pi \mu \underline{X}. G \upharpoonright d_i \Rightarrow \text{error} \notin \varphi'$$

**Case**  $G = d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle$

$$\begin{aligned} d_1 \rightarrow d_2 : r \langle \widetilde{T \mapsto G} \rangle_M \upharpoonright d_1 &\stackrel{\text{def}}{=} \oplus_N ((c! \langle T_n \rangle, \text{UNIT}); G_n \upharpoonright d_1 \& (c! \langle T_n \rangle, l_\emptyset); G_{l_\emptyset} \upharpoonright d_1) \\ &\quad \emptyset \neq N \cup l_\emptyset \mapsto G_{l_\emptyset} \subseteq M \\ d_1 \rightarrow d_2 : r \langle \widetilde{T \mapsto G} \rangle \upharpoonright d_2 &\stackrel{\text{def}}{=} \&_M (c? (T_m); G_m \upharpoonright d_1) \\ d_1 \rightarrow d_2 : r \langle \widetilde{T \mapsto G} \rangle \upharpoonright d &\stackrel{\text{def}}{=} G_1 \upharpoonright d \end{aligned}$$

By induction we know that:

$$\Phi' \Pi \varphi'_i \equiv \Pi G_m \upharpoonright d_i \Rightarrow \text{error} \notin \varphi'$$

Hence:

$$\varphi_i^a \equiv \oplus_N ((c! \langle T_n \rangle, \text{UNIT}); G_n \upharpoonright d_1 \& (c! \langle T_n \rangle, l_\emptyset); G_{l_\emptyset} \upharpoonright d_1) \Rightarrow \text{error} \notin \varphi_i^a$$

$$\varphi_i^n \equiv \&_M (c? (T_m); G_m \upharpoonright d_1) \Rightarrow \text{error} \notin \varphi_i^n$$

Therefore we have that:

$$\Phi \equiv \Pi d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle \upharpoonright d_i \Rightarrow \text{error} \notin \Phi$$

**Case**  $n > 0$

**Case: type in queue**

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_b^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\delta'_k = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \varphi_m \quad k \in J$$

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^j \parallel \dots \parallel \varphi_a^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\delta'_k = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \parallel \varphi_a^k \parallel \dots \varphi_m \quad k \in J$$

By induction we know that:

$$[\Sigma] \Phi' = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_b^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n \Rightarrow \text{error} \notin \Phi'$$

Hence we have that:

$$\text{error} \notin \&_J(c?(T_j); \varphi_b^j)$$

Therefore we can show that:

$$[\_] \Phi = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \varphi_m \Rightarrow \text{error} \notin \Phi$$

**Case: type hole in queue**

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\Sigma[c \mapsto q_{[]} \_] \quad \delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^j \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\begin{aligned}\delta'_k = & [c \mapsto []; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^{l_{\square}} \parallel \dots \parallel \\ & \oplus_N((c!(\langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c!(\langle T_k \rangle, l_{\square}); \varphi_a^{l_{\square}}) \parallel \dots \parallel \varphi_m) \quad k \in J\end{aligned}$$

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\Sigma[c \mapsto q_{\square}] \quad \delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_b^j \parallel \dots \parallel \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\begin{aligned}\delta'_k = & [c \mapsto []; \Sigma] \varphi_1 \parallel \dots \parallel \oplus_N((c!(\langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c!(\langle T_k \rangle, l_{\square}); \varphi_a^{l_{\square}}) \parallel \dots \parallel \\ & \varphi_b^{l_{\square}} \parallel \dots \parallel \varphi_m) \quad k \in J\end{aligned}$$

By induction we know that:

$$[\Sigma] \Phi' = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_b^j \parallel \dots \parallel \varphi_m \quad \delta_j \in \Delta_n \Rightarrow \text{error} \notin \Phi'$$

Hence we have that:

$$\text{error} \notin \oplus_N((c!(\langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c!(\langle T_k \rangle, l_{\square}); \varphi_a^{l_{\square}})$$

Therefore we can show that:

$$\begin{aligned}[-] \Phi = & [c \mapsto []; \Sigma] \varphi_1 \parallel \dots \parallel \oplus_N((c!(\langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c!(\langle T_k \rangle, l_{\square}); \varphi_a^{l_{\square}}) \parallel \dots \parallel \varphi_b^{l_{\square}} \parallel \dots \parallel \varphi_m \\ & \Rightarrow \text{error} \notin \Phi\end{aligned}$$

□

## C.4 Valid States Liveness Lemma

Lemma 3.17

$$[\Sigma] \Phi \in \Delta \Rightarrow [\Sigma] \Phi \rightarrow [\Sigma'] \Phi' \vee \Phi \equiv \Pi \epsilon \vee \Phi \equiv \Pi \underline{X}$$

*Proof.* By induction over  $n$ .

**Case**  $n = 0$

By case analysis of  $G$ .

If  $[\Sigma_\emptyset] \Phi \in \Delta_0$  and  $\Phi \equiv \Phi'$  then  $[\Sigma_\emptyset] \Phi' \in \Delta_0$  and hence it is sufficient to consider direct reduction.

**Case**  $G = d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle$

If

$$d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle \upharpoonright d_a \equiv \oplus_N((c! \langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c! \langle T_k \rangle, l_\square); \varphi_a^{l_\square})$$

then:

$$\begin{aligned} [\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel \oplus_N((c! \langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c! \langle T_k \rangle, l_\square); \varphi_a^{l_\square}) \parallel \dots \parallel \varphi_o \rightarrow \\ [\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel ((c! \langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c! \langle T_k \rangle, l_\square); \varphi_a^{l_\square}) \parallel \dots \parallel \varphi_o \end{aligned}$$

If

$$d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle \upharpoonright d_a \equiv ((c! \langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c! \langle T_k \rangle, l_\square); \varphi_a^{l_\square})$$

then:

$$\begin{aligned} [\Sigma_\emptyset] \varphi_1 \parallel \dots \parallel ((c! \langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c! \langle T_k \rangle, l_\square); \varphi_a^{l_\square}) \parallel \dots \parallel \varphi_o \rightarrow \\ [c \mapsto T_n] \varphi_1 \parallel \dots \parallel G_n \upharpoonright d_a \parallel \dots \parallel \varphi_o \end{aligned}$$

**Case**  $G = 0$

$$\Phi \equiv \Pi \epsilon$$

**Case**  $G = X$

$$\Phi \equiv \Pi \underline{X}$$

**Case**  $n > 0$

**Case: type in queue**

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_b^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\delta'_k = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \varphi_m \quad k \in J$$

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^j \parallel \dots \parallel \varphi_a^j \parallel \dots \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\delta'_k = [c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \parallel \varphi_a^k \parallel \dots \varphi_m \quad k \in J$$

We can show that:

$$[c \mapsto T_k; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \&_J(c?(T_j); \varphi_b^j) \parallel \dots \varphi_m \rightarrow [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^k \parallel \dots \parallel \varphi_b^j \parallel \dots \varphi_m$$

**Case: type hole in queue**

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\Sigma[c \mapsto q_\square] \quad \delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^j \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\begin{aligned} \delta'_k = & [c \mapsto []; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^{l_\square} \parallel \dots \parallel \\ & \oplus_N((c!(\langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c!(\langle T_k \rangle, l_\square); \varphi_a^{l_\square}) \parallel \dots \parallel \varphi_m) \quad k \in J \end{aligned}$$

$\forall \varphi_i, c = d_a d_b \forall j \in J$  choose some  $\varphi_a^j, \varphi_b, \Sigma$  such that:

$$\Sigma[c \mapsto q_\square] \quad \delta_j = [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_a^j \parallel \dots \parallel \varphi_b^j \parallel \dots \parallel \varphi_m \quad \delta_j \in \Delta_n$$

then  $\delta'_k \in \Delta_{n+1}$  where:

$$\begin{aligned} \delta'_k = & [c \mapsto []; \Sigma] \varphi_1 \parallel \dots \parallel \oplus_N((c!(\langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c!(\langle T_k \rangle, l_\square); \varphi_a^{l_\square}) \parallel \dots \parallel \\ & \varphi_b^{l_\square} \parallel \dots \parallel \varphi_m) \quad k \in J \end{aligned}$$

We can show that:

$$\begin{aligned} [c \mapsto []; \Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^{l_{[]}^1} \parallel \dots \parallel \oplus_N((c! \langle T_k \rangle, \text{UNIT}); \varphi_a^k) \& ((c! \langle T_k \rangle, l_{[]}^1); \varphi_a^{l_{[]}^2}) \parallel \dots \parallel \varphi_m \\ \rightarrow [\Sigma] \varphi_1 \parallel \dots \parallel \varphi_b^{l_{[]}^1} \parallel \dots \parallel \varphi_a^{l_{[]}^2} \parallel \dots \parallel \varphi_m \end{aligned}$$

□



## Appendix D

# Global Typability and Local Update

In this chapter we provide the auxiliary proofs used to show that Global Typability and Local Update imply General Update Compatibility of dynamic programs.

## D.1 Participants Relation as Strict Partial Order

**Lemma D.1.**  $<_G$  is a partial order.

*Proof.* A partial order is anti-symmetric and transitive. We have transitivity straightforwardly from the rules. We then need to prove anti-symmetry: that  $d_1 <_G d_2$  implies that  $d_2 <_G d_1$  is not the case. We proceed by induction on the last derivation rule used for  $d_1 <_G d_2$ .

Consider the case where the rule used is that for  $f(G')$ ,  $\mu \underline{X}.G'$ , or  $\text{dmod}_F(G')$ . The conclusions then follow directly by application of the induction hypothesis.

Consider the case where the last rule used is:

$$\frac{d_1 <_{G_k} d_2 \quad d_2 \neq d, d'}{d_1 <_{d \rightarrow d':c(\widetilde{T \mapsto G})} d_2}$$

By induction we know that it is not the case that  $d_2 <_{G_k} d_1$ , for some  $G_k \in \widetilde{G}$ . Suppose for contradiction that  $d_2 <_G d_1$  also holds. In that case, we also have  $d_2 <_{G'_k} d_1$  for some  $G'_k \in \widetilde{G}$ . We know that  $G_k \upharpoonright d_2 = G'_k \upharpoonright d_2$  and by Lemma D.2 we also have  $d_2 <_{G_k} d_1$ . This contradicts the fact that  $d_2 <_{G_k} d_1$  does not hold.

Consider the case where the last rule used is:

$$\overline{d_1 <_{d_1 \rightarrow d_2:c(\widetilde{T \mapsto G})} d_2}$$

as the only possible rule that could derive  $d_2 <_G d_1$  is an application of

$$\frac{d_1 <_{G_k} d_2 \quad d_2 \neq d, d'}{d_1 <_{d \rightarrow d':c(\widetilde{T \mapsto G})} d_2}$$

on some  $d' <_{G_k} d_1$  (to allow for transitivity) for some  $G_k \in \widetilde{G}$ . The side-condition  $d_1 \neq d, d'$  cannot, however, hold in this case, so this is not possible.  $\square$

## D.2 Equivalent Projection Give Equivalent Orderings

**Lemma D.2.** *If  $d_1 <_G d_2$  and  $G \upharpoonright d_2 \equiv G' \upharpoonright d_2$  then  $d_1 <_{G'} d_2$*

*Proof.* By lexicographic induction on the structure of  $G$  then the structure of  $G'$ .

Consider the cases where  $G$  is 0 or  $\underline{X}$ . There are no derivation rule for  $d_1 <_G d_2$  for such  $G$ , and hence we can disregard these cases.

Consider the cases where  $G$  is  $\mu \underline{X}. G_1$ ,  $f(G_1)$ ,  $\text{dmod}_F(G_1)$ . Each case follows directly by induction.

Consider the case where  $G$  is  $d \rightarrow d' : c \langle \widetilde{T \mapsto G} \rangle$ . We proceed by a case split on the last rule used to derive  $d_1 <_G d_2$ .

If the last rule used is

$$\overline{d_1 <_{d_1 \rightarrow d_2 : c \langle \widetilde{T \mapsto G} \rangle} d_2}$$

Recall that here  $d = d_1$  and  $d' = d_2$ . We then proceed by a case split over the structure of  $G'$ . In the case where  $G' = \mu \underline{X}. G$ , 0 or  $\underline{X}$ , then we have a contradiction, as  $G' \upharpoonright d_2 \not\equiv G \upharpoonright d_2$ . Consider the case where  $G' = d_3 \rightarrow d_4 : c' \langle \widetilde{T' \mapsto G'} \rangle$ . If we have that  $d_4 = d_2$ , then by the hypotheses we have that  $G' \upharpoonright d_2 = G \upharpoonright d_2$ . By principle channel allocation, and the fact that  $c = d_1, d_2$  we also know that  $d_1 = d_3$ . Then by

$$\overline{d_1 <_{d_1 \rightarrow d_2 : c \langle \widetilde{T' \mapsto G'} \rangle} d_2}$$

we have that  $d_1 <'_G d_2$ . If we have that  $d_4 \neq d_2$ , then by the definition of projection we know that  $G \upharpoonright d_2 = G'_k \upharpoonright d_2$ , for some, for some  $G_k \in \widetilde{G}$  and for some  $G'_k \in \widetilde{G'}$ . Then, by lexicographic induction we then can show that  $d_1 <_{G'_k} d_2$ , and hence that  $d_1 <_{G'} d_2$ .

If the last rule used is

$$\frac{d_1 <_{G_k} d_2 \quad d_2 \neq d, d'}{d_1 <_{d \rightarrow d' : c \langle \widetilde{T \mapsto G} \rangle} d_2}$$

Here we know that  $d_2 \neq d, d'$ . By the definition of projection we know that  $G \upharpoonright d_2 = G_k \upharpoonright d_2$ , for some  $G_k \in \widetilde{G}$ . By applying the inductive hypothesis we know that  $d_1 <_{G'} d_2$ . as required.

□

### D.3 Global Session Type Updatability Lemma

Lemma 5.4

If

$$\Phi \equiv \Pi_I G \upharpoonright d_i \quad u = \widetilde{f_i \mapsto t_i}_I \quad \varphi_i \wr \emptyset \vdash t_i : \text{UNIT} \quad G' \upharpoonright d_i \equiv \varphi_i$$

Then

$$\text{upd}(\Phi, u, b) \equiv \Pi_I G'' \upharpoonright d_i$$

*Proof.* By induction over  $G$ .

**Case**  $G = 0$

By the definition of  $\upharpoonright$  we have that:

$$0 \upharpoonright d = \epsilon$$

Hence we have that:

$$\Phi \equiv \Pi_I \epsilon$$

By Lemma A.4 and the definition of  $\text{upd}$  we know that:

$$\text{upd}(\Phi, u, b) \equiv \Pi_I \epsilon$$

Hence we can say that:

$$\text{upd}(\Phi, u, b) \equiv \Pi_I 0 \upharpoonright d_i$$

**Case**  $G = X$

Similar to case:  $G = 0$ .

**Case**  $G = d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle$

By the definition of  $\upharpoonright$  we know that:

$$\Pi_I d_a \rightarrow d_b : c \langle \widetilde{T \mapsto G} \rangle \upharpoonright d_i \Pi \varphi_i$$

where:

$$\varphi_a \equiv \oplus(\text{snd}(\text{Res } r, T_1), \text{UNIT}); G_j \upharpoonright d_1 \quad \varphi_b \equiv \&(\text{rcv}(\text{Res } r), T_i); G_j \upharpoonright d_2$$

By Lemma A.4 and the definition of  $\text{upd}$  we know that:

$$\text{upd}((\alpha(\tilde{v}), T); \varphi, u, b) \equiv (\alpha(\tilde{v}), T); \text{upd}(\varphi, u, b)$$

By the well formedness of Global Session Types we know that:

$$G \upharpoonright d_i \equiv G_j \upharpoonright d_i$$

where  $i \in I \setminus a, b$ . Using the inductive hypothesis we have that:

$$\Phi' \equiv \Pi_I G_j \upharpoonright d_i \Rightarrow \text{upd}(\Phi', u, b) \equiv \Pi_I G'_j \upharpoonright d_i$$

Hence we have that:

$$\text{upd}(\Pi_I d_a \rightarrow d_b : c\langle T \mapsto \widetilde{G_j} \rangle \upharpoonright d_i, u, b) \equiv \Pi_I d_a \rightarrow d_b : c\langle T \mapsto \widetilde{G'_j} \rangle \upharpoonright d_i$$

**Case**  $G = \mu X.G_1$

By the definition we know that:

$$\mu \underline{X}.G_1 \upharpoonright d \equiv \mu \underline{X}.G_1 \upharpoonright d$$

By induction we know that:

$$\Pi \text{upd}(G_1 \upharpoonright d, u, b) \equiv \Pi G''' \upharpoonright d$$

By Lemma A.4 and the definition of  $\text{upd}$  we know that:

$$\Pi \text{upd}(\mu \underline{X}.G_1 \upharpoonright d, u, b) \equiv \Pi \mu \underline{X}.G''' \upharpoonright d$$

**Case**  $G = f(G_1)$

**Case**  $b = \text{true}$

By the definition of  $\upharpoonright$  we know that:

$$\Pi f(G_1) \upharpoonright d \equiv \Pi f_d(G_1 \upharpoonright d)$$

By Lemma A.4 and the definition of  $\text{upd}$  we know that:

$$\begin{aligned} \Pi \text{upd}(f(G_1) \upharpoonright d, u, \text{true}) &\equiv \Pi \text{upd}(\varphi'_i, u, \text{true}) \\ &\equiv \Pi \text{upd}(G' \upharpoonright d, u, \text{true}) \end{aligned}$$

**Case  $b = \text{false}$**

By Lemma A.4 and the definition of  $\text{upd}$  we know that:

$$\Pi \text{upd}(f(G_1) \upharpoonright d, u, \text{false}) \equiv \Pi f_d(\text{upd}(G_1 \upharpoonright d, u, \text{false}))$$

By induction we know that:

$$\Pi \text{upd}(G_1 \upharpoonright d, u, \text{false}) \equiv \Pi G''' \upharpoonright d$$

Hence we have that:

$$\Pi \text{upd}(f(G_1) \upharpoonright d, u, \text{false}) \equiv \Pi f_d(\text{upd}(G''' \upharpoonright d, u, \text{false}))$$

$$G'' = f(G''')$$

**Case  $G = \text{dmod}_F(G_1)$**

By Lemma A.4 and the definition of  $\text{upd}$  we know that:

$$\Pi \text{upd}(\text{dmod}_F(G_1) \upharpoonright d, u, b) \equiv \Pi \text{dmod}_{F_d}(\text{upd}(G_1 \upharpoonright d, u, \text{true}))$$

where  $\text{regions}(G_1) \cap F = \emptyset$  and

$$\Pi \text{upd}(\text{dmod}_F(G_1) \upharpoonright d, u, b) \equiv \Pi \text{dmod}_{F_d}(\text{upd}(G_1 \upharpoonright d, u, \text{false}))$$

where  $\text{regions}(G_1) \cap F \neq \emptyset$ . By induction we know that:

$$\Pi \text{upd}(G_1 \upharpoonright d, u, b) \equiv \Pi G''' \upharpoonright d$$

Hence we have that:

$$\Pi \text{upd}(\text{dmod}_F(G_1) \upharpoonright d, u, b) \equiv \Pi (\text{dmod}_F(G''')) \upharpoonright d$$

□

# References

Ajmani, S. (2004). Automatic Software Upgrades for Distributed Systems. *Artificial Intelligence*, (Cm):1–23.

Ajmani, S., Liskov, B., and Shrira, L. (2006). Modular Software Upgrades for Distributed Systems. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, pages 452–476. Massachusetts Institute of Technology.

Altekar, G., Bagrak, I., Burstein, P., and Schultz, A. (2005). OPUS : Online Patches and Updates for Security. In *Proceedings of the 14th USENIX Security Symposium*, pages 287–302, Baltimore, Maryland, USA. USENIX Association.

Anderson, G. (2011). Dynamic Software Update for Behavioural Properties of Concurrent Programs. In *Proceedings of the 11th Grace Hopper Celebration of Women in Computing*, Portland, Oregon, USA.

Anderson, G. and Rathke, J. (2009). Migrating Protocols in Multi-threaded Message-passing Systems. In *Proceedings of the 2nd International Workshop on Hot Topics in Software Upgrades*, pages 8:1–8:5, Orlando, Florida, USA. ACM.

Anderson, G. and Rathke, J. (2011). Resource Access with Variably Typed Return. In *Proceedings of the 4th International Workshop on Programming Language Approaches to Concurrency and Communication-cCentric Software*, pages 51–57, Saarbrucken.

Anderson, G. and Rathke, J. (2012). Dynamic Software Update for Message Passing Programs. In *The Tenth International Asian Symposium on Programming Languages and Systems*, page (to appear), Kyoto, Japan. LNCS.

Appel, A. (1994). Hot-sliding in ML. Technical report, Princeton University, Princeton, New Jersey, USA.

Arthur, L. (1988). *Software Evolution: the Software Maintenance Challenge*. Wiley-Interscience, New York, New York, USA.

Banker, R., Datar, S., Kemerer, C., and Zweig, D. (1993). Software Complexity and Maintenance Costs. *Communications of the ACM*, 36(11):81–94.

Bartoletti, M., Degano, P., Ferrari, G., and Zunino, R. (2007). Types and effects for resource usage analysis. In *In: Proceedings of the 10th international conference on Foundations of software science and computational structures*, volume vol, pages 4423pp32–47.

Baumann, A., Appavoo, J., Wisniewski, R., Silva, D. D., Krieger, O., and Heiser, G. (2007). Reboots are for Hardware: Challenges and Solutions to Updating an Operating System on the Fly. In *Proceedings of the 18th USENIX Annual Technical Conference*, pages 26:1–26:14, Santa Clara, California, USA. USENIX Association.

Baumann, A., Heiser, G., Appavoo, J., Silva, D. D., Krieger, O., Wisniewski, R. W., and Kerr, J. (2005). Providing dynamic update in an operating system. pages 32–32, Anaheim, CA. USENIX Association.

Bettini, L., Coppo, M., Dezani-Ciancaglini, M., Luca, M. D., and Yoshida, N. (2008). Global Progress in Dynamically Interleaved Multiparty Sessions. In *Proceedings of the 19th International Conference on Concurrency Theory*, pages 418–433, Toronto, Ontario, Canada. Springer-Verlag.

Bierman, G., Hicks, M., Sewell, P., and Stoyle, G. (2003). Formalizing Dynamic Software Updating. In *Proceedings of the 2nd International Workshop on Unanticipated Software Evolution*, Warsaw, Poland.

Bierman, G., Parkinson, M., and Noble, J. (2008). UpgradeJ : Incremental Typechecking for Class Upgrades. In *Proceedings of the 22nd European Conference on Object-Oriented Programming*, pages 235–259, Paphos, Cypress. Springer-Verlag.

Boyapati, C., Liskov, B., Shrira, L., Moh, C.-h., and Richman, S. (2003). Lazy Modular Upgrades In Persistent Object Stores. In *Proceedings of the 18th Annual ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications*, pages 403–417, Anaheim, California, USA. ACM.

Carbone, M., Honda, K., and Yoshida, N. (2007). Structured Communication-Centred Programming for Web Services. In *Proceedings of the 16th European Conference on Programming*, pages 2–17, Braga, Portugal. Springer-Verlag.

Cardelli, L. (1997). Program Fragments, Linking, and Modularization. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 266–277, Paris, France. ACM.

Cook, R. and Lee, I. (1983). DYMOS: a Dynamic Modification System. *SIGSOFT Software Engineering Notes*, 8(4):201–202.

Deniélou, P.-m. and Yoshida, N. (2010). Buffered Communication Analysis in Distributed Multiparty Sessions. In *Proceedings of the 21st International Conference on Concurrency Theory*, pages 343–357, Paris, France. Springer-Verlag.

Dezani-Ciancaglini, M. and Liguoro, U. D. (2010). Sessions and Session Types : an Overview. In *Proceedings of the 6th International Conference on Web Services and Formal Methods*, pages 1–28, Bologna, Italy. Springer-Verlag.

Dmitriev, M. (2001). Towards Flexible and Safe Technology for Runtime Evolution of Java Language Applications. In *Proceedings of the Workshop on Engineering Complex Object-Oriented Systems for Evolution, in association with OOPSLA 2001 International Conference*, Tampa, Florida, USA.

Drossopoulou, S. and Eisenbach, S. (2002). Manifestations of Java Dynamic Linking - an Approximate Understanding at Source Language Level. In *Proceedings of the 1st International Workshop on Unanticipated Software Evolution*, Málaga, Spain. Springer-Verlag.

Duggan, D. (2005). Type-based Hot Swapping of Running Modules. In *Proceedings of the 6th ACM SIGPLAN International Conference on Functional Programming*, volume 41, pages 181–220, Florence, Italy. ACM.

Eaddy, M. and Feiner, S. (2005). Multi-Language Edit-and-Continue for the Masses. Technical report, Columbia University, New York, New York, USA.

Erlang (2010). Appup Cookbook - Typical Cases of Upgrades/Downgrades Done in Runtime. [http://www.erlang.org/doc/design\\_principles/appup\\_cookbook.html](http://www.erlang.org/doc/design_principles/appup_cookbook.html).

Flanagan, C. and Freund, S. (2000). Type-based Race Detection for Java. In *Proceedings of the ACM SIGPLAN 2000 Conference on Programming Language Design and Implementation*, volume 35, pages 219–232, Vancouver, British Columbia, Canada. ACM.

Freeman, P. (1986). Programming Without Tears. *High Technology*, 6(4):38–45.

Gay, S. and Hole, M. (2005). Subtyping for session types in the pi calculus. *Acta Informatica*, 42(2):191–225.

Gay, S. and Vasconcelos, V. (2007). Asynchronous Functional Session Types. Technical report 2007–251, Department of Computing, University of Glasgow, Glasgow, Scotland, UK.

Gay, S. and Vasconcelos, V. (2010). Linear type theory for asynchronous session types. *Journal of Functional Programming*, 20:19–50.

Gay, S., Vasconcelos, V., and Ravara, A. (2003). Session Types for Inter-Process Communication. Technical report, University of Glasgow, Glasgow, Scotland, UK.

Gifford, D. K. and Lucassen, J. M. (1986). Integrating Functional and Imperative Programming. In *Proceedings of the ACM Conference on LISP and Functional Programming*, pages 28–38, Cambridge, Massachusetts, USA. ACM.

Gilmore, S., Kirli, D., and Walton, C. (1998). Dynamic ML without Dynamic Types. Technical report.

Gregersen, A. and Jorgensen, B. (2009). Dynamic Update of Java Applications - Balancing Change Flexibility vs Programming Transparency. *Journal of Software Maintenance and Evolution: Research and Practice*, 21(2):81–112.

Gregersen, A. and Jorgensen, B. (2011). Run-time Phenomena in Dynamic Software Updating : Causes and Effects. In *Proceedings of the 12th International Workshop on Principles on Software Evolution*, pages 6–15, Szeged, Hungary. ACM.

Gupta, D., Jalote, P., and Barua, G. (1996). A Formal Framework For Online Software Version Change. *IEEE Transactions on Software Engineering*, 22(2):120–131.

Gurtner, D. (2011). *Safe Dynamic Software Updates in Multi-Threaded Systems with ActiveContext*. PhD thesis, University of Bern.

Hayden, C., Hardisty, E., Hicks, M., and Foster, J. (2009). Efficient Systematic Testing For Dynamically Updatable Software. In *Proceedings of the 2nd International Workshop on Hot Topics in Software Upgrades*, pages 9:1–9:5, New York, New York, USA. ACM.

Hayden, C., Magill, S., Hicks, M., Foster, N., and Foster, J. (2011a). Specifying and Verifying the Correctness of Dynamic Software Updates. In *Proceedings of the International Conference on Verified Software: Theories, Tools, and Experiments*, Philadelphia, Pennsylvania, USA. Springer-Verlag.

Hayden, C., Smith, E. K., Denchev, M., Hicks, M., and Foster, J. (2011b). Kitsune : Efficient , General-purpose Dynamic Software Updating for C.

Hayden, C. M., Smith, E. K., Hardisty, E. A., Hicks, M., and Foster, J. S. (2011c). Evaluating Dynamic Software Update Safety Using Systematic Testing. *IEEE Transactions on Software Engineering*, pages 1–41.

Hayden, C. M., Smith, E. K., Hicks, M., and Foster, J. S. (2011d). State Transfer for Clear and Efficient Runtime Upgrades. In *Proceedings of the 3rd International Workshop on Hot Topics in Software Upgrades*, Hannover, Germany. IEEE Press.

Hicks, M. (2005). *Dynamic Software Updating*. PhD thesis, University of Pennsylvania.

Hjálmtýsson, G. and Gray, R. (1998). Dynamic C++ Classes: A Lightweight Mechanism to Update Code in a Running Program. In *Proceedings of the 9th USENIX Annual Technical Conference*, pages 6–6, New Orleans, Louisiana, USA. USENIX Association.

Honda, K. (1993). Types for Dyadic Interaction. In *Proceedings of the 4th International Conference on Concurrency Theory*, number CONCUR’93, pages 509–523, Hildesheim, Germany. Springer-Verlag.

Honda, K., Kubo, M., and Vasconcelos, V. (1998). Language Primitives and Type Discipline for Structured Communication-Based Programming. In *Proceedings of the 7th European Symposium on Programming Languages and Systems*, volume 171, pages 122–138, Lisbon, Portugal. Springer-Verlag.

Honda, K., Yoshida, N., and Carbone, M. (2008). Multiparty Asynchronous Session Types.

Kaashoek, F. and Arnold, J. (2009). Ksplice: Automatic Rebootless Kernel Updates. In *Proceedings of the 4th ACM European Conference on Computer Systems*, pages 187–198, Nuremberg, Germany. ACM.

Lea, D. (1999). *Concurrent Programming in Java: Design Principles and Patterns*. Addison-Wesley Longman Publishing Co., Inc., 2nd edition.

Makris, K. and Bazzi, R. (2009). Immediate Multi-Threaded Dynamic Software Updates Using Stack Reconstruction. In *Proceedings of the 19th USENIX Annual technical conference*, pages 31–44, San Diego, California, United States. USENIX Association.

Marino, D. and Millstein, T. (2008). A Generic Type-and-effect System. In *Proceedings of the 4th International Workshop on Types in Language Design and Implementation*, pages 39–50, Savannah, Georgia, USA. ACM.

Mostrou, D., Yoshida, N., and Honda, K. (2009). Global Principal Typing in Partially Commutative Asynchronous Sessions. In *Proceedings of the 18th European Symposium on Programming Languages and Systems*, pages 316–322, York, UK. Springer-Verlag.

Neamtiu, I., Foster, J., and Hicks, M. (2005). Understanding Source Code Evolution Using Abstract Syntax Tree Matching. *SIGSOFT Software Engineering Notes*, 30(4):1–5.

Neamtiu, I. and Hicks, M. (2009). Safe and Timely Dynamic Updates to Multi-threaded Programs. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 13–24, Dublin, Ireland. ACM.

Neamtiu, I., Hicks, M., Foster, J., and Pratikakis, P. (2008). Contextual Effects For Version-Consistent Dynamic Software Updating And Safe Concurrent Programming. In *Proceedings of the 35th annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 37–49, San Francisco, California, USA. ACM.

Neamtiu, I., Hicks, M., Stoyle, G., and Oriol, M. (2006). Practical Dynamic Software Updating for C. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 72–83, Ottawa, Ontario, Canada. ACM.

Nicoara, A., Alonso, G., and Roscoe, T. (2008). Controlled, systematic, and efficient code replacement for running java programs. *Proceedings of the 3rd ACM SIGOPS/EuroSys European Conference on Computer Systems 2008 - Eurosyst '08*, (5005):233.

Nielson, F. and Nielson, H. R. (1999). Type and Effect Systems. pages 114–136. Springer-Verlag.

Orso, A., Rao, A., and Harrold, M.-J. (2002). A Technique for Dynamic Updating of Java Software. In *Proceedings of the International Conference on Software Maintenance*, pages 649–658, Montréal, Canada. IEEE Computer Society.

Scott, D. (1998). Assessing the Costs of Application Downtime.

Sewell, P. (2001). Modules, Abstract Types, and Distributed Versioning. In *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 236–247, London, UK. ACM.

Soules, C. A. N., Appavoo, J., Hui, K., Wisniewski, R. W., Silva, D. D., Auslander, M., Krieger, O., Stumm, M., Ganger, G. R., Ostrowski, M., Rosenburg, B., and Xenidis, J. (2003). System Support for Online Reconfiguration. *Online*.

Stewart, D. and Chakravarty, M. (2005). Dynamic Applications from the Ground Up. In *Proceedings of the 9th ACM SIGPLAN Workshop on Haskell*, pages 27–38, Tallinn, Estonia. ACM.

Stoyle, G. (2006). *A Theory of Dynamic Software Updates*. Phd, University of Cambridge.

Stoyle, G., Hicks, M., Bierman, G., Sewell, P., and Neamtiu, I. (2005). Mutatis Mutandis: Safe and predictable dynamic software updating. *ACM Trans. Program. Lang. Syst.*, 29(4):183–194.

Subramanian, S. (2010). *Dynamic Software Updates : A VM-Centric Approach*. PhD thesis, University of Texas at Austin.

Subramanian, S., Hicks, M., and Mckinley, K. (2009). Dynamic Software Updates : A VM-centric Approach. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 1–12, Dublin, Ireland. ACM.

Sun-Microsystems (2006). Java Request for Enhancement: Hot Code swapping. [http://bugs.sun.com/view\\_bug.do?bug\\_id=4910812](http://bugs.sun.com/view_bug.do?bug_id=4910812).

Tempero, E., Bierman, G., Noble, J., and Parkinson, M. (2008). From Java to UpgradeJ: an empirical study. pages 1–5, Nashville, Tennessee. ACM.

Thomas, W., Wimmer, C., and Stadler, L. (2010). Dynamic Code Evolution for Java. In *Proceedings of the 8th International Conference on the Principles and Practice of Programming in Java*, pages 10–19, Vienna, Austria. ACM.

Wernli, E., Gurtner, D., and Nierstrasz, O. (2011). Using First-class Contexts to realize Dynamic Software Updates. In *Proceedings of the 3rd International Workshop on Smalltalk Technologies*, pages 21–31, Edinburgh, Scotland. ACM.

Wikstrom, R. and Williams, M. (1993). *Concurrent Programming in Erlang*. Prentice Hall International (UK) Ltd., 2nd edition.

Yoshida, N. and Vasconcelos, V. (2006). Language Primitives and Type Discipline for Structured Communication-Based Programming Revisited : Two Systems for Higher-Order Session Communication. *Electronic Notes in Theoretical Computer Science*, 171(4):73–93.

ZeroTurnaround (2011). JRebel. <http://zeroturnaround.com/jrebel/>.