On Confidentiality and Formal Methods

Michael James Banks

Submitted for the degree of Doctor of Philosophy

The University of York Department of Computer Science

March 2012

Abstract

The contemporary challenge of engineering verifiably secure software has motivated various techniques for measuring and regulating the flow of confidential data from systems to their users. Unfortunately, these techniques suffer from a lack of integration with modern formal methods for software development, which inhibits their application in practice.

This thesis proposes a novel approach for integrating information flow security concerns with formal methods. Working in the Unifying Theories of Programming (UTP), this thesis presents a generic framework for modelling interactions between users and systems. This framework can be applied to encode information flow about a system's activities to its users. It thereby allows confidentiality properties to be formalised in the UTP as upper bounds on information flow to users.

The main contribution of this thesis is a unified platform for designing software that is not only functionally correct, but also secure by design. This platform specialises the information flow encoding to the *Circus* formal method, making it possible to specify confidentiality properties within *Circus* processes. In this setting, conflicts between functionality and confidentiality are represented as miracles, rendering insecure functionality infeasible.

The platform provides techniques for verifying that functionality and confidentiality properties are mutually consistent. These techniques can be applied to develop a process through a series of feasibility-preserving refinement steps, to achieve a system implementation that does not leak secret information to untrusted users. These techniques are evaluated with a brief case study.

1	Intro	oduction	1
	1.1	Confidentiality	1
	1.2	Formal Methods and Security	2
	1.3	A Platform for Secure Software Development	4
	1.4	Structure of this Thesis	5
2	Bac	kground	7
	2.1	Unifying Theories of Programming	7
	2.2	Programs as Predicates	8
	2.3	Relations	8
		2.3.1 Refinement	9
		2.3.2 Defining Program Constructs	10
	2.4	Designs	11
	2.5	Reactive Processes	13
	2.6	Reactive Designs	14
	2.7	<i>Circus</i>	16
	2.8	Circus Actions	17
	2.9	The <i>Circus</i> Family	20
3	Mult	i-User Systems	21
	3.1	Introduction	21
		3.1.1 Motivation	21
		3.1.2 Contributions of this Chapter	22
	3.2	Modelling Users	23
		3.2.1 Views	24
		3.2.2 Healthiness Conditions for Views	25

	3.2.3	Conjoining Views	27			
	3.2.4	Functional Views	28			
3.3	Calcu	lating Interactions	29			
	3.3.1	Localisation	29			
	3.3.2	Globalisation	31			
	3.3.3	A Galois Connection	32			
	3.3.4	Inference	33			
3.4	Multi	-User Designs	35			
	3.4.1	Views for Designs	35			
	3.4.2	Projecting Designs	35			
	3.4.3	Worked Example: A Byte Register	37			
3.5	Multi	-User Reactive Designs	40			
	3.5.1	Reactive Views	40			
	3.5.2	Projecting a Reactive Process	43			
	3.5.3	Lazy Abstraction	45			
3.6	Refining Multi-User Systems					
	3.6.1	Local Refinement	48			
	3.6.2	Co-operating and Independent Refinement	49			
3.7	Concl	usion	51			
•	.	l'he Breezeller	50			
		lity Properties	53			
4.1	Introd	luction	53			
	4.1.1	Access Control	53			
	4.1.2	Inference Control	54			
		Contributions of this Chapter	55			
4.2	An In	distinguishability Semantics	56			
	4.2.1	Lifting Relations	56			
	4.2.2	Modelling Inference	58			
4.3	Encod	ling Confidentiality	61			
	4.3.1	Secrets and Cover Stories	61			
	4.3.2	Obligations	62			
	4.3.3	Conformance	63			
	4.3.4	The Lattice of Obligations	64			

4

		4.3.5	Multi-Obligations	65
	4.4	Deriv	ing Secure Specifications	67
		4.4.1	Reconciling Obligations	67
		4.4.2	Fixed Points of Confidentiality	69
		4.4.3	Reconciling Multi-Obligations	71
	4.5	Refine	ement and Confidentiality	71
	4.6	Specif	fying Noninterference Properties	73
		4.6.1	Noninterference	73
		4.6.2	Worked Example: Basic Security Predicates	75
	4.7	Relate	ed Work	78
		4.7.1	The Limitations of Noninterference	78
		4.7.2	Security Specifications	79
		4.7.3	Other Approaches	80
		4.7.4	Refinement and Confidentiality	81
	4.8	Concl	usion	83
5	Spe	cifvina	Confidentiality in <i>Circus</i>	85
	5.1		luction	85
	5.2			-)
		LIIIII	g the Circus Semantics	86
			g the <i>Circus</i> Semantics	86 86
		5.2.1	Lifting <i>Circus</i> Actions	86
		5.2.1 5.2.2	Lifting Circus Actions	86 88
	5.3	5.2.1 5.2.2 5.2.3	Lifting Circus Actions	86 88 91
	5.3	5.2.1 5.2.2 5.2.3 Comp	Lifting Circus Actions	86 88 91 91
	5.3	5.2.1 5.2.2 5.2.3 Comp 5.3.1	Lifting Circus Actions	86 88 91 91 93
	5.3	5.2.1 5.2.2 5.2.3 Comp 5.3.1 5.3.2	Lifting Circus Actions	86 88 91 91 93 94
	5.3	5.2.1 5.2.2 5.2.3 Comp 5.3.1 5.3.2 5.3.3	Lifting Circus Actions	86 88 91 91 93 94 94
	5.3	5.2.1 5.2.2 5.2.3 Comp 5.3.1 5.3.2 5.3.3 5.3.4	Lifting Circus Actions	86 88 91 91 93 94 94 95
	5.3	5.2.1 5.2.2 5.2.3 Comp 5.3.1 5.3.2 5.3.3 5.3.4 5.3.5	Lifting Circus Actions	86 88 91 93 93 94 95 95 96
	5.3	5.2.1 5.2.2 5.2.3 Comp 5.3.1 5.3.2 5.3.3 5.3.4 5.3.5 5.3.6	Lifting Circus Actions	 86 88 91 93 94 94 95 96 97
		5.2.1 5.2.2 5.2.3 Comp 5.3.1 5.3.2 5.3.3 5.3.4 5.3.5 5.3.6 5.3.7	Lifting Circus Actions	 86 88 91 91 93 94 94 95 96 97 99
	5.3 5.4	5.2.1 5.2.2 5.2.3 Comp 5.3.1 5.3.2 5.3.3 5.3.4 5.3.5 5.3.6 5.3.7	Lifting Circus Actions	 86 88 91 93 94 94 95 96 97

	5.4.3	Blocks	102
	5.4.4	On Equivalence Laws	104
5.5	Confi	dentiality Annotations	105
	5.5.1	Defining Confidentiality Annotations	105
	5.5.2	Properties of Confidentiality Annotations	106
	5.5.3	Passive Confidentiality Annotations	108
	5.5.4	Active Confidentiality Annotations	109
	5.5.5	Discussion	110
5.6	Event	-Based Confidentiality Annotations	112
	5.6.1	Inputs and Outputs	113
	5.6.2	Occurrence of Events	113
	5.6.3	Superposition	115
	5.6.4	Discussion	117
5.7	Declas	ssification	117
	5.7.1	Relaxing Confidentiality Annotations	118
	5.7.2	A Declassification Construct	119
5.8	Relate	ed Work	120
5.0	iterate		120
5.9			120
5.9	Concl	usion	122
5.9 Seci	Concl ure Soft	usion	122 123
5.9 Secu 6.1	Concl ure Soft Introd	usion	122 123 123
5.9 Seci	Concl ure Soft Introd Confid	usion	122 123 123 124
5.9 Secu 6.1	Concl ure Soft Introd Confid 6.2.1	usion	122 123 123 124 124
5.9 Secu 6.1	Concl ure Soft Introd Confid 6.2.1 6.2.2	usion	 122 123 123 124 124 126
5.9 Secu 6.1 6.2	Concl ure Soft Introd Confid 6.2.1 6.2.2 6.2.3	usion	122 123 123 124 124
5.9 Secu 6.1	Concl Jre Soft Introd Confid 6.2.1 6.2.2 6.2.3 Backw	usion	 122 123 124 124 126 128 128
5.9 Secu 6.1 6.2	Concl Jre Soft Introd Confid 6.2.1 6.2.2 6.2.3 Backw 6.3.1	usion	 122 123 124 124 126 128 128 129
5.9 Secu 6.1 6.2	Concl Jre Soft Introd Confid 6.2.1 6.2.2 6.2.3 Backw 6.3.1 6.3.2	usion	 122 123 124 124 126 128 128 129 129
5.9 Secu 6.1 6.2	Concl Jre Soft Introd Confid 6.2.1 6.2.2 6.2.3 Backw 6.3.1	usion	 122 123 124 124 126 128 128 129 132
5.9 Secu 6.1 6.2	Concl Jre Soft Introd Confid 6.2.1 6.2.2 6.2.3 Backw 6.3.1 6.3.2 6.3.3 6.3.4	usion	122 123 124 124 126 128 128 129 129 132 134
5.9 Secu 6.1 6.2	Concl Jre Soft Introd Confid 6.2.1 6.2.2 6.2.3 Backw 6.3.1 6.3.2 6.3.3 6.3.4	usion	122 123 124 124 126 128 129 129 132 134 135
5.9 Secu 6.1 6.2	Concl Jre Soft Introd Confid 6.2.1 6.2.2 6.2.3 Backw 6.3.1 6.3.2 6.3.3 6.3.4 Laws	usion	122 123 124 124 126 128 128 129 129 132 134

6

	6.5	Verify	ving Consistency	141
		6.5.1	Propagated Confidentiality Annotations	142
		6.5.2	Left Justification	143
		6.5.3	Full Justification	149
		6.5.4	Resolving Inconsistency	149
		6.5.5	Discussion	150
	6.6	A Sec	cure Refinement Strategy	151
		6.6.1	Action Refinement	152
		6.6.2	Data Refinement	153
		6.6.3	Process Refinement	156
	6.7	Imple	ementation	157
		6.7.1	Translating to Circus	157
		6.7.2	Compilation to Code	158
	6.8	Concl	lusion	160
7	Case	e Study	,	163
	7.1	Introc	duction	163
	7.2	A Sea	led-Bid Auction	164
	7.3	Mode	elling Contestants' Inferences	164
		7.3.1	Declaration	166
		7.3.2	Submission	167
		7.3.3	Initialisation	169
	7.4	Specif	fying Confidentiality Properties	171
		7.4.1	Comparing Bids	171
		7.4.2	Value of Bids	172
		7.4.3	Order of Bids	174
	7.5	Poten	tial Refinement Steps	177
		7.5.1	Bidding Order	179
		<i>/....</i>	8	
		7.5.2		179
			Deterministic Declaration	179 180

8	Con	onclusions				
	8.1	Summary of the Thesis				
		Modelling User Interactions	184			
		8.1.2	Encoding Confidentiality	184		
		8.1.3	Integrating Confidentiality with Circus	185		
		8.1.4	Verifying Security	186		
	8.2	Direct	ions for Future Research	186		
		8.2.1	Tool Support	187		
		8.2.2	Mechanising the Theory	187		
		8.2.3	Case Studies	188		
		8.2.4	A Confidentiality Policy Language	189		
		8.2.5	Probabilistic Information Flow	189		
		8.2.6	Other Notions of Security	190		
	8.3	Coda		191		
Pro	oofs			193		
Lis	st of F	Reference	ces	229		

Acknowledgements

I could not have wished for a more patient and supportive supervisor than Jeremy Jacob. I am deeply grateful to him for giving me the freedom to seek out my own course of research, and for guiding me to surmount the obstacles along the way. Our discussions have challenged me to craft my ideas with logical rigour and to present those ideas precisely and fluently, in speech and in writing.

I have been fortunate to find myself supported by such an intellectually stimulating research community in the Department of Computer Science. Special thanks are owed to all of my officemates through the years — especially Peter Booth, Matthew Naylor, Ankur Patel, Marco Perez, Chris Poskitt and Jason Reich — for their camaraderie, their enthusiasm and for encouraging me to keep going when all seemed lost.

The direction of my research has been influenced by engaging — and often enlightening — discussions with Ana Cavalcanti, Will Harwood, Colin Runciman, Jim Woodcock and Frank Zeyda. Thanks to James Wood (NHS CfH) and Jonathan Lewis (Phoenix Interactive) for taking the time to propose topics for case studies to me, which I regrettably was unable to carry out owing to time constraints.

I also owe thanks to the Engineering and Physical Sciences Research Council, for providing financial support for the duration of my doctoral studies; and to the Department of Computer Science, where I have spent a lively, productive and enriching seven years of study.

Declaration

Unless stated otherwise, all contributions of this thesis are my own original work. Parts of this thesis are based on published papers:

- On Modelling User Observations in the UTP. In 3rd International Symposium on Unifying Theories of Programming (UTP 2010). Chapter 3 is an extended version of this paper.
- 2. Unifying Theories of Confidentiality. In *3rd International Symposium on Unifying Theories of Programming (UTP 2010).*

Chapter 4 incorporates some parts of this paper.

3. Specifying Confidentiality in Circus. In FM 2011: Formal Methods.

Chapter 5 and Chapter 6 build on work presented in this paper. Some material is also reproduced in Chapter 2.

4. Confidentiality Annotations for *Circus*. In *Proceedings of the Fourth York Doctoral Symposium on Computer Science (YDS 2011).*

Some ideas in this paper are used in Chapter 4 and Chapter 5. The case study in Chapter 7 also extends work in this paper.

The List of References contains full details of these papers.

1 Introduction

1.1 Confidentiality

Any organisation which handles private or sensitive information assets is obliged to protect the *confidentiality* of those assets. The U.S. National Institute of Standards and Technology defines confidentiality as:

"Preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information." (NIST, 2006)

The choice of security mechanisms for protecting confidential data should be proportionate to the sensitivity of those data (Anderson, 2003). In particular, the cost of implementing security measures should be balanced against the risk of disclosing those data to unauthorised persons:

- For governmental or military establishments, the compromise of national secrets could have catastrophic consequences.
- Likewise, a business may possess valuable trade secrets, where unauthorised exposure may result in financial or legal repercussions.
- The revelation of sensitive data about individuals such as their banking, employment or medical records — may compromise their privacy and facilitate identity theft or blackmail against them.

In practice, software systems which store confidential data are commonly guarded with firewalls, access control and encryption. Each of these measures are important, but they do not guarantee security: all too often, a determined adversary can bypass them by identifying and

1 Introduction

exploiting flaws in their design and implementation. For this reason, it is widely accepted that bolting security mechanisms onto a software system is insufficient to obtain reliable assurances that the system is really secure (McLean, 1987).

With the advent of ubiquitous Internet connectivity, the demand for software systems that are *secure by design* is greater than ever before. In the words of Santen et al. (2002):

"The consent is growing that secure systems cannot be built by adding security features ex post to an existing implementation, but that "security-aware" engineering of systems and software must take security concerns into account, starting from requirements engineering through architectural and detailed design to coding, testing and deployment."

This sentiment is echoed by Jones (2009), reporting the outcomes of an international meeting of security experts drawn from academia, commerce and government. These experts envisage:

"The development and procurement of software and systems which are resilient and sustainable by design, where requirements such as security and privacy are, as a matter of course, defined at project initiation and implemented and assured throughout in risk-based, whole-life processes."

In circumstances where the confidentiality of information is critical, it is therefore imperative that software systems are developed to minimise the possibility of undesirable leakages of confidential information. These developments demand the application of precise and robust software engineering techniques.

1.2 Formal Methods and Security

An influential school of thought in software engineering holds that programs should be treated as mathematical objects (Hoare, 1969; Scott and Strachey, 1971; Dijkstra, 1976; Gries, 1981; Morgan, 1994). This philosophy has motivated the production of a broad spectrum of deductive techniques, collectively known as *formal methods*, for constructing and analysing models of computer systems with logical rigour. The field of formal methods is surveyed by Wing (1990), Saiedian (1996) and Woodcock et al. (2009), among others.

Software engineers can apply formal methods to provide a high level of assurance that software products are *correct by construction*. These methods emphasise the value of creating abstract models to specify the desired characteristics of systems. These models can be subjected to rigorous consistency checks to determine whether they faithfully represent the customer's requirements for the system. A satisfactory model can then be developed — through a series of correctness-preserving *refinement steps* — into an implementation of the system (Wirth, 1971).

In the study and application of formal methods, correctness is predominantly interpreted in terms of *functionality properties*, which dictate the relation between the inputs and outputs of systems. Functionality places a lower bound on the information that a system reveals to its users. In contrast, confidentiality is a *non-functional* property: in fact, it imposes an upper bound on the disclosure of information to untrusted users. Since these bounds oppose each other, standard methods for developing systems to achieve functional correctness may fail to achieve a secure system (Jacob, 1992).

The advantages of applying formal methods in the development of security-critical systems are well documented by Wing (1998) and Ryan (2001). Indeed, it is widely recognised that formalisation is essential for designing trustworthy cryptographic protocols (Burrows et al., 1989). However, much of the theory associated with formal notions of confidentiality has been developed independently from formal methods and other software engineering practices. As a consequence, well-established formal development platforms — such as VDM (Jones, 1990) and the B-Method (Abrial, 1996; Schneider, 2001) — do not incorporate facilities for encoding confidentiality requirements within system designs. Hence,

1 Introduction

practitioners of formal methods need specialised knowledge to integrate confidentiality properties within the development of systems successfully.

1.3 A Platform for Secure Software Development

This thesis describes a formal platform for designing software systems that satisfy a joint specification of confidentiality and functionality properties. By promoting confidentiality to the status of a first-class citizen, we make it accessible to formal methods practitioners who do not possess specialised security knowledge.

At the heart of our platform is a generic framework for modelling systems with multiple users, presented in Chapter 3. Based on the *Unify-ing Theories of Programming* (UTP) (Hoare and He, 1998), this framework is suitable for modelling a user's interactions across different semantic models of systems. In Chapter 4, we study how the framework can be applied to formalise users' knowledge about a system's behaviour, providing a medium for the specification of confidentiality properties.

We elevate the framework to a fully fledged formal development platform by embedding it within *Circus*. The *Circus* notation, introduced by Woodcock and Cavalcanti (2001), combines the process algebra known as CSP (Hoare, 1985a; Roscoe, 1997; Schneider, 1999) and the state-based style of Z (Spivey, 1992; Woodcock and Davies, 1996) to provide facilities for modelling state-rich concurrent and reactive systems.

In Chapter 5, we extend the UTP theory of *Circus* processes to support new constructs for expressing confidentiality attributes over the state and behaviour of *Circus* processes. This construct integrates seamlessly with the *Circus* syntax; it can be mixed with regular *Circus* constructs to specify a wide range of confidentiality properties.

It is possible to specify a process where its functionality and confidentiality properties conflict with each other. Such conflicts between properties are undesirable, because no implementation can possibly satisfy a contradictory specification. Chapter 6 presents a technique for detecting these conflicts, which draws inspiration from Dijkstra's weakest precondition

1.4 Structure of this Thesis

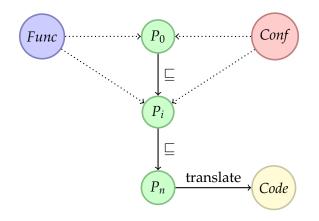


Figure 1.1: A depiction of a formal system development strategy incorporating both functionality and confidentiality requirements. Both functionality and confidentiality properties can be added to the development at each design step. These correctness properties are inherited under refinement, leading to a final design that may be implemented as code.

calculus. Moreover, this technique can be applied in combination with the *Circus* refinement calculus, to develop processes with the assurance that refinement steps do not induce conflicts between functionality and confidentiality.

To evaluate the effectiveness of these techniques, we present a case study development in Chapter 7.

1.4 Structure of this Thesis

Following this introduction, this thesis is structured as follows:

- Chapter 2 provides an overview of the UTP and Circus.
- Chapters 3 through 6 present the main contributions of the thesis.
- Chapter 7 delivers a case study that applies the techniques delivered in previous chapters.

1 Introduction

- Chapter 8 summarises the contributions of the thesis and identifies topics for future work.
- The Appendix contains formal proofs of all theorems and lemmas given in the thesis.

This thesis is written with a formal methods audience in mind, owing to the nature of the material it covers. In particular, the reader should be familiar with the first-order predicate calculus and the algebra of CSP.

A discussion of the ethical issues surrounding computer security and information confidentiality is beyond the scope of this thesis. The reader interested in these issues may wish to consult Stoll (1988) and Schneier (2004) for various perspectives on this topic.

2 Background

2.1 Unifying Theories of Programming

This chapter provides a whistle-stop tour of Hoare and He's Unifying Theories of Programming (UTP). The UTP framework imparts a uniform semantics to a range of programming paradigms, spanning imperative, object-oriented and functional languages, across sequential and concurrent models of computation in timed and untimed domains.

The UTP strives for simplification: it achieves a *separation of concerns* by building semantic models of programming on top of more elementary programming models. The power of this modular approach is its ability to reveal connections between different programming paradigms, enabling results from one paradigm to be translated to other paradigms.

This chapter does not cover all aspects of the UTP, but focuses only on aspects that are relevant to this thesis. The reader who is familiar with the UTP may wish to proceed directly to Chapter 3. Likewise, the reader who desires a more comprehensive account of the UTP is directed to:

- the tutorial by Woodcock and Cavalcanti (2004), which focuses on the UTP theory of designs;
- the tutorial by Cavalcanti and Woodcock (2006), which constructs a UTP semantics for CSP processes; and
- the textbook by Hoare and He (1998), which is the definitive reference work for UTP.

2.2 Programs as Predicates

The semantic foundation of the UTP is the first-order predicate calculus, augmented with fixed point constructs from second-order logic.

Following the tradition of Hehner (1984a,b), *programs are predicates* in the UTP, and there is no distinction between programs and specifications at the semantic level. In the words of Hoare (1984b):

"A computer program is identified with the strongest predicate describing every relevant observation that can be made of the behaviour of a computer executing that program."

In UTP parlance, a *theory* is a model of a particular programming paradigm. A UTP theory is composed of three ingredients:

- an *alphabet*, which is a set of variable names denoting the attributes of the paradigm that can be observed by an external entity;
- a *signature*, which is the set of programming language constructs intrinsic to the paradigm.
- a collection of *healthiness conditions*, which define the space of programs that fit within the paradigm.

In the UTP, healthiness conditions are typically defined as monotonic¹ idempotent predicate transformers. Given a healthiness condition **F**, a predicate *P* is said to be **F**-healthy if and only if $P = \mathbf{F}(P)$.

2.3 Relations

The most basic UTP theory is the alphabetised predicate calculus, which has no alphabet restrictions or healthiness conditions.

The theory of relations is slightly more specialised. It requires the alphabet of a relation to consist of only:

¹**F** is monotonic if and only if $P_1 \sqsubseteq P_2$ implies **F** $(P_1) \sqsubseteq$ **F** (P_2) , for all P_1 and P_2 .

- undecorated variables (*a*, *b*, *c*, . . .), modelling an observation of the program at the start of its execution; and
- primed variables (*a*', *b*', *c*', . . .), modelling an observation of the program at a later stage of its execution.

Hereafter, we write x to denote the list of all undecorated variables in a relation's alphabet, and x' to denote all primed variables in the alphabet. The alphabet is *homogeneous* if the names of its primed variables match the names of its undecorated variables.

A relation models all observable behaviours of a program by mapping initial states (over x) to final states (over x').

Example 2.1. The relation $F2C \triangleq c' = (f - 32) \times \frac{5}{9}$ models a program that converts temperature measurements from the Fahrenheit scale (*f*) to the Celsius scale (*c'*). By renaming *c'* and *f* to *c* and *f'* respectively:

$$C2F \triangleq F2C[c,f'/c',f] = f' = (c \times 9/5) + 32$$

we derive a program that performs the inverse conversion.

 \Diamond

2.3.1 Refinement

Program P_1 is refined by program P_2 — i.e. every observation of P_2 is an observation of P_1 — if and only if:

$$P_1 \sqsubseteq P_2 \quad \triangleq \quad [P_2 \Rightarrow P_1]$$

where the square brackets $[\cdot]$ denote universal quantification over all variables in the alphabet (Dijkstra and Scholten, 1990).

Refinement succinctly captures the notion of program correctness: if $P_1 \sqsubseteq P_2$ holds, then P_2 correctly implements P_1 . This notion of refinement is the same across all UTP theories.

The \sqsubseteq relation is a partial order and induces a complete lattice over the space of relations. At the bottom of the lattice, the least refined

П		x' = x	(skip)
a := E		$a' = E \wedge u' = u$	(assignment)
$P_1; P_2$	$\underline{\underline{\frown}}$	$\exists x_0 \bullet P_1[x_0/x'] \land P_2[x_0/x]$	(sequence)
$P_1 \sqcap P_2$	$\underline{\underline{\frown}}$	$P_1 \vee P_2$	(non-determinism)
$P_1 \lhd C \rhd P_2$		$(C \wedge P_1) \vee (\neg C \wedge P_2)$	(conditional)
$\mu X \bullet F(X)$		$\prod \{X \mid F(X) \sqsubseteq X\}$	(recursion)

Figure 2.1: The relational semantics of program constructs.

program **true** (sometimes called **abort**) has no constraints whatsoever on its behaviour. The lattice top is **false**, which admits no behaviour at all and cannot be implemented. Yet this "program" satisfies all specifications, which makes it a *miracle* (Nelson, 1989; Morgan, 1994). Miracles are conceptually very useful, because they give meaning to contradictory specifications.

Sometimes we write $P_1 \sqsubset P_2$ to specify P_2 is a *strict* refinement of P_1 ; i.e. P_1 and P_2 are not equivalent:

$$P_1 \sqsubset P_2 \triangleq P_1 \sqsubseteq P_2 \land P_2 \not\sqsubseteq P_1$$

2.3.2 Defining Program Constructs

The language constructs in the signature of the theory of relations are defined in Figure 2.1. Each construct is explained as follows:

- The skip statement, which does not alter the program state in any way, is modelled as the relational identity *Π*.
- The assignment of value *E* to a state variable *a* is defined by setting *a*' to *E* and keeping all other variables (denoted by *u*) constant.
- The sequential composition of two programs is just relational composition over their intermediate state.

- Non-deterministic choice between programs is their greatest lower bound.
- Conditional choice between programs is specified using a ternary operator, _ < C ▷ _, where C is a relation.²
- A semantics for recursion is given by the weakest fixed point μF of a monotonic predicate transformer *F*. In particular, we have:

$$\mu F = \mu X \bullet F(X)$$

A corollary of the definition of μF is that an infinite recursion $\mu X \bullet X$ is equivalent to **abort**:

$$\mu X \bullet X = \bigcap \{X \mid X \sqsubseteq X\} = \text{true}$$

2.4 Designs

Hoare and He (1998) identify an inconsistency between the theory of relations and our intuitive understanding of programs. Consider the relation ($\mu X \bullet X$); *a* := 1. By calculation, we find that:

$(\mu X \bullet X); a := 1$					
=	true ; <i>a</i> := 1	[as above]			
=	$\exists x_0 \bullet \mathbf{true}[x_0/x'] \land a := 1[x_0/x]$	[def ;]			
=	$\exists x \bullet a := 1$	[renaming]			
=	a' = 1	[def :=]			

This result implies that a program can break out of a non-terminating loop, which is inconsistent with the reality of programming. To resolve this inconsistency, a theory of programming needs to distinguish between terminating and non-terminating programs. This goal is met by the theory of *designs*.

²The notation $P_1 \triangleleft C \triangleright P_2$ should be read " P_1 if *C* else P_2 " (Hoare, 1985b).

2 Background

A design models a program as a precondition and a postcondition. It also introduces two special Boolean observational variables into the alphabet:

- *ok* holds if the program has started properly; while
- *ok'* holds if the program has terminated.

The theory of designs is characterised by two idempotent monotonic healthiness conditions, **H1** and **H2**:

$$\mathbf{H1}(P) \quad \triangleq \quad ok \Rightarrow P \\ \mathbf{H2}(P) \quad \triangleq \quad P; ((ok \Rightarrow ok') \land v' = v)$$

where *v* denotes the list of state variables excluding *ok*.

H1 specifies that a design makes no predictions about its behaviour before it is started. H2 specifies that a design cannot insist on non-termination by forcing ok' to be **false**.³

Provided *P*'s alphabet includes *ok* and *ok'*, **H1** and **H2** are commutative. Any relation that is **H1** and **H2**-healthy can be written in the form:

$$Pre \vdash Post \triangleq (ok \land Pre) \Rightarrow (ok' \land Post)$$

A design $Pre \vdash Post$ reflects a contract between the customer and the programmer: if the design is started in a state satisfying *Pre*, then it must terminate in a state such that *Post* is satisfied. By this construction, refinement of a design corresponds to weakening its precondition and strengthening its postcondition (Hoare and He, 1998, Theorem 3.1.2).

Some program constructs in the theory of relations are not designs in their own right. These constructs need to be redefined as designs:

 $\Pi_D \triangleq \mathbf{true} \vdash \Pi \qquad (\text{design-skip})$ $a := E \triangleq \mathbf{true} \vdash a' = E \land u' = u \qquad (\text{design-assignment})$

³The definition of **H2** presented here is taken from Cavalcanti and Woodcock (2006). Its semantics is equivalent to Hoare and He's definition, but reformulated as an idempotent monotonic predicate transformer.

2.5 Reactive Processes

A reactive process communicates with its environment by engaging in a series of events. Therefore, a model of a reactive process needs to admit intermediate observations at points during its execution.

The theory of reactive processes is defined over an extended alphabet, featuring four pairs of distinguished observational variables. The Boolean variables *ok* and *wait* model the process status at the initial observation:

- *ok* ∧ *wait*: the process is waiting to start its execution;
- $ok \land \neg$ *wait*: the process has started its execution properly;
- \neg *ok*: the process has not been started properly;

Likewise, the primed Boolean variables ok' and wait' model the process status at subsequent observations:

- *ok'* ∧ *wait'*: the process has reached a stable intermediate point in its execution and is awaiting interaction with the environment;
- $ok' \land \neg$ *wait'*: the process has terminated in a stable point.
- ¬ ok': the process has neither reached a stable intermediate point nor has terminated: it is said to have *diverged*.

The variables tr and tr' record the sequence of events observed by the environment. tr records all events prior to the process starting, while tr' records events that occur up to the point when the subsequent observation is made. Hence, the process trace is given by tr' - tr.

The *ref* and *ref*' variables record the events that may be refused by the process. They are unconstrained by the theory of reactive processes.

Reactive processes are characterised by three healthiness conditions:

2 Background

where the reactive identity Π_{rea} is defined as follows:

$$\Pi_{rea} \triangleq \begin{pmatrix} \neg ok \wedge tr \leq tr' \\ \lor ok' \wedge wait' = wait \wedge tr' = tr \wedge ref' = ref \wedge v' = v \end{pmatrix}$$

A **R1**-healthy process can only extend the trace; in other words, it has no control over events prior to starting. A **R2**-healthy process is oblivious to events that occurred before it starts. A **R3**-healthy process has no effect on the observation until it commences execution.

Since **R1**, **R2** and **R3** are idempotent and commute with each other, they can be combined into a single healthiness condition **R**:

$$\mathbf{R}(P) \triangleq \mathbf{R1} \circ \mathbf{R2} \circ \mathbf{R3}(P)$$

2.6 Reactive Designs

The theories of designs and reactive processes each have valuable qualities. Designs are useful for modelling programs in terms of preconditions and postconditions, while reactivity enables the intermediate behaviour of programs to be captured. However, the healthiness conditions **H1** and **R1** do not commute with each other:

$$\begin{array}{lll} \mathbf{H1} \circ \mathbf{R1}(P) & = & ok \Rightarrow (P \land tr \leq tr') \\ & \neq & (ok \Rightarrow P) \land tr \leq tr' \\ & = & \mathbf{R1} \circ \mathbf{H1}(P) \end{array}$$

Intuitively, a **H1**-healthy relation allows any observation whatsoever when *ok* is **false**, whereas a **R1**-healthy relation enforces $tr \le tr'$ in all circumstances. This non-commutativity between **H1** and **R1** implies that no design is a reactive process, and vice versa.

The space of *reactive designs* is a sub-space of the reactive processes, which is derived by applying **R** to the space of designs.⁴ Reactive designs

⁴Contrary to their name, reactive designs are *not* designs, since they are not H1-healthy.

have the shape:

$$P = \mathbf{R} \left(Pre \vdash Post \right)$$

where *Pre* specifies the conditions in which *P* does not diverge; and *Post* specifies all possible non-diverging outcomes of executing *P* in any state satisfying *Pre*. In this way, reactive designs facilitate the definition of process semantics in a design-like manner.

In particular, the bottom of the lattice of reactive designs is:

$$\mathsf{R}\left(\mathsf{false}\vdash\mathsf{true}\right) \quad = \quad \mathsf{R}(\mathsf{true})$$

which diverges whenever it is permitted to start execution. The top of the lattice of reactive designs is:

$$\mathsf{R}(\mathsf{true} \vdash \mathsf{false}) \quad = \quad \mathsf{R}(\neg ok)$$

which would achieve the impossible (postcondition **false**) if it could ever be executed. This process is the *reactive design miracle* (Woodcock, 2010).

The theory of reactive designs is extended by Hoare and He (1998) and Cavalcanti and Woodcock (2006) to give a UTP semantics to CSP processes. Two healthiness conditions characterise the sub-space of reactive designs corresponding to CSP processes:

CSP1 specifies that, if a process diverges, then only the extension of the trace is guaranteed. **CSP2** recasts **H2**: it specifies a process may not require non-termination.

Every CSP process is expressible in the form of a reactive design.

2 Background

Theorem 2.2 (from Hoare and He (1998)). For each CSP process *P*:

$$P \quad = \quad \mathbf{R} \left(\neg P_f^f \vdash P_f^t \right)$$

where P_c^b abbreviates P[b, c/ok', wait].

The theory of CSP processes serves as a baseline for the semantics of *Circus* actions. For brevity, we skip over this theory and jump straight to the closely related theory of *Circus* actions in Section 2.8.

2.7 Circus

Circus is a formal specification language which fuses the CSP process algebra with a notion of state, to achieve a cohesive platform for modelling state-rich concurrent and reactive systems.

The specification constructs of *Circus* are called *actions*. *Circus* actions can be grouped into three classes:

- CSP constructs, modelling interaction with the environment over a collection of named channels;
- Z schema expressions and specification statements (Morgan, 1994), representing state operations; and
- guarded commands (Dijkstra, 1976), orchestrating control flow according to the state.

A *Circus* process specifies an internal state (hidden from the environment), a state invariant and a collection of named actions. The behaviour of a *Circus* process is defined by a distinguished nameless main action, which follows the declarations of the other actions.

Example 2.3. Figure 2.2 presents a *Circus* process modelling a memory cell that stores an integer value from the *in* channel. The cell can be switched between two modes. In public mode, the value currently stored in the cell is broadcast on the *out* channel; whereas in private mode, an arbitrary value is broadcast on *out*.

```
MODE == \{PUB, PRV\}

channel on, off

channel in, out : \mathbb{N}

process Cell \triangleq begin

state Mem \triangleq [val : \mathbb{N}, m : MODE]

Init \triangleq [Mem' |val' = 0 \land m' = PUB]

Read \triangleq \begin{pmatrix} m = PUB \& out!val \rightarrow Skip \\ \Box m = PRV \& \Box n : \mathbb{N} \bullet out!n \rightarrow Skip \end{pmatrix}

Write \triangleq in?n \rightarrow val := n?

Switch \triangleq (on \rightarrow m := PRV) \Box (off \rightarrow m := PUB)

\bullet Init ; \mu X \bullet (Read \Box Write \Box Switch) ; X

end
```

Figure 2.2: An example *Circus* process: a memory cell.

2.8 Circus Actions

A UTP theory of *Circus* actions is defined by Oliveira et al. (2009), based on work by Woodcock and Cavalcanti (2002) and Oliveira (2005). This theory enriches the theory of CSP processes with facilities for accessing and manipulating state variables.

The primitive actions of *Circus* are inherited from CSP. The action *Skip* terminates immediately, while *Stop* deadlocks and *Chaos* diverges:⁵

$$Skip \triangleq \mathbf{R} (true \vdash tr' = tr \land \neg wait' \land v' = v)$$

$$Stop \triangleq \mathbf{R} (true \vdash tr' = tr \land wait')$$

$$Chaos \triangleq \mathbf{R} (false \vdash true)$$

A specification statement (Morgan, 1994) is a construct for specifying operations on state, in terms of a precondition, a postcondition and a *frame* w of variables, whose values the operation may change. The *Circus*

⁵Notice the *Chaos* here is that of Hoare (1985a), rather than Roscoe (1997).

2 Background

semantics of specification statements is:

$$w : [Pre, Post] \triangleq \mathbf{R} (Pre \vdash Post \land \neg wait' \land tr' = tr \land u' = u)$$

where *u* denotes all state variables outside the frame *w*. Normalised Z schemata can be translated to specification statements using the basic conversion rules given by Cavalcanti and Woodcock (1998).

Following Morgan (1994), assignments, assumptions and coercions⁶ are defined in terms of specification statements:

$$a := E \triangleq a : [true, a' = E]$$
$$\{C\} \triangleq : [C, true]$$
$$[C] \triangleq : [true, C]$$

It follows from the semantics of the specification statement that:

$$a := E = \mathbf{R} (\mathbf{true} \vdash a' = E \land \neg wait' \land tr' = tr \land u' = u)$$

$$\{C\} = \mathbf{R} (C \vdash \neg wait' \land tr' = tr \land v' = v)$$

$$[C] = \mathbf{R} (\mathbf{true} \vdash C \land \neg wait' \land tr' = tr \land v' = v)$$

Assumptions and coercions are annotation constructs, specifying conditions that are expected to be satisfied by the process state. The assumption { *C* } behaves as *Skip* if the process state satisfies *C*, and diverges otherwise. The coercion [C] *forces* the process state to satisfy *C*.

Example 2.4. The coercion [*C*] behaves as *Skip* if *C* holds; for instance:

$$h := 1; [h = 1] = h := 1; [true] = h := 1$$

 $h := 0; [h = 1] = h := 0; [false] = R(\neg ok)$

but behaves as the reactive design miracle otherwise.

 \Diamond

⁶The terms "assumption" and "coercion" are adopted by Morgan (1994). In the design space, { · } and [·] are called "assertions" and "assumptions" respectively by Hoare and He (1998). To avoid confusion, we use Morgan's terms throughout this thesis.

As Example 2.3 shows, *Circus* actions can be composed using the CSP operators. As before, sequential composition (;) is relational composition, internal choice (\Box) is disjunction and recursion is defined in terms of the weakest fixed point operator. Other CSP operators are given a UTP semantics by Oliveira et al. (2009). We now outline the semantics of three *Circus* operators.

Prefixing A prefixing $c.e \rightarrow Skip$ never diverges and awaits synchronisation on channel *c*. Once it has made that synchronisation, it terminates:

$$c.e \to Skip \triangleq \mathbf{R} \left(\mathbf{true} \vdash \left(\begin{array}{c} tr' = tr \land (c,e) \notin ref' \\ \lhd wait' \rhd \\ tr' = tr \frown \langle (c,e) \rangle \end{array} \right) \land v' = v \right)$$

Here, an event is represented in the trace as a pair of a channel name and the value transmitted on that channel. The semantics of a CSP prefixing $c \rightarrow Skip$ (where channel *c* does not communicate values) is defined as $c.Sync \rightarrow Skip$, where *Sync* denotes a generic synchronisation event.

The prefixing $c.e \rightarrow A$ can be rewritten as $c.e \rightarrow Skip$; A.

External choice An external choice $A_1 \square A_2$ offers the environment the ability to select between the events offered by A_1 and A_2 . Its semantics is:

$$A_{1} \Box A_{2} \triangleq \mathbf{R} \left(\neg A_{1f}^{f} \land \neg A_{2f}^{f} \vdash \left(\begin{array}{c} A_{1f}^{t} \land A_{2f}^{t} \\ \lhd tr' = tr \land wait' \rhd \\ A_{1f}^{t} \lor A_{2f}^{t} \end{array} \right) \right)$$

This choice does not diverge provided that neither A_1 nor A_2 diverge. Provided that A_1 and A_2 do not terminate instantaneously, then the environment can select from the initial events offered by either action. The prefixing operator specifies its refusals negatively, so the refusal set of $A_{1f}^{t} \wedge A_{2f}^{t}$ cannot contain any event that is accepted by either A_1 or A_2 .

Once an event is performed, the choice between A_1 and A_2 is resolved.

2 Background

Guards A guarded command *g* & *A* behaves as *A* if *g* holds, and as *Stop* otherwise:

$$g \& A \triangleq \mathbf{R} \left((g \Rightarrow \neg A_f^f) \vdash (A_f^t \lhd g \rhd tr' = tr \land wait') \right)$$

Definitions of the remaining *Circus* operators — including parallel composition and hiding — are given by Oliveira et al. (2009). Their details are not necessary to follow the rest of this thesis.

The UTP semantics of *Circus* actions facilitates the definition of a refinement calculus (Oliveira, 2005), featuring a wide variety of laws proved directly from the semantics. This refinement calculus enables software engineers to transform abstract process designs into concrete implementations, without exposure to the underlying *Circus* semantics.

2.9 The Circus Family

Building on the reference *Circus* semantics, several versions of *Circus* have been defined to address specialised domains in software engineering. These domains include (but are not limited to) object-oriented programming (Cavalcanti et al., 2005), real-time specification (Wei et al., 2010) and synchronously-clocked systems (Gancarski and Butterfield, 2009).

In the body of this thesis, we describe how notions of information confidentiality can be captured in the UTP and, from there, embedded into the reference *Circus* semantics. We speculate that our work could be replicated across the other versions of *Circus*, giving rise to languages in each domain with a confidentiality-sensitive semantics.

3 Multi-User Systems

3.1 Introduction

This chapter presents techniques for formally reasoning about shared computing systems that offer services to multiple end-users. This class of *multi-user* systems encompasses a wide variety of software products, ranging from operating systems and database servers to telecommunications networks and cloud computing facilities.

The word "user" has many connotations in computing but, in this chapter, we use the term to mean any entity that can interact with a computer system or be influenced by its behaviour. Thus, a user may denote a client that issues requests to the system and receives timely responses; an administrator responsible for monitoring the system; or an entity that observes the system's behaviour in a passive way.

3.1.1 Motivation

When designing a multi-user system, it is important to characterise the relationship between the system and its users, to determine whether the system will satisfy the expectations of its users.

A multi-user system may offer different services to different users. Hence, it is usual for the system to supply each user with its own dedicated interface. These interfaces impose structure upon the system's environment, enabling the system to distinguish between its users.

Many formal methods for software development treat the environment of a system as a single uniform entity. In these formalisms, the users of a system are conventionally modelled as components of the specification

3 Multi-User Systems

of the system itself. For instance, in state-based formalisms such as B and Z, a user's interface to a system is typically defined as a collection of operations that are considered to be accessible to the user. Likewise, when working with process algebras such as CSP, it is usual to model users as individual processes that synchronise with a process representing the operations offered by the system. Using a tool such as FDR (Roscoe, 1994), this approach can be used to verify that a process delivers the functionality expected by its users, by analysing the synchronisations between these processes for the absence of deadlock or divergence.

Modelling users as implicit components of a system works well for analysing the functionality delivered by the system to its environment. However, by failing to distinguish the users from the environment, this approach does not distinguish the interactions of users from the behaviours of the system. This shortcoming is inconvenient when dealing with correctness properties — such as confidentiality properties — where the relation between behaviours and interactions has intrinsic significance.

3.1.2 Contributions of this Chapter

In this chapter, we investigate an alternative approach for modelling users of systems. The core of this approach is outlined in Section 3.2: we specify a user's interface separately from a system, in terms of the elements of a system's behaviour that are visible to that user. In this way, we become able to analyse the interactions that users may perform with the system, and thence verify that a system design upholds the users' expectations about their interactions.

Section 3.3 presents a method for calculating the space of interactions that a user may perform with a given system. We explore how this approach can be applied to the theories of designs in Section 3.4 and reactive processes in Section 3.5.

We also investigate how knowledge of the interfaces of users can be used to justify certain kinds of functional design decisions in system developments. In Section 3.6, we outline a weakened notion of refinement applicable to multi-user systems that is framed by the observational abilities of isolated users or groups of collaborating users.

The work in this chapter constitutes the foundations of later chapters of this thesis, which build on the model of users to formalise the notion of information flow in the UTP and in *Circus*.

3.2 Modelling Users

Recall from Chapter 2 that, in the UTP theory of relations, a system is modelled in terms of all possible observations that an external entity could make of the system.

When reasoning about multi-user systems modelled as relations, we draw a distinction between two kinds of observations:

Behaviours are observations of a system made by the global environment.

Interactions are local observations of a system made by a single user.¹

We assume that users cannot monitor the system directly, but can only observe the system through their respective interfaces. We capture this assumption by introducing a secondary observation space to model interactions separately from behaviours. To do this, we introduce two new lists of observational variables, x_V and x'_V , to encode interactions as alphabetised predicates:

- the *x_V* variables model the state of the user's interface before an interaction takes place; whereas
- the *x*[']_V variables model the state of the interface after the interaction takes place.

¹Although the term "interaction" suggests that users actively engage with the system and influence its behaviour, this need not be so.

3.2.1 Views

In general, a user's interface to a system relays only partial information about the system's behaviour to that user; after all, users need not be concerned with every detail of the system's behaviour. Hence, a user's interaction is a *projection* of the behaviour through that user's interface.

A *view* is a predicate expressing a relation between observation spaces. More precisely, we define a view as a relation from the space of behaviours to the space of interactions. Hence, a view formalises a user's interface to a system, because it determines which aspects of the system's behaviour show through the user's interface.

Example 3.1. Suppose that *x* and *y* are integer-valued observational variables. Consider the following views:

$$V1 \triangleq a_1 = x \land b_1 = y$$

$$V2 \triangleq c_2 = \max(x, y)$$

$$V3 \triangleq d_3 = x \land (e_3 = 0 \lhd x \lt y \triangleright e_3 = 1)$$

*V*1 provides the exact values of *x* and *y*, since there is a one-to-one correspondence between these variables and a_1 and b_1 . *V*2 provides the value of the larger of *x* and *y*, but no indication of whether x < y, x = y or x > y. *V*3 provides the value of *x*, but offers only partial information about *y*'s value in relation to *x*'s value.

When dealing with multiple views, each representing a different user, we require that each view's alphabet of interaction variables is disjoint from those of the other views. (All of the views listed in Example 3.1 are pairwise disjoint.)

Definition 3.2 (Disjoint views). A pair of views V_1 and V_2 are *disjoint* if and only if their alphabets have no interaction variables in common:

$$\alpha V_1 \cap \alpha V_2 \subseteq x \cup x'$$

The abstractness of views allows us to describe the observational abilities of users in exact relation to the behaviour of systems. We note in passing that a view's structure should conform to the UTP theory in which it is applied; we return to this point in Sections 3.4 and 3.5.

3.2.2 Healthiness Conditions for Views

To ensure that all views represent a viable mapping between behaviours and interactions, it is necessary to impose some healthiness conditions on the structure of views.

Our first healthiness condition ensures that views do not restrict the behaviour of systems in any way.

Definition 3.3 (VH1). A view *V* is **VH1**-healthy if and only if *V* maps each behaviour to *at least one* interaction:

$$\mathsf{VH1}(V) \triangleq (\exists x_V, x'_V \bullet V) \Rightarrow V$$

VH1 may be more comprehensible when presented in a form that asserts the **VH1**-healthiness of a view, as given by Lemma 3.4.

Lemma 3.4 (VH1 reformulated).

VH1
$$(V) = V$$
 if and only if $[\exists x_V, x'_V \bullet V]$

A **VH1**-healthy view is a total relation from behaviours to interactions. If *V* does not map a behaviour ϕ to at least one interaction, then **VH1**(*V*) maps ϕ to *every* interaction. Hence, **VH1**-healthy views do not restrict the behaviour of systems in any way.

Remark 3.5. An alternative formulation of **VH1** would transform views by mapping behaviours associated with zero interactions to a nominated "null" interaction. However, our definition of **VH1** is simpler, because it avoids the need to specify what that null interaction should be.

Example 3.6. Predicates such as $a = b \land c_V > 0$ and **false** are not **VH1**-healthy, because they restrict the behaviour space. For instance:

$$VH1(a = b \land c_V > 0)$$

$$= (\exists x_V, x'_V \bullet a = b \land c_V > 0) \Rightarrow (a = b \land c_V > 0)$$
 [def VH1]
$$= a = b \Rightarrow (a = b \land c_V > 0)$$
 [pred calc]
$$= a = b \Rightarrow c_V > 0$$
 [prop calc]

which does not equal $a = b \land c_V > 0$.

We could write a "precognitive" **VH1**-healthy predicate — for example, $a_V = a'$ — where the initial variables of an interaction depend upon the values of the final behavioural variables. Such a predicate does not correspond to any user interface in reality, because no interface can predict the behaviour of an unspecified system before that system starts its execution. We exclude these unrealistic predicates from the space of views by defining a second healthiness condition **VH2**.

Definition 3.7 (VH2 healthiness condition). A view *V* is **VH2**-healthy if and only if its alphabet contains only unprimed variables:

$$\mathsf{VH2}(V) \triangleq \exists x', x'_V \bullet V$$

The definition of a healthiness condition to rule out precognitive predicates over the initial and final observational space would be complicated. With **VH2**, we cut this Gordian knot by assuming that a user's interface does not change during the system's execution.

By assuming that a user's interface is constant, we are free to encode views over the unprimed variables alone. The notation:

$$\Delta V \triangleq V \wedge V'$$

denotes the interface corresponding to the **VH2**-healthy view V, where V' denotes V with all unprimed variables renamed with a prime.

 \Diamond

We now present some useful properties of **VH1** and **VH2**. Both **VH1** and **VH2** are idempotent.

Law 3.8 (VH1 idempotent). VH1 \circ VH1(V) = VH1(V)

Law 3.9 (VH2 idempotent). VH2 \circ VH2(V) = VH2(V)

Law 3.9 is a trivial consequence of the definition of **VH2**. Unfortunately, **VH1** and **VH2** do not commute with each other.

Law 3.10 (VH1-VH2 non-commutative). VH2 \circ VH1 $(V) \sqsubseteq$ VH1 \circ VH2(V)

We say that a view that is both VH1-healthy and VH2-healthy is VHhealthy. In light of Law 3.10, we define VH as the stronger combination of VH1 and VH2.

Definition 3.11 (VH healthiness condition).

$$\mathsf{VH}(V) \triangleq \mathsf{VH1} \circ \mathsf{VH2}(V)$$

Corollary 3.12 (VH idempotent). By Law 3.8 and the definition of VH:

 $VH1 \circ VH(V) = VH(V)$

Likewise, since the x' and x'_V variables are free in **VH**-healthy views:

$$\mathbf{VH2} \circ \mathbf{VH}(V) \quad = \quad \mathbf{VH}(V)$$

It follows from these results that **VH** is idempotent.

All of the views in Example 3.1 are **VH**-healthy, as are all the views that we consider in the following sections.

3.2.3 Conjoining Views

Given two disjoint views V_1 and V_2 , the view $V_1 \wedge V_2$ models the combined interactions of a pair of users. This view represents those users sharing knowledge of their interactions with the system with each

other: conjoining V_1 and V_2 ensures that pairs of V_1 -interactions and V_2 -interactions are consistent with the same behaviour.

Lemma 3.13 (Conjunction preserves VH). If V_1 and V_2 are disjoint VH-healthy views, then $V_1 \wedge V_2$ is also VH-healthy.

3.2.4 Functional Views

While **VH1** insists that a view describes a total mapping from behaviours to interactions, this mapping need not be functional. A non-functional view would correspond to a user's interface that can yield *different* interactions for the *same* behaviour.

Example 3.14. A lossy communications channel can be modelled as a **VH**-healthy view. For instance, the view $a_V \in [a - k, a + k]$ could represent a signal that fluctuates by up to *k* units. This view could serve as a crude model of a physical wire carrying the signal.

However, the interfaces to a system are typically designed to yield a deterministic projection of the system's behaviour, at a reasonable level of abstraction.² We define the space of these *functional* views in terms of a new healthiness condition **VH3**.

Definition 3.15 (**VH3** healthiness condition). A view *V* is **VH3**-healthy if and only if *V* maps each behaviour to *exactly one* interaction or *indefinitely many* interactions:

$$\mathbf{VH3}(V) \triangleq (\exists_1 x_V, x'_V \bullet V) \Rightarrow V$$

Every VH3-healthy view is also VH1-healthy.

Example 3.16. Provided k > 0, the **VH1**-healthy view $a_V \in [a - k, a + k]$ from Example 3.14 is not **VH3**-healthy:

VH3 $(a_V \in [a - k, a + k])$

²Quantum mechanics suggests that no physical interface can be truly deterministic at the lowest level, but this is generally not a concern for system designers in practice!

3.3 Calculating Interactions

$$= (\exists_1 x_V, x'_V \bullet a_V \in [a - k, a + k]) \Rightarrow a_V \in [a - k, a + k] \quad [def VH3]$$

$$= false \Rightarrow a_V \in [a - k, a + k] \quad [not deterministic]$$

$$= true \quad [prop calc]$$

The resulting view **true** models a user whose interface is disconnected from the system, so the user's interaction bears no resemblance whatsoever to the system's behaviour.

From here onwards, the views we consider in examples are **VH3**-healthy, to aid intuition. However, unlike **VH1** and **VH2**, the special functional property of **VH3** is not essential to justify the theorems and lemmas presented in the following sections. For this reason, we do not restrict the framework to **VH3**-healthy views; rather, we appeal to **VH3** only when it is needed.

3.3 Calculating Interactions

In this section, we formalise procedures for translating a model of a system's behaviour into a complementary model representing a user's interactions with that system, and vice versa.

3.3.1 Localisation

Given a view *V* and (a behavioural model of) a system *P*, the predicate L(V, P) is the *image* of *P* projected through *V*. In other words, this predicate encodes all interactions with *P* that can be made through *V*.

Definition 3.17 (L predicate transformer).

$$\mathsf{L}(V,P) \triangleq \exists x, x' \bullet \Delta V \land P$$

Example 3.18. Consider the condition *Ex*:

$$Ex \quad \triangleq \quad x \ge 0 \land y \ge 0 \land x + y = 10$$

The images of *Ex* corresponding to each view listed in Example 3.1 are:

$$L(V1, Ex) = a_1 \ge 0 \land b_1 \ge 0 \land a_1 + b_1 = 10$$

$$L(V2, Ex) = c_2 \ge 5 \land c_2 \le 10$$

$$L(V3, Ex) = (d_3 \ge 0 \land d_3 < 5 \land e_3 = 0) \lor (d_3 \ge 5 \land d_3 \le 10 \land e_3 = 1)$$

Notice that **L** hides the behavioural variables of Ex, modelling a user's inability to observe those variables directly.

We now describe three properties of L that we use later in this chapter.

Distributivity Each interaction with a system that behaves either as P_1 or P_2 is an interaction either with P_1 or with P_2 .

Law 3.19 (L(V) is disjunctive). $L(V, P_1 \lor P_2) = L(V, P_1) \lor L(V, P_2)$

Monotonicity If $P_1 \sqsubseteq P_2$, then every *V*-interaction with P_2 (where *V* is **VH**-healthy) will also be a *V*-interaction with P_1 .

Law 3.20 (L(V) is monotonic). Provided P_1 and P_2 share an alphabet:

$$P_1 \sqsubseteq P_2$$
 implies $L(V, P_1) \sqsubseteq L(V, P_2)$

Splitting views If a view *V* is divided into two disjoint views V_1 and V_2 , such that $V = (V_1 \land V_2)$, then $L(V_1, P) \land L(V_2, P)$ is potentially weaker than L(V, P) itself. Intuitively, this is because the interactions given by $L(V_1, P)$ and $L(V_2, P)$ are not synchronised with each other. We generalise this result to an arbitrary number of views in Law 3.21.

Law 3.21 (Splitting views weakens projections). Provided V_1, \ldots, V_n are pairwise disjoint views:

$$\bigwedge_{i\in 1..n} \bullet \mathsf{L}(V_i, P) \quad \sqsubseteq \quad \mathsf{L}\left(\left(\bigwedge_{i\in 1..n} V_i\right), P\right)$$

3.3.2 Globalisation

Given an image U of interactions formed by projecting a system P through view V, some knowledge of the structure of P can be recovered by substituting V back into U. This procedure is formalised by the **G** predicate transformer.

Definition 3.22 (G predicate transformer).

$$\mathbf{G}(V,U) \triangleq \forall x_V, x'_V \bullet \Delta V \Rightarrow U$$

The relation **G** (V, U) is the weakest specification of a system such that each behaviour of **G** (V, U) projects through *V* to an interaction of *U*.

Example 3.23. Continuing from Example 3.18, the conditions recovered by applying **G** to each **L**-projection of *Ex* are:

$$\begin{array}{lll} {\bf G} \left(V1, {\bf L} \left(V1, Ex \right) \right) & = & x \ge 0 \, \land \, y \ge 0 \, \land \, x + y = 10 \\ \\ {\bf G} \left(V2, {\bf L} \left(V2, Ex \right) \right) & = & \max(x, y) \ge 5 \, \land \, \max(x, y) \le 10 \\ \\ {\bf G} \left(V3, {\bf L} \left(V3, Ex \right) \right) & = & (x \ge 0 \, \land \, x < 5) \, \lhd \, x < y \vartriangleright \, (x \ge 5 \, \land \, x \le 10) \end{array}$$

Observe that $\mathbf{G}(V1, \mathbf{L}(V1, Ex)) = Ex$, because view V1 preserves the values of x and y in a_1 and b_1 . However, both $\mathbf{G}(V2, \mathbf{L}(V2, Ex))$ and $\mathbf{G}(V3, \mathbf{L}(V3, Ex))$ are weaker than Ex, because information about the values of x and y is discarded by applying $\mathbf{L}(V2)$ or $\mathbf{L}(V3)$ to Ex.

The predicate $\mathbf{G}(V, U)$ features all behaviours that *V* only maps to interactions in *U*. However, if *V* is **VH3**-healthy, then *V* maps each behaviour to exactly one interaction. This property of *V* enables the definition of $\mathbf{G}(V, U)$ to be reshaped, as Lemma 3.24 shows.

Lemma 3.24 (G and VH3). Provided V is VH3-healthy:

$$\mathbf{G}(V, U) = \neg \mathbf{G}(V, \neg U) = \exists x_V, x'_V \bullet \Delta V \wedge U$$

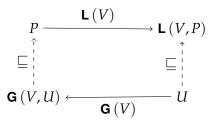


Figure 3.1: The Galois connection between **L** and **G**.

3.3.3 A Galois Connection

As Example 3.23 demonstrates, the **L** and **G** predicate transformers are not inverses of each other, since a view need not define a one-to-one correspondence between behaviours and interactions. They do, however, form a Galois connection under the refinement ordering.

Theorem 3.25 (**L** and **G** form a Galois connection). The **L** and **G** predicate transformers form a Galois connection between the space of behaviours and the space of interactions linked by a given view. Thus:

$$U \sqsubseteq \mathbf{L}(V, P)$$
 if and only if $\mathbf{G}(V, U) \sqsubseteq P$

Theorem 3.25 enables us to apply existing results on Galois connections from the literature. For instance, Corollary 3.26 is an immediate consequence of Theorem 3.25.

Corollary 3.26. Substituting L(V, P) for *U* in Theorem 3.25 yields:

$$\mathbf{G}(V, \mathbf{L}(V, P)) \sqsubseteq P$$

Corollary 3.26 justifies the intuition that applying L(V) to a system *P* may discard information about *P*'s behaviours that cannot be recovered by applying **G**(*V*) to the result.

In Chapter 4 of the UTP textbook, Hoare and He emphasise the value of Galois connections for linking UTP theories. Theorem 3.25 links UTP theories with counterparts of those theories defined over the space of user interactions. We explore these counterpart theories in Section 3.4 and Section 3.5.

Remark 3.27. A view *V* can be interpreted as a linking predicate between two data types, where behaviours are instances of the concrete data type, while interactions are instances of the abstract data type. In this setting, L(V) is an abstraction function and G(V) is a concretisation function.

This insight suggests how *abstract interpretation* (Cousot and Cousot, 1977, 1992) — a theory for approximating the behaviour of programs to simplify proofs of correctness properties, by abstracting away irrelevant detail — could be formulated in the UTP. By defining a view *V* to retain only information about a program *P* essential for a property, it would be more efficient to prove the property over L(V, P) rather than *P*.

3.3.4 Inference

Suppose a user knows the structure of a system *P* and its own view *V*. On making a interaction ψ with *P*, the user can calculate the space of all behaviours of *P* that *may* have generated ψ . This calculation is specified by Definition 3.28.

Definition 3.28 (Inference function).

infer
$$(P, V, \psi) \triangleq P \land \neg \mathbf{G} (V, \neg \psi)$$

Corollary 3.29. Lemma 3.24 implies that whenever *V* is **VH3**-healthy:

infer
$$(P, V, \psi) = P \land \exists x_V, x'_V \bullet \Delta V \land \psi$$

A function similar to that presented in Corollary 3.29 is called the *inference function* by Jacob (1988, 1989a). It expresses the maximum information that a user at *V* can deduce about the behaviour of *P* from its interaction ψ alone.

Example 3.30. Suppose Alice plays a simple guessing game $GUESS_0$. In this game, a number $n \in 1..10$ is chosen and concealed from Alice. Alice

makes a guess g of the value of n, whereupon she is informed whether her guess is correct or greater or smaller than n, via the response variable r'. We specify the guessing game as follows:

$$GUESS_0 \quad \triangleq \quad n \in 1..10 \land g \in 1..10 \land \begin{pmatrix} g > n \Rightarrow r' > 0 \\ \land g = n \Rightarrow r' = 0 \\ \land g < n \Rightarrow r' < 0 \end{pmatrix}$$

Alice's view allows her to observe the values of *g* and *r*, but not *n*:

ALICE
$$\triangleq g_A = g \wedge r_A = r$$

If Alice guesses g = 7 and she is informed that r' > 0, we find:

infer
$$(GUESS_0, ALICE, g_A = 7 \land r'_A > 0)$$

= $GUESS_0 \land \exists g_A, r_A, g'_A, r'_A, \bullet \Delta ALICE \land g_A = 7 \land r'_A > 0$ [def infer]
= $n \in 8..10 \land g = 7 \land r' > 0$ [pred calc]

 \Diamond

Hence, Alice can infer that $n \in 8..10$ from her interaction.

From the perspective of the user at *V*, each behaviour of infer (P, V, ψ) is consistent with ψ . The stronger infer (P, V, ψ) is, the *fewer* behaviours of *P* are consistent with ψ , enabling the user at *V* to draw stronger inferences about the behaviour of *P*.

Remark 3.31. The inference function can also be used to calculate what information a user at V_1 can infer about the possible interactions through a disjoint view V_2 belonging to another user. For instance:

$$L(V_2, infer(P, V_1, \psi))$$

models all *V*₂-interactions with *P* that comply with a *V*₁-interaction ψ .

3.4 Multi-User Designs

So far, we have studied multi-user systems in the basic setting of alphabetised relations. We now investigate how the concepts we have developed can be specialised for the theory of designs.

3.4.1 Views for Designs

We interpret a UTP design *D* as starting in an initial state consisting of user inputs and terminating (or aborting) in a final state output to users.

It is reasonable to expect that the users of *D* can observe whether *D* has started or has terminated. Since the behavioural variables ok and ok' are hidden from users, it is necessary to extend the alphabet of a view *V* by introducing a new interaction variable ok_V to mirror ok. We formalise the expectation that $ok_V = ok$ (and $ok'_V = ok'$ by extension) with a new healthiness condition.

Definition 3.32 (OK healthiness condition).

$$OK(V) \triangleq V \wedge ok_V = ok$$

Definition 3.33 (VHD). A view is **VHD**-healthy if and only if it is both **VH**-healthy and **OK**-healthy.

Henceforth, we require that views applied to designs are **VHD**-healthy.

3.4.2 Projecting Designs

The predicate that results from projecting a design through a **VHD**-healthy view is not itself a design, because it is defined over interactions, not behaviours. Yet this predicate is isomorphic to designs, so it possesses the special properties of designs in a renamed observational space. For this reason, we introduce a notation for projected designs in Definition 3.34.

Definition 3.34 (Local design). Let *V* denote a view and *Pre*, *Post* denote predicates with alphabet x_V, x'_V :

$$Pre \vdash_V Post \triangleq ok_V \land Pre \Rightarrow ok'_V \land Post$$

The **L** predicate transformer can be applied to a **VHD**-healthy view V and design D, to obtain an interface design that expresses the behaviours of D as projected through V.

Lemma 3.35 (L and designs). Provided *V* is **VHD**-healthy, $L(V, Pre \vdash Post)$ can always be expressed as a local design:

$$\mathsf{L}(V, Pre \vdash Post) = \neg \mathsf{L}(V, \neg Pre) \vdash_{V} \mathsf{L}(V, Post)$$

Moreover, if *Pre* is a condition (i.e. its alphabet contains only undecorated variables), then:

$$\mathsf{L}(V, Pre \vdash Post) = (\forall x \bullet V \Rightarrow Pre) \vdash_V \mathsf{L}(V, Post)$$

The precondition $PreL = (\forall x \bullet V \Rightarrow Pre)$ is the weakest predicate over interactions such that *all* behaviours of *D* that are consistent with each interaction of *PreL* satisfy *Pre*. Therefore, if a user's interaction with *D* satisfies $ok_V \land PreL$, then that user can be certain that $ok \land Pre$ holds. In turn, the user can ascertain that *Post* (and its projection L(V, Post)) will hold upon *D* terminating.

Remark 3.36. If no projection of *D* through *V* guarantees that *Pre* holds, then the precondition of L(V,D) collapses to **false**. It follows that:

$$false \vdash_{V} L(V, Post)$$

$$= ok_{V} \land false \Rightarrow ok'_{V} \land L(V, Post)$$
[Definition 3.34]
$$= true$$
[prop calc]

and so a user at *V* can infer nothing about the final state of L(V, D).

3.4.3 Worked Example: A Byte Register

We now consider how the theory developed in the previous sections can be applied to a simple multi-user system. The aim of this example is to demonstrate how the specification and design of multi-user systems may be guided by knowledge of user interactions.

We focus on a register capable of storing a single (8-bit) byte. We model the register's value by an integer variable *reg* with domain 0..255. The register also features a binary-valued flag *err* that indicates numeric overflow when raised. The state space of the register is described by:

$$ST \triangleq reg \in 0..255 \land err \in 0..1$$

Consider an operation that doubles the value stored in *reg*, provided that the initial value of *reg* lies in the range 0..127 and the overflow flag is not raised. We model this operation as a UTP design as follows:

$$DBL \triangleq reg \in 0..127 \land err = 0 \vdash reg' = reg \times 2 \land err' = 0$$

Suppose that two users can observe the register: the first user can observe the values of the higher four bits of the value of *reg*, and the second user can observe the lower four bits of *reg*. The overflow flag *err* is visible to both users. The views of these users are defined as follows:

$$H \triangleq \mathbf{OK} \left(reg_H = \left\lfloor \frac{reg}{16} \right\rfloor \land err_H = err \right)$$
$$L \triangleq \mathbf{OK} \left(reg_L = reg \mod 16 \land err_L = err \right)$$

Effectively, the *H* view masks out the lower four bits of the register from the first user's perspective, while the *L* view hides the higher four bits from the second user. Both of these views are **VHD**-healthy.

Example 3.37. Consider the instantiation $\phi = (reg = 60 \land err = 0)$ of *ST*. This state projects through *H* and *L* as follows:

$$\mathbf{L}(H,\phi) = reg_H = 3 \wedge err_H = 0$$

$$\mathbf{L}(L,\phi) = reg_L = 12 \wedge err_L = 0$$

Notice the binary representation of *reg* (00111100) is given by concatenating reg_H (0011) with reg_L (1100).

Using Lemma 3.35, we now calculate the projections of *DBL*'s behaviour through the views *H* and *L*, constrained by the register's state space *ST*.

This calculation indicates the user at *H* can be sure that *DBL*'s precondition is satisfied if and only if $reg_H \in 0..7$ (i.e. $reg \in 0..127$), since any other value of reg_H corresponds to a value of reg that violates the precondition of *DBL*.

In the postcondition of *DBL*, the value of the fifth most significant bit of *reg* determines whether $reg'_H = reg_H \times 2$ or $reg'_H = (reg_H \times 2) + 1$. Since that bit of *reg* cannot be observed through *H*, the user at *H* can only be sure of the value of reg'_H once the operation is complete.

It is instructive to consider what can be observed of *DBL* through *L*:

$$L (L \land ST, DBL)$$

$$= \begin{pmatrix} \forall x \bullet (L \land ST) \Rightarrow reg \in 0..127 \land err = 0 \\ \vdash_L \exists x, x' \bullet \Delta(L \land ST) \land reg' = reg \times 2 \land err' = 0 \end{pmatrix}$$

$$= false \vdash_L (reg'_L = (reg_L \times 2) \mod 16 \land err'_L = 0)$$
 [pred calc]
$$= true$$
 [Remark 3.36]

Since an observation at *L* provides no information regarding the most significant bit of *reg*, a user at *L* cannot distinguish between $reg \in 0..127$ or $reg \in 128..255$. Hence, that user cannot determine in any circum-

stances whether the precondition of *DBL* is satisfied. The outcome of this calculation reflects that, from that user's perspective, nothing can be guaranteed about *DBL*'s behaviour.

Depending on the context in which the register is used, this limitation may be unacceptable. Hence, we relax the precondition of *DBL* to cover all values of *reg* that can be stored by the register and, when $reg \ge 128$, to assign an arbitrary value from the range 0..255 to *reg'* and set *err'* to 1:

$$DBL2 \triangleq reg \in 0..255 \land err = 0 \vdash \left(\begin{array}{c} reg' = reg \times 2 \land err' = 0 \\ \lhd reg \in 0..127 \vartriangleright \\ reg' \in 0..255 \land err' = 1 \end{array} \right)$$

Observe that $DBL \sqsubseteq DBL2$, because the postcondition of DBL2 reduces to the postcondition of DBL when the precondition of DBL is satisfied.

The projection of *DBL*2 through *L* is as follows:

After the *DBL*2 operation completes, the user at *L* can verify the doubling operation was successful by checking that its interaction satisfies the precondition of L(L, DBL2) and that $err'_L = 0$.

Remark 3.38. If the state of *DBL*2 is data-refined, then the corresponding data refinements must also be made to the *H* and *L* views. For instance, if we were to data-refine *reg* into eight binary variables r7, ..., r0 to

represent the bits of the register, we would replace *H* with the view:

OK
$$(r7_H = r7 \land r6_H = r6 \land r5_H = r5 \land r4_H = r4 \land err_H = err)$$

which could be calculated by projecting a linking predicate through *H*.

3.5 Multi-User Reactive Designs

We now apply our framework to the UTP theory of reactive processes. This work allows us to formalise users' interactions with CSP or *Circus* processes, since these formalisms have a reactive design semantics.

3.5.1 Reactive Views

Our first task is to formalise how users can interact with reactive processes, by defining a family of *reactive views*. It is natural to interpret a user's interface in terms of the events that a reactive process communicates with its environment.

Definition 3.39 (Window). A user's *window* is the set of events communicated by a reactive process which are visible to the user.

We now consider which observational elements of a reactive process are relayed to a user with window W.

- **State** A process does not expose its internal state to the external environment. Hence, a view should not convey any information about the state variables v to the user.
- **Stability** We suppose the user can observe whether the process has reached a stable point in its execution by requiring its view to be **OK**-healthy, for the same reasons as those given in Section 3.4.
- **Progress** We enable the user to determine whether the process is waiting to start, is waiting for any interaction with the environment or has terminated, by expanding the user's view with a local variable $wait_V$ corresponding to *wait*.

Trace Since events take place sequentially, we model a user as perceiving their order of occurrence. Hence, a user's observation of the global trace *tr* is given by restricting *tr* to only those events in *W*:

$$tr_V = tr \upharpoonright \mathcal{W}$$

where $tr \upharpoonright W$ denotes tr, but with all events outside W erased.

Refusals A user who interacts with the process directly may discover that, at some point, the events it wishes to perform with the process are refused. We model the projection of a process's refusal set as:

$$ref_V \cap \mathcal{W} \subseteq ref$$

This expression only reveals information about the events of *ref* in W to the user.

Remark 3.40. At first glance, it would seem more natural to define:

$$ref_V \subseteq ref \cap \mathcal{W}$$

but this expression implies that, from the user's perspective, the process can never refuse any event outside W. When combined with the projection functions for the other observational elements, this expression leads to projected processes that violate the failures model of CSP: they would not engage in events outside W, but would not refuse those events either.

We impose these expectations on views by defining another healthiness condition **VHR**.

Definition 3.41 (VHR healthiness condition).

$$\mathbf{VHR}(\mathcal{W}, V) \triangleq \begin{pmatrix} \exists v \bullet \mathbf{VH}(V) \\ \land & ok_V = ok \land wait_V = wait \\ \land & tr_V = tr \upharpoonright \mathcal{W} \land ref_V \cap \mathcal{W} \subseteq ref \end{pmatrix}$$

where v denotes the list of state variables, excluding ok, wait, tr, ref.

Following **VHD**, the inclusion of *ok* in **VHR** is vital to support any meaningful model of a user's interactions with a process. The *wait* variable, while not essential, is still important if we wish to study a user's interactions with a process in terms of its component actions. (Indeed, compositional reasoning plays a key role in Chapter 5.)

The inclusion of the trace and refusal variables is less important. If we take $W = \emptyset$, then we derive a minimal interface representing a user who cannot engage in activity with the process, but who could still perceive divergence and termination via $\neg ok'_V$ and $\neg wait'_V$ respectively.

Remark 3.42. We may supply **VHR** with views modelling other facets of a user's interface to a process. For instance, an interface that indicates whether the process engages in an event in \mathcal{Y} before an event in \mathcal{Z} (where \mathcal{Y} and \mathcal{Z} are disjoint from \mathcal{W}) is modelled by the view:

$$\mathsf{VHR}\left(\mathcal{W}, xy_V \Leftrightarrow \exists \alpha, \beta \bullet \left(\begin{array}{c} tr = \alpha \cap \beta \land \alpha \upharpoonright \mathcal{Y} \neq \langle \rangle \\ \land \alpha \upharpoonright \mathcal{Z} = \langle \rangle \land \beta \upharpoonright \mathcal{Z} \neq \langle \rangle \end{array}\right)\right)$$

These more elaborate views inhibit compositional reasoning about users' interactions with processes, which would be a hindrance in later chapters. For this reason, we do not consider these views further.

Recall from Section 2.5 that a **R2**-healthy process *P* is insensitive to the value of *tr*. If we limit our attention to **R2**-healthy processes, then the projection tr_V of *tr* given by **VHR** is redundant:

R2 $(tr_V = tr \upharpoonright \mathcal{W} \land tr'_V = tr' \upharpoonright \mathcal{W})$		
=	$(tr_V = tr \upharpoonright \mathcal{W} \land tr'_V = tr' \upharpoonright \mathcal{W})[\langle \rangle, tr' - tr/tr, tr']$	[def R2]
=	$tr_V = \langle angle \upharpoonright \mathcal{W} \wedge tr'_V = (tr' - tr) \upharpoonright \mathcal{W}$	[substitution]
=	$tr_V = \langle angle \wedge tr'_V = (tr' - tr) \upharpoonright \mathcal{W}$	[property of ↾]
=	$tr_V = \langle \rangle \wedge tr'_V - tr_V = (tr' - tr) \upharpoonright \mathcal{W}$	[property of –]

We adopt a slight variation of Δ **VHR** (W, **true**) as our model of a user's interface to a reactive process, given in Definition 3.43. This variation

omits the $tr_V = \langle \rangle$ conjunct.

Definition 3.43 (Reactive interface).

$$\mathcal{R}(\mathcal{W}) \triangleq \begin{pmatrix} ok_{V} = ok \land ok'_{V} = ok' \\ \land wait_{V} = wait \land wait'_{V} = wait' \\ \land tr'_{V} - tr_{V} = (tr' - tr) \upharpoonright \mathcal{W} \\ \land ref_{V} \cap \mathcal{W} \subseteq ref \land ref'_{V} \cap \mathcal{W} \subseteq ref' \end{pmatrix}$$

To some extent, the decision to work with interfaces of the $\mathcal{R}(\mathcal{W})$ form is arbitrary. Nevertheless, there is a close relationship between this form and existing work in the literature, which is revealed in Subsection 3.5.3.

Remark 3.44. While $\mathcal{R}(\mathcal{W})$ includes a projection of *ref*, that variable has no significance when dealing with **CSP3**-healthy processes, such as *Circus* actions (Oliveira et al., 2009).

3.5.2 Projecting a Reactive Process

For convenience, we define a specialised version of the **L** predicate transformer that encodes $\mathcal{R}(\mathcal{W})$ in place of a view.

Definition 3.45 (**LR** predicate transformer). Provided *P* is **R2**-healthy:

$$\mathsf{LR}(\mathcal{W}, P) \triangleq \mathsf{L}(\mathsf{VHR}(\mathcal{W}, \mathsf{true}), P)$$
$$= \exists x, x' \bullet \mathcal{R}(\mathcal{W}) \land P$$

Applying **LR** to a **R**-healthy reactive process does not yield a **R**-healthy process, but a predicate that is isomorphic to a **R**-healthy process.

Definition 3.46 (Local **R**). Let $\mathbf{R}_V(P) \triangleq \mathbf{R}\mathbf{1}_V \circ \mathbf{R}\mathbf{2}_V \circ \mathbf{R}\mathbf{3}_V(P)$ and:

$$\mathbf{R1}_{V}(P) \triangleq P \wedge tr_{V} \leq tr'_{V}$$

$$\mathbf{R2}_{V}(P) \triangleq P[\langle\rangle, tr'_{V} - tr_{V}/tr_{V}, tr'_{V}]$$

$$\mathbf{R3}_{V}(P) \triangleq ((\exists v, v' \bullet \Pi_{rea})[x_{V}, x'_{V}/x, x'] \lhd wait_{V} \rhd P)$$

If *P* is **R**-healthy, then LR(W, P) is R_V -healthy. Since a reactive design is expressible as a design made **R**-healthy, Lemma 3.47 presents a result similar to Lemma 3.35.

Lemma 3.47 (LR and reactive designs). Provided *P* is a reactive design:

$$\mathsf{LR}(\mathcal{W}, P) = \mathsf{R}_{V}\left(\left(\forall x, x' \bullet \mathcal{R}(\mathcal{W}) \Rightarrow \neg P_{f}^{f}\right) \vdash_{V} \mathsf{LR}\left(\mathcal{W}, P_{f}^{t}\right)\right)$$

where P_c^b abbreviates P[b, c/ok', wait] (as used in Theorem 2.2).

The predicate $\forall x, x' \bullet \mathcal{R}(\mathcal{W}) \Rightarrow \neg P_f^f$ describes all interactions for which a user at \mathcal{W} can be certain that P does not diverge. The same effect manifests in the preconditions of multi-user designs, as described in Section 3.4.

Example 3.48. The reactive design semantics of $a \rightarrow b \rightarrow Skip$ is:

$$\mathbf{R}\left(\mathbf{true} \vdash \left(\begin{array}{c} (tr' = tr \land a \notin ref') \lor (tr' = tr \frown \langle a \rangle \land b \notin ref') \\ \lhd wait' \rhd \\ tr' = tr \frown \langle a, b \rangle \end{array}\right)\right)$$

The **LR**-projection of $a \rightarrow b \rightarrow Skip$ through the window $\{a\}$ is:

$$\mathbf{LR} \left(\{a\}, a \to b \to Skip \right)$$

$$= \mathbf{R}_{V} \left(\begin{pmatrix} (\forall x, x' \bullet \mathcal{R} (\{a\}) \Rightarrow \neg (a \to b \to Skip)_{f}^{f} \\ \vdash_{V} \mathbf{LR} (\{a\}, a \to b \to Skip)_{f}^{t} \end{pmatrix} \right)$$
[Lemma 3.47]
$$= \mathbf{R}_{V} \left(\mathbf{true} \vdash_{V} \mathbf{LR} \left(\{a\}, (a \to b \to Skip)_{f}^{t} \right) \right)$$
[no divergence]

Focusing on the postcondition of the interior local design, we obtain:

$$\mathsf{LR}\left(\left\{a\right\}, (a \to b \to Skip)_{f}^{t}\right)$$

$$= \mathsf{LR}\left(\left\{a\right\}, \left(\begin{array}{cc} tr' = tr \land a \notin ref' \\ \lor tr' = tr \land a \land b \notin ref' \\ \lor wait' \rhd \\ tr' = tr \land a, b \rangle \end{array}\right)\right)$$
[semantics]

$$= \begin{pmatrix} \left(\begin{array}{c} \mathbf{LR}(\{a\}, tr' = tr \land a \notin ref') \\ \lor \ \mathbf{LR}(\{a\}, tr' = tr \land \langle a \rangle \land b \notin ref') \\ \lhd wait'_{V} \rhd \\ \mathbf{LR}(\{a\}, tr' = tr \land \langle a, b \rangle) \end{pmatrix} \quad [Law \ 3.19]$$

$$= \begin{pmatrix} (tr'_{V} = tr_{V} \land a \notin ref'_{V}) \lor (tr'_{V} = tr_{V} \land \langle a \rangle) \\ \lhd wait'_{V} \rhd \\ tr'_{V} = tr_{V} \land \langle a \rangle \end{pmatrix} \quad [def \ \mathbf{LR}]$$

Hence, **LR** ({*a*}, *a* \rightarrow *b* \rightarrow *Skip*) is equal to:

$$\mathbf{R}_{V}\left(\mathbf{true}\vdash_{V}\left(\begin{array}{c} (tr'_{V}=tr_{V}\wedge a\notin ref'_{V})\vee(tr'_{V}=tr_{V}\frown\langle a\rangle)\\ \lhd wait'_{V}\rhd\\ tr'_{V}=tr_{V}\frown\langle a\rangle\end{array}\right)\right)$$

which is isomorphic to the semantics of the process $a \rightarrow (Skip \sqcap Stop)$.

Likewise, the projection of $a \rightarrow b \rightarrow Skip$ through the window $\{b\}$ is:

$$\mathbf{R}_{V}\left(\mathbf{true}\vdash_{V}\left(tr'_{V}=tr_{V}\lhd \mathit{wait}'_{V}\vartriangleright tr'_{V}=tr_{V}\land \langle b\rangle\right)\right)$$

which is isomorphic to the process *Stop* \sqcap *b* \rightarrow *Skip*.

 \Diamond

3.5.3 Lazy Abstraction

In two papers (Roscoe et al., 1994; Roscoe, 1995), Roscoe introduced the notion of the *lazy abstraction* of a CSP process. The lazy abstraction of a process *P*, denoted by Lazy (W, *P*), is a process describing the interactions that a user with window W can make of *P*.

Definition 3.49 reproduces Roscoe's definition of lazy abstraction.

Definition 3.49 (Lazy abstraction). The lazy abstraction of a CSP process P to the user at window W is:

 $\mathsf{Lazy}\,(\mathcal{W}, P) \quad \triangleq \quad (P \, [\![\, \mathcal{H}\,]\!] \, (\mu \, X \bullet Stop \sqcap (\sqcap e : \mathcal{H} \bullet e \to X))) \setminus \mathcal{H}$

where \mathcal{H} denotes $(\Sigma \setminus \mathcal{W})$, the set of all events outside \mathcal{W} .

The process³ $\mu X \bullet Stop \sqcap (\sqcap e : \mathcal{H} \bullet e \to X)$ represents the activities of *P* that are controlled by the other users of *P*. Placing this process in parallel with *P* and hiding the \mathcal{H} events models the behaviours of *P* as observed by the user at \mathcal{W} .

Example 3.50. The lazy abstractions of $a \rightarrow b \rightarrow Skip$ through the windows $\{a\}$ and $\{b\}$ respectively are:

$$\begin{aligned} \mathsf{Lazy}\left(\left\{a\right\}, a \to b \to Skip\right) &= a \to (Skip \sqcap Stop) \\ \mathsf{Lazy}\left(\left\{b\right\}, a \to b \to Skip\right) &= Stop \sqcap b \to Skip \end{aligned}$$

These processes mirror those derived in Example 3.48.

Roscoe (1997) presents multiple forms of lazy abstraction defined over the various semantic models of CSP and proves their equivalence. We build on this work to establish a correspondence between the failuresdivergences version of lazy abstraction and our **LR** predicate transformer.

Theorem 3.51 (Lazy abstraction as projection). For any *divergence-free* CSP process *P*, Lazy (W, *P*) is isomorphic to LR (W, P_R), where P_R denotes the reactive design formulation of *P*.

Theorem 3.51 is useful, because it allows us to characterise **LR** algebraically. Hence, existing tools and techniques designed for calculating and analysing lazy abstractions can be ported to our framework. Theorem 3.51 also validates our reactive interface \mathcal{R} (Definition 3.43). In the opposite direction, Theorem 3.51 suggests that lazy abstraction could be generalised by adopting a different interface specification instead of \mathcal{R} .

3.6 Refining Multi-User Systems

In previous sections, we have described how the UTP can be applied to calculate a user's interactions with systems. Following Jacob (1987), this

 \Diamond

³This process is called $CHAOS(\mathcal{H})$ by Roscoe (1997). It should not be confused with the *Chaos* action of *Circus*: while $CHAOS(\mathcal{H})$ can accept or refuse any sequence of events drawn from \mathcal{H} , it never diverges, whereas *Chaos* can diverge.

section argues that it is permissible to weaken the canonical notion of refinement when developing multi-user systems. We also define multiuser notions of refinement in terms of the canonical refinement relation.

Intuitively, $P_1 \sqsubseteq P_2$ formalises the following *correctness criterion*:

CC1 P_2 may replace P_1 without its environment detecting the change.

This criterion assumes the existence of an entity which can monitor the entire environment of the system, and thereby observe its behaviour completely. However, in a multi-user system, the environment is structured by the interfaces of its users. If no user's interface spans the entirety of the environment, then the criterion is arguably too strong.

If we are developing a multi-user system, we may wish to weaken our correctness criterion to:

CC2 P_2 may replace P_1 without its users detecting the change.

Depending on the users' views, there are instances of P_1 and P_2 which uphold this criterion, but for which $P_1 \sqsubseteq P_2$ fails to hold.

Example 3.52. Returning to the byte register (Subsection 3.4.3), consider the following variant of the *DBL*2 operation that doubles the lower four bits of *reg* and the upper four bits of *reg* separately:

$$INDBL2 \triangleq reg \in 0..255 \land err = 0 \vdash \begin{pmatrix} HDBL \land LDBL \land err' = 0 \\ \lhd reg \in 0..127 \triangleright \\ reg' \in 0..255 \land err' = 1 \end{pmatrix}$$
$$HDBL = \lfloor reg'/16 \rfloor = (\lfloor reg/16 \rfloor \times 2)$$
$$LDBL = reg' \mod 16 = (reg \times 2) \mod 16$$

INDBL2 allows the users at *H* and *L* to access separate registers without needing to keep those registers synchronised, which means that an implementation of *INDBL2* may provide each user with their own local instance of the register.

Notice that *DBL*2 $\not\sqsubseteq$ *INDBL*2, because when $reg \in 0..127$ and err = 0, *INDBL*2 always sets the fourth most significant bit of reg' to 0 regardless

of the value of the fifth most significant bit of *reg*. However, the users at *H* and *L* cannot individually distinguish *INDBL2* from *DBL2*, because *INDBL2* provides the same interactions to each user as does *DBL2*. \diamond

This section outlines a class of refinement relations — based on work by Jacob (1987, 1989b, 1992) — formulated over user interactions with a system, rather than the behaviours of the system. These refinement relations are more relaxed than behavioural notions of refinement: they satisfy **CC2** but not necessarily **CC1**.

3.6.1 Local Refinement

We call P_2 a *local refinement* of P_1 with respect to a view V if and only if each V-interaction with P_2 is also a V-interaction with P_1 .

Definition 3.53 (Local refinement).

$$P_1 \sqsubseteq_V P_2 \triangleq \mathsf{L}(V, P_1) \sqsubseteq \mathsf{L}(V, P_2)$$

Local refinement allows new behaviours to be *added* to a system, provided these new behaviours do not induce new interactions through *V*. Hence, if $P_1 \sqsubseteq_V P_2$ holds, then a user at *V* cannot detect the replacement of P_1 with P_2 .

Corollary 3.54. Definition 3.53 and Law 3.20 imply, for all V:

$$P_1 \sqsubseteq P_2$$
 implies $P_1 \sqsubseteq_V P_2$

Unlike \sqsubseteq , the \sqsubseteq_V ordering is not antisymmetric in general, because two different systems offering the same interactions through *V* are equivalent under \sqsubseteq_V . Nevertheless, \sqsubseteq_V is always a pre-order.

Remark 3.55. If P_1 and P_2 are CSP processes and V = VHR(W, true) then, by Theorem 3.51, $P_1 \sqsubseteq_V P_2$ is equivalent to the condition:

$$Lazy(\mathcal{W}, P_1) \subseteq Lazy(\mathcal{W}, P_2)$$

This result implies $P_1 \sqsubseteq_V P_2$ can be tested with a CSP model checker.

Since local refinement is formulated in terms of a single user, it needs to be extended if we wish to reason about multiple users.

3.6.2 Co-operating and Independent Refinement

Jacob (1987) proposed two notions of refinement, known as *co-operating refinement* and *independent refinement*, intended for the development of multi-user systems.

- **Co-operating refinement** assumes that users can exchange information with each other about their interactions with a system. This means the users could reconstruct more information about the system behaviour than they could infer from their individual interactions.
- **Independent refinement** assumes users cannot communicate with each other; instead, the only information that each user can obtain about the behaviour of a system is from their own interaction.

Jacob defined co-operating and independent refinement in the form:

$$P_1 \sqsubseteq_{\mathcal{A}}^{co} P_2 \triangleq \forall \phi_2 : \omega(P_2) \bullet \exists \phi_1 : \omega(P_1) \bullet \forall U : \mathcal{A} \bullet \phi_1 \cong_U \phi_2$$
$$P_1 \sqsubseteq_{\mathcal{A}}^{ind} P_2 \triangleq \forall U : \mathcal{A} \bullet \forall \phi_2 : \omega(P_2) \bullet \exists \phi_1 : \omega(P_1) \bullet \phi_1 \cong_U \phi_2$$

where \mathcal{A} denotes a set of users, $\omega(P)$ denotes the space of behaviours of P (with respect to some semantic model), and $\phi_1 \cong_U \phi_2$ holds if and only if the behaviours ϕ_1 and ϕ_2 yield the same interaction to the user U.

We express co-operating and independent refinement in the UTP by extending the definition of local refinement to a set of pairwise disjoint views *VS* (representing multiple users) in different ways.

Definition 3.56 (Co-operating and independent refinement).

$$P_1 \sqsubseteq_{VS}^{co} P_2 \triangleq P_1 \sqsubseteq_{\wedge VS} P_2$$
$$P_1 \sqsubseteq_{VS}^{ind} P_2 \triangleq \bigwedge_{V \in VS} P_1 \sqsubseteq_V P_2$$

The UTP style obscures the essential equivalence between these definitions and Jacob's. For co-operating refinement, $P_1 \sqsubseteq_{\Lambda VS} P_2$ demands that for each behaviour ϕ_2 of P_2 , there exists a behaviour ϕ_1 of P_1 such that ϕ_1 and ϕ_2 yield the same (ΛVS)-interaction, where the view ΛVS represents the combined interactions made by the users (Subsection 3.2.3). A similar comparison can be made for independent refinement.

As with \sqsubseteq_V , the \sqsubseteq_{VS}^{co} and \sqsubseteq_{VS}^{ind} orderings are pre-orders but not partial orders (Jacob, 1989b). When $VS = \{V\}$, both \sqsubseteq_{VS}^{co} and \sqsubseteq_{VS}^{ind} reduce to \sqsubseteq_V .

Co-operating refinement is no stronger than behavioural refinement:

$$P_1 \sqsubseteq P_2$$
 implies $P_1 \sqsubseteq_{VS}^{co} P_2$

which follows from Law 3.20 and the definition of local refinement. Similarly, it follows from Law 3.21 that independent refinement is no stronger than co-operating refinement:

$$P_1 \sqsubseteq_{VS}^{co} P_2$$
 implies $P_1 \sqsubseteq_{VS}^{ind} P_2$

Even though cooperating and independent refinement are generally weaker than behavioural refinement, they are strong enough to preserve the functionality inherent in a system's specification *from the perspectives of the users*, provided their assumptions about users are upheld.

Example 3.57. Returning to Example 3.52, it is the case that:

$$DBL2 \sqsubseteq_{\{H,L\}}^{ind} INDBL2$$
 but not $DBL2 \sqsubseteq_{\{H,L\}}^{co} INDBL2$

Provided the register can only be accessed through H and L — and the users at H and L do not co-operate — the replacement of DBL2 with INDBL2 is justified by independent refinement. However, this replacement is not justified by co-operating refinement, because if the users at H and L combine their interactions, they can identify behaviours of INDBL2 that are not behaviours of DBL2.

We do not study co-operating or independent refinement further in this

thesis. Nevertheless, as Example 3.57 demonstrates, they provide system developers with an extra degree of flexibility in making design choices than behavioural refinement can offer. Indeed, they would be particularly appropriate in the design of distributed multi-user systems, because they offer the opportunity to distribute a system's workload across multiple processors without needing to keep those processors synchronised.

3.7 Conclusion

This chapter has presented a generic framework for studying the interactions between systems and their users. By separating the concern of modelling users from designing systems, we can specify directly which elements of a system's behaviour are viewable by users. With predicate transformers, these specifications of views can be applied to calculate a user's interactions with a system, as well as a user's inferences about the system's behaviour from those interactions.

This framework fits seamlessly within the predicate semantics of the UTP. Hence, it can be applied to derive a multi-user interpretation of any UTP theory. We have demonstrated its application to the theories of designs and reactive processes and, in the latter case, have identified a link with Roscoe's lazy abstraction.

We have also applied the framework to capture weaker notions of refinement that are suitable for multi-user system development. The connections between our framework, Roscoe's lazy abstraction and the notions of co-operating and independent refinement together imply there is scope for using model checkers for CSP for automatically verifying these weaker refinement relations hold between processes.

The notions of views and inferences are taken forward to the next chapter, where we study information flow security in the UTP context.

4 Confidentiality Properties

4.1 Introduction

Confidentiality demands that a system only releases secret information — such as cryptographic keys or classified documents — to users who are authorised to access that information. Since the advent of time-sharing computer systems in the late 1960s, much research has focused on protecting the confidentiality of data stored within multi-user systems. The concepts of *access control* and *inference control* have proved influential, in different ways.

4.1.1 Access Control

Access control models — most famously, the Bell–La Padula (BLP) model (Bell and La Padula, 1973, 1976) — describe an abstract mechanism for controlling the operations on the system's state that each user of the system is entitled to perform. It is intended that, if a system is implemented in accordance with this mechanism, then the mechanism will prevent a user from invoking an operation (or sequence of operations) to access data that it is not authorised to access.

Access control is readily comprehensible by software engineers and is widely deployed in software products. Yet despite its intuitive appeal, its adequacy for describing security policies has been criticised on both theoretical and practical grounds:

Covert channels A *covert channel* allows users to communicate data in a manner unanticipated by the system's designers (Lampson, 1973). These channels often emerge from shared resources: for instance, a

4 Confidentiality Properties

high-level user (or a Trojan horse program invoked with high-level privileges) may modulate its consumption of a system resource (such as processor time or memory) to signal data to a low-level user. Covert channels lie outside the domain of access control, because they are an emergent property of the system's resources.

- **Narrow focus** Access control models are formulated in terms of active subjects (users) operating over a collection of passive objects, such as files and database records. However, confidentiality requirements may apply to other facets of a system, such as whether a high-level user has engaged in particular activities.
- **Semantic ambiguity** It is the job of software engineers to determine how an access control model should be mapped onto a system design. For instance, models such as BLP are specified in terms of "read" and "write" operations, but the semantics of these operations is not made precise. Such ambiguities can lead to insecurity: for instance, McLean (1987) outlined a hypothetical "System Z" which implements BLP, but downgrades the security level of all data objects for the duration of each operation it performs, enabling low-level users to access any object whatsoever.

4.1.2 Inference Control

A radically different philosophy for protecting the confidentiality of information is to regulate the *information flow* from systems to their users. We observed in Subsection 3.3.4 that a user can infer information about the system's behaviour by analysing its interaction with the system. By extension, that user can calculate information about the system's state, or the activities of other users (Denning, 1982).

A confidentiality property specifies which aspects of a system's behaviour are secret (in some way) and prescribes an upper bound on the flow of information about those aspects to low-level users. The meaning of confidentiality properties is unambiguous and is independent of the system under consideration. Unlike access control models, they do not prescribe any mechanism for safeguarding confidential data. Hence, confidentiality properties grant system designers the flexibility to choose how to implement security measures as they see fit (McLean, 1994b).

Various specialised techniques for verifying a system design against particular confidentiality properties have been proposed in the literature. However, the topic of integrating these techniques with existing highintegrity software engineering practices has remained under-explored.

4.1.3 Contributions of this Chapter

Building on the framework of user interactions developed in Chapter 3, this chapter outlines a style for specifying confidentiality properties over any aspect of the observation space defined by a UTP theory. This formulation of confidentiality properties provides a starting point for integrating notions of confidentiality into specification and programming languages with a UTP semantics.

Section 4.2 applies the predicate semantics of the UTP — coupled with the notion of views described in Chapter 3 — to formalise a user's inferences about the behaviour of systems. This work provides a semantic foundation for our definition of confidentiality properties, which is presented in Section 4.3.

Since confidentiality properties are inherently non-functional, they are typically kept separate from the functional specification of a system. In Section 4.4, we show how a specification of a *secure* system may be derived by merging a functionality specification with confidentiality properties. Then, in Section 4.5, we investigate why confidentiality properties are not preserved by refinement, but also show how this problem is resolved by merging functionality and confidentiality together.

In Section 4.6, we show how various information flow properties defined in the literature can be expressed in our confidentiality framework. We also compare our approach for specifying confidentiality with the body of existing work on information flow security in Section 4.7.

4 Confidentiality Properties

4.2 An Indistinguishability Semantics

Information flow is a product of the relationship between a system's behaviours and a user's inferences about those behaviours. In this section, we describe a method for extending UTP theories to encode information flow: in effect, we model a user's inferences about the behaviour of a system *in terms of the system's behaviours themselves*.

The cornerstone of our method is a mapping from the behaviours of a system to an isomorphic space of behaviours that we call the *fog space*. Intuitively, if a behaviour ϕ is mapped to a fog behaviour $\tilde{\phi}$, then $\tilde{\phi}$ represents an *alternative explanation* for ϕ in some respect. It is important to stress that our fog space is a fiction: we do not expect fog behaviours to manifest in any real system.

The fog space is conceptually very useful, because it enables us to formulate relational properties between behaviours that cannot be defined in the observable space alone. In particular, we employ the fog space to capture two different security concerns. In this section, we study how a user's inferences about the process's behaviour can be encoded in the fog. In Section 4.3, we investigate how the fog can serve as the medium underlying the specification of confidentiality properties.

4.2.1 Lifting Relations

We derive a fog alphabet $\tilde{x}, \tilde{x'}$ by renaming the observation variables x, x'. Given a relation *P*, the *lift* of *P* relates the observable behaviours of *P* to a renaming of those behaviours over the fog variables. Definition 4.1 presents a predicate transformer **U** ("up") for lifting relations.

Definition 4.1 (U predicate transformer). Let \tilde{P} denote $P[\tilde{x}, \tilde{x'}/x, x']$:

$$\mathbf{U}(P) \triangleq P \wedge \widetilde{P}$$

From this point onwards, we say that a predicate is a *lifted relation* if its alphabet contains the x, x' variables and the $\tilde{x}, \tilde{x'}$ fog variables.

Example 4.2. To understand the structure of lifted relations, it is helpful to expand them into a form where observable behaviours are paired with fog behaviours. For instance:

$$\begin{aligned} \mathbf{U} \left(h = 0 \lor h = 1 \right) \\ = & (h = 0 \lor h = 1) \land (\tilde{h} = 0 \lor \tilde{h} = 1) \\ = & \begin{pmatrix} (h = 0 \land \tilde{h} = 0) \lor (\tilde{h} = 0 \land \tilde{h} = 1) \\ \lor & (h = 1 \land \tilde{h} = 0) \lor (\tilde{h} = 1 \land \tilde{h} = 1) \end{pmatrix} \end{aligned} \qquad [prop calc] \end{aligned}$$

Here, each observable behaviour is paired with each fog behaviour. \Diamond

Remark 4.3. The **U** function is related to the *separating simulation* functions defined by Hoare and He (1998). Separating simulations rename the alphabet of a relation, to facilitate reasoning about updates to shared variables. Here, U(P) models two completely separate instances of a system with no shared variables between them.

Theorem 4.4 highlights an important property of **U**: it preserves the *structure* of a UTP theory.

Theorem 4.4 (U is an order-embedding). The **U** predicate transformer is both monotonic and order-reflecting under the refinement ordering:

$$P_1 \sqsubseteq P_2$$
 if and only if $U(P_1) \sqsubseteq U(P_2)$

A consequence of this order-embedding is that, by hiding the fog variables of a lifted relation U(P), we can recover the relation P. The **D** ("down") predicate transformer does just that.

Definition 4.5 (D predicate transformer).

$$\mathsf{D}(Q) \triangleq \exists \widetilde{x}, \widetilde{x'} \bullet Q$$

Lemma 4.6 is a simple consequence of Definition 4.5.

4 Confidentiality Properties

Lemma 4.6 (**D** inverse of **U**). Provided *P* is a relation:

$$\mathsf{D}\left(\mathsf{U}\left(P\right)\right) = P$$

Since **U** is defined in terms of conjunction, it is conjunctive. Likewise, **D** is disjunctive, because existential quantification is disjunctive.

Law 4.7 (U is conjunctive). $U(P_1 \land P_2) = U(P_1) \land U(P_2)$

Law 4.8 (D is disjunctive). $D(Q_1 \vee Q_2) = D(Q_1) \vee D(Q_2)$

However, **U** is not disjunctive, and **D** is not conjunctive.

Law 4.9 (U and disjunction). $U(P_1 \vee P_2) \sqsubseteq U(P_1) \vee U(P_2)$

Law 4.10 (D and conjunction). $\mathbf{D}(Q_1) \wedge \mathbf{D}(Q_2) \sqsubseteq \mathbf{D}(Q_1 \wedge Q_2)$

Since the conditional is defined as a disjunction (Subsection 2.3.2), Law 4.11 is a direct consequence of Law 4.9.

Law 4.11 (U conditional). $U(P_1 \triangleleft C \triangleright P_2) \sqsubseteq U(C \land P_1) \lor U(\neg C \land P_2)$

We established in Theorem 4.4 that U is monotonic. Likewise, D is monotonic, because existential quantification is monotonic.

Law 4.12 (D is monotonic). $Q_1 \sqsubseteq Q_2$ implies $D(Q_1) \sqsubseteq D(Q_2)$

Finally, negation does not distribute through **U** or **D**.

Law 4.13 (U and negation). \neg **U**(*P*) \sqsubseteq **U**(\neg *P*)

Law 4.14 (**D** and negation). $\mathbf{D}(\neg Q) \sqsubseteq \neg \mathbf{D}(Q)$

4.2.2 Modelling Inference

A long-standing principle for evaluating the security of a system is commonly known as Shannon's maxim (Shannon, 1949):

"the enemy knows the system being used"

According to this principle, we should assume this "enemy" — in our case, an adversarial user of a system P — possesses complete knowledge of the design of P. The user's complete knowledge of P is reflected by the definition of **U**: the space of fog behaviours in **U** (P) is no larger than the space of observable behaviours in **U** (P).

It follows from Shannon's maxim that a user can analyse its interaction ψ with *P* in tandem with its knowledge of *P* (and its interface) to infer all behaviours of *P* that are consistent with ψ . In Chapter 3, we modelled these deductions with the inference function (Definition 3.28). We now formalise a user's inferences about *P*'s behaviour in the lifted space.

Recall from Subsection 3.2.1 that a view *V* maps behaviours to interactions. We say two behaviours ϕ_1 and ϕ_2 are *indistinguishable* through a view *V* if and only if ΔV maps both ϕ_1 and ϕ_2 to the same interactions.

Definition 4.15 (Indistinguishability). ϕ_1, ϕ_2 are *V*-indistinguishable if and only if $L(V, \phi_1) = L(V, \phi_2)$

We capture a user's inability to distinguish between behaviours by defining an *indistinguishability relation* over the lifted space. An indistinguishability relation \mathcal{I} relates the observable behaviour ϕ_1 to the fog behaviour $\tilde{\phi_2}$ (and ϕ_2 to $\tilde{\phi_1}$) if and only if ϕ_1 and ϕ_2 yield the same interaction to the user. In this way, \mathcal{I} partitions the space of behaviours into equivalence classes.

Definition 4.16 presents a predicate transformer for constructing a (symmetric) indistinguishability relation from V, by reflecting the x_V variables back on themselves.

Definition 4.16 (IR predicate transformer).

$$\mathsf{IR}(V) \triangleq \Delta(\exists x_V \bullet V \land V[\widetilde{x}/x])$$

By imposing IR(V) upon the lifted space of U(P), we model a user's inferences about the behaviour of P in terms of the Low-indistinguishable fog behaviours of \tilde{P} . Definition 4.17 introduces a special form of U that incorporates IR in this way.

Definition 4.17 (UI predicate transformer).

$$\mathsf{UI}(V,P) \triangleq \mathsf{U}(P) \land \mathsf{IR}(V)$$

UI (*V*, *P*) relates each observable behaviour ϕ of *P* to *only* those fog behaviours of *P* that are indistinguishable to ϕ through *V*. Therefore, it reflects the user's ability to infer all behaviours of *P* that are consistent with its interaction with *P*. It works in a similar way to the inference function, but abstracts away from the user's interactions with the process.

The **UI** predicate transformer inherits the characteristics of **U**; it too is an order-embedding from relations to lifted relations.

Example 4.18. Recall the guessing game described in Example 3.30. The indistinguishability relation corresponding to Alice's view is:

$$\mathsf{IR}(ALICE) = g = \tilde{g} \land r = \tilde{r} \land g' = g' \land r' = r'$$

The lifted form of *GUESS*₀ with respect to $g = \tilde{g} \wedge r' = \tilde{r'}$ is:

$$\begin{pmatrix} g \in 1..10 \\ \land n \in 1..10 \land \widetilde{n} \in 1..10 \\ \land g = \widetilde{g} \land r' = \widetilde{r'} \end{pmatrix} \land \begin{pmatrix} g > n \Rightarrow r' > 0 \land \widetilde{g} > \widetilde{n} \\ \land g = n \Rightarrow r' = 0 \land \widetilde{g} = \widetilde{n} \\ \land g < n \Rightarrow r' < 0 \land \widetilde{g} < \widetilde{n} \end{pmatrix}$$

This lifted relation indicates that Alice can distinguish between runs where g > n, g = n and g < n.

Remark 4.19. We could have defined **UI** to accommodate multiple fog alphabets, in order to represent behavioural indistinguishability with respect to multiple views. However, a single fog alphabet suffices to illustrate the main topics in this chapter, so we avoid cluttering the presentation with multiple fog alphabets.

In the sections that follow, we use views only to construct indistinguishability relations. This enables us to dispense with user interactions and work exclusively in the lifted space of behaviours.

4.3 Encoding Confidentiality

Some of a system's behaviours may entail the communication of secret data. The confidentiality of those data is violated if an adversarial low-level user Low can infer significant information about that data. Hence, regulating Low's inferences about the system's behaviour is crucial to safeguard the confidentiality of those data.

This section presents a scheme for encoding confidentiality properties that fits naturally with the semantics of lifted relations. We first outline the philosophy of this scheme, before explaining how these properties can be applied to evaluate the security of a system design.

4.3.1 Secrets and Cover Stories

Let ϕ denote a behaviour of a system *P* which is secret in some respect. If *P* behaves as ϕ , then it is imperative that Low cannot establish that fact, in order to maintain the confidentiality of ϕ .

Our scheme for specifying that ϕ is confidential is based on *cover stories*. Intuitively, a cover story behaviour for ϕ is a designated alternative behaviour of *P* that lacks the secret aspect of ϕ . The role of these cover story behaviours is to *confuse* Low; they prevent Low from determining when the system's actual behaviour is ϕ . Exactly which behaviours of a system are classed as secret — and which behaviours should be designated as cover stories — depends on the system's functionality and its security policy.

Example 4.20. Consider a military commander who issues the order "attack at dawn". If the attack is a secret, then we may interpret an order to "do not attack" as an appropriate cover story. However, if we suppose that an enemy (Low) knows an attack is inevitable (by observing troop movements) and the secret is the time of the attack, then "attack at noon" or "attack at dusk" may serve as cover stories instead.

For a behaviour to be an effective cover story for ϕ , it must be Lowindistinguishable from ϕ . It follows that, if a system *P* features both ϕ and

a Low-indistinguishable cover story for ϕ , then Low cannot determine (with certainty) from its interaction when *P* behaves as ϕ . Hence, if *P* does behave as ϕ , then the presence of Low-indistinguishable cover story behaviours in *P* maintains the confidentiality of ϕ , because Low cannot rule out the possibility that a cover story behaviour took place instead. We now formalise these concepts in the lifted space.

4.3.2 Obligations

*Obligations*¹ are our basic unit for expressing confidentiality properties. An obligation is a lifted relation, where the observable space represents secret behaviours and the fog space represents cover story behaviours.

Example 4.21. The obligation:

$$h = 0 \Rightarrow h > 0$$

specifies that each behaviour where $\tilde{h} > 0$ serves as a cover story for each behaviour where h = 0.

An obligation may relate a behaviour to itself, meaning that behaviour is not confidential. For instance, Example 4.21 maps behaviours where $h \neq 0$ to themselves.

Example 4.22. Consider a program which accepts an integer h as its input and produces an output l' that is computed from h (in an unspecified manner). Given that Low views l' but not h, we may define some obligations to model certain confidentiality properties over the value of h:

$$\begin{array}{rcl} \theta_1 & \triangleq & \widetilde{h} \neq h \\ \theta_2 & \triangleq & \widetilde{h} \mod 2 \neq h \mod 2 \\ \theta_3 & \triangleq & h \mod 2 = 1 \Rightarrow \widetilde{h} \mod 2 = 0 \\ \theta_4 & \triangleq & h = 0 \Rightarrow \widetilde{h} \neq h \end{array}$$

¹The term "obligation" is borrowed from Seehusen and Stølen (2007, 2009), but we ascribe a different meaning to the term.

 θ_1 specifies that every value of *h* is secret, while every non-equal value of \tilde{h} is a cover story. In short, θ_1 specifies that Low must be unable to establish the exact value of *h*. θ_2 specifies that Low cannot determine the parity of *h*. θ_3 specifies that any even value of *h* is a cover story for any odd value of *h*. Finally, θ_4 specifies that Low must be unable to determine that h = 0.

Remark 4.23. In this chapter, we focus on specifying confidentiality properties over the whole behaviour of the system. However, it is sometimes desirable to limit Low's inferences about a high-level (High) user's interactions with a system. Given a relation θ_H — defined over High's interaction variables x_H, x'_H and fog variables $\widetilde{x_H}, \widetilde{x'_H}$ — the obligation:

$$\forall x_H, x'_H, \widetilde{x_H}, x'_H \bullet \Delta (H \land H[\widetilde{x}, \widetilde{x_H}/x, x_H]) \Rightarrow \theta_H$$

projects θ_H through *H* to yield an obligation over *x*, *x'* and $\tilde{x}, \tilde{x'}$.

4.3.3 Conformance

Let *P* denote a relation and θ denote an obligation. Informally, *P* conforms to θ with respect to a view *L* if, for each behaviour ϕ exhibited by *P*, there exists *at least one* behaviour $\tilde{\phi}$ of \tilde{P} such that:

- $\tilde{\phi}$ is indistinguishable through *L* from ϕ ; and
- θ prescribes $\tilde{\phi}$ as a cover story for ϕ .

If *P* fails to conform to θ , then Low could rule out all cover stories for a particular behaviour of *P*. In this case, we interpret *P* as being able to leak secret information to Low.

Definition 4.24 codifies this conformance condition.

Definition 4.24 (Conformance).

$$P \propto_L \theta \triangleq [P \Rightarrow \mathbf{D}(\mathbf{UI}(L, P) \land \theta)]$$

Example 4.25. The following table evaluates a selection of candidate programs against the obligations listed by Example 4.22 according to the conformance condition, assuming Low's view *LV* is defined as $l_L = l$.

Program	θ_1	θ_2	θ_3	$ heta_4$
l' = h	×	×	Х	×
l'=0	\checkmark	\checkmark	\checkmark	\checkmark
$l' = h \operatorname{mod} 2$	\checkmark	×	×	\checkmark
$l' = \lfloor h/2 \rfloor$	\checkmark	\checkmark	\checkmark	\checkmark
$l'=h\vee l'=0$	×	×	\checkmark	\checkmark

(Each \checkmark indicates $P \propto_L \theta$ and each \times indicates otherwise.)

The program l' = h allows Low to establish the exact value of h, so it violates each of the obligations. Conversely, the program l' = 0 reveals no information about h to Low, so it conforms to all the obligations.

The program $l' = h \mod 2$ enables Low to distinguish between $h \in \{0, 2\}$ and $h \in \{1, 3\}$, so Low can determine the parity of h, in violation of θ_2 and θ_3 . The program $l' = \lfloor h/2 \rfloor$ enables Low to distinguish between $h \in \{0, 1\}$ and $h \in \{2, 3\}$, but this does not violate any of the obligations.

The program $l' = h \lor l' = 0$ gives Low perfect information about *h* when $l' \in \{1, 2, 3\}$. However, when l' = 0, Low cannot determine any information about the value of *h*.

Obligations need to be defined with respect to the UTP theory in which a system is modelled. However, they may be defined without making presumptions about the system's functionality.

4.3.4 The Lattice of Obligations

The ordering over obligations is simply the refinement ordering. If $\theta_1 \sqsubseteq \theta_2$, then for each behaviour ϕ , θ_2 prescribes no more cover stories for ϕ as does θ_1 . Therefore, if a relation *P* conforms to θ_2 , then *P* must also conform to θ_1 .

Lemma 4.26 (\sqsubseteq is closed under \propto_L).

$$\theta_1 \sqsubseteq \theta_2 \land P \propto_L \theta_2$$
 implies $P \propto_L \theta_1$

Example 4.27. Recalling Example 4.22, we have $\theta_3 \sqsubseteq \theta_2$ and $\theta_4 \sqsubseteq \theta_1$. It follows that each program in Example 4.25 conforming to θ_2 also conforms to θ_3 , and each program conforming to θ_1 conforms to θ_4 . \diamond

The refinement ordering induces a complete lattice over the space of obligations, with a unique top and bottom. Every system conforms to the weakest obligation **true**, but only the miracle **false** conforms to the strongest obligation **false**.

We denote a set of obligations by Θ . The least upper bound and greatest lower bound of Θ are standard.

Definition 4.28 (obligation lub). $\Box \Theta \triangleq \land \Theta$

Definition 4.29 (obligation glb). $\square \Theta \triangleq \lor \Theta$

Even if a system *P* conforms to θ_1 and θ_2 individually, *P* need not conform to $\theta_1 \wedge \theta_2$, as Example 4.30 shows.

Example 4.30. Suppose θ_1 and θ_2 prescribe mutually exclusive cover stories: for instance, $\theta_1 = \tilde{h} = 1$ and $\theta_2 = \tilde{h} = 2$. Then $\theta_1 \wedge \theta_2 =$ **false**. \diamond

It follows from Lemma 4.26 that, if a system conforms to $\theta_1 \wedge \theta_2$, then it conforms to θ_1 and θ_2 separately.

Law 4.31 (lub conformance). $P \propto_L (\theta_1 \wedge \theta_2)$ implies $P \propto_L \theta_1 \wedge P \propto_L \theta_2$

There is a dual law for the greatest lower bound of obligations, which also follows from Lemma 4.26.

Law 4.32 (glb conformance). $P \propto_L \theta_1 \vee P \propto_L \theta_2$ implies $P \propto_L (\theta_1 \vee \theta_2)$

4.3.5 Multi-Obligations

The notion of conformance is rather weak: it demands that Low is unable to rule out at least one cover story behaviour — out of possibly many

cover stories that an obligation prescribes — for each behaviour of a system. However, confidentiality properties often demand that Low cannot rule out *any* cover story drawn from a set of behaviours, which is a more severe imposition.

We now define an extension of obligations to encode properties such as these. Following the style of Mantel (2003), we specify a *multi-obligation* (σ , τ) as a pair:

- The *obligation template* τ is an obligation with alphabet including *auxiliary variables*.
- The restriction σ is a predicate transformer; given a system design *P*, $\sigma(P)$ is a predicate over the auxiliary variables.

The purpose of σ is to specify a range of values to fill the auxiliary variables of the obligation template. Hence, a multi-obligation (σ , τ) encodes a set of obligations, where each valuation of the auxiliary variables satisfying $\sigma(P)$ is substituted into τ to instantiate a separate obligation.

Example 4.33. The multi-obligation $(\lambda P \bullet i \in S, \tilde{h} = i)$ specifies that Low must be unable to rule out each of the individual values in *S* for *h*. It encodes the set of obligations $\{\tilde{h} = i \mid i \in S\}$.

Definition 4.34 extends the conformance condition to multi-obligations. A relation *P* conforms to (σ, τ) if and only if *P* conforms to each σ instantiation of τ . We assume the alphabet of *P* is disjoint from the
auxiliary variables named by σ and τ ; if they are not, then we may
rename those auxiliary variables to resolve such clashes.

Definition 4.34 (Multi-obligation conformance).

$$P \propto_L (\sigma, \tau) \triangleq [P \wedge \sigma(P) \Rightarrow \mathsf{D}(\mathsf{UI}(L, P) \wedge \tau)]$$

Typically, we specify σ to be a constant function. Nevertheless, some kinds of confidentiality properties are sensitive to the details of a system's

behaviours. These properties can be encoded by defining σ as a predicate transformer over the space of systems.

Definition 4.35 presents one possible formulation of the least upper bound of an indexed set $\{(\sigma_1, \tau_1), \dots, (\sigma_n, \tau_n)\}$ of multi-obligations.

Definition 4.35 (Multi-obligation least upper bound).

$$\bigsqcup i \bullet (\sigma_i, \tau_i) \triangleq \left(\lambda P \bullet \bigwedge i \bullet j = i \Rightarrow \sigma_i(P), \bigwedge i \bullet j = i \Rightarrow \tau_i\right)$$

where *j* is a fresh variable that is free in each σ_i and θ_i .

Since the indexing variable is part of the alphabet of the obligation template and restriction of the multi-obligation $\bigsqcup i \bullet (\sigma_i, \tau_i)$, it is quantified by the universal closure of the multi-obligation conformance condition (Definition 4.34). Hence, we have:

$$P \propto_L | |i \bullet (\sigma_i, \tau_i)|$$
 if and only if $\forall i \bullet P \propto_L (\sigma_i, \tau_i)$

4.4 Deriving Secure Specifications

So far, we have treated obligations (i.e. specifications of confidentiality) separately from specifications of the functionality of systems. In this section, we show how confidentiality properties can be combined with a functional specification, to realise a *secure* specification of a system.

4.4.1 Reconciling Obligations

Let *FC* denote the combination of a lifted functionality specification of the form **UI** (*L*, *P*) and an obligation θ :

$$FC = \mathbf{UI}(L, P) \wedge \theta$$

FC maps each behaviour ϕ of *P* to only those fog behaviours of **UI** (*L*, *P*) that θ lists as cover stories for ϕ . If *P* conforms to θ , then each behaviour of *P* is present in *FC*. On the other hand, if *P* does not conform to θ , then

P features at least one behaviour ϕ that **UI** (*L*, *P*) associates with no fog behaviour that θ prescribes as a cover story for ϕ . Hence, a contradiction is induced in *FC* which renders the behaviour ϕ infeasible.

Example 4.36. Recall the program $l' = h \lor l' = 0$ from Example 4.25 and the obligation $\theta_1 \triangleq \tilde{h} \neq h$ from Example 4.22. By lifting the program (with respect to Low's view *LV*) and conjoining θ_1 , we derive:

$$\begin{aligned} & \textbf{UI} (LV, l' = h \lor l' = 0) \land \tilde{h} \neq h \\ &= (l' = h \lor l' = 0) \land (\tilde{l}' = \tilde{h} \lor \tilde{l}' = 0) \land l' = \tilde{l}' \land \tilde{h} \neq h \qquad [\text{def UI}] \\ &= ((l' = h \land l' = \tilde{h}) \lor l' = 0) \land l' = \tilde{l}' \land \tilde{h} \neq h \qquad [\text{prop calc}] \\ &= (\textbf{false} \lor l' = 0) \land l' = \tilde{l}' \land \tilde{h} \neq h \qquad [\text{contradiction}] \\ &= l' = 0 \land l' = \tilde{l}' \land \tilde{h} \neq h \qquad [\text{prop calc}] \end{aligned}$$

Here, each behaviour of the original program where $l' \neq 0$ — revealing the value of *h* to Low — is made infeasible.

The purpose of combining a functional system design (a lifted relation) with a confidentiality specification (an obligation) is to disable functionality that potentially reveals a secret to Low. A remarkable consequence of this combination is that it *makes the system design secure*.

In extreme cases, *P*'s functionality may be irreconcilable with θ . Then, *FC* admits *no behaviour whatsoever*, as Example 4.37 demonstrates.

Example 4.37. Applying θ_1 to the lifted form of program l' = h yields:

$$UI (LV, l' = h) \land \tilde{h} \neq h$$

= $l' = h \land \tilde{l}' = \tilde{h} \land l' = \tilde{l}' \land \tilde{h} \neq h$ [def UI]
= false [contradiction]

Since l' = h always reveals h to Low, no cover story for l' = h prescribed by θ_1 remains in the fog space. With no supporting fog behaviours, l' = h itself is rendered infeasible.

Indeed, it is entirely appropriate for the combination of irreconcilable

specifications of functionality and confidentiality to be a miracle, because only a miracle could ever satisfy both specifications.

4.4.2 Fixed Points of Confidentiality

An anomaly induced by conjoining **UI** (L, P) with θ is that observable behaviours made infeasible by θ are not necessarily made infeasible in the fog space. This anomaly is illustrated by Example 4.38.

Example 4.38. Consider a very simple system $h = 0 \lor h = 1$ and the obligation $\theta = h \neq \tilde{h} \land (h = 0 \Rightarrow \tilde{h} = 2)$. We have:

$$\begin{aligned} & \textbf{UI} (\textbf{true}, h = 0 \lor h = 1) \land h \neq \tilde{h} \land (h = 0 \Rightarrow \tilde{h} = 2) \\ &= \textbf{U} (h = 0 \lor h = 1) \land h \neq \tilde{h} \land (h = 0 \Rightarrow \tilde{h} = 2) \\ &= (h = 0 \land \tilde{h} = 1 \lor h = 1 \land \tilde{h} = 0) \land (h = 0 \Rightarrow \tilde{h} = 2) \\ &= h = 1 \land \tilde{h} = 0 \end{aligned} \qquad [def U] \\ &= h = 1 \land \tilde{h} = 0 \end{aligned}$$

Here, the behaviour h = 0 is infeasible, but its counterpart $\tilde{h} = 0$ remains in the fog space. It is as though Low cannot rule out the infeasible h = 0, which is incompatible with our expectation that Low knows the system's construction (Subsection 4.2.2).

To rectify this anomaly, we define a new predicate transformer **C** over the space of lifted relations.

Definition 4.39 (C predicate transformer).

$$\mathbf{C}(Q) \triangleq Q \land \widetilde{\mathbf{D}(Q)}$$

The result of applying **C** to Q is a lifted relation that allows all observable behaviours of Q, but makes infeasible all fog behaviours that do not correspond to any of those observable behaviours.

Lemma 4.40 (C is monotonic). $Q_1 \sqsubseteq Q_2$ implies $C(Q_1) \sqsubseteq C(Q_2)$

C is not idempotent. Its application to Q has the potential to make yet other observable behaviours of Q infeasible, if the fog counterparts of those behaviours are excluded from the fog space of **C** (Q).

Example 4.41. Let $Q = h > 0 \land \tilde{h} > 0$ and $\theta = \tilde{h} = h - 1$. Then:

$$FC = Q \wedge \theta = h > 1 \wedge \tilde{h} = h - 1$$

Repeatedly applying **C** to *FC* reveals a pattern:

$$\mathbf{C}(FC) = h > 1 \land \tilde{h} = h - 1 \land \tilde{h} > 1 = h > 2 \land \tilde{h} = h - 1$$

$$\mathbf{C}^{2}(FC) = h > 2 \land \tilde{h} = h - 1 \land \tilde{h} > 2 = h > 3 \land \tilde{h} = h - 1$$

$$\mathbf{C}^{n}(FC) = h > n + 1 \land \tilde{h} = h - 1$$

The application of **C** to *FC* does not converge because $\mathbf{C}^n(FC) \sqsubset \mathbf{C}^{n+1}(FC)$ for $n \ge 0$. In the limit, we obtain **false**: every behaviour is made miraculous.

Let $\mathbf{CC}(Q)$ denote the fixed point of **C** applied to *Q*.

Definition 4.42 (CC predicate transformer).

$$\mathbf{CC}(Q) \triangleq \mu X \bullet \mathbf{C}(Q \wedge X)$$

Since **C** is monotonic (Lemma 4.40) and the lifted space is a complete lattice, the Knaster-Tarski theorem (Tarski, 1955) implies **CC** (Q) is well-defined for every point Q in the lifted space. Unlike **C**, **CC** is idempotent.

Lemma 4.43 (CC is idempotent). CC (CC (Q)) = CC (Q)

The idempotence of **CC** is significant. Let $Y = CC(UI(L, P) \land \theta)$. By definition, the fog space of *Y* is no larger than its observable space. In particular, the observable space of *Y* contains only behaviours that θ maps to at least one cover story in the fog space of *Y*. It follows that projecting *Y* back to the unlifted space yields a relation **D**(*Y*) that conforms to θ .

Theorem 4.44 (Fixed point of confidentiality).

$$D(CC(UI(L,P),\theta)) \propto_L \theta$$

4.4.3 Reconciling Multi-Obligations

Since multi-obligations are not defined as predicates, they cannot be conjoined with lifted system designs in the same manner as obligations. The procedure for conjoining obligations can be extended to multi-obligations, by treating each instantiation of a multi-obligation as an obligation in its own right.

Given a multi-obligation (σ , τ) and a relation *P*, we can instantiate the obligation template τ with a valuation of the auxiliary variables emitted by $\sigma(P)$. The instantiated obligation θ can be reconciled against the lift of *P*, to identify a lifted relation *Q* (where $P \sqsubseteq \mathbf{D}(Q)$) such that $\mathbf{D}(Q)$ conforms to θ . Then, the procedure can be repeated with $\mathbf{D}(Q)$ in place of *P* and instantiations of τ generated from $\sigma(\mathbf{D}(Q))$.

This procedure needs to be applied until a relation P_{ω} is identified that conforms to each $\sigma(P_{\omega})$ -instantiation of τ . In the worst case, the derivation of P_{ω} may be extremely long-winded. For reasons described in Section 4.5, $P_1 \propto_L \theta$ and $P_1 \sqsubseteq P_2$ do not together imply $P_2 \propto_L \theta$, so the procedure needs to be restarted whenever P_1 is replaced by P_2 .

4.5 Refinement and Confidentiality

Refinement steps that reduce non-determinism may fail to preserve confidentiality properties. Example 4.45 demonstrates how refinement can introduce insecurity into a system design.

Example 4.45. Recall from Example 4.25 the program $l' = h \lor l' = 0$ and the obligation $\theta_4 \triangleq h = 0 \Rightarrow \tilde{h} \neq h$, where $l' = h \lor l' = 0$ conforms to θ_4 . While $l' = h \lor l' = 0$ is refined by l' = h, the program l' = h does not conform to θ_4 (Example 4.25), because Low can establish h = 0 if it observes l' = 0.

This problem was first noted in print by Jacob (1989a) and was subsequently termed the "refinement paradox" by Roscoe (1995). This term is perhaps misleading, because the "paradox" is readily explained. By resolving non-determinism, refinement resolves uncertainty about the behaviours of a system, so Low's inferences about those behaviours are *never weakened* but *may be strengthened*. Therefore, naïve refinement steps may increase Low's inferences beyond the upper limit on information flow imposed by a confidentiality property, thus violating the property.

Roscoe (1995) argues the root of the so-called paradox lies in the fact that formal specification notations typically draw no distinction between two different roles that non-determinism may serve:

- **Under-specification** is "don't-care" non-determinism, providing software engineers with freedom to choose how to implement a specification.
- **Unpredictability** is non-determinism intended to limit Low's inferences about secret information.

It is safe to refine away under-specification non-determinism within a specification. However, care is needed when reducing the unpredictability of a system design, in order not to jeopardise confidentiality properties.

Our formulation of confidentiality in the lifted space naturally distinguishes between under-specification and unpredictability. By introducing a secondary observation space to model Low's inferences, we become able to reason about unpredictability at the semantic level. Furthermore, the effect of conjoining an obligation θ to a lifted relation Q — in the style of Section 4.4 — is to specify that *P*'s behaviour can be no more predictable (to Low) than θ allows.

Example 4.46. Returning to the insecure refinement described in Example 4.45, we now attempt the same refinement in the lifted space.

$$UI(LV, l' = h \lor l' = 0) \land (h = 0 \Rightarrow \tilde{h} \neq h)$$

$$\subseteq UI(LV, l' = h) \land (h = 0 \Rightarrow \tilde{h} \neq h) \qquad [order embedding]$$

4.6 Specifying Noninterference Properties

=	$UI\left(LV,l'=h\right)\wedge\left(h=0\Rightarrow\widetilde{h}\neq h\wedge\widetilde{h}=h\right)$	[def UI]
=	$UI\left(\mathit{LV}, \mathit{l'}=h\right) \land \left(h=0 \Rightarrow \mathbf{false}\right)$	[contradiction]
=	UI $(LV, l'=h) \wedge h eq 0$	[prop calc]

Since the program fails to provide the cover story $\tilde{h} \neq h$ when h = 0, the obligation excludes behaviours where h = 0 from taking place.

The so-called refinement paradox evaporates in the consistent lifted space, because refining a lifted system with embedded obligations cannot make the system behave in a way that violates those obligation. Since obligations are never weakened by refinement (by Lemma 4.26), they remain within the system design at each step of its development.

4.6 Specifying Noninterference Properties

This section examines how many information flow security properties defined in the literature can be formulated within the framework. By doing so, we demonstrate that our framework is sufficiently expressive for capturing these recognised notions of confidentiality.

A side-effect of this work is that we link our UTP formulation of confidentiality properties to the body of existing work on information flow properties. This link could potentially be used to port existing results from the literature to our framework.

4.6.1 Noninterference

The canonical information flow property is *noninterference* (Goguen and Meseguer, 1982, 1984). Noninterference mandates that High's inputs to a system must have no effect on the system's outputs to Low; in other words, Low must be unable to rule out any possible High interaction with the system (or even that High has interacted with the system at all). Therefore, a system satisfying noninterference does not disclose to Low any information about High's interaction.

In Goguen and Meseguer's original definition of noninterference, a system is modelled as a deterministic state machine, where outputs to Low are a function of the inputs from High and Low. This model allows analysis of whether High's choice of inputs can influence the outputs visible to Low. However, this model does not support reasoning about non-deterministic systems. Moreover, this model also requires the system to be *input total*: in every state, the system must be willing to accept every input provided by the environment. This requirement is unrealistic in practice, since a system with a finite memory can only buffer a finite series of inputs pending computation.

Despite the limitations of its original definition, the notion of noninterference has been highly influential in theoretical studies of information flow security. A multitude of noninterference-like properties have been defined to replicate the intent of noninterference in a non-deterministic setting. Typically, these properties are formulated in a trace semantics, where the interactions between a system and its users are recorded as events within the trace. These properties include, but are not limited to, the following:

- **Noninference** (O'Halloran, 1990) requires that for every trace of *P* featuring a high-level event, there is a Low-indistinguishable trace featuring no high-level events.
- **Generalised noninference** (McLean, 1994a) is similar to noninference, but is weaker: it requires only that for every trace of *P* featuring a high-level input event, there is a Low-indistinguishable trace featuring no high-level input events.
- **Generalised noninterference** (McCullough, 1987) requires that perturbing a trace by changing high-level inputs does not imply a change in subsequent low-level events.
- **Forward correctability** (Johnson and Thayer, 1988) requires that a perturbation of a trace that inserts or deletes a high-level input event can be rectified by inserting or deleting high-level output events.

Separability (McLean, 1994a) is an extremely strict property, which prohibits any information flow between High and Low. A system which enforces separability between High and Low is equivalent to two physically separate ("air-gapped") sub-systems, where one sub-system interacts with High alone and the other with Low alone.

Separability is based on a proposal by Rushby (1981) for verifying security, which involves cutting known communication channels between different components of the system and establishing the resulting system can be split into separate components.²

4.6.2 Worked Example: Basic Security Predicates

Several frameworks for expressing a range of information flow properties in a uniform manner have been presented in the security literature (Jacob, 1991; McLean, 1994a; Focardi and Gorrieri, 1995; Zakinthinos and Lee, 1997; Mantel, 2000b). The objective of these frameworks is to consolidate the existing definitions of noninterference-like properties in the literature, in order to evaluate and compare these properties rigorously and to enable new information flow properties to be formalised.

These frameworks formulate confidentiality properties as closure conditions over an observation space (typically trace sets). We hypothesise that multi-obligations are sufficiently expressive to encode these closure conditions. We aim to substantiate this hypothesis by demonstrating how the closure conditions from one such framework can be encoded in our own framework.

The *Modular Framework for Information Flow Properties* (MAKS), introduced by Mantel (2000b) (as the *Modular Assembly Kit for Security*) and later elaborated by Mantel (2003), is capable of expressing a wide range of noninterference-like properties from the security literature in a uniform trace-based style. It defines a collection of *basic security predicates* (BSPs) to express a variety of transformations on high-level components of sys-

²Jacob (1990) observed this proposal is subtly flawed, because cutting known channels may hide the presence of undesired covert channels.

tem traces, capturing low-level users' uncertainty about high-level inputs and outputs. These BSPs are closure conditions that share a common structure that is similar to our structuring of multi-obligations. Therefore, we can translate the BSPs defined by Mantel (2003) into multi-obligations, in the context of the UTP theory of reactive processes.

In the MAKS, the events communicated between a system and its users are partitioned into three sets:

- *V*, denoting non-confidential events visible to Low;
- *N*, denoting non-confidential events not visible to Low; and
- *C*, denoting confidential events not visible to Low.

We identify \mathcal{V} with Low's window and identify \mathcal{C} with High's window. The BSPs assume that Low can observe a projection of the system trace, so a suitable choice of Low's view would be a **VHR**-healthy view (Definition 3.41) constructed from window \mathcal{V} .

Example 4.47 presents multi-obligation encodings of a selection of BSPs from the MAKS. We summarise the meaning of these BSPs as follows:

- *R* (removal) demands that each trace of *P* that contains confidential events can be *perturbed* by deleting all confidential events to reach another trace of *P*.
- *BSD* (backwards strict deletion) demands that each trace *tr* of *P* that contains a confidential event can be perturbed by deleting the final confidential event *c* in *tr* but not changing the events leading up to *c* to reach another trace of *P*.
- *BSI* (backwards strict insertion) is the dual of *BSD*. It demands that each trace *tr* of *P* can be perturbed by inserting any confidential event *c* after all other confidential events in *tr*, but not changing the events leading up to the inserted *c*.
- *BSIA^{pc}* (backwards strict insertion of admissible events) is a weaker form of *BSI*, which applies only to those traces of *P* that can be

Example 4.47 (Reformulations of BSPs as multi-obligations). Let
$$\alpha \sim \gamma \, \widetilde{\alpha} \triangleq \alpha \restriction \mathcal{V} = \widetilde{\alpha} \restriction \mathcal{V}$$
 in:

$$R(\mathcal{V}, \mathcal{N}, \mathcal{C}) \triangleq \left(\lambda P \bullet \mathbf{true}, \left(\widetilde{t}^{\gamma} - \widetilde{t}^{\gamma} \right) \restriction \mathcal{C} = \langle \rangle \right)$$

$$BSD(\mathcal{V}, \mathcal{N}, \mathcal{C}) \triangleq \left(\lambda P \bullet \left(\begin{array}{c} c \in \mathcal{C} \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \end{array} \right), \exists \widetilde{\alpha} \bullet \left(\begin{array}{c} \left(\widetilde{t}^{\gamma} - \widetilde{t}^{\gamma} \right) = \beta \cap \widetilde{\alpha} \\ \wedge \widetilde{\alpha} \land \gamma \alpha \end{array} \right) \right)$$

$$BSI(\mathcal{V}, \mathcal{N}, \mathcal{C}) \triangleq \left(\lambda P \bullet \left(\begin{array}{c} c \in \mathcal{C} \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \end{array} \right), \exists \widetilde{\alpha} \bullet \left(\begin{array}{c} \left(\widetilde{t}^{\gamma} - \widetilde{t}^{\gamma} \right) = \beta \cap \langle c \rangle \cap \widetilde{\alpha} \end{array} \right) \right)$$

$$BSI(\mathcal{V}, \mathcal{N}, \mathcal{C}) \triangleq \left(\lambda P \bullet \left(\begin{array}{c} c \in \mathcal{C} \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \end{array} \right), \exists \widetilde{\alpha} \bullet \left(\begin{array}{c} \left(\widetilde{t}^{\gamma} - \widetilde{t}^{\gamma} \right) = \beta \cap \langle c \rangle \cap \widetilde{\alpha} \end{array} \right) \right)$$

$$BSI(\mathcal{V}, \mathcal{N}, \mathcal{C}) \triangleq \left(\lambda P \bullet \left(\begin{array}{c} c \in \mathcal{C} \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \end{array} \right), \exists \widetilde{\alpha} \bullet \left(\begin{array}{c} \left(\widetilde{t}^{\gamma} - \widetilde{t}^{\gamma} \right) = \beta \cap \langle c \rangle \cap \widetilde{\alpha} \end{array} \right) \right)$$

$$BSIA^{p}(\mathcal{V}, \mathcal{N}, \mathcal{C}) \triangleq \left(\lambda P \bullet \left(\begin{array}{c} c \in \mathcal{C} \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \\ \wedge \alpha \restriction \mathcal{C} = \langle \rangle \end{array} \right), \exists \widetilde{\alpha} \bullet \left(\begin{array}{c} \left(\widetilde{t}^{\gamma} - \widetilde{t} \right) = \beta \cap \langle c \rangle \cap \widetilde{\alpha} \\ \wedge \widetilde{\alpha} \restriction \mathcal{C} = \langle \rangle \end{array} \right) \right)$$

Information flow property	Multi-obligation
Noninference	R(H)
Generalised noninference	R(HI)
Generalised noninterference	$BSD(HI) \sqcup BSI(HI)$
Separability (over <i>P</i>)	$BSD(H) \sqcup BSIA^{pc}(H)$

Table 4.1: MAKS definition of noninterference properties from Mantel (2003), where $H = (\mathcal{L}, \emptyset, \mathcal{H})$, $HI = (\mathcal{L}, \mathcal{H} \setminus \mathcal{I}, \mathcal{H} \cap \mathcal{I})$ and \mathcal{I} is a subset of $\mathcal{V} \cup \mathcal{N} \cup \mathcal{C}$ denoting input events.

constructed by appending a confidential event to a trace in *P*. This BSP needs to be initialised with the specification of *P*, to filter the traces satisfying this condition.

The definition of conformance implies the perturbations of each trace *tr* must themselves be Low-indistinguishable from *tr*.

The BSPs of MAKS can be conjoined to express many (but not all) of the noninterference-like information flow properties defined in the literature. Table 4.1 reproduces the MAKS definitions of the noninterference-like properties described in Subsection 4.6.1, but in terms of the least upper bounds of the multi-obligations in Example 4.47. Full justifications of these definitions are given by Mantel (2003).

4.7 Related Work

Since the advent of the noninterference property, many researchers have studied the problem of restricting information flow from systems to their users. This section surveys some of their work, and compares it with the approach presented in this chapter.

4.7.1 The Limitations of Noninterference

Information flow properties based on noninterference typically require that a system does not leak any information about High's activities to Low. This blanket security policy often conflicts with a system's functional requirements, which may require High to be able to communicate non-confidential data to Low through the system. In these cases, noninterference from High to Low is simply not appropriate.

No consensus has been reached in the literature about which nondeterministic formulations of noninterference are most suitable for expressing the security requirements of realistic information systems (Ryan, 2001). Indeed, the strictness of noninterference-like properties has led researchers such as Ryan et al. (2001) to question whether these properties are relevant to secure software development.

The position elaborated in this thesis is that software engineers should be able to specify custom confidentiality properties that are tailored to the intricacies of the system domain, rather than choosing from a limited range of ready-made information flow security properties.

4.7.2 Security Specifications

Our notion of obligations for specifying confidentiality has much in common with the *security specifications* defined by Jacob (1988).

A security specification *f* is a function from Low's interactions to sets of system behaviours. A system *P* satisfies *f* if, for each interaction ψ_L with *P* that Low can make, $f(\psi_L)$ is a subset of Low's inference set for ψ_L .

There is a duality between obligations and security specifications: *f* specifies the maximum information that Low *may establish* about the system's behaviour; whereas an obligation specifies the minimum information that Low *cannot rule out* about the system's behaviour.

Security specifications do not incorporate the notion of cover story behaviours. Hence, if we wish to specify that $\tilde{\phi}$ is a cover story for ϕ using a security specification, then it is necessary that ϕ serves as a cover story for $\tilde{\phi}$ as well. Moreover, we can define obligations that offer a choice of which cover stories are present in a system, whereas security specifications provide no such choice.

4.7.3 Other Approaches

This chapter is primarily concerned with one flavour of confidentiality properties. Nevertheless, a variety of alternative approaches towards information flow security have been studied extensively in the literature. These approaches have advantages over our confidentiality framework, but they also have drawbacks.

- Active and Passive Flows Jacob (1991) and Roscoe (1997) identify two scenarios for information leakage between users. These are *passive flows*, where Low tries to infer information about the system's behaviour without the help of other users; and *active flows*, where a treacherous High user tries to leak secret information to Low by interacting with the system according to a pre-arranged protocol. The framework set out in this chapter is intended to counter passive flows, but it does not guarantee the prevention of active flows. Nevertheless, if such a protocol could be established by High and Low then, in many situations, High could pass secrets to Low outside the boundaries of the system anyway.
- **Quantifying Flow** Small leaks of data from high-level users to low-level users are often tolerable (and sometimes unavoidable) in realistic systems, provided that a low-level user cannot infer any significant details about confidential data from such a leak. However, our notion of confidentiality is *qualitative*: an obligation is violated if the slightest amount of data deemed secret by a confidentiality property is leaked to Low. This effect motivates defining confidentiality properties in terms of a *quantitative* measure of information flow between users. Recent research has addressed quantitative confidentiality properties using information theory (Shannon, 1948) as a foundation; Mu (2008) surveys this work and its applications.
- **Probability** Another limitation of our treatment of confidentiality properties is that we do not address the probability distribution of a

system's behaviours. This shortcoming could lead to serious security breaches, because if Low has knowledge of this distribution, then it may be able to deduce confidential data with near certainty but without violating the confidentiality property.

Some research has investigated *probabilistic* confidentiality properties which account for the likelihood of alternative high-level activities (Santen, 2006, 2008). While these properties are attractive in theory, their application in practice is not without difficulty: the probability distribution of a system's behaviours may be unknown. Even if that distribution is known, it may be intractable to determine whether a non-trivial system model satisfies a probabilistic confidentiality property (Ryan, 2001).

The security of confidential information need not rest on technical measures alone. Measures such as access control and inference control are complemented by physical security mechanisms for computer terminals, such as barriers and surveillance systems (Anderson, 2001).

Sociological factors also have a role to play in system security: for instance, ethics training or the threat of punishment may deter users from attempting to access confidential information without the appropriate clearance (Jacob, 1989a; Workman and Gathegi, 2007).

4.7.4 Refinement and Confidentiality

Most notions of confidentiality-preserving refinement in the literature are realised in two ways:

- by limiting the space of confidentiality properties to those that are *refinement closed* — that is, if a system satisfies such a property, then all refinements of the system will also satisfy that property; or
- 2. by strengthening the refinement relation to ensure it preserves confidentiality properties in specifications.

Obligations are refinement closed in the lifted space. The method proposed in Section 4.4 for deriving a secure specification of a system *P* is

refinement-closed. If *P* is refined in an insecure fashion with respect to θ , then that insecurity is manifested when the lifted form of *P* is conjoined with θ . Therefore, the notion of confidentiality-preserving refinement described in Section 4.5 is just the standard notion of refinement in the UTP, but over the space of lifted relations.

Roscoe et al. (1994) describe a noninterference-like property over CSP processes in the line of the first approach. The property requires Low's interface to a process — formulated using the concept of lazy abstraction (described in Subsection 3.5.3) — to be deterministic. If this property is met — as can be verified with a model checker — then noninterference is assured, because High's interaction with the process cannot influence Low's interactions in any way. Furthermore, no (failures-divergences) refinement of the process can induce new sources of information flow from High to Low, because Low's interface must remain deterministic. While this property is theoretically appealing, Ryan (2001) argues that it is unsuitable for a large class of systems which contain unavoidable non-determinism in their outputs to Low.

Seehusen and Stølen (2007, 2009) have proposed a further instance of the first approach. They advocate structuring a system specification as a set of trace sets. Each trace set represents under-specification, while the presence of multiple trace sets within a specification represents unpredictability. Traces can be refined away from the specification, so long as doing so makes no trace set in the specification empty. Indeed, this notion of refinement generalises the structure of a confidentiality refinement relation defined by Jacob (1992).

Instances of the second approach that are specialised to particular confidentiality properties are detailed by Mantel (2001) and Alur et al. (2006), among others. A generic approach for strengthening refinement to preserve any given confidentiality property is described by Banks and Jacob (2010a). In that work, the standard refinement ordering is combined with a confidentiality ordering, which dictates that refinement does not strengthen Low's inferences about behaviours of a system marked as secret. This notion of refinement is very strong indeed: it

prohibits some refinement steps that do not jeopardise security. To make it tractable, it may be necessary to adopt a multi-user refinement relation (see Section 3.6) in place of the canonical notion of refinement.

4.8 Conclusion

This chapter has presented an abstract framework for defining confidentiality properties in a UTP style. The framework cleanly captures the essence of confidentiality properties in an abstract manner, without specialising to any particular UTP theory. By specialising the framework to a particular UTP theory, it could be deployed to formulate confidentiality properties over the program state (in the theory of designs); the trace generated by a process (in the theory of reactive processes); or the execution period of operations (in timed theories).

A further novelty of the framework lies in how it enables confidentiality properties to be integrated with a functional specification of a system. To our knowledge, our approach for uniting functionality and confidentiality concerns with a single semantics is original. In the next chapter, this novelty is applied to incorporate both functionality and confidentiality attributes within a single specification language.

Not all conceivable confidentiality properties can be expressed as obligations. The notion of multi-obligations allows us to widen the space of confidentiality properties that can be expressed within the framework. As Section 4.6 demonstrates, the addition of multi-obligations gives expressive power that suffices to capture many of the information flow properties detailed in the literature.

The abstractness of the framework undermines its practicality as a platform for developing secure software. Some significant obstacles that hinder the practical application of the framework are as follows:

Compositionality Compositionality is essential for understanding a system in terms of its individual components. But the framework is defined over the totality of a system's behaviour. To make the

framework practical, it is necessary to identify how it can be applied to individual components of a system.

- **Localisation** Confidentiality requirements are often directly related to particular aspects of a system, principally those involving interactions with high-level users. To formulate these requirements within our framework, they must be encoded *indirectly* over the space of behaviours of a system. This indirection adds undesirable complexity to specifying confidentiality properties in the framework.
- **Verification** The conformance condition is defined purely semantically. Without proof rules, it is necessary to translate a system into its underlying semantics in order to verify it against an obligation, which is cumbersome and inefficient.

The next chapter overcomes these obstacles with a tighter integration of the confidentiality framework with the syntax and semantics of *Circus*. It demonstrates how the framework can be made tractable, but at the expense of sacrificing some of the framework's generality.

5 Specifying Confidentiality in Circus

5.1 Introduction

In this chapter, we specialise the confidentiality framework presented in Chapter 4 to express confidentiality properties over *Circus* processes. This work paves the way towards an *integrated* language for specifying the functionality and confidentiality attributes of systems alongside each other. The philosophy of this integrated language is to exploit the facilities provided by *Circus* as far as practicable, in order to maintain compatibility with the existing laws and refinement calculus of *Circus*.

In Section 5.2 and Section 5.3, we lift the UTP semantics of *Circus* actions and operators to model Low's inferences about the execution of those constructs in a compositional fashion. Section 5.4 presents a simple extension of the *Circus* syntax that is needed to ensure processes in our lifted language have an unambiguous semantics.

The central novelty of this chapter is a specification construct for expressing confidentiality properties within the body of a *Circus* process. We define *confidentiality annotations* in Section 5.5, in terms of the lifted semantics of *Circus* and the obligations described in Chapter 4. These annotations can be used in tandem with the specification facilities of *Circus* to formulate a variety of confidentiality properties over the state and behaviour of *Circus* processes.

In Section 5.6, we show how confidentiality annotations can be applied to express a diverse variety of confidentiality properties in terms of a process's interactions with its environment. This work culminates with a technique for *superposing* confidentiality annotations with existing *Circus* processes, enabling the specification of noninterference-like properties in

5 Specifying Confidentiality in Circus

our integrated language.

In Section 5.7, we discuss how the integrated language may be extended to accommodate the declassification of secret information. Finally, in Section 5.8, we draw comparisons between our integrated language and recent work on accommodating security concerns within formal methods.

This chapter emphasises the theoretical aspects of the integrated language. We defer the study of a tractable method for developing systems from specifications expressed in the language until the next chapter.

5.2 Lifting the Circus Semantics

The first step towards defining our integrated language is to codify Low's inferences about the behaviour of *Circus* actions at the semantic level. We accomplish this step by lifting the UTP semantics of *Circus* actions, following the approach described in Subsection 4.2.1. This lifted semantics supplies the foundation we need for defining confidentiality annotations in Section 5.5.

5.2.1 Lifting Circus Actions

Since *Circus* actions are reactive designs, we model Low's interactions (and hence inferences) with a *Circus* process in terms of the reactive views defined in Subsection 3.5.1. The indistinguishability relation formed by applying the **IR** predicate transformer (Definition 4.16) to the reactive interface $\mathcal{R}(\mathcal{L})$ (Definition 3.43) corresponding to Low's window \mathcal{L} is equivalent to:

$$\begin{pmatrix} ok = \widetilde{ok} \land ok' = \widetilde{ok'} \\ \land wait = \widetilde{wait} \land wait' = \widetilde{wait'} \\ \land (tr' - tr) \upharpoonright \mathcal{L} = (\widetilde{tr'} - \widetilde{tr}) \upharpoonright \mathcal{L} \\ \land ref \cap \mathcal{L} = \widetilde{ref} \cap \mathcal{L} \land ref' \cap \mathcal{L} = \widetilde{ref'} \cap \mathcal{L} \end{pmatrix}$$

We can simplify this indistinguishability relation in the context of *Circus* actions. Each *Circus* action *A* is **CSP3**- and **CSP4**-healthy; that is, *A* does

not depend upon *ref*, nor does it restrict *ref* on termination (Oliveira et al., 2009). Hence, our *Circus* indistinguishability relation need only constrain the relation between *ref* and \widetilde{ref} on non-termination.

Definition 5.1 (Circus indistinguishability relation).

$$\mathcal{I}(\mathcal{L}) \triangleq \begin{pmatrix} ok = \widetilde{ok} \land ok' = \widetilde{ok'} \\ \land wait = \widetilde{wait} \land wait' = \widetilde{wait'} \\ \land (tr' - tr) \upharpoonright \mathcal{L} = (\widetilde{tr'} - \widetilde{tr}) \upharpoonright \mathcal{L} \\ \land wait' \Rightarrow ref' \cap \mathcal{L} = \widetilde{ref'} \cap \mathcal{L} \end{pmatrix}$$

The $\mathcal{I}(\mathcal{L})$ relation codifies the expectation that Low's observation of the behaviour of a *Circus* process amounts to a projection of its trace through \mathcal{L} , together with knowledge of whether the process has started correctly, is waiting for interaction from the environment, or has terminated. Henceforth, we adopt $\mathcal{I}(\mathcal{L})$ as our model of Low's observational abilities of *Circus* actions and processes.

By lifting the semantics of *Circus* actions and conjoining $\mathcal{I}(\mathcal{L})$, we can model Low's inferences about an action's behaviour in terms of the fog behaviours. Definition 5.2 presents a specialised form of the lifting function **UI** (Definition 4.17) which incorporates $\mathcal{I}(\mathcal{L})$.

Definition 5.2 (UC predicate transformer).

$$\mathsf{UC}(\mathcal{L},A) \triangleq \mathsf{U}(A) \land \mathcal{I}(\mathcal{L})$$

We call $UC(\mathcal{L}, A)$ the lift of the *Circus* action A with respect to \mathcal{L} . Reassuringly, there exists an order-embedding between *Circus* actions and **UC**-lifted actions, which is a special form of Theorem 4.4.

Lemma 5.3 (UC order-embedding).

 $A_1 \sqsubseteq A_2$ if and only if **UC** $(\mathcal{L}, A_1) \sqsubseteq$ **UC** (\mathcal{L}, A_2)

Furthermore, Lemma 4.6 applies to lifted Circus actions too.

Law 5.4 (Recovering actions). For all Circus actions,

$$\mathbf{D}\left(\mathbf{UC}\left(\mathcal{L},A\right)\right) = \mathbf{D}\left(\mathbf{U}\left(A\right)\right) = A$$

5.2.2 Lifting Reactive Designs

While *Circus* actions are reactive designs, their lifted counterparts are not reactive, because **UC** does not commute with **R3**:

$$\begin{aligned} \mathsf{UC}\left(\mathcal{L},\mathsf{R3}(P)\right) &= & \mathsf{U}\left(\varPi_{rea} \lhd wait \rhd P\right) \land \mathcal{I}(\mathcal{L}) \\ &\neq & \varPi_{rea} \lhd wait \rhd \left(\mathsf{U}\left(P\right) \land \mathcal{I}(\mathcal{L})\right) \\ &= & \mathsf{R3} \circ \mathsf{UC}\left(\mathcal{L},P\right) \end{aligned}$$

Nevertheless, the **UC**-lifts of *Circus* actions can be reshaped into a form with the character of reactive designs. Given an action *A*, its fog counterpart \widetilde{A} (with alphabet $\widetilde{x}, \widetilde{x'}$) satisfies the following healthiness conditions:

$$\begin{split} \widetilde{\mathbf{R1}}(A) &= A \wedge \widetilde{tr} \leq t\widetilde{r'} \\ \widetilde{\mathbf{R2}}(A) &= A[\langle \rangle, t\widetilde{r'} - t\widetilde{r}/t\widetilde{r}, t\widetilde{r'}] \\ \widetilde{\mathbf{R3}}(A) &= (\widetilde{II_{rea}} \lhd \widetilde{wait} \triangleright A) \end{split}$$

Naturally, $\widetilde{\mathbf{R1}}$, $\widetilde{\mathbf{R2}}$ and $\widetilde{\mathbf{R3}}$ are idempotent and commute with each other. Moreover, $\widetilde{\mathbf{R1}}$ and $\widetilde{\mathbf{R2}}$ commute with $\mathbf{R1}$ and $\mathbf{R2}$. However, $\mathbf{R3}$ and $\widetilde{\mathbf{R3}}$ do not commute with each other in general, because $\mathbf{R3}(A)$ behaves as \mathcal{I}_{rea} when *wait* holds, but $\widetilde{\mathbf{R3}}(A)$ behaves as $\widetilde{\mathcal{II}_{rea}}$ when *wait* holds. To sidestep this mismatch, we insist that *wait* = *wait* and define a compound healthiness condition:

$$\widehat{\mathbf{R3}}(A) \triangleq \mathbf{R3}(A) \land \widetilde{\mathbf{R3}}(A) \land wait = \widetilde{wait} \\ = ((\Pi_{rea} \land \widetilde{\Pi_{rea}}) \lhd wait \triangleright A) \land wait = \widetilde{wait}$$

 $\widehat{\mathbf{R3}}$ is idempotent and commutes with $\mathbf{R1}$, $\widehat{\mathbf{R1}}$, $\mathbf{R2}$ and $\widehat{\mathbf{R2}}$. We call the composition of these healthiness conditions $\widehat{\mathbf{R}}$.

Definition 5.5 (**R** healthiness condition).

$$\widehat{\mathbf{R}}(A) \triangleq \mathbf{R1} \circ \widetilde{\mathbf{R1}} \circ \mathbf{R2} \circ \widetilde{\mathbf{R2}} \circ \widehat{\mathbf{R3}}(A)$$

Since the exterior $\mathcal{I}(\mathcal{L})$ term of UC (\mathcal{L}, A) implies wait = wait, we can express UC (\mathcal{L}, A) — where A is **R**-healthy — in terms of **R**.

Lemma 5.6 (Lifted reactive process). Provided *A* is **R**-healthy:

$$\mathsf{UC}(\mathcal{L}, A) = \widehat{\mathsf{R}}(\mathsf{U}(A)) \wedge \mathcal{I}(\mathcal{L})$$

Theorem 5.7 builds on Lemma 5.6 by reformulating the interior **U**(*A*) term as a design, which is justified because $\mathcal{I}(\mathcal{L})$ implies the condition $ok = \widetilde{ok} \wedge ok' = \widetilde{ok'}$.

Theorem 5.7 (Lifted reactive design). Provided A is a reactive design:

$$\mathbf{UC}(\mathcal{L}, A) = \widehat{\mathbf{R}}\left(\neg \mathbf{U}(A)_{ff}^{ff} \vdash \mathbf{U}(A)_{ff}^{tt}\right) \wedge \mathcal{I}(\mathcal{L})$$

where B_{de}^{bc} denotes $B[b, c, d, e/ok', \widetilde{ok'}, wait, \widetilde{wait}]$.

The semantics of **UC** (\mathcal{L} , A) from Theorem 5.7 is clarified by Lemma 5.8, which expresses A in the form **R** ($Pre \vdash Post$).

Lemma 5.8 (Unfolded lifted reactive design). Provided Pre is a condition:

$$\mathsf{UC}\left(\mathcal{L}, \mathsf{R}\left(\operatorname{Pre} \vdash \operatorname{Post}\right)\right) = \widehat{\mathsf{R}}\left(\operatorname{Pre} \lor \widetilde{\operatorname{Pre}} \vdash \mathsf{U}\left(\operatorname{Pre} \Rightarrow \operatorname{Post}\right)\right) \land \mathcal{I}(\mathcal{L})$$

Lemma 5.8 can be instantiated with the semantics of the *Circus* specification statement, to derive a lifted specification statement presented in Lemma 5.9. The lifted forms of assignments, assumptions and coercions are all derivable from Lemma 5.9.

Lemma 5.9 (Lifted specification statement).

$$\begin{aligned} & \mathsf{UC}\left(\mathcal{L}, w : [Pre, Post]\right) = \\ & \widehat{\mathsf{R}}\left(Pre \lor \widetilde{Pre} \vdash \mathsf{U}\left(Pre \Rightarrow (Post \land \neg wait' \land tr' = tr \land u' = u\right)\right) \land \mathcal{I}(\mathcal{L}) \end{aligned}$$

5 Specifying Confidentiality in Circus

where u denotes all variables outside of the frame w.

Reasoning about Low's inferences about a *Circus* action's behaviour can be simplified by studying the lifted postcondition of that action. Example 5.10 demonstrates how this principle can be applied.

Example 5.10. Consider the action $OO = on \rightarrow h := 0 \square off \rightarrow h := 1$. Its semantics is:

$$\mathbf{R}\left(\mathbf{true} \vdash \left(\begin{array}{c} v' = v \wedge tr' = tr \wedge (on, Sync), (off, Sync) \notin ref' \\ \lhd wait' \rhd \\ u' = u \wedge \left(\begin{array}{c} h' = 0 \wedge tr' = tr \cap \langle (on, Sync) \rangle \\ \lor h' = 1 \wedge tr' = tr \cap \langle (off, Sync) \rangle \end{array}\right)\right)\right)$$

where *u* denotes the list of state variables excluding *h*; and $s_1, s_2 \notin S$ abbreviates $s_1 \notin S \land s_2 \notin S$.

Suppose $\mathcal{L} = \{(on, Sync)\}$. Unfolding the postcondition of **UC** (\mathcal{L}, OO) shows how its fog behaviours are reduced in the presence of $\mathcal{I}(\mathcal{L})$:

$$\begin{aligned} \mathbf{U} \left(\begin{array}{c} v' = v \wedge tr' = tr \wedge (on, Sync), (off, Sync) \notin ref' \\ \lhd wait' \rhd \\ u' = u \wedge \left(\begin{array}{c} h' = 0 \wedge tr' = tr \cap \langle (on, Sync) \rangle \\ \lor h' = 1 \wedge tr' = tr \cap \langle (off, Sync) \rangle \end{array} \right) \right) \wedge \mathcal{I}(\mathcal{L}) \end{aligned} \\ = \left[\mathcal{I}(\mathcal{L}) \text{ implies } wait' = \widetilde{wait'} \right] \\ \left(\begin{array}{c} \mathbf{U} \left(v' = v \wedge tr' = tr \wedge (on, Sync), (off, Sync) \notin ref') \\ \lhd wait' \rhd \\ \lor wait' \rhd \\ \lor h' = 1 \wedge tr' = tr \cap \langle (off, Sync) \rangle \end{array} \right) \right) \right) \wedge \mathcal{I}(\mathcal{L}) \end{aligned} \\ = \left[\mathcal{I}(\mathcal{L}) \text{ implies } (tr' - tr) \upharpoonright \mathcal{L} = (t\widetilde{r'} - t\widetilde{r}) \upharpoonright \mathcal{L} \right] \\ \left(\begin{array}{c} \mathbf{U} \left(v' = v \wedge tr' = tr \wedge (on, Sync), (off, Sync) \notin ref') \\ \lhd wait' \rhd \\ \lor wait' \succ \\ \lor wait' \succ \\ \mathbf{U} \left(v' = v \wedge tr' = tr \wedge (on, Sync), (off, Sync) \notin ref') \\ \lhd wait' \succ \\ \mathbf{U} \left(u' = u \wedge h' = 0 \wedge tr' = tr \cap \langle (on, Sync) \rangle \right) \right) \right) \wedge \mathcal{I}(\mathcal{L}) \end{aligned} \\ \left(\begin{array}{c} \mathbf{U} \left(u' = u \wedge h' = 0 \wedge tr' = tr \cap \langle (on, Sync) \rangle \right) \\ \lhd wait' \rhd \\ \lor \mathbf{U} \left(u' = u \wedge h' = 1 \wedge tr' = tr \cap \langle (off, Sync) \rangle \right) \end{array} \right) \wedge \mathcal{I}(\mathcal{L}) \end{aligned}$$

90

At the semantic level, this result models Low as being able to distinguish between the labelled alternative behaviours and therefore being able to infer information about the state: In case (1), Low observes deadlock, so it infers h' = h. In (2), Low observes termination and the *on* event, so it infers h' = 0. In (3), Low observes termination without observing the *off* event, so it infers the occurrence of *off* and h' = 1.

5.2.3 Addressing Divergence

Although *A* and \widetilde{A} may diverge individually, Lemma 5.8 makes clear that **UC** (\mathcal{L} , A) may only diverge in its own right — i.e. $\neg ok' \land \neg \widetilde{ok'}$ — if neither of the preconditions of A or \widetilde{A} are satisfied, or if the lifted process has already diverged. If **UC** (\mathcal{L} , A) does diverge, the exterior $\mathcal{I}(\mathcal{L})$ term forces the diverging behaviour of \widetilde{A} to be Low-indistinguishable from the behaviour of A. This effect is not problematic, because all diverging behaviours of A are present in **UC** (\mathcal{L} , A) (via the order embedding).

Of course, the possibility of divergence evaporates if *A* has **true** as its precondition. In practice, divergence is undesirable (Hoare, 1985a) and tools (such as FDR) for establishing that CSP constructs are divergence-free are readily available.

5.3 Composing Lifted Actions

To analyse Low's inferences about a *Circus* process, it is helpful to *divide* the process into multiple parts, so that each part can be *conquered* by analysing it separately. However, we need to compose those separate analyses, in order to study Low's inferences about the whole process.

Circus actions can be composed using the *Circus* operators. However, the definitions of the *Circus* operators given by Oliveira et al. (2009) are unsuitable for composing lifted actions, because they do not include the fog variables. In this section, we rectify this problem by lifting the definitions of these operators to make them suitable for composing lifted actions. Moreover, these lifted operators are also needed for composing

5 Specifying Confidentiality in Circus

confidentiality properties with lifted actions, as we see later.

For each *Circus* operator \oplus , we denote its lifted counterpart by $\widehat{\oplus}$. We now describe three principal conditions to guide the definition of $\widehat{\oplus}$. These conditions ensure that $\widehat{\oplus}$ can be used to compose lifted actions in the same way that \oplus composes *Circus* actions.

Condition 5.11. Lifted operators respect the order-embedding between *Circus* actions and lifted actions:

$$\bigoplus i \bullet A_i = \mathbf{D}\left(\widehat{\bigoplus} i \bullet \mathbf{UC}\left(\mathcal{L}, A_i\right)\right)$$

This condition ensures that applying $\widehat{\oplus}$ to the lifts of actions A_1 and A_2 yields a lifted action with exactly the same observable behaviours as the *Circus* action $A_1 \oplus A_2$.

Condition 5.12. Decomposing a lifted action into parts and combining those parts with lifted operators respects the refinement ordering:

$$\mathsf{UC}\left(\mathcal{L},\bigoplus i\bullet A_i\right) \quad \sqsubseteq \quad \widehat{\bigoplus}\,i\bullet\mathsf{UC}\left(\mathcal{L},A_i\right)$$

Splitting a lifted composite action $UC(\mathcal{L}, A_1 \oplus A_2)$ into its constituents and composing those lifted constituents with $\widehat{\oplus}$ preserves correctness, but not necessarily equivalence. The construct $\widehat{\bigoplus} i \bullet UC(\mathcal{L}, A_i)$ models Low as being able to observe the execution of each sub-action A_i individually. This means $\widehat{\bigoplus} i \bullet UC(\mathcal{L}, A_i)$ represents an over-approximation of the information Low can deduce about the behaviour of $\bigoplus i \bullet A_i$.

This condition is a compromise. It would be preferable to insist on the stronger condition that $UC(\mathcal{L}, \bigoplus i \bullet A_i) = \bigoplus i \bullet UC(\mathcal{L}, A_i)$. However, the definitions of lifted operators respecting that condition would be excessively complicated and intractable to analyse.

Condition 5.13. Lifted operators are monotonic with respect to \sqsubseteq :

$$\begin{array}{l} \mathsf{UC}\left(\mathcal{L},A_{i}\right) \sqsubseteq \mathsf{UC}\left(\mathcal{L},A_{i}'\right) \quad \text{for each } i \\ \text{implies} \quad \left(\widehat{\bigoplus} i \bullet \mathsf{UC}\left(\mathcal{L},A_{i}\right)\right) \sqsubseteq \left(\widehat{\bigoplus} i \bullet \mathsf{UC}\left(\mathcal{L},A_{i}'\right)\right) \end{array}$$

This condition is essential to support piecewise refinement of lifted specifications. It enables us to refine a lifted action *B* in isolation, with the assurance that any context in which *B* appears is also refined.

Theorem 5.14. Each lifted operator defined in this section satisfies Conditions 5.11, 5.12 and 5.13.

5.3.1 Sequential Composition

Closed forms of some lifted operators can be derived by unfolding the lifts of the respective *Circus* operators. For instance, Lemma 5.15 presents the outcome of unfolding the sequence operator.

Lemma 5.15 (Unfolded sequence).

$$\mathsf{UC}(\mathcal{L}, A_1; A_2) \quad \sqsubseteq \quad \exists x_0, \widetilde{x_0} \bullet \begin{pmatrix} \mathsf{UC}(\mathcal{L}, A_1) [x_0, \widetilde{x_0} / x', \widetilde{x'}] \\ \land \quad \mathsf{UC}(\mathcal{L}, A_2) [x_0, \widetilde{x_0} / x, \widetilde{x}] \end{pmatrix}$$

We obtain a definition of lifted sequential composition by generalising Lemma 5.15: we replace $U(A_1)$ and $U(A_2)$ with arbitrary points in the lifted space B_1 and B_2 .

Definition 5.16 (Lifted sequential composition).

$$B_1; B_2 \triangleq \exists x_0, \widetilde{x_0} \bullet B_1[x_0, \widetilde{x_0}/x', \widetilde{x'}] \land B_2[x_0, \widetilde{x_0}/x, \widetilde{x}]$$

Since the \hat{i} operator is just relational composition over the $x', \tilde{x'}$ and x, \tilde{x} variables, it respects the three conditions. Moreover, Law 5.17 shows that the laws of sequential composition (such as associativity) also apply to the lifted sequential composition operator.

Law 5.17 (U and \hat{i}). **U** $(A_1; A_2) =$ **U** $(A_1) \hat{i}$ **U** (A_2)

5 Specifying Confidentiality in Circus

Whenever a **UC**-lifted construct comprises an action in sequence with a coercion, assignment or *Skip*, it can be restructured as Lemma 5.18 shows.

Lemma 5.18. Provided that either A_1 or A_2 terminate instantaneously, leaving the trace unchanged:

$$UC(\mathcal{L}, A_1; A_2) = UC(\mathcal{L}, A_1); UC(\mathcal{L}, A_2)$$

5.3.2 Internal Choice

A closed form for lifted internal choice is easily derived by unfolding.

Lemma 5.19 (Unfolded internal choice).

$$\mathsf{UC}\left(\mathcal{L}, A_1 \sqcap A_2\right) \quad \sqsubseteq \quad \mathsf{UC}\left(\mathcal{L}, A_1\right) \lor \mathsf{UC}\left(\mathcal{L}, A_2\right)$$

Here, the refinement step cuts down the fog behaviours associated with A_1 and A_2 individually, but all behaviours of A_1 and A_2 are retained. Again, we generalise Lemma 5.19 to define a lifted internal choice operator which, in this case, is identical to the *Circus* internal choice.

Definition 5.20 (Lifted internal choice).

$$B_1 \sqcap B_2 \triangleq B_1 \lor B_2$$

5.3.3 Prefixing

A prefixed action $c \rightarrow A$ is identical to $(c \rightarrow Skip)$; *A*. This leads to a straightforward definition of the lifted forms of prefixing and input prefixing, based on the lifted sequence operator.

Definition 5.21 (Lifted prefixing).

$$c.e \xrightarrow{\frown} B \triangleq UC(\mathcal{L}, c.e \to Skip) ; B$$
$$c?e \xrightarrow{\frown} B \triangleq UC(\mathcal{L}, c?e \to Skip) ; B$$

The lifted reactive design form of lifted prefixing — presented in Lemma 5.22 — is derived by instantiating Theorem 5.7.

Lemma 5.22 (Unfolded prefixing).

$$\begin{aligned} & \mathsf{UC} \left(\mathcal{L}, c.e \to Skip \right) &= \\ & \widehat{\mathsf{R}} \left(\mathsf{true} \vdash \mathsf{U} \left(v' = v \right) \land \left(\begin{array}{c} \mathsf{U} \left(tr' = tr \land c.e \notin ref' \right) \\ \lhd wait' \rhd \\ \mathsf{U} \left(tr' = tr \land \langle (c,e) \rangle \right) \end{array} \right) \right) \land \mathcal{I}(\mathcal{L}) \end{aligned}$$

Lemma 5.22 indicates that, if Low could observe $c.e \rightarrow Skip$ in isolation, then Low could infer whether the event *c.e* has taken place — from *wait* and *wait'* — even if that event is absent from Low's window.

If Low can observe the event *c.e*, then the lifted form of $c.e \rightarrow A$ can be decomposed by separating the event instance from the action, as Lemma 5.23 shows.

Lemma 5.23 (Decomposing prefixing). Provided $(c, e) \in \mathcal{L}$:

$$UC(\mathcal{L}, c.e \rightarrow A) = UC(\mathcal{L}, c.e \rightarrow Skip); UC(\mathcal{L}, A)$$

5.3.4 Guarded Actions

Lemma 5.24 reformulates lifted guarded actions as lifted reactive designs.

Lemma 5.24 (Unfolded guard).

$$\mathbf{UC}(\mathcal{L}, g \& A) \subseteq \widehat{\mathbf{R}}\left(\left(\mathbf{U}(g) \Rightarrow \neg \mathbf{U}(A)_{ff}^{ff}\right) \vdash \left(\begin{array}{c} \mathbf{U}(g) \land \mathbf{U}(A)_{ff}^{tt} \\ \lor \mathbf{U}(\neg g \land tr' = tr \land wait') \end{array}\right)\right) \land \mathcal{I}(\mathcal{L})$$

Here, the postcondition of $UC(\mathcal{L}, g \& A)$ is strengthened by refinement, to impose the constraint $U(g) \lor U(\neg g)$ upon the fog space. This refinement is necessary to obtain a closed form purely in terms of U(A).

Definition 5.25 (Lifted guarded action).

$$g \widehat{\&} B \triangleq \widehat{\mathsf{R}} \left(\left(\mathsf{U}(g) \Rightarrow \neg B_{ff}^{ff} \right) \vdash \left(\begin{array}{c} \mathsf{U}(g) \land B_{ff}^{tt} \\ \lor \mathsf{U}(\neg g \land tr' = tr \land wait') \end{array} \right) \right)$$

Under this definition, we are effectively modelling Low as being able to determine whether *g* is **true** or **false**. Depending on the process in which the guarded action is embedded, this may over-approximate Low's inferences about the process state.

5.3.5 External Choice

Lemma 5.26 shows how lifted external choice can be refined to derive an expression in terms of the lifted actions $U(A_1)$ and $U(A_2)$.

Lemma 5.26 (Unfolded external choice).

$$\begin{aligned} & \operatorname{UC} \left(\mathcal{L}, A_1 \Box A_2 \right) & \sqsubseteq \\ & & \operatorname{R} \begin{pmatrix} \neg \operatorname{U} \left(A_{1_f}^{f} \right) \land \neg \operatorname{U} \left(A_{2_f}^{f} \right) \\ & \vdash \begin{pmatrix} \operatorname{U} \left(tr' = tr \land wait' \right) \land \operatorname{U} \left(A_1 \right)_{f\!f}^{tt} \land \operatorname{U} \left(A_2 \right)_{f\!f}^{tt} \\ & \lor \operatorname{U} \left(\neg \left(tr' = tr \land wait' \right) \right) \land \left(\operatorname{U} \left(A_1 \right)_{f\!f}^{tt} \lor \operatorname{U} \left(A_2 \right)_{f\!f}^{tt} \right) \\ & \land \mathcal{I}(\mathcal{L}) \end{aligned} \right) \end{aligned}$$

In summary, the consequences of the refinement in Lemma 5.26 are:

- Weakening the precondition means that both A₁ and A₁ (or A₂ and A₂) must diverge for the lifted external choice to diverge. This contrasts with UC (L, A₁ □ A₂), where divergence in either the observable space or the fog space suffices to make the entire lifted action divergent. Since divergence is to be avoided (Subsection 5.2.3), this point is unlikely to have importance in practice.
- Strengthening the postcondition models Low as being able to determine whether the external choice has resolved. Since Low can

perceive deadlock (by definition of $\mathcal{I}(\mathcal{L})$), this means that Low can detect whether any events have been performed, even if those events are outside its window. This over-approximation of Low's inferences, while perhaps counter-intuitive, has little impact on the application of lifted constructs.

Lemma 5.26 motivates the following definition for the lifted external choice operator.

Definition 5.27 (Lifted external choice).

$$B_1 \stackrel{\frown}{\Box} B_2 \triangleq \widehat{\mathbf{R}} \begin{pmatrix} \neg B_{1ff}^{ff} \land \neg B_{2ff}^{ff} \\ & \left(\mathbf{U} \left(tr' = tr \land wait' \right) \land B_{1ff}^{tt} \land B_{2ff}^{tt} \\ & \lor \mathbf{U} \left(\neg \left(tr' = tr \land wait' \right) \right) \land \left(B_{1ff}^{tt} \lor B_{2ff}^{tt} \\ & \left(B_{1ff}^{tt} \lor B_{2ff}^{tt} \right) \\ & \left(B_{1ff}^{tt} \lor B_{2ff}^{tt} \right) \end{pmatrix} \end{pmatrix}$$

5.3.6 Parallel Composition

In CSP and *Circus*, a parallel composition $A_1 | [cs] | A_2$ diverges whenever A_1 or A_2 diverge, provided those actions agree on their previous synchronsiations. Lemma 5.28 lifts Oliveira's definition of the precondition of a parallel composition.

Lemma 5.28 (Unfolded parallel precondition).

$$PREPAR(\mathbf{U}(A_1),\mathbf{U}(A_2)) \subseteq \mathbf{U}\left(\neg (A_1 | [cs] | A_2)_f^f\right)$$

where $PREPAR(\mathbf{U}(A_1), \mathbf{U}(A_2))$ abbreviates:

$$\neg \exists 1.tr, 2.tr, \widetilde{1.tr}, \widetilde{2.tr} \bullet \begin{pmatrix} \mathsf{U} (A_1)_{ff}^{ff} \, \widehat{;} \, \mathsf{U} (1.tr' = tr) \\ \land \quad \mathsf{U} (A_2)_{ff} \, \widehat{;} \, \mathsf{U} (2.tr' = tr) \\ \land \quad \mathsf{U} (1.tr' \upharpoonright cs = 2.tr' \upharpoonright cs) \end{pmatrix}$$

$$\land \quad \neg \exists 1.tr, 2.tr, \widetilde{1.tr}, \widetilde{2.tr} \bullet \begin{pmatrix} \mathsf{U} (A_1)_{ff} \, \widehat{;} \, \mathsf{U} (1.tr' = tr) \\ \land \quad \mathsf{U} (A_2)_{ff}^{ff} \, \widehat{;} \, \mathsf{U} (2.tr' = tr) \\ \land \quad \mathsf{U} (A_2)_{ff}^{ff} \, \widehat{;} \, \mathsf{U} (2.tr' = tr) \\ \land \quad \mathsf{U} (1.tr' \upharpoonright cs = 2.tr' \upharpoonright cs) \end{pmatrix}$$

As with lifted external choice, weakening the precondition is necessary to derive an expression in terms of $U(A_1)$ and $U(A_2)$. Here, this weakening ensures the lifted parallel composition diverges only if both the observable trace and the fog trace of either lifted action leads to divergence, and the other lifted action is prepared to synchronise on the projections of those traces. Again, this point is of no importance if the lifted actions are divergence-free.

Lemma 5.28 is a stepping stone for the following derivation of a lifted reactive design modelling the parallel composition of two lifted actions:

$$\begin{aligned} & \mathsf{UC} \left(\mathcal{L}, A_1 \left| [ns_1 \mid cs \mid ns_2] \right| A_2 \right) \\ & \sqsubseteq \quad \widehat{\mathsf{R}} \begin{pmatrix} PREPAR(\mathsf{U} \left(A_1 \right), \mathsf{U} \left(A_2 \right)) \\ & \vdash \quad \mathsf{U} \left(\begin{pmatrix} A_1_f^t ; U1(out\alpha A_1) \\ & \land A_2_f^t ; U2(out\alpha A_2) \end{pmatrix} ; M_{\parallel}(cs) \end{pmatrix} \right) \land \mathcal{I}(\mathcal{L}) \end{aligned}$$

[Theorem 5.7; Lemma 5.28 (weaken precondition)]

$$= \widehat{\mathbf{R}} \begin{pmatrix} PREPAR(\mathbf{U}(A_1), \mathbf{U}(A_2)) \\ \vdash \begin{pmatrix} \mathbf{U}(A_1)_{ff}^{tt} ; \mathbf{U}(U1(out\alpha A_1)) \\ \wedge \mathbf{U}(A_2)_{ff}^{tt} ; \mathbf{U}(U2(out\alpha A_2)) \end{pmatrix} ; \mathbf{U}(M_{\parallel}(cs)) \end{pmatrix} \wedge \mathcal{I}(\mathcal{L})$$
[Law 4.7; Law 5.17]

We interpret the relabelled variables i.x, $\tilde{i.x}$ as members of x, \tilde{x} respectively, so **U** extends over them. By this interpretation, lifting the relabelling functions *U*1 and *U*2 is straightforward:

$$\mathbf{U}(Ui(out\alpha A)) = x.i = x' \land \widetilde{x.i} = \widetilde{x'}$$

Expanding the lifted form of the merge function $M_{\parallel}(cs)$ is not illuminating, so we refrain from doing so.

Finally, Definition 5.29 presents a parallel operator for lifted actions.

Definition 5.29 (Lifted parallel composition).

$$B_{1} |[ns_{1} | cs | ns_{2}]| B_{2} \triangleq$$

$$\widehat{\mathbf{R}} \left(PREPAR(B_{1}, B_{2}) \vdash \begin{pmatrix} B_{1}_{ff}^{tt} ; U1(out\alpha B_{1}) \\ \wedge B_{2}_{ff}^{tt} ; U2(out\alpha B_{2}) \end{pmatrix} ; \mathbf{U} (M_{\parallel}(cs)) \right)$$

5.3.7 Hiding

Hiding is the final operator we consider. Again, we derive a lifted form by unfolding **UC** (\mathcal{L} , $A \setminus cs$).

Lemma 5.30 (Unfolded hiding).

$$\mathsf{UC}\left(\mathcal{L}, A \setminus cs\right) \quad = \quad \left(HID(\mathsf{U}\left(A\right)) \land \mathcal{I}(\mathcal{L})\right) \,\widehat{;} \, \mathsf{UC}\left(\mathcal{L}, Skip\right)$$

where *HID*(*B*) denotes:

$$\widehat{\mathbf{R}}\left(\exists s, \widetilde{s} \bullet \left(\begin{array}{c} B[s, \widetilde{s}, (cs \cup ref'), (cs \cup \widetilde{ref'})/tr', t\widetilde{r'}, ref', \widetilde{ref'}] \\ \land \quad \mathbf{U}\left((tr' - tr) = (s - tr) \upharpoonright (EVENT - cs)\right) \end{array}\right)\right)$$

As usual, we adapt Lemma 5.30 to yield Definition 5.31.

Definition 5.31 (Lifted hiding).

$$B \widehat{\setminus} cs \triangleq (HID(B) \land \mathcal{I}(\mathcal{L})) ; UC(\mathcal{L}, Skip)$$

Lifted hiding distributes through lifted operators in the same manner that *Circus* hiding distributes through *Circus* operators.

5.4 Lifting Circus Processes

This section addresses two issues. First, we identify a method for specifying the windows of users in a *Circus* specification. Second, we extend the *Circus* syntax to enable the specification of lifted processes.

5.4.1 Encoding Windows

It is natural to model Low's window as a subset of the channels of the process. Within a *Circus* specification, we can use the *Circus* keyword **channelset** to assign a label to that window.

Following Oliveira et al. (2009), a channel set specification can be readily translated to the corresponding set of all events visible to Low. Let $\delta(c)$ denote the type of channel *c*; if *c* does not communicate typed values, then we take $\delta(c) = \{Sync\}$. We then have:

$$\mathcal{L} = \{(c, e) \mid c \in L \land e \in \delta(c)\}$$

Example 5.32. To specify that Low can observe events on the *on* and *off* channels of a process, we write:

channel *on*, *off* **channelset** $Low \triangleq \{ | on, off | \}$

This channel set corresponds to the window $\{(on, Sync), (off, Sync)\}$.

5.4.2 Lifting and Closure

In the UTP, it is desirable for program operators defined in some theory to be closed with respect to the theory: for instance, if P_1 and P_2 are objects of the theory, then so should be $P_1 \oplus P_2$. However, by applying the lifted *Circus* operators defined in Section 5.3 to **UC**-lifted actions, we can formulate lifted constructs that are not closed with respect to **UC**.

Example 5.33. The construct $UC(\mathcal{L}, h := 0) \cap UC(\mathcal{L}, h := 1)$ cannot be derived by applying UC to any single *Circus* action. As Figure 5.1 shows, it is stronger than $UC(\mathcal{L}, h := 0 \sqcap h := 1)$, but weaker than each of $UC(\mathcal{L}, h := 0)$ and $UC(\mathcal{L}, h := 1)$.

Constructs such as $UC(\mathcal{L}, h := 0) \cap UC(\mathcal{L}, h := 1)$ are not part of the UC-image of the space of *Circus* actions. Instead, they inhabit a larger space of predicates with the joint alphabet of observable and fog variables.

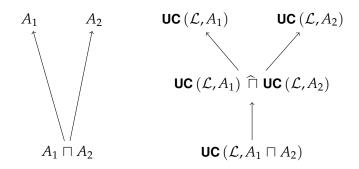


Figure 5.1: A comparison of the *Circus* refinement lattice (left side) with the refinement lattice of the lifted space (right side).

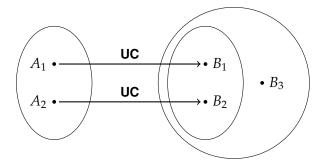


Figure 5.2: An illustration of the relationship between the *Circus* space (left side) and the lifted space (right side).

Figure 5.2 visualises the relationship between the *Circus* space, the **UC** space and this lifted space.

Example 5.34. In Figure 5.2, B_1 and B_2 are the **UC**-projections of A_1 and A_2 . However, B_3 (which could denote $B_1 \oplus B_2$) lies outside the **UC**-image of the space of *Circus* actions.

Results which apply to the lifted **UC** space may fail to hold over the larger lifted space, in general. However, we shall restrict our attention to constructs that can be formed by combining **UC**-lifted actions with the lifted operators. This restriction guarantees the three principal conditions of the lifted operators are maintained for the constructs that we work with. It also ensures these constructs are expressible in the form $\widehat{\mathbf{R}}$ (*Pre* \vdash *Post*)

via the definitions of the lifted operators. A new UTP theory would be needed to define this restricted space of acceptable lifted constructs formally, but this topic is left for future work.

For our purposes, the restriction on the form of lifted constructs is sufficient to avoid practical difficulties stemming from the lack of **UC**-closure. In particular, the techniques presented in Chapter 6 can be applied to lifted constructs both inside and outside the **UC** space.

An important consequence of moving outside the **UC** space is that a compound *Circus* construct — such as $A_1 \sqcap A_2$ — can be translated to the lifted space in multiple ways. For instance:

- UC (L, A₁ ⊓ A₂) gives an exact representation of Low's inferences about A₁ ⊓ A₂; while
- UC (L, A₁) ∩ UC (L, A₂) models Low's inferences about A₁ and A₂ separately. Segmenting lifted actions in this way makes it easier to analyse Low's inferences, albeit at the cost of over-approximation.

The difference between these forms has important consequences for verification and refinement, which are described in detail in Chapter 6. Hence, when we interpret a *Circus* specification in the lifted space, we need to pay special attention to distinguishing between these forms.

5.4.3 Blocks

We now present the *block*, which is a syntactic construct for specifying how *Circus* actions should be translated to the lifted space. Blocks are little more than concrete syntax for **UC**; we apply them to *Circus* actions to encode the granularity of lifted actions within the body of a process. By delineating the boundaries between lifted actions explicitly, blocks describe the translation of *Circus* processes into the lifted space.

Definition 5.35 (Block).

$$\langle L : A \rangle \triangleq (\mathbf{UC}(\mathcal{L}, A) \lhd \ell = L \triangleright A)$$

where \mathcal{L} denotes the window declared by the channel set *L*.

Each block is labelled with the name of a channel set, denoting a user's window. The metavariable ℓ can be instantiated with a window label, to focus attention on a specific user. This means that blocks can be used to encode the inferences of multiple users within a single process specification.

We say that blocks that are labelled with the same channel set are members of the same *family*. In the presentation that follows, we typically consider only a single Low user and omit an explicit window declaration from blocks. As a shorthand, we define:

$$\langle A \rangle \triangleq \langle \ell : A \rangle$$

Taking Low's window \mathcal{L} to be implicit, Definition 5.35 gives us:

$$\langle A \rangle = \mathbf{UC}(\mathcal{L}, A)$$

The block construct sets the language apart from standard *Circus*. To maintain a level of compatibility with *Circus* specifications, we adhere to a list of syntactic conventions:

- 1. Each action of a process must be nested within a block.
- 2. Blocks of the same family cannot be nested within each other.
- 3. Outside of blocks, we use the syntax of *Circus* operators to denote the lifted forms of those operators.

Example 5.36. Recalling Example 5.33, we have:

$$\langle A_1 \sqcap A_2 \rangle = \mathbf{UC} (\mathcal{L}, A_1 \sqcap A_2) \langle A_1 \rangle \sqcap \langle A_2 \rangle = \mathbf{UC} (\mathcal{L}, A_1) \widehat{\sqcap} \mathbf{UC} (\mathcal{L}, A_2)$$

Henceforth, we tend to use the block notation in preference to UC, to aid readability. \diamond

When dealing with processes that interact with multiple users, it may be convenient to model the inferences of different users at different levels of granularity. The atomicity of lifted constructs — from the perspective of different users — may be specified compactly by nesting blocks of different families.

Example 5.37. Let *Alice* label Alice's window, and *Bob* label Bob's window, within the following specification:

$$\langle Alice : \langle Bob : A_1 \rangle \sqcap \langle Bob : A_2 \rangle \rangle$$

If we are reasoning about Alice's inferences (i.e. $\ell = Alice$), we read this specification as $\langle Alice : A_1 \sqcap A_2 \rangle$. Whereas if we are reasoning about Bob, we read it as $\langle Bob : A_1 \rangle \sqcap \langle Bob : A_2 \rangle$.

5.4.4 On Equivalence Laws

As a consequence of the order embedding between *Circus* actions and lifted actions (Lemma 5.3), all equivalence laws over *Circus* actions apply also to actions nested within blocks. However, Condition 5.12 implies that standard equivalence laws between *Circus* actions do not necessarily apply across blocks. For instance, while sequence distributes through internal choice within a *Circus* action:

$$\langle (A_1 \sqcap A_2); A_3 \rangle = \langle (A_1; A_3) \sqcap (A_2; A_3) \rangle$$

it does not distribute across blocks, in general:

$$\langle A_1 \sqcap A_2 \rangle$$
; $\langle A_3 \rangle \neq \langle A_1; A_3 \rangle \sqcap \langle A_2; A_3 \rangle$

Nevertheless, many algebraic properties of the *Circus* operators also apply to the lifted operators; for instance:

$$(\langle A_1 \rangle \sqcap \langle A_2 \rangle); \langle A_3 \rangle = (\langle A_1 \rangle; \langle A_3 \rangle) \sqcap (\langle A_2 \rangle; \langle A_3 \rangle)$$

104

The task of identifying a general collection of equivalence laws is left as future work.

5.5 Confidentiality Annotations

This section introduces confidentiality annotations — or CAs for short — which are algebraic constructs for incorporating obligations within lifted processes. A CA may be embedded at any point in the body of a lifted process, alongside lifted actions, to express a confidentiality property over the process state when the CA is invoked. By specifying confidentiality properties in this way, they become amenable to compositional reasoning: we can consider the effect of CAs on a process in terms of their effect on the behaviour of individual *Circus* actions.

5.5.1 Defining Confidentiality Annotations

The fog variables of a **UC**-lifted action track the information that Low can infer about the action's behaviour. The relation between the observable variables and fog variables can be employed for a secondary purpose: to capture the meaning of CAs within the lifted semantics directly.

Definition 5.38 formalises CAs with respect to the lifted semantics.

Definition 5.38 (Confidentiality annotation).

$$\langle L : \theta \rangle \triangleq \mathbf{UC}(\mathcal{L}, Skip) \land (\ell = L \Rightarrow Conf(\theta))$$

where Conf $(\theta) \triangleq (ok \land \neg wait \Rightarrow \theta)$.

As with blocks, we abbreviate $\langle \ell : \theta \rangle$ to $\langle \theta \rangle$ if we wish to consider only a single Low user:

$$\langle \theta \rangle = \mathbf{UC} (\mathcal{L}, Skip) \wedge \text{Conf} (\theta)$$

With respect to the observable behaviour of processes, a CA behaves as the lifted *Skip*: it terminates instantaneously, leaving the process trace and state unchanged.

Example 5.39. When invoked, the CA $\langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$ specifies that the process must not reveal to Low that the state satisfies the condition h = 0, with states where h > 0 serving as cover stories.

We can combine CAs with lifted *Circus* constructs using the lifted operators defined in Section 5.3. Hence, we can use CAs to impose confidentiality properties over the process state at intermediate points of the process's execution. Even though CAs are localised to a specific region of the process body, their effects diffuse throughout the process.

Typically, we place CAs in sequence with blocks, to impose confidentiality properties on the process state at specific points of its execution. For instance, to specify the obligation θ applies to the process state when the action *A* terminates, we would write:

$$\ldots$$
; $\langle A \rangle$; $\langle \theta \rangle$; \ldots

CAs may be used in other ways. For instance, a process may be offered a choice between performing the action *A* or satisfying a CA:

$$\ldots$$
; $(\langle A \rangle \sqcap \langle \theta \rangle)$; \ldots

5.5.2 Properties of Confidentiality Annotations

CAs inherit many of the properties of obligations. This is because the semantics of CAs induces an order isomorphism between the \Box -ordering over obligations and the \Box -ordering over CAs.

Theorem 5.40 (CAs and \sqsubseteq). Provided θ_1, θ_2 reference only state variables:

$$\theta_1 \sqsubseteq \theta_2$$
 if and only if $\langle \theta_1 \rangle \sqsubseteq \langle \theta_2 \rangle$

Hence, in the lifted semantics, the refinement ordering represents better confidentiality as well as better functionality.

From Definition 4.28, it follows that the sequential composition of CAs can be simplified to a CA embedding the least upper bound of their respective obligations.

Law 5.41 (CA least upper bound).

$$\langle \theta_1 \rangle; \ldots; \langle \theta_n \rangle = \langle \bigsqcup \{ \theta_1, \ldots, \theta_n \} \rangle$$

Law 5.41 implies that CAs in sequence are commutative.

Law 5.42 (CA commutativity). $\langle \theta_1 \rangle$; $\langle \theta_2 \rangle = \langle \theta_2 \rangle$; $\langle \theta_1 \rangle$

Following Definition 4.29, a choice between multiple CAs is equivalent to the greatest lower bound of their respective obligations.

Law 5.43 (CA greatest lower bound).

$$\langle \, heta_1 \,
angle \, \sqcap \, \lor \sqcap \, \langle \, heta_n \,
angle \quad = \quad \langle \, igcap_{\{ eta_1, \ldots, eta_n \}} \,
angle$$

The semantics of CAs is closely connected with the semantics of lifted *Circus* coercions. Indeed, the CA is a generalisation of the lifted coercion, as Law 5.44 shows.

Law 5.44 (CA and coercion). $\langle [C] \rangle = \langle \mathbf{U}(C) \rangle$

Applying **D** to a CA yields a *Circus* coercion, as Law 5.45 shows.

Law 5.45 (CA and D). $D(\langle \theta \rangle) = [D(\theta)]$

There is a kind of duality between CAs and *Circus* coercions. A lifted coercion $\langle [C] \rangle$ demands that both the process state and the fog state satisfy *C*, thus placing a *lower bound* on Low's inferences about the state. In contrast, CAs place an *upper bound* on Low's inferences.

It is often convenient to rewrite the sequential composition of a CA and a *Circus* action as a single compound action. Law 5.46 and Law 5.47 are concerned with folding a CA with a block that immediately precedes or succeeds it.

Law 5.46 (CA left sequence). $\langle \theta \rangle$; $\langle A \rangle = \langle A \rangle \land Conf(\theta)$

Law 5.47 (CA right sequence). $\langle A \rangle$; $\langle \theta \rangle = \langle A \rangle \land Conf'(\theta)$

where Conf' $(\theta) \triangleq (ok' \land \neg wait' \Rightarrow \theta')$

5.5.3 Passive Confidentiality Annotations

A CA $\langle \theta \rangle$ is *passive* in a state ψ if $[\psi \Rightarrow \theta]$ holds; that is, θ does not constrain the fog space associated with ψ in any way. When passive, a CA's behaviour is just $\langle Skip \rangle$.

If a CA within a process body is passive in every state in which it can be invoked by the process, then we call that CA *innocuous* in the context of that process. CAs can also be innocuous if they are never invoked: for instance, $\langle Stop \rangle$ and $\langle Chaos \rangle$ are left zeros (with respect to sequential composition) for CAs, just as they are for lifted *Circus* actions.

An innocuous CA has no effect whatsoever upon lifted actions. Hence, innocuous CAs can be freely inserted or removed within a process, without changing its meaning.

Example 5.48. Suppose the CA from Example 5.39 follows the lifted command $\langle h := 1 \rangle$. Let:

$$B1 \quad riangleq \quad \langle h := 1 \rangle \ ; \langle h = 0 \Rightarrow \widetilde{h} > 0 \rangle$$

The CA is innocuous in this context, as calculation shows:

$$\langle h := 1 \rangle ; \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$$

$$= \langle h := 1 \rangle ; \langle [h = 1] \rangle ; \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$$
 [coercion introduction]
$$= \langle h := 1 \rangle ; \langle h = 1 \land \tilde{h} = 1 \rangle ; \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$$
 [Law 5.44]
$$= \langle h := 1 \rangle ; \langle h = 1 \land \tilde{h} = 1 \land (h = 0 \Rightarrow \tilde{h} > 0) \rangle$$
 [lub (Law 5.41)]
$$= \langle h := 1 \rangle ; \langle [h = 1] \rangle$$
 [Law 5.44]
$$= \langle h := 1 \rangle ; \langle [h = 1] \rangle$$
 [Law 5.44]

 \Diamond

5.5.4 Active Confidentiality Annotations

A CA $\langle \theta \rangle$ is *active* in a state ψ if θ constrains the fog space associated with ψ . Operationally, the Conf (θ) predicate prunes the fog of ψ , leaving only those fog states prescribed as cover stories by θ for ψ . This pruning commits the process design to supporting at least one (Low-indistinguishable) fog behaviour that passes through the CA in a cover story state.

The commitment made by a CA is broken if none of the cover story states prescribed by θ are present in the fog when the CA is invoked. *Breaking the commitment symbolises a conflict between the functionality of the process and its confidentiality properties.* This conflict becomes manifest in the following example.

Example 5.49. Suppose *B*0 names the CA from Example 5.39 following the command h := 0:

B0
$$\triangleq \langle h := 0 \rangle; \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$$

Intuitively, the demands of functionality ($\tilde{h} = 0$) and confidentiality ($\tilde{h} > 0$) that *B*0 places on the fog space are irreconcilable. The consequences of this conflict are brought into stark relief by calculation:

This CA prunes every fog state.

In a very real sense, the CA \langle **false** \rangle specifies the impossible. By Law 5.44 and the definition of coercions, we have:

$$\langle \text{ false } \rangle = \langle [\text{ false }] \rangle = \langle \mathsf{R}(\neg ok) \rangle$$

Recall from Section 2.6 that $\mathbf{R}(\neg ok) = \mathbf{R} (\text{true} \vdash \text{false})$ is the reactive design miracle. As in Section 4.4, miracles give meaning to pro-

 \Diamond

cess specifications where confidentiality properties are inconsistent with functionality properties. *By behaving miraculously, a CA disables insecure functionality.*

Example 5.50. The specification $\langle h := 0 \rangle$; $\langle false \rangle$ from Example 5.49 can be simplified as follows:

$$\langle h := 0 \rangle; \langle false \rangle$$

$$= \langle h := 0 \rangle; \langle [false] \rangle$$

$$= \langle h := 0; [false] \rangle$$

$$[Lemma 5.18]$$

$$= \langle [false] \rangle$$

$$[property of coercions]$$

which, of course, is unimplementable.

The commitment made by a CA can also be broken later in a process's execution, again inducing miraculous behaviour, as Example 5.51 shows.

Example 5.51. Suppose that Low can observe values transmitted on the *out* channel. Consider the process fragment:

$$\langle h = 0 \Rightarrow \widetilde{h} > 0 \rangle$$
; $\langle out!h \to Skip \rangle$

Were Low to observe *out*!0, Low could infer h = 0 with certainty. Hence, this specification behaves miraculously if h = 0.

5.5.5 Discussion

Example 5.48 and Example 5.49 represent the extreme ends of a scale of the effects of CAs. The interplay between *Circus* actions and active CAs is often more subtle.

Example 5.52. Suppose a process is offered the choice between the (infeasible) *B*0 and (feasible) *B*1. Its behaviour is given by calculation:

 $B0 \sqcap B1$ $= UC(\mathcal{L}, \mathbf{R}(\neg ok)) \sqcap UC(\mathcal{L}, h := 1) \quad [Example 5.48, Example 5.49]$

 \Diamond

$$= UC (\mathcal{L}, h := 1)$$
[lattice top]
$$= B1$$

The miraculous *B*0 is bypassed by the internal choice, because *B*1 is strictly weaker than *B*0. \Diamond

As Example 5.52 shows, the presence of miracles within a process body does not necessarily render the whole process miraculous. While a miracle may disable part of a process body, the process may remain feasible, so long as the process can "backtrack" from the miracle by resolving non-determinism. For instance, the *Circus* action $c \rightarrow [$ **false**]cannot refuse to engage in an event c. Yet it can never actually accept c, because it would need to satisfy the reactive design miracle [**false**] were it to do so (Woodcock, 2010).

As Woodcock (2010) and Wei et al. (2010) point out, $c \rightarrow [$ **false**] is a *Circus* action that lies outside the space of CSP processes, because it violates the axioms of the failures-divergences model of CSP. The same effect in the lifted space may be induced by CAs. For instance, $\langle c \rightarrow Skip \rangle$; \langle **false** \rangle is equivalent to $\langle c \rightarrow [$ **false** $] \rangle$, so *c* is prevented from occurring.

We may interpret the meaning of miracles within process designs — whether induced by CAs or otherwise — in two ways:

- A miraculous process specification cannot be implemented as a program; therefore, it should be deemed invalid. This interpretation reflects the philosophy of the conformance condition defined in Chapter 4: a process that fails to conform to the confidentiality properties in its specification does not satisfy its specification.
- 2. Alternatively, we may embrace specifications with miraculous elements, since they provide the designers of a process with greater flexibility in making implementation choices. Provided these miraculous elements can be avoided by the process, or are finessed away by refinement, their presence can be tolerated.

We adopt the second interpretation, because it is more conducive to the spirit of the UTP (Hoare, 1984a) and refinement calculi (Morgan, 1994). Furthermore, if we were to reject miraculous specifications as invalid, then we could only treat CAs as a form of documentation.

Remark 5.53. An interesting application of miraculous CAs is to direct the process's behaviour according to Low's inferences about its state. Let *Normal* stand for an arbitrary action in the specification:

$$\langle C \Rightarrow \neg \widetilde{C} \rangle$$
; $\langle Normal \rangle \square \langle C \& alarm \rightarrow Skip \rangle$

If Low can establish the state satisfies *C*, then the CA $\langle C \Rightarrow \neg \tilde{C} \rangle$ is violated and this specification reduces to:

$$\langle$$
 false $\rangle \square \langle$ *alarm* \rightarrow *Skip* \rangle

which forces the *alarm* to be raised *urgently*, because the miraculous \langle **false** \rangle forces the external choice to be resolved instantly. (This effect is a consequence of the semantics of external choice; it is explained in more detail by Woodcock (2010).) Alternatively, if the state satisfies \neg *C*, then the process behaves as $\langle Normal \rangle$. However, this specification does not prevent "false alarms": the environment may choose to raise the alarm in a state satisfying *C*, even if the CA is satisfied.

5.6 Event-Based Confidentiality Annotations

Up to now, we have studied how CAs can specify confidentiality properties over the state of a process. However, it is often the interactions that a process performs with its environment (or its other users) which are confidential. This section presents techniques for using CAs in order to express confidentiality properties over those interactions.

5.6.1 Inputs and Outputs

Be they inputs or outputs, the data values that a process communicates with its environment may be secret. It is straightforward to specify confidentiality properties over those values using CAs.

The semantics of an input prefixing $c?i \rightarrow A$ declares *i*? as a local variable in the context of *A*. This variable can be incorporated into a CA within the scope of the input prefixing.

Example 5.54. The following process fragment specifies that Low cannot deduce if a value transmitted on the *c* channel is a member of the set *S*:

$$\langle c?i \rightarrow h := i? \rangle$$
; $\langle h \in S \Rightarrow \tilde{h} \notin S \rangle$

This CA permits Low to infer the occurrence of the event.

 \Diamond

With an output event *c*!*E*, the variables within the expression *E* can just be referenced by a CA directly.

5.6.2 Occurrence of Events

Other kinds of confidentiality properties — such as the noninterference properties discussed in Section 4.6 — are concerned with restricting Low's inferences about whether a process has (or has not) engaged in particular events. In the literature, these properties are typically formulated over the traces over a process. Since the process trace is a product of the whole process, we cannot encode trace-based properties using CAs directly.

Fortunately, the symbiosis between the state and behaviour of *Circus* processes allows information about the control flow of the process to be recorded with state variables. In essence, we can use a local variable to record whether a particular interaction is performed by the process. This variable can then be referenced by a CA to indirectly encode what Low is entitled to infer about those interactions. This scheme is illustrated briefly by Example 5.55.

Example 5.55. Recall the specification *OO* given in Example 5.10, joined with the CA from Example 5.39:

$$\langle on \rightarrow h := 0 \Box off \rightarrow h := 1 \rangle$$
; $\langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$

This specification stipulates that Low cannot be sure that *on* has occurred: it demands that, if *on* does occur, then Low must be unable to rule out the occurrence of *off* instead. \diamond

A process need not record its control flow in terms of its state variables, but by introducing local variables, it can always be transformed into an equivalent process that does. In this way, we can formulate various confidentiality properties over the interactions a process performs.

Example 5.56. Consider the following choice construct:

$$\langle sec \rightarrow A_1 \Box cov \rightarrow A_2 \rangle$$

Suppose that *sec* is a high-level event whose occurrence is deemed to be secret. If we regard the occurrence of *cov* (implying the non-occurrence of *sec*) as a suitable cover story, then we can formulate a confidentiality property by extending the choice construct with a local variable as follows:

$$\operatorname{var} h \bullet \left(\begin{array}{c} \langle \sec \to h := 0 \ \Box \ \cos \to h := 1 \rangle \ ; \\ \langle h = 0 \Rightarrow \widetilde{h} > 0 \rangle \ ; \\ \langle h = 0 \& A_1 \ \Box \ h = 1 \& A_2 \rangle \end{array} \right)$$

It is important to invoke the CA immediately after *sec* takes place, for otherwise it would not take effect if A_1 deadlocks.

With this scheme, it is necessary to extend the process with a separate CA wherever a confidential interaction may take place. Moreover, in a realistic *Circus* process, the secret and cover story events need not be located within the same choice construct, so determining how the process should be extended with a CA may require some creativity.

We now propose a compositional method for structuring the specification of confidentiality properties over the events of a process.

5.6.3 Superposition

CAs over the events of a process can be packaged together in a *confidentiality specification*. The purpose of a confidentiality specification is to monitor the activity of a process and to invoke CAs whenever the process engages in a confidential interaction.

We formalise a confidentiality specification as a recursive construct that is always ready to engage in a set of high-level events. Here, we focus our attention on confidentiality specifications expressible in the form:

$$\mu X \bullet \operatorname{var} h : E \bullet \langle \Box e : E \bullet c.e \to h := e \rangle ; \langle \theta \rangle ; X$$

where *E* denotes the type of channel *c*; *h* is a variable of type *E*; and θ is an obligation over *h* and \tilde{h} .

Remark 5.57. More sophisticated kinds of confidentiality specifications could be defined with a local state, to record the activity of the process in detail and to impose CAs when particular patterns arise.

Suppose we place the confidentiality specification above in parallel composition (synchronising on channel c) with a lifted process. Then, whenever the process accepts an event *c.e.*, the same event invokes the CA in the confidentiality specification with h set to e. In this way, the confidentiality property defined by the CA is imposed upon the process.

We call this method *superposition*, as it is closely related to the superposition method for program structuring proposed by Chandy and Misra (1988), Back and Sere (1996) and others. Superposition promotes separation of concerns: it enables us to specify the confidentiality requirements of a system without regard to its functionality, and vice versa.

The following two examples show how the confidentiality properties described in Example 5.54 and Example 5.56 can be reformulated as confidentiality specifications.

Example 5.58. Continuing from Example 5.54, a confidentiality specification over the values communicated over channel *c* is as follows:

$$\mu X \bullet \operatorname{var} h \bullet \langle c?i \to h := i? \rangle ; \langle h \in S \Rightarrow \widetilde{h} \notin S \rangle ; X$$

Again, this specification does not prohibit Low from inferring that events on *c* have taken place. \Diamond

Example 5.59. The confidentiality specification for the choice construct described in Example 5.56 is:

$$CS = \mu X \bullet \operatorname{var} h \bullet \left\langle \begin{array}{c} \operatorname{sec} \to h := \operatorname{sec} \\ \Box & \operatorname{cov} \to h := \operatorname{cov} \end{array} \right\rangle; \langle h = \operatorname{sec} \Rightarrow \tilde{h} = \operatorname{cov} \rangle; X$$

The superposed specification is:

$$\langle sec \rightarrow A_1 \square cov \rightarrow A_2 \rangle \parallel \{ | sec, cov \mid \} \parallel CS$$

This superposed specification is not identical to the specification given in Example 5.56. If A_1 or A_2 engage in *sec* events, then the confidentiality specification will apply to those events as well.

We can apply superposition to express noninterference-like confidentiality properties (Subsection 4.6.1), where the occurrence of a high-level event is secret, and the non-occurrence of that event is the cover story. Example 5.60 illustrates this technique, in a manner reminiscent of Roscoe's lazy abstraction (Subsection 3.5.3).

Example 5.60. Let *Q* denote a lifted process which communicates *sec* events. The process:

$$QQ = Q \parallel \langle \mu X \bullet cov \to X \rangle$$

behaves as *Q*, but in addition may perform a *cov* event at any time. Taking the confidentiality specification from Example 5.59, the superposition:

 $(QQ |[\{ | sec, cov |\}]| CS) \setminus \{ | cov |\}$

specifies that Low cannot infer that QQ — and therefore Q — has performed a *sec* event. In this way, we impose the noninference property (described in Subsection 4.6.1) on *sec* events of Q.

5.6.4 Discussion

Superposing a confidentiality specification to a process achieves the same effect as modifying the process by attaching a separate CA to each highlevel event. Hence, a superposition of a confidentiality specification over a lifted process can be re-expressed as a single integrated process specification, by applying the step laws of parallel composition.

It would be worthwhile to create a library of reusable confidentiality specifications, to encapsulate common "design patterns" for encoding confidentiality requirements. Without needing to understand their technical detail, software engineers could compose these confidentiality specifications with *Circus* processes, to realise secure process designs.

5.7 Declassification

Up to now, we have treated the confidentiality of information as a static property. However, the secrecy of information may change over time. Information ceases to be secret when it is *declassified*.

Declassification may ease the implementation of the system. For instance, if information is deemed to be valueless to an adversary — such as old passwords or expired cryptographic keys — then it may be acceptable to reveal that information to the adversary. In other cases, the ability to declassify data is often essential to avoid conflicts with a system's functionality. For instance, the questions of an examination paper should be declassified to candidates at the moment when the exam starts, while sample answers may be declassified after the exam ends.

Declassification is not an issue for the confidentiality framework presented in Chapter 4. In that framework, if an obligation does not constrain the fog space associated with a particular behaviour of a system, then

that behaviour is implicitly unclassified. However, once a CA is invoked, the confidentiality property it imposes remains in force in perpetuity.

In this section, we discuss a simple approach for declassifying confidentiality properties within lifted processes. This approach relies on expressing CAs in a way such that they can be disabled by other constructs within a process specification. However, we have not solved the problem of defining a semantics for these constructs that is consistent with the reactive design model of *Circus*, which is a significant shortcoming of the approach.

5.7.1 Relaxing Confidentiality Annotations

Recall from Subsection 4.2.1 the dual purposes of the fog space: to model Low's inferences; and to capture confidentiality properties. As we have modelled them, these concerns are not separable from each other, which impedes the declassification of confidentiality properties.

When a CA is invoked, it cuts away fog behaviours that do not represent acceptable cover stories. By cutting down the fog space, the CA destroys information about Low's inferences encoded in the semantics of lifted constructs. Once lost, this information cannot be put back into the lifted space at any later point in the process's execution. For this reason, the semantics of CAs rules out a general scheme for declassifying them once they have taken effect.

It is possible to prevent CAs from taking effect when they are invoked. Consider a CA of the form $\langle c \Rightarrow \theta \rangle$ where *c* denotes a Boolean variable. This CA is innocuous if *c* is **false**, but is equivalent to $\langle \theta \rangle$ if *c* is **true**. Hence, the CA is *contingent* on *c*, because *c* controls whether the CA imposes θ upon the process state.

Making a CA contingent on c does not solve the problem of relaxing that CA after its invocation. If c is an ordinary state variable, then setting its value after the CA is invoked does not affect the CA. However, if c is not a state variable, but a variable whose value is determined by the control flow of the process, then it could be used to declassify the CA

once the process has reached a particular point in its execution.

5.7.2 A Declassification Construct

We postulate a special specification construct, **dec** *c*, which can coerce the value of *c* according to whether the construct is invoked by the process.

Definition 5.61 (dec statement).

 $\operatorname{dec} c \quad \triangleq \quad Skip \wedge c \Leftrightarrow wait'$

The statement **dec** c specifies that the truth of c is contingent on whether the statement is invoked by the process. Here, c is not an ordinary state variable: c is **false** if the process invokes **dec** c during its execution; and cis **true** if the process halts without executing **dec** c. Hence, the invocation of **dec** c retroactively declassifies any CA contingent on c.

Example 5.62. Here is a fragment of a specification for a (very simple) examination management system:

$$\mathsf{var} dq, ds \bullet \begin{pmatrix} set_paper?q.s \rightarrow \\ \langle Student : dq \Rightarrow \theta Questions(q?) \rangle; \\ \langle Student : ds \Rightarrow \theta Solutions(s?) \rangle; \\ \dots \\ \langle \mathsf{dec} dq; start_exam!q \rightarrow Skip \rangle; \\ \langle Run_Exam \rangle; \\ \langle \mathsf{dec} ds; end_exam!s \rightarrow Skip \rangle \end{pmatrix}$$

Here, a confidentiality property is applied to the question paper (*q*) and solution sheet (*s*) upon creation. The question paper is declassified immediately before *start_exam*, while the solution is declassified on *end_exam*. Without declassification, the CAs would prevent *start_exam* and *end_exam* from taking place.

Unfortunately, the dec construct is not a reactive design, because it is

not **R3**-healthy:

 $\mathbf{R3}(\operatorname{dec} c) = (\amalg_{rea} \lhd wait \rhd \operatorname{dec} c) \neq \operatorname{dec} c$

Furthermore, the sequential composition of **dec** *c* and a reactive design *A* is not a reactive design, because the constructs:

$$A; \mathbf{dec} c = A \wedge c = wait'$$
$$\mathbf{dec} c; A = A \wedge c = wait$$

are again not **R3**-healthy.

Since **dec** breaks the reactive design model of *Circus*, we lose the assurance that the laws of *Circus* actions continue to hold in the presence of **dec** constructs. In order to deploy **dec**-like constructs in a rigorous manner, it would be necessary to redefine the semantic model of *Circus* and to re-prove existing results. Hence, we must abandon **dec**, because to do otherwise would force us to sacrifice the semantic basis of the language. We leave the task of defining a workable semantics for a declassification construct to future work.

An alternative strategy for accommodating declassification would be to extend the lifted semantics with another alphabet of fog variables, for the sole purpose of tracking Low's inferences about a process. Then, to declassify a preceding CA, the fog variables can be overwritten with these backup variables, to nullify the effect of the CA. However, this approach also has its drawbacks; most notably, the extended alphabet would necessitate redefining the lifting functions, as well as the semantics of lifted operators and CAs. In turn, tracking the extra variables would complicate analysis of specifications at the semantic level.

5.8 Related Work

The lifted semantics for *Circus* presented in this chapter has much in common with Morgan's *shadow semantics* (Morgan, 2009, 2012), which

extends the refinement calculus for sequential programs (Morgan, 1994). The shadow semantics has been applied by McIver (2009) and McIver and Morgan (2010) to derive secure implementations of security-critical algorithms and communication protocols by stepwise refinement.

In the shadow semantics, the program state is divided into a secret variable h and a non-secret variable v. It is assumed that Low can monitor the value of v (but not h) at each point in the program's execution. In addition, a special "shadow set" state variable H is defined to contain all values of h that are consistent with Low's observation of the program's execution. The semantics of each program construct is defined to update v, h and H as appropriate.

In some respects, the lifted semantics of *Circus* actions is a generalisation of the shadow semantics. An analogue of the shadow set is achieved (over all the observational variables) by incorporating the $\mathcal{I}(\mathcal{L})$ relation within the lifting function for *Circus* actions. While Morgan (2009) defines a shadow semantics for sequential programs, our approach is applicable to the wider domain of *Circus* actions. This makes our approach extensible to other languages in the *Circus* family, or indeed to any language with a UTP semantics.

We make an important deviation from the philosophy of the shadow semantics by not differentiating between secret and non-secret variables at the semantic level. On the contrary, confidentiality annotations allow software engineers to specify exactly which properties of the state and by extension, the behaviour — of processes are secret. This is arguably a more flexible (albeit more involved) approach for integrating confidentiality properties within formal methods.

McIver (2009) has formulated a relative of CAs for the shadow semantics. A *visibility annotation* cuts down the shadow associated with *a* to only the actual value of *a*, thereby modelling that value as revealed to Low. Visibility annotations can be formulated in terms of CAs:

reveal
$$a \triangleq \langle a = \widetilde{a} \rangle$$

We hypothesise that specialised annotation constructs may be useful for specifying other kinds of security properties besides confidentiality. Indeed, a family of annotation-like operators to specify a class of availability properties over CSP and timed *Circus* processes has been proposed by Woodcock (2010) and Wei et al. (2010). However, such prospects are beyond the scope of this thesis, and we do not pursue them further.

5.9 Conclusion

This chapter has presented an extension of the *Circus* language which integrates confidentiality properties with *Circus* processes. This extended language inherits the semantic basis of the confidentiality framework of Chapter 4. In particular, we have retained the notion of using miracles to denote inconsistencies between functionality and confidentiality. Hence, processes in the extended language are incapable of engaging in functionality that would (directly or indirectly) reveal information classed as secret to low-level users.

The confidentiality annotation is a versatile and elegant construct for specifying confidentiality properties over both the state and behaviour of *Circus* processes. This construct sets the language apart from the framework of Chapter 4 by the compositional manner in which confidentiality properties are joined with the functional specification facilities of *Circus*. In turn, the specification of confidentiality properties can be localised to security-critical regions of the process body. In this way, the language overcomes two of the problems described in Section 4.8.

While we have taken *Circus* as the formal foundation of our language, we are confident that its underlying principles could be translated to other formalisms. Nevertheless, the generality of *Circus* makes our language (in its current form) suitable for embedding confidentiality within sequential programs and concurrent processes alike.

There remains the problem, described in Section 4.8, of verifying that the functionality of a process in the extended language is consistent with its embedded CAs. We address this final problem in the next chapter.

6 Secure Software Development

6.1 Introduction

Chapter 5 focused on extending *Circus* with a syntax and semantics for specifying confidentiality properties. In this chapter, we turn our attention towards an engineering method to guide and support the development of process designs expressed in this extended language. This method builds upon the *Circus* refinement strategy, but with additions to maintain faithfulness with the lifted semantics.

Recall from Section 4.5 that, in lifted UTP theories, refinement preserves confidentiality properties encoded as obligations, because unfulfilled obligations induce miraculous behaviour. Section 6.2 explores this issue in the context of lifted *Circus* processes. It argues that confidentiality annotations deserve special attention during the development of a process, to avoid jeopardising the implementability of the process.

Before embarking upon the development of a process in our extended language, it is expedient to verify that its functionality is consistent with its embedded CAs. Section 6.3 presents a procedure for calculating Low's inferences about the state of the process at each point of its execution. In Section 6.4, this procedure is distilled into laws defined over lifted *Circus* actions. In Section 6.5, we describe how this procedure can be used to verify that a *Circus* process satisfies its CAs.

In Section 6.6, we tailor the *Circus* refinement strategy to accommodate verified processes in the lifted language. This *secure refinement strategy* enables software engineers to apply the laws of the *Circus* refinement calculus to lifted processes, safe in the knowledge that refinement steps maintain the consistency of the process.

6 Secure Software Development

The final development step is to translate a concrete-level process design into executable code. However, there are important caveats which mean that assurances about the security of the process design may not be applicable to the code. We discuss these issues in Section 6.7.

In this chapter, the development strategy is primarily concerned with safeguarding confidential information from a single adversarial user. The techniques presented here are readily extendible to cover multiple low-level users, who may be encoded within a process specification using blocks of different families (Subsection 5.4.3). Indeed, the case study of Chapter 7 involves a development where the confidentiality properties apply to multiple users' inferences.

6.2 Confidentiality-Preserving Refinement

Recall from Section 5.5 that a CA makes miraculous those behaviours of a process that would disclose secret information to Low. Since a process is compelled not to behave miraculously, the process is forced to satisfy the confidentiality property embodied by the CA.

Suppose a lifted process Q is refined in a way that removes all cover story behaviours prescribed by a CA embedded in Q. The refined process Q' violates the CA and, by definition, the CA becomes miraculous. In turn, by behaving miraculously, the CA disables all behaviours of Q'that would leak secret information. The refinement of Q to Q' is still legitimate, because miracles occupy the top of the lattice of (lifted) reactive designs. However, such refinements *cannot make a process insecure*, because processes cannot behave miraculously. For the reasons described in Section 4.5, it follows that *refinement is confidentiality-preserving* in the lifted semantics.

6.2.1 Refinement and Feasibility

While refinement is confidentiality-preserving, it must be used with care. In the worst case, a refinement step may come at the price of sacrificing the feasibility of the process, making it impossible to implement (and therefore useless). This phenomenon is made clear by Example 6.1.

Example 6.1. Recall from Example 5.48 and Example 5.49 that:

$$\begin{array}{rcl} B0 & \triangleq & \langle h := 0 \rangle \ \text{;} \ \langle h = 0 \Rightarrow \widetilde{h} > 0 \rangle & = & \langle \mathbf{R}(\neg ok) \rangle \\ B1 & \triangleq & \langle h := 1 \rangle \ \text{;} \ \langle h = 0 \Rightarrow \widetilde{h} > 0 \rangle & = & \langle h := 1 \rangle \end{array}$$

Now consider the specification:

B01
$$\triangleq \langle h := 0 \sqcap h := 1 \rangle; \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$$

B01 is feasible in isolation: Low cannot distinguish between h := 0 and h := 1 by observing the behaviour of B01 alone. Both B0 and B1 are refinements of B01, but only B1 is a feasible refinement of B01. \Diamond

Refinement is not guaranteed to maintain the feasibility of processes containing CAs. This is because removing cover stories may induce conflicts between CAs and their surrounding contexts, making those contexts miraculous. On the other hand, refinement may also induce miracles into any process with partial specification constructs such as specification statements and coercions (Nelson, 1989; Morgan, 1994). The next example replicates the phenomenon shown in Example 6.1, but using *Circus* coercions in place of CAs.

Example 6.2 (adapted from Zeyda et al. (2003)). Working in the *Circus* semantics, the action h := 1 can be rewritten using a coercion:

$(h := 0 \sqcap h := 1); [h = 1]$		
=	$(h:=0 ; [h=1])\sqcap (h:=1 ; [h=1])$	[distributivity]
=	${f R}(eg ok) \sqcap h := 1$	[Example 2.4]
=	h := 1	[lattice top]

While refining the $h := 0 \sqcap h := 1$ action to h := 0 is feasible in isolation, the resulting specification — h := 0; [h = 1] — is miraculous, as we saw

in Example 2.4.

To support confidentiality-preserving refinement with a practical development method, we need to employ special measures to verify whether a refinement step, applied to one part of a process, maintains the feasibility of the process as a whole. These measures — which are the main topic of this chapter — depend on analysing the information that Low can infer about the process state at each point during the process's execution.

6.2.2 Atomic and Composite Non-Determinism

Information flow to Low can be modelled at different levels of granularity. We introduce these levels of granularity by drawing an analogy with Morgan's shadow semantics, which was discussed in Section 5.8.

The shadow semantics grants Low the ability to monitor the control flow of a program. Non-deterministic specification constructs in the shadow semantics are either *atomic* or *composite*:

- **Atomic constructs** are considered to execute instantaneously, so Low cannot observe how non-determinism is resolved during their execution. These constructs are needed for specifying security-critical elements of a program, such as conditional branches that depend on the values of secret variables.
- **Composite constructs** enable Low to monitor how they resolve nondeterminism. Hence, refining away non-determinism within a composite construct does not reveal more information about the construct's behaviour to Low.

We can express atomic and composite non-deterministic constructs in the lifted semantics. The atomic construct denoting non-deterministic choice between A_1 and A_2 is represented by:

 $\langle A_1 \sqcap A_2 \rangle$

while the corresponding composite construct is represented by:

$$\langle A_1 \rangle \sqcap \langle A_2 \rangle$$

Morgan (2009) defines *ignorance-preserving refinement* over the shadow semantics, which mandates the shadow set *H* can never be decreased by refining any component of a program. In other words, a refinement is ignorance-preserving if it reveals no more information to Low about the secret variable *h*. Under ignorance-preserving refinement, composite constructs can be refined into atomic constructs, because doing so can only increase Low's ignorance about a program's behaviour. However, Morgan forbids the refinement of atomic constructs in any way.

Owing to ignorance-preserving refinement, a program design in the shadow semantics *implicitly* specifies the maximal information about *h* that is permitted to flow to Low. On the contrary, CAs provide an *explicit* specification of the maximal information flow to Low. Therefore, we do not insist that refinement is ignorance-preserving, because CAs disable insecure functionality of a process. It follows that atomic choice can be refined to composite choice in our semantics:

$$\langle A_1 \sqcap A_2 \rangle \subseteq \langle A_1 \rangle \sqcap \langle A_2 \rangle$$

which is a consequence of Condition 5.12.

By making this refinement, we sever the indistinguishability relation between the observable behaviours of A_1 and the fog behaviours of A_2 (and vice versa). Intuitively, the specification $\langle A_1 \rangle \sqcap \langle A_2 \rangle$ discloses to Low how the process resolves the choice between A_1 and A_2 . (It is as though Low can observe the "program counter" to determine which of A_1 and A_2 is executed.)

A consequence of this refinement is that $\langle A_1 \rangle \sqcap \langle A_2 \rangle$ reveals at least as much information to Low as does $\langle A_1 \sqcap A_2 \rangle$. For this reason, this refinement may violate a CA elsewhere within a process, so it is not guaranteed to preserve the feasibility of the process.

6 Secure Software Development

6.2.3 Equivalence

There are some lifted constructs whose atomic and composite forms are equivalent, as Example 6.3 shows.

Example 6.3. Recall from Example 5.10 that a lifted action can be unfolded according to the $\mathcal{I}(\mathcal{L})$ relation. By extension, given that *L* includes the *on* or *off* channels, we have:

$$\langle L: \textit{on} \to h := 0 \sqcap \textit{off} \to h := 1 \rangle = \left(\begin{array}{c} \langle L: \textit{on} \to h := 0 \rangle \\ \sqcap \langle L: \textit{off} \to h := 1 \rangle \end{array} \right)$$

Since Low can distinguish an *on* event from an *off* event (by observing termination), it follows that Low can always determine how the choice is resolved within this specification. \Diamond

Example 6.3 motivates the following law.

Law 6.4 (Split choice). Provided *L* includes the channels c_1 and c_2 :

$$\langle L : c_1 \to A_1 \sqcap c_2 \to A_2 \rangle = \langle L : c_1 \to A_1 \rangle \sqcap \langle L : c_2 \to A_2 \rangle$$

Where possible, it is advantageous to refine an atomic construct into a composite construct, because we can analyse the components of the composite construct in isolation from each other. Moreover, in cases where atomic and composite constructs are semantically equivalent, substituting one for the other cannot change the feasibility of the process.

6.3 Backwards Propagation

As we observed in Chapter 4, confidentiality properties are not closed under composition. Thus, it can be tricky to determine if a CA is violated within a process body, because the CA imposes a confidentiality constraint that extends over the whole of a process.

This section describes a technique for *compacting* our model of Low's inferences about the behaviour of lifted *Circus* actions into indistinguishab-

ility relations over the process state. These indistinguishability relations have the same form as obligations, so they can be combined with a CA to identify which cover stories prescribed by the CA are not discarded by the action's behaviour. In turn, we can *propagate* this information through a process, to calculate the confidentiality constraints that apply over the process state at each point in the process's execution.

Section 6.5 explains how propagation can help to warn software engineers if the functionality and confidentiality attributes of a process specification are mutually inconsistent.

6.3.1 Weakest Reactive Preconditions

In predicate transformer semantics (Dijkstra, 1976), **wp** (P, Post) denotes the *weakest precondition* which guarantees that program P terminates in a state satisfying *Post*. Following Hoare and He (1998), **wp** can be defined in the UTP theory of relations as follows:

$$\mathsf{wp}(P, Post) \triangleq \forall x' \bullet P \Rightarrow Post$$

The following predicate transformer describes all initial states of a *Circus* action *A* such that *every* normal execution of *A* started in any such state — regardless of whether *A* diverges, deadlocks or terminates — is guaranteed to reach an observable state satisfying the postcondition *Post*:

$$wrp(A, Post) \triangleq wp(A, (ok \land \neg wait \Rightarrow Post))$$
$$= \forall x' \bullet (ok \land \neg wait \land A) \Rightarrow Post$$

wrp (*A*, *Post*) is called the *weakest reactive precondition* (for normal behaviour) by Cavalcanti and Woodcock (2003); we adopt that name here.

6.3.2 Masking States

Propagation involves the translation of Low-indistinguishable process behaviours into a notion of Low-indistinguishable process states. We

6 Secure Software Development

say that a fog state $\tilde{\psi}$ masks a state ψ through A if and only if, for each behaviour of A started in ψ , there is a counterpart Low-indistinguishable fog behaviour of \tilde{A} started in $\tilde{\psi}$.

Definition 6.5 (Masking). $\tilde{\psi}$ masks ψ through *A* (with respect to \mathcal{L}) if:

$$\left[A \land \psi \Rightarrow \exists \, \widetilde{x}, \widetilde{x'} \bullet \mathsf{UC} \, (\mathcal{L}, A) \land \widetilde{\psi} \,\right]$$

Given that $\tilde{\psi}$ masks ψ , if Low's interaction is consistent with *A* starting in ψ , then Low cannot rule out the possibility that *A* started in $\tilde{\psi}$ instead.

Example 6.6. Given that Low can observe on and off events, consider:

$$\langle g \& (on \rightarrow Skip \sqcap off \rightarrow Skip) \Box \neg g \& on \rightarrow Skip \rangle$$

Here, each state satisfying the condition *g* masks each state satisfying \neg *g*, because *on* may be performed in either case. Notice that symmetry does not apply here, because the occurrence of *off* would reveal to Low that *g* holds. \Diamond

We can adapt **wrp** to encode Low's inferences about the behaviour of a *Circus* action in terms of its initial state. The predicate:

$$\mathsf{wrp}\left(A, \exists \, \widetilde{x'} \bullet \mathsf{UC}\left(\mathcal{L}, A
ight)
ight)$$

describes the *weakest masking relation* over the initial states of action *A*: it maps ψ to $\tilde{\psi}$ if and only if $\tilde{\psi}$ masks ψ .

We now define a specialised form of the weakest masking relation over lifted constructs. In order to simplify the presentation of the results to follow, it hides the ok, ok, wait, wait variables.

Definition 6.7 (wmr predicate transformer).

wmr (B)
$$\triangleq (\forall x' \bullet \mathbf{D} (B!) \Rightarrow \exists \widetilde{x'} \bullet B!)$$

where *B*! denotes B[**true**, **true**, **false**, **false**/*ok*, \widetilde{ok} , *wait*, \widetilde{wait}].

wmr (B) encodes the maximum information that Low can deduce about the initial state of B, regardless of the interaction that Low makes with B. Hence, **wmr** (B) may potentially over-approximate Low's inferences: not every behaviour of B may reveal all the information to Low modelled by **wmr** (B). Example 6.8 clarifies the nature of this over-approximation.

Example 6.8. Consider the block $ST = \langle s!a \rightarrow Skip \sqcap t!b \rightarrow Skip \rangle$. Given that Low's window includes the *s* and *t* channels, we have:

wmr (ST) =
$$a = \widetilde{a} \wedge b = \widetilde{b}$$

Intuitively, Low may learn by observing ST either the value of a (on channel s) or the value of b (on channel t), depending on how the non-determinism in ST is resolved. However, before ST starts, there is no way of knowing how that non-determinism is resolved. Hence, **wmr** over-approximates Low's inferences about the initial state of ST.

Now consider the block $TT = \langle t!a \rightarrow Skip \sqcap t!b \rightarrow Skip \rangle$. We have:

wmr
$$(TT) = a = \widetilde{a} \lor b = \widetilde{b}$$

By observing a value transmitted on channel t, Low learns only that either a or b have that value, in the absence of other information. \Diamond

Since **wmr** ($\langle A \rangle$) expresses an upper bound on Low's inferences about *A*'s initial state, it is not monotonic with respect to the \sqsubseteq ordering. This non-monotonicity of **wmr** is demonstrated by the next example.

Example 6.9. The *ST* block (Example 6.8) is refined by $\langle s!a \rightarrow Skip \rangle$. This block reveals the value of *a* to Low, but tells Low nothing about *b*:

wmr
$$(\langle s!a \rightarrow Skip \rangle) = a = \tilde{a}$$

We have $ST \sqsubseteq \langle s!a \rightarrow Skip \rangle$, but **wmr** $(ST) \not\sqsubseteq$ **wmr** $(\langle s!a \rightarrow Skip \rangle)$, so **wmr** is not monotonic.

6.3.3 Propagating Obligations

We now turn our attention to propagating obligations backwards through blocks. Consider a block placed in sequence with a CA:

B;
$$\langle \theta \rangle$$

Should *B* terminate, then $\langle \theta \rangle$ takes effect and constrains the fog space associated with the final state of *B* according to the obligation θ .

We extend the **wmr** predicate transformer to propagate θ backwards through *B*. Let **bwQ**(*B*, θ) denote the masking relation over the initial states of *B*, but strengthened with an obligation parameter.

Definition 6.10 (bwQ predicate transformer).

$$\mathbf{bwQ}(B,\theta) \triangleq \forall x' \bullet \left(\mathbf{D}(B!) \Rightarrow \exists \widetilde{x'} \bullet B! \land \operatorname{Conf}'(\theta) \right)$$

If a *Circus* action *A* starts in a state ψ and terminates in a state ψ' satisfying the postcondition **D** (θ'), then **bwQ** ($\langle A \rangle$, θ) relates ψ to only those initial fog states that mask ψ through $\langle A \rangle$ and lead to Low-indistinguishable behaviours terminating in fog states marked as cover stories for ψ' by θ' . In fact, **bwQ** is more general: it can be applied to arbitrary lifted constructs to identify an indistinguishability relation between initial states.

The masking relation **bwQ**(B, θ) is itself an obligation, which fuses Low's inferences about the behaviour of B with the confidentiality property imposed upon those inferences by $\langle \theta \rangle$ whenever B terminates.

Example 6.11. Suppose *ST* is named in sequence with a CA:

$$\langle s!a \rightarrow Skip \sqcap t!b \rightarrow Skip \rangle$$
; $\langle a = 0 \Rightarrow \tilde{a} > 0 \rangle$

Intuitively, the CA prevents the *s*!*a* event from taking place if a = 0. This intuition is justified by calculation:

bwQ ((
$$s!a \rightarrow Skip \sqcap t!b \rightarrow Skip$$
), ($a = 0 \Rightarrow \tilde{a} > 0$))
= $a = \tilde{a} \land b = \tilde{b} \land (a = 0 \Rightarrow \tilde{a} > 0)$ [by Example 6.8]

$$= a = \tilde{a} \wedge b = \tilde{b} \wedge a \neq 0$$
 [prop calc]

No cover stories are prescribed by this back-propagated obligation when a = 0.

Theorem 6.12 shows the result of applying **bwQ** to a lifted reactive design.

Theorem 6.12 (Unfolding bwQ).

$$\begin{aligned} \mathsf{bwQ}\left(\mathsf{UC}\left(\mathcal{L},\mathsf{R}\left(\operatorname{Pre}\vdash\operatorname{Post}\right)\right),\theta\right) &= \\ &\forall x' \bullet \left(\begin{array}{c} \operatorname{Pre!}\wedge ok' \wedge \operatorname{Post!}\wedge tr \leq tr' \\ &\Rightarrow \exists \, \widetilde{x'} \bullet (\widetilde{\operatorname{Pre!}} \Rightarrow \widetilde{\operatorname{Post!}}) \wedge \widetilde{tr} \leq \widetilde{tr'} \wedge \mathcal{I}(\mathcal{L})! \wedge \operatorname{Conf'}(\theta) \end{array}\right) \\ &\wedge \quad \widetilde{\operatorname{Pre!}} \Rightarrow \operatorname{Pre} \end{aligned}$$

If *B* always exhibits miraculous behaviour in some initial states, then Low can rule out the possibility that *B* starts in those states. The initial states where execution of *B* is feasible are given by **bwR** (*B*).

Definition 6.13 (bwR predicate transformer).

bwR (B)
$$\triangleq \exists x', \widetilde{x'} \bullet B!$$

Whenever *B* is a non-miraculous construct, we have **bwR** (*B*) = **true**. Theorem 6.14 shows the result of applying **bwR** to a lifted reactive design.

Theorem 6.14 (Unfolding bwR).

bwR (UC (\mathcal{L} , R ($Pre \vdash Post$))) = $\exists x', \tilde{x'} \bullet U$ ($tr \leq tr'$) \land (($Pre! \lor \widetilde{Pre!}$) $\Rightarrow ok' \land U$ ($Pre! \Rightarrow Post!$)) $\land \mathcal{I}(\mathcal{L})!$

We define a predicate transformer **bw**, which conjoins **bwQ** with **bwR** to exclude initial states which lead to miraculous behaviour from back-propagated obligations.

Definition 6.15 (bw predicate transformer).

$$\mathbf{bw}\left(B,\theta\right) \triangleq \mathbf{bwR}\left(B\right) \wedge \mathbf{bwQ}\left(B,\theta\right)$$

Informally, the obligation **bw** (B, θ) encodes the maximum information that Low can infer about the initial state of *B* via its interaction with *B*. Moreover, if *B* terminates, then **bw** (B, θ) incorporates Low's inferences about the final state reached by *B* derived from the subsequent behaviour of the process, as encoded by θ .

bw is monotonic in its second argument with respect to the \sqsubseteq ordering.

Lemma 6.16 (Monotonicity of **bw**). $\theta_1 \subseteq \theta_2$ implies **bw** $(B, \theta_1) \subseteq$ **bw** (B, θ_2)

However, **bw** is not necessarily monotonic in its first argument with respect to \sqsubseteq , because **wmr** is not monotonic with respect to \sqsubseteq (Example 6.9).

6.3.4 Forwards Propagation

An alternative to backwards propagation would be to propagate CAs and Low's inferences *forwards* through blocks. Forwards propagation is analogous to calculating the *strongest postcondition* that a program satisfies if started in any initial state satisfying a specified precondition.

Example 6.17. The counterpart to wmr for forwards propagation is:

$$\forall x \bullet \mathbf{D}(B!) \land \neg wait' \Rightarrow \exists \widetilde{x} \bullet B!$$

Applying this predicate transformer to ST (Example 6.8) yields:

$$\begin{array}{c} ok' \wedge \neg \ wait' \Rightarrow \\ \left(\begin{array}{c} \widetilde{ok'} \wedge \neg \ \widetilde{wait'} \\ \wedge \ \left(\begin{array}{c} \mathsf{last} \left(tr' \right) = (s,a) \wedge \mathsf{last}(\widetilde{tr'}) = (s,\widetilde{a}) \wedge a = \widetilde{a} \\ \vee \ \mathsf{last} \left(tr' \right) = (t,b) \wedge \mathsf{last}(\widetilde{tr'}) = (t,\widetilde{b}) \wedge b = \widetilde{b} \end{array} \right) \end{array} \right)$$

Unlike wmr(ST), this relation distinguishes the cases where Low learns

a or *b*, since the tr' variable retains information about how the non-determinism is resolved. \diamond

As Example 6.17 shows, forwards propagation retains more detailed information about Low's inferences than does backwards propagation. However, forwards propagations are far more laborious to calculate. This result correlates with Dijkstra and Scholten's observation that weakest preconditions are simpler to calculate than strongest postconditions (Dijkstra and Scholten, 1990, §12).

It is pragmatic to adopt backwards propagation as the basis of the verification method presented in subsequent sections. However, there is no fundamental reason why this verification method could not instead be defined in the style of forwards propagation.

6.4 Laws of Backwards Propagation

We can apply **bw** to any lifted construct, be it a lifted *Circus* action, a CA, or a composite construction of lifted actions, lifted operators and CAs.

As should be expected, back-propagating an obligation θ_1 through a CA embedding θ_2 yields the least upper bound of θ_1 and θ_2 .

Law 6.18 (bw CA). bw $(\langle \theta_1 \rangle, \theta_2) = \theta_1 \wedge \theta_2$

Calculating the backwards propagation of obligations through lifted actions can be difficult, especially with actions with a complex semantics. In this section, we propose tactics for dividing this calculation into manageable parts. We also identify laws for back-propagating obligations through the basic *Circus* constructs.

6.4.1 Heuristics

Calculating **bw** (B, θ) from first principles can be tedious. Lemma 6.19 shows this calculation can be simplified by breaking B into parts and applying **bw** to each part.

Lemma 6.19 (Over-approximating bw).

$$\mathbf{bw}\,(B,\theta) \quad \sqsubseteq \quad \mathbf{bw}\,(B \land C,\theta) \land \mathbf{bw}\,(B \land \neg C,\theta)$$

Lemma 6.19 may over-approximate **bw** (B, θ) , but it is often desirable to calculate **bw** (B, θ) exactly. An exact calculation can be obtained by partitioning *B* according to Low's abilities to distinguish between the behaviours of *B*. Theorem 6.20 specifies the scheme of these calculations.

Theorem 6.20 (Splitting bw).

bw (B, θ) = **bw** $(B \land C, \theta) \land$ **bw** $(B \land \neg C, \theta)$

provided the following conditions hold:

(SC1)
(SC2)
(SC3)
(SC4)

Side-conditions 1 and 2 of Theorem 6.20 require that at least one behaviour of *B* satisfies *C* and at least one behaviour of *B* does not satisfy *C*. Side-condition 3 stipulates that *B* is constructed from **UC**-lifted actions; this condition is trivially satisfied if *B* is a block or composed from blocks using the lifted *Circus* operators. Side-condition 4 demands *C* holds for each behaviour ϕ if and only if *C* holds for all behaviours that are Low-indistinguishable from ϕ .

We specialise Theorem 6.20 in three ways by instantiating *C*.

Traces If the projections of the traces of two behaviours through Low's window \mathcal{L} are different, then Low can distinguish between those behaviours. This insight motivates the following lemma.

Lemma 6.21 (bw traces). Let $W \subseteq \mathcal{L}$ and $HT = (tr' - tr) \upharpoonright W \in S$, where *S* denotes a set of traces drawn from events in W. Then, provided

side-conditions 1, 2 and 3 of Theorem 6.20 are upheld:

$$\mathbf{bw}(B,\theta) = \mathbf{bw}(B \wedge HT,\theta) \wedge \mathbf{bw}(B \wedge \neg HT,\theta)$$

Refusals If the process deadlocks, it *may* refuse any of the events in Low's window. If Low attempts to engage in those events and the process refuses them, then Low will perceive that refusal. By observing these refusals, Low may be able to infer some information about the state of the process at deadlock and, consequently, its state prior to deadlocking.

Lemma 6.22 (bw refusals). Let $\mathcal{W} \subseteq \mathcal{L}$ and $HR = \mathcal{W} \subseteq ref'$. Then, provided side-conditions 1,2 and 3 of Theorem 6.20 are upheld:

$$\mathbf{bw}\,(B,\theta) \quad = \quad \mathbf{bw}\,(B \wedge HR,\theta) \wedge \mathbf{bw}\,(B \wedge \neg HR),\theta)$$

Moreover, if *B* is a lifted *Circus* action, we also have (by **CSP4**-healthiness):

$$\mathbf{bw}(B,\theta) = \mathbf{bw}(B \land (wait' \Rightarrow HR), \theta) \land \mathbf{bw}(B \land (wait' \Rightarrow \neg HR), \theta)$$

Termination Since $\mathcal{I}(\mathcal{L})$ models Low as being able to distinguish terminating from non-terminating behaviours, we also have Lemma 6.23. Again, Low may exploit knowledge of (non-)termination to infer information about the initial state of a block.

Lemma 6.23 (bw termination). Provided side-conditions 1,2 and 3 of Theorem 6.20 are upheld:

$$\mathsf{bw}\,(B,\theta) \quad = \quad \mathsf{bw}\,\big(B \land \mathit{wait}',\theta\big) \land \mathsf{bw}\,\big(B \land \neg \mathit{wait}',\theta\big)$$

Naturally, these lemmas can be used together, to partition lifted constructs more finely than on traces, refusals or termination alone.

6.4.2 Propagating through Circus Constructs

bw inherits many of the properties of the **wp** predicate transformer. Following the style of the weakest precondition calculus, we now define laws for back-propagating obligations through lifted *Circus* constructs.

Since these laws are defined over individual lifted actions, they cannot be applied to more complex blocks. Nevertheless, by Condition 5.12, such a block can be broken apart (by refinement) into a composite construct, featuring simpler blocks that are amenable to these laws.

Primitives As should be expected, the lifted *Skip* is an identity for backwards propagation.

Law 6.24 (bw *Skip*). bw $(\langle Skip \rangle, \theta) = \theta$

Likewise, the lifted *Stop* and the lifted *Chaos* are zeros for backwards propagation.

Law 6.25 (bw *Stop*). bw $(\langle Stop \rangle, \theta) =$ true

Law 6.26 (bw *Chaos*). bw ($\langle Chaos \rangle, \theta$) = true

Scoping Backwards propagation distributes through a declaration of variable *a* simply by hiding *a* and \tilde{a} .

Law 6.27 (bw scope).

bw (var $a: T \bullet B, \theta$) = $\forall a: T \bullet \exists \tilde{a}: T \bullet bw (B, \forall a: T \bullet \exists \tilde{a}: T \bullet \theta)$

Commands Law 6.28 provides a general account of back-propagating obligations through specification statements.

Law 6.28 (bw specification statement).

bwR $(\langle w : [Pre, Post] \rangle) = \mathbf{U} (Pre \Rightarrow \exists v' \bullet Post \land u' = u)$

138

6.4 Laws of Backwards Propagation

$$\mathbf{bwQ}\left(\langle w : [Pre, Post] \rangle, \theta\right) = \begin{pmatrix} \forall v' \bullet (Pre \land Post \land u' = u) \Rightarrow \\ \exists \widetilde{v'} \bullet (\widetilde{Pre} \Rightarrow \widetilde{Post} \land \widetilde{u'} = \widetilde{u}) \land \theta' \\ \land \widetilde{Pre} \Rightarrow Pre \end{pmatrix}$$

where u denotes all state variables not listed in w.

Law 6.29 follows by specialising Law 6.28.

Law 6.29 (bw assumption).

 $\mathsf{bw}\left(\left\langle \left\{ \,C \,\right\} \,\right\rangle, \theta\right) \quad = \quad (C \Rightarrow \exists \, \widetilde{v'} \, \bullet \, (\widetilde{C} \Rightarrow \widetilde{v'} = \widetilde{v}) \, \land \, \theta'[v/v']) \, \land \, (\widetilde{C} \Rightarrow C)$

As CAs are a generalisation of coercions, Law 6.30 is a direct consequence of Law 5.44 and Law 6.18.

Law 6.30 (bw coercion). bw $(\langle [C] \rangle, \theta) = \theta \land U(C)$

The law for assignment is typical for a weakest precondition calculus.

Law 6.31 (bw assignment). bw $(\langle a := E \rangle, \theta) = \theta[E, \widetilde{E}/a, \widetilde{a}]$

Sequence As lifted sequential composition is just relational composition, back-propagating an obligation through a sequence of blocks can be performed block by block.

Law 6.32 (bw sequence).

$$\mathbf{bwR}\left(B_{1}\,\widehat{\mathbf{j}}\,B_{2}\right) = \mathbf{bwR}\left(B_{1}\wedge \operatorname{Conf}'\left(\mathbf{bwR}\left(B_{2}\right)\left[x',\widetilde{x'}/x,\widetilde{x}\right]\right)\right)$$
$$\mathbf{bwQ}\left(B_{1}\,\widehat{\mathbf{j}}\,B_{2},\theta\right) = \mathbf{bwQ}\left(B_{1},\mathbf{bwQ}\left(B_{2},\theta\right)\right)$$

Guard Recall from Definition 5.25 that the lifted guard g & B reveals the truth of the condition g to Low. If g is **false**, this construct deadlocks and reveals no extra information to Low. If g is **true**, this construct behaves as B and reveals whatever information to Low that B reveals.

Law 6.33 (bw guard). **bw** $(g \& B, \theta) = (\mathbf{bw} (B, \theta) \land \mathbf{U} (g)) \lor \mathbf{U} (\neg g)$

Prefixing A prefixing $c!E \rightarrow Skip$ which outputs the value of an expression *E* to the environment on channel *c* may reveal to Low information about the process state in terms of *E*, provided Low can observe *c*.

Law 6.34 (bw output prefix).

$$\mathsf{bw}\left(\left\langle L : c : E \to Skip \right\rangle, \theta\right) = \begin{cases} \theta \land E = \widetilde{E} & \text{if } c \in L \\ \theta & \text{otherwise} \end{cases}$$

A prefixing $c.e \rightarrow Skip$ which communicates a constant e to the environment does not change the state or reveal state information to Low. It is therefore an identity for backwards propagation, as Law 6.35 shows. This law can be derived from Law 6.34 by taking E = e.

```
Law 6.35 (bw prefix). bw (\langle c.e \rightarrow Skip \rangle, \theta) = \theta
```

A prefixing which accepts an input value *e*? from the environment on channel *c* reveals the exact value of *e*? to Low, provided Low can observe *c*. Conversely, if Low cannot observe *c*, Low can still infer that *e*? has the type $\delta(c)$ and that P(e?) holds.

Law 6.36 (bw input prefix).

$$\mathbf{bw} \left(\left\langle L : c?e : P \to e := e? \right\rangle, \theta \right) \\ = \begin{cases} \forall e : \delta(c) \bullet P(e) \Rightarrow \theta[e/\tilde{e}] & \text{if } c \in L \\ \forall e : \delta(c) \bullet P(e) \Rightarrow \exists \tilde{e} : \delta(c) \bullet P(\tilde{e}) \land \theta & \text{otherwise} \end{cases}$$

Choice To back-propagate θ through a composite internal choice, it is sufficient to back-propagate θ through each construct of the choice separately and take the least upper bound.

Law 6.37 (bw internal choice).

$$\mathsf{bw}\left(B_{1} \sqcap B_{2}, \theta\right) \quad \sqsubseteq \quad \mathsf{bw}\left(B_{1}, \theta\right) \land \mathsf{bw}\left(B_{2}, \theta\right)$$

It is harder to derive a general law for back-propagating θ through a composite external choice. Should B_1 and B_2 both deadlock without performing any event (or diverging), then $B_1 \square B_2$ behaves as $B_{1ff}^{tt} \land B_{2ff}^{tt}$. This term may be miraculous if B_{1ff}^{tt} and B_{2ff}^{tt} restrict the state space in different ways — for instance, if they contain CAs — even if B_{1ff}^{tt} and B_{2ff}^{tt} are individually feasible. It follows that:

bw
$$(B_1 \Box B_2, \theta) \sqsubseteq$$
bw $(B_1, \theta) \land$ **bw** (B_2, θ)

is *not* guaranteed in all cases. We side-step this problem by defining a backwards propagation law for a more specialised form of external choice. In this form, the choice is resolved by the environment before any state change can take place.

Law 6.38 (**bw** guarded prefixed external choice). Provided at least one of the guards g_1, g_2, \ldots, g_n is satisfied by each state:

$$\mathsf{bw} (\Box i \bullet g_i \& c_i \to B_i, \theta) \sqsubseteq \bigwedge i \bullet (\mathsf{U}(g_i) \land \mathsf{bw}(B_i, \theta)) \lor \mathsf{U}(\neg g_i)$$

Law 6.39 follows from Law 6.38 by taking true for each guard.

Law 6.39 (bw prefixed external choice).

$$\mathsf{bw} \ (\Box \ i \bullet c_i \to B_i, \theta) \quad \sqsubseteq \quad \bigwedge i \bullet \mathsf{bw} \ (B_i, \theta)$$

6.5 Verifying Consistency

Owing to the non-compositional nature of confidentiality properties, a CA in a process specification may make other components of the process unimplementable, if the functionality of those components is inconsistent with the CA. Hence, it is expedient to ascertain that a process is feasible before proceeding with its development, to avoid unpleasant surprises if the process is indeed unimplementable.

This section describes a procedure for dividing the task of verifying that a process satisfies a CA into multiple cases, and then to conquer

each case individually. This procedure is based on propagating the CA through the blocks of the process.

Propagation makes infeasible blocks readily identifiable, so this procedure enables the implementability (and therefore the security) of processes to be verified in a piecewise manner. Should the process be found to be unimplementable, then the requirements will need to be revised to resolve the conflict between functionality and confidentiality.

6.5.1 Propagated Confidentiality Annotations

Recall from Subsection 6.3.3 that backwards propagation can be applied to specification constructs of the form $\langle A \rangle$; $\langle \theta \rangle$ to translate θ into an obligation over the initial state of $\langle A \rangle$. This back-propagated obligation can itself be embedded within a CA preceding $\langle A \rangle$:

$$\langle \mathsf{bw}(\langle A \rangle, \theta) \rangle; \langle A \rangle; \langle \theta \rangle$$

Inserting this CA into the specification refines the specification:

$$\begin{array}{l} \langle A \rangle ; \langle \theta \rangle \\ = & \langle \mathbf{true} \rangle ; \langle A \rangle ; \langle \theta \rangle & [\text{property of CA}] \\ \sqsubseteq & \langle \mathbf{bw} (\langle A \rangle, \theta) \rangle ; \langle A \rangle ; \langle \theta \rangle & [\text{Theorem 5.40}] \end{array}$$

Example 6.40 shows this refinement is non-trivial if $\langle A \rangle$ may deadlock or diverge.

Example 6.40. By Lemma 6.23, back-propagating θ through $\langle Skip \sqcap Stop \rangle$ just yields θ . Therefore:

$$\langle Skip \sqcap Stop \rangle; \langle \theta \rangle \sqsubseteq \langle \theta \rangle; \langle Skip \sqcap Stop \rangle; \langle \theta \rangle$$

Provided θ is not vacuously **true**, the specification is strictly refined because $\langle \theta \rangle$ is enforced even if the non-deterministic choice resolves to $\langle Stop \rangle$.

Suppose that, when $\langle A \rangle$ is started in a lifted state ψ , it terminates in a final state that violates θ . Then $\langle \mathbf{bw} (\langle A \rangle, \theta) \rangle$ is miraculous in ψ .

Example 6.41. Recall the specification from Example 6.11:

$$\langle s!a \rightarrow Skip \sqcap t!b \rightarrow Skip \rangle ; \langle a = 0 \Rightarrow \widetilde{a} > 0 \rangle$$

We insert a CA encapsulating the back-propagation of $a = 0 \Rightarrow \tilde{a} > 0$ through the block (see Example 6.11) into the specification, to obtain:

$$\langle a = \widetilde{a} \land b = b \land a \neq 0 \rangle$$
; $\langle s!a \rightarrow Skip \sqcap t!b \rightarrow Skip \rangle$; $\langle a = 0 \Rightarrow \widetilde{a} > 0 \rangle$

Now, the feasibility of the specification is made explicit: if it starts in a state where a = 0, then miraculous behaviour is assured.

6.5.2 Left Justification

We now describe a procedure for transforming a lifted process design into a form where inconsistencies between the functionality and confidentiality attributes of a process can be detected by studying the CAs within the process alone. The procedure is based on inserting propagated CAs between every lifted specification construct within the process body. These CAs are inserted to record Low's inferences about the process state at each point in the process's execution, based on Low's observations of the subsequent behaviour of the process.

To formalise this procedure, we define a special format for lifted constructs, where each construct is surrounded by a pair of CAs.

Definition 6.42 (Left justification). A lifted construct is *left-justified* if and only if it has the form:

$$\langle \, \theta^{\alpha} \, \rangle$$
 ; B ; $\langle \, \theta^{\omega} \, \rangle$

such that **bw** $(B, \theta^{\omega}) \sqsubseteq \theta^{\alpha}$.

Of course, we could left-justify any construct by setting $\theta^{\alpha} =$ **false**. That would not be helpful for our broader goal of verifying the process

is consistent with its CAs. Rather, we are interested in identifying the *weakest* obligation θ^{α} that left-justifies the construct.

Every CA is left-justified in its own right:

$$\langle \theta \rangle = \langle \text{true} \rangle; \langle \theta \rangle; \langle \text{true} \rangle = \langle \theta \rangle; \langle \text{true} \rangle; \langle \text{true} \rangle$$

It is straightforward to left-justify primitive constructs — such as specification statements and guards — with the laws of **bw** (Section 6.4).

To back-propagate a CA through a composite lifted construct, we leftjustify each of its components individually, in order to derive a CA to left-justify the composite construct as a whole. We now consider the details of this procedure for specific composite constructs.

Sequence We propagate a CA through a sequence of blocks in a pattern reminiscent of the right fold operator in functional programming (Hutton, 1999). For instance, given:

$$\langle A_1 \rangle$$
;...; $\langle A_n \rangle$; $\langle \theta_n \rangle$

we left-justify each lifted action as follows:

$$\langle \theta_0 \rangle$$
; $\langle A_1 \rangle$; $\langle \theta_1 \rangle$; ...; $\langle \theta_{n-1} \rangle$; $\langle A_n \rangle$; $\langle \theta_n \rangle$

where, for each $0 \le i < n$, **bw** $(\langle A_{i+1} \rangle, \theta_{i+1}) \sqsubseteq \theta_i$ holds.

Remark 6.43. It is most efficient to left-justify sequences of blocks from right to left. Consider:

$$\langle A_1 \rangle$$
; $\langle \theta_1 \rangle$; $\langle A_2 \rangle$; $\langle \theta_2 \rangle$

If **bw** $(\langle A_2 \rangle, \theta_2) \not\sqsubseteq \theta_1$, then it is efficient to left-justify $\langle A_2 \rangle$ first, so that $\langle A_1 \rangle$ can be left-justified with respect to $\theta_1 \wedge \mathbf{bw} (\langle A_2 \rangle, \theta_2)$.

Choice When propagation reaches a composite choice construct, we distribute the CA through the choice as follows:

$$\begin{array}{lll} (B_1 \sqcap B_2); \langle \theta \rangle &=& B_1; \langle \theta \rangle \sqcap B_2; \langle \theta \rangle \\ (B_1 \sqcap B_2); \langle \theta \rangle &=& B_1; \langle \theta \rangle \sqcap B_2; \langle \theta \rangle \end{array}$$

We left-justify B_1 and B_2 (with respect to θ) separately. Following Law 6.37 and Law 6.38, we then left-justify the choice construct as a whole with a CA containing the least upper bound of the obligations from each branch:

$$\begin{array}{ccc} (B_1 \sqcap B_2); \langle \theta \rangle & \sqsubseteq & \langle \mathsf{bw} (B_1, \theta) \land \mathsf{bw} (B_2, \theta) \rangle; (B_1 \sqcap B_2); \langle \theta \rangle \\ (B_1 \sqcap B_2); \langle \theta \rangle & \sqsubseteq & \langle \mathsf{bw} (B_1, \theta) \land \mathsf{bw} (B_2, \theta) \rangle; (B_1 \sqcap B_2); \langle \theta \rangle \end{array}$$

Parallel Parallel-by-merge is defined so that, upon termination, the final values of the state variables are selected from the final states reached by the concurrent actions (Hoare and He, 1998; Oliveira et al., 2009). Therefore, a CA which succeeds a parallel construct can be split into parts, depending on the merge specified by the parallel operator.

Let ns_1 and ns_2 denote disjoint subsets (but not necessarily a partition) of the state variables of a process. Then, let θ_0 , θ_1 and θ_2 denote obligations that satisfy the conditions:

$$\begin{aligned} \alpha \theta_0 \cap (ns_1 \cup \widetilde{ns_1} \cup ns_2 \cup \widetilde{ns_2}) &= \emptyset \\ \alpha \theta_1 \subseteq ns_1 \cup \widetilde{ns_1} \\ \alpha \theta_2 \subseteq ns_2 \cup \widetilde{ns_2} \end{aligned}$$

These conditions demand the alphabets of θ_0 , θ_1 and θ_2 are disjoint. From this point onwards, we consider parallel constructs of the form:

$$(B_1 | [ns_1 | cs | ns_2] | B_2); \langle \theta_0 \land \theta_1 \land \theta_2 \rangle$$

Remark 6.44. Some obligations cannot be structured in this form. An example obligation would be $\tilde{g} \neq \tilde{h}$, where $ns_1 = \{g\}$ and $ns_2 = \{h\}$.

However, this obligation can be refined to $\tilde{g} \ge 0 \land \tilde{h} < 0$ to achieve an obligation structured in our required form.

Since a parallel composition terminates only if both its components terminate, we can distribute θ_1 and θ_2 according to the rule:

$$(B_{1} | [ns_{1} | cs | ns_{2}] | B_{2}); \langle \theta_{0} \land \theta_{1} \land \theta_{2} \rangle \\ \sqsubseteq \\ (B_{1}; \langle \theta_{1} \rangle | [ns_{1} | cs | ns_{2}] | B_{2}; \langle \theta_{2} \rangle); \langle \theta_{0} \land \theta_{1} \land \theta_{2} \rangle$$

Since B_1 and B_2 must synchronise on events in *cs*, each of these components may constrain the other's space of behaviours. In some cases, the synchronisation may prevent B_1 or B_2 from engaging in cover story behaviours mandated by their respective CAs.

Example 6.45. Consider the specification, where $c_1, c_2 \notin L$:

$$\left(\begin{array}{c} \langle c_1 \to h := 0 \Box c_2 \to h := 1 \rangle \\ \| [\{h\} \mid \{ \mid c_1, c_2 \mid \} \mid \emptyset] \| \\ \langle c_1 \to Skip \rangle \end{array}\right); \langle h = 0 \Rightarrow \widetilde{h} > 0 \rangle$$

Since the parallel branches can only synchronise on c_1 , this specification can be rewritten as:

$$\langle c_1 \to h := 0 \rangle ; \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$$
 [semantics of parallel]
= $\langle c_1 \to Skip \rangle ; \langle h := 0 \rangle ; \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$ [Lemma 5.18]
= $\langle c_1 \to Skip \rangle ; \langle h := 0 \rangle ; \langle false \rangle$ [by Example 5.49]

We would fail to detect this insecurity in the original specification if we were to analyse the two parallel branches without accounting for the synchronisation between them. \diamond

We can sidestep the problem highlighted by Example 6.45 by enlarging Low's window on B_1 and B_2 to include all channels listed by *cs*. In effect, we assume *pessimistically* that Low can observe all synchronisation events

performed by B_1 and B_2 individually, even if they are hidden from the environment.

Example 6.46. Continuing from Example 6.45, interpreting the left branch of the parallel construct under the assumption that Low *can* observe synchronisation events over c_1 and c_2 gives:

$$\langle c_{1} \rightarrow h := 0 \Box c_{2} \rightarrow h := 1 \rangle ; \langle h = 0 \Rightarrow h > 0 \rangle$$

$$\Box \quad (\langle c_{1} \rightarrow h := 0 \rangle \Box \langle c_{2} \rightarrow h := 1 \rangle); \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle$$

$$[Condition 5.12]$$

$$= \quad \begin{pmatrix} \langle c_{1} \rightarrow h := 0 \rangle; \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle \\ \Box \quad \langle c_{2} \rightarrow h := 1 \rangle; \langle h = 0 \Rightarrow \tilde{h} > 0 \rangle \end{pmatrix}$$

$$= \quad \langle c_{1} \rightarrow h := 0 \rangle; \langle \mathbf{false} \rangle \Box \langle c_{2} \rightarrow h := 1 \rangle; \langle \mathbf{true} \rangle$$

$$[as Example 5.48, Example 5.49]$$

Now, the inconsistency hidden in the original specification is exposed. \Diamond

Remark 6.47. Assuming that Low can observe synchronisations between parallel processes is incompatible with the superposition technique described in Subsection 5.6.3. As superposition is used to impose confidentiality properties over the synchronisation events, it would be self-defeating to assume Low can observe those events. This problem could be solved by reforming a superposed process into a single process (with CAs embedded alongside blocks) before applying back-propagation.

We then left-justify B_1 ; $\langle \theta_1 \rangle$ and B_2 ; $\langle \theta_2 \rangle$ separately:

$$B_{1}; \langle \theta_{1} \rangle \sqsubseteq \langle \mathsf{bw} (B_{1}^{+}, \theta_{1}) \rangle; B_{1}; \langle \theta_{1} \rangle$$
$$B_{2}; \langle \theta_{2} \rangle \sqsubseteq \langle \mathsf{bw} (B_{2}^{+}, \theta_{2}) \rangle; B_{2}; \langle \theta_{2} \rangle$$

where B_1^+ and B_2^+ denote B_1 and B_2 respectively, but with the events in *cs* modelled as visible to Low.

Finally, the parallel composition construct can be left-justified by taking

the least upper bound of the left-justified CAs:

$$\langle \theta_0 \wedge \mathbf{bw} (B_1^+, \theta_1) \wedge \mathbf{bw} (B_2^+, \theta_2) \rangle; (\langle \mathbf{bw} (B_1^+, \theta_1) \rangle; B_1; \langle \theta_1 \rangle |[ns_1 | cs | ns_2]| \langle \mathbf{bw} (B_2^+, \theta_2) \rangle; B_2; \langle \theta_2 \rangle); \langle \theta_0 \wedge \theta_1 \wedge \theta_2 \rangle$$

Since θ_0 references state variables that are left unchanged by the parallel construct, it is important that θ_0 is included within the left-justifying CA.

Recursion In *Circus*, a recursive construct can take the form $\mu X \bullet F(X)$, or other forms that can be transformed to this form (Oliveira et al., 2009). In order to left-justify the lifted loop body $\langle F(Skip) \rangle$ with respect to θ , our strategy is to identify an *invariant obligation* θ_I (where $\theta \sqsubseteq \theta_I$) that is maintained by $\langle F(Skip) \rangle$.

Definition 6.48 (Invariant obligation). The obligation θ_I is an invariant obligation for $\langle F(Skip) \rangle$ if back-propagating θ_I through $\langle F(Skip) \rangle$ yields an obligation no stronger than θ_I ; that is:

bw
$$(\langle F(Skip) \rangle, \theta_I) \subseteq \theta_I$$

We can take the obligation θ_0 from the CA following the recursive construct and calculate:

$$\theta_{i+1} = \theta_i \wedge \mathsf{bw}\left(\langle F(Skip) \rangle, \theta_i\right)$$

until we identify an θ_n such that $\theta_n = \theta_{n+1}$, whereupon (by Definition 6.48) we know that θ_n is an invariant for the loop body.

Since **bw** is not monotonic in its first argument, the conjunction of θ_i in θ_{i+1} ensures that θ_{i+1} is no weaker than θ_i . However, this method is not guaranteed to identify a finite *n* such that $\theta_n = \theta_{n+1}$ in the general case. In such cases, we may resort to intuition in order to identify an invariant obligation, by looking for patterns in the sequence of obligations $\theta_0, \theta_1, \theta_2, \ldots$ Nor is θ_n necessarily the weakest (and hence most desirable) invariant obligation for the loop body. We leave the problem of devising

more sophisticated techniques for identifying invariant obligations for future work.

6.5.3 Full Justification

We say that a process is *fully justified* if every lifted construct in its body is left-justified. Full justification entails that each component within a composition is left-justified and, moreover, that each composite construct is also left-justified in its own right.

The purpose of fully justifying a process specification is to reveal miraculous behaviour stemming from inconsistencies between functionality and confidentiality. These inconsistencies are indicated by the presence of the miraculous CA \langle **false** \rangle within a fully justified process. This CA is generated by backwards propagation if execution of the constructs that follow it could leak information to Low in violation of a subsequent CA.

Example 6.49. Assuming Low can observe events on channel *s*, left-justifying the specification fragment:

$$ST1 = \langle s!a \to Skip \sqcap t!b \to Skip \rangle; \langle a \neq \widetilde{a} \rangle$$

yields the infeasible specification $\langle false \rangle$; *ST*1.

Fully justifying a process is a sound method for verifying its feasibility, because left-justification is a sound method for verifying the feasibility of individual constructs. Hence, the absence of the CA \langle **false** \rangle in a fully justified process body implies the process is feasible. However, full justification is only complete as a verification method if no approximations are made: if Low's inferences are over-approximated too conservatively, then fully justifying a process could generate spurious inconsistencies.

6.5.4 Resolving Inconsistency

Each back-propagated CA reflects Low's inferences about the behaviour of the constructs executed after that CA is invoked. Therefore, whenever

 \Diamond

backwards propagation generates the CA \langle **false** \rangle , we should focus our attention on the constructs following that CA, in order to understand the conflict between functionality and confidentiality.

The presence of miraculous CAs within a process does not necessarily imply the process development is doomed. For instance, if a miraculous CA is local to a non-deterministic construct, then it can be eliminated by refinement, as Example 6.50 shows.

Example 6.50. A specification fragment which refines ST1 is:

$$ST2 = (\langle s!a \to Skip \rangle \sqcap \langle t!b \to Skip \rangle); \langle a \neq \widetilde{a} \rangle$$

Left-justifying each block of ST2 separately gives:

$$\langle \text{ false } \rangle ; \langle s!a \to Skip \rangle ; \langle a \neq \widetilde{a} \rangle$$

 \sqcap
 $\langle a \neq \widetilde{a} \land b = \widetilde{b} \rangle ; \langle t!b \to Skip \rangle ; \langle a \neq \widetilde{a} \rangle$

Since \langle **false** \rangle is the lattice top, this specification is equivalent to:

$$\langle a \neq \widetilde{a} \land b = b \rangle$$
; $\langle t!b \rightarrow Skip \rangle$; $\langle a \neq \widetilde{a} \rangle$

which is itself a fully-justified specification.

In practice, it is expedient to verify the feasibility of back-propagated CAs when calculating them. If a miraculous CA is detected, then further backwards propagation is futile if the inconsistency it reveals cannot be resolved by refining away insecure components of the process.

6.5.5 Discussion

Our propagation procedure is related to the *unwinding* technique first proposed by Goguen and Meseguer (1984), which aims to simplify the task of verifying that a system satisfies noninterference. Modern formulations of unwinding transform a global confidentiality property (expressed

 \diamond

in terms of trace sets) over a system into a set of conditions over its individual state transitions (Ryan and Schneider, 1999; Mantel, 2000a). These conditions can then be discharged using standard proof methods.

Propagation, in a sense, is a dual of unwinding. Propagation transforms a CA — a confidentiality property located in one part of a process but applying to the whole process — into a collection of CAs distributed throughout the process. The procedure for distributing these CAs entails a proof of the process's feasibility, or a result that suggests (but does not necessarily imply) the process is inconsistent.

It is useful for propagated CAs to be retained in a process design, to save the effort of re-calculating them after performing each refinement step. As we show in the next section, they provide useful guidance for ensuring that refinement steps do not violate CAs.

6.6 A Secure Refinement Strategy

For stepwise refinement to be useful, it must lead to process designs that are implementable. Therefore, once a process has been verified to be consistent with its CAs, it is highly desirable that each refinement step maintains the consistency of the process. However, we established in Section 6.2 that applying the *Circus* refinement laws to lifted processes containing CAs is not guaranteed to preserve their feasibility.

This section extends the *Circus* refinement strategy proposed by Cavalcanti et al. (2003) to accommodate process designs featuring CAs. Given a consistent process design, we justify each refinement of the process design with respect to the CAs embedded within the process, in order to maintain the process's consistency.

The secure refinement strategy set out in this section addresses action refinement, data refinement and process refinement. It is closely related to the *Circus* refinement strategy, but it describes additional checks that need to be made to ensure the consistency of processes is preserved by refinement steps.

The application of full justification to a process design greatly reduces

the effort needed to justify that refinement steps maintain the consistency of the process. Using the propagated CAs within a process, we can verify the consistency of the refined process without needing to re-verify the entire process at each refinement step.

6.6.1 Action Refinement

Given a fully justified process *Q*, we are free to refine lifted *Circus* actions nested within blocks. Consider the following left-justified fragment of *Q*:

$$ig\langle$$
 bw $(\langle A_1
angle$, $heta) ig
angle$; $\langle A_1
angle$; $\langle heta
angle$

Let $A_1 \sqsubseteq A_2$. If **bw** ($\langle A_2 \rangle, \theta$) $\not\sqsubseteq$ **bw** ($\langle A_1 \rangle, \theta$), then replacing $\langle A_1 \rangle$ with $\langle A_2 \rangle$ means the specification fragment is no longer left-justified. Hence, the resulting process is not fully justified.

To verify the refinement is feasibility-preserving, we return the refined process to its fully justified form, so that insecurities can be identified. First, we re-apply left-justification to the block to obtain:

$$\langle \mathbf{bw}(\langle A_2 \rangle, \theta) \rangle; \langle A_2 \rangle; \langle \theta \rangle$$

Then, we back-propagate $\langle \mathbf{bw}(\langle A_2 \rangle, \theta) \rangle$ to the preceding blocks and apply left-justification until a fully justified process is reached.

As before, if propagation yields the CA \langle **false** \rangle , then it indicates the refinement step is not feasibility-preserving. In that case, it is necessary either to rectify that infeasibility by refining the process further, or to discard the refined process and try a different development step instead.

In other cases, where **bw** $(\langle A_2 \rangle, \theta) \sqsubseteq$ **bw** $(\langle A_1 \rangle, \theta)$, replacing $\langle A_1 \rangle$ with $\langle A_2 \rangle$ will maintain a fully justified process. There is no need to back-propagate **bw** $(\langle A_2 \rangle, \theta)$, because **bw** is monotonic in its second argument (Lemma 6.16). Therefore, no additional verification is necessary. This observation facilitates a useful shortcut, as Example 6.51 shows.

Example 6.51. The left-justified form of the specification fragment from Example 6.41 can be refined by resolving the non-deterministic choice

between the *s* and *t* actions:

$$\langle a = \widetilde{a} \land b = \widetilde{b} \land a \neq 0 \rangle$$
; $\langle s!a \rightarrow Skip \rangle$; $\langle a = 0 \Rightarrow \widetilde{a} > 0 \rangle$

This block remains left-justified, because:

bw
$$(s!a \rightarrow Skip, a = 0 \Rightarrow \tilde{a} > 0) \sqsubseteq a = \tilde{a} \land b = \tilde{b} \land a \neq 0$$

and so no further backwards propagation is necessary.

 \Diamond

6.6.2 Data Refinement

Data refinement is conventionally conducted by identifying a retrieve relation linking a concrete state space to an abstract state space (He et al., 1986). To ensure that a concrete process simulates the abstract process correctly, this relation must satisfy particular soundness conditions. Cavalcanti et al. (2003) and Oliveira (2005) define the soundness conditions that a retrieve relation must satisfy to be a *forwards simulation* over *Circus* actions.

Definition 6.52 (Forwards simulation (Oliveira, 2005)). A forwards simulation between actions A_1 and A_2 of processes P_1 and P_2 , with state spaces v_1 and v_2 respectively and local state L, is a relation R between v_1 , v_2 and L satisfying:

- 1. (Feasibility) $\forall v_2, L \bullet (\exists v_1 \bullet R)$
- 2. (Correctness) $\forall v_1, v_2, v'_2, L \bullet R \land A_2 \Rightarrow (\exists v'_1, L \bullet A_1 \land R')$

A lifted forwards simulation between lifted actions has essentially the same structure as a forward simulation between *Circus* actions.

Definition 6.53 (Lifted forwards simulation). A lifted forwards simulation between lifted constructs B_1 and B_2 of processes P_1 and P_2 , with state spaces $v_1, \tilde{v_1}$ and $v_2, \tilde{v_2}$ respectively and local state L, \tilde{L} , is a relation R between $v_1, \tilde{v_1}, v_2, \tilde{v_2}$ and L, \tilde{L} satisfying:

1. (Feasibility)
$$\forall v_2, \widetilde{v_2}, L, \widetilde{L} \bullet (\exists v_1, \widetilde{v_1} \bullet R)$$

2. (Correctness)
$$\begin{array}{c} \forall v_1, \widetilde{v_1}, v_2, \widetilde{v_2}, v'_2, \widetilde{v'_2}, L, \widetilde{L} \bullet \\ R \land B_2 \Rightarrow (\exists v'_1, \widetilde{v'_1}, L, \widetilde{L} \bullet B_1 \land R') \end{array}$$

Corollary 6.54 is a consequence of the order-embedding between the space of *Circus* actions and the space of lifted actions (Lemma 5.3).

Corollary 6.54. The relation **U**(*R*) is a lifted forwards simulation between $\langle A_1 \rangle$ and $\langle A_2 \rangle$ exactly if *R* is a forwards simulation between A_1 and A_2 .

With Corollary 6.54, we can apply the existing simulation laws given by Oliveira et al. (2009) for data-refining lifted *Circus* actions. Of course, if these data refinements entail the reduction of non-determinism, they need not yield a fully justified process.

When data refinement is applied to a abstract process featuring CAs, those CAs need to be translated into CAs over the concrete state space. Given an abstract obligation θ_1 , we derive a concrete obligation θ_2 by projecting θ_1 through the simulation relation *R*:

$$\theta_2 = \forall v_1, \widetilde{v_1} \bullet \mathbf{U}(R) \Rightarrow \theta_1$$

 θ_2 is the weakest obligation such that, for each state ψ_2 and fog state $\widetilde{\psi}_2$ that satisfy θ_2 , *all* projections of $\psi_2 \wedge \widetilde{\psi}_2$ through **U**(*R*) satisfy θ_1 .

Example 6.55. Recall the process fragment B01 from Example 6.1:

$$B01 \quad \triangleq \quad \langle h := 0 \sqcap h := 1 \rangle ; \langle h = 0 \Rightarrow h > 0 \rangle$$

If we data-refine each component of this fragment, with the simulation $(g = 0 \lhd h > 0 \rhd g = 1)$, then we obtain the concrete process fragment:

$$\langle \, g := 0 \sqcap g := 1 \, \rangle$$
 ; $\langle \, g = 0 \Rightarrow \widetilde{g} = 1 \, \rangle$

 \Diamond

which is feasible.

Inconsistencies may arise if data refinement is applied to CAs in isolation from the functionality surrounding them, as Example 6.56 shows. **Example 6.56.** Taking *B*01 as before, let $R = (g = 0 \lhd h < 10 \triangleright g = 1)$. Projecting the CA of *B*01 through *R* yields:

$$\begin{array}{l} \forall h, \widetilde{h} \bullet \mathbf{U} \left(g = 0 \lhd h < 10 \vartriangleright g = 1 \right) \Rightarrow \left(h = 0 \Rightarrow \widetilde{h} > 0 \right) \\ = & \forall h, \widetilde{h} \bullet \left(\mathbf{U} \left(g = 0 \lhd h < 10 \vartriangleright g = 1 \right) \land h = 0 \right) \Rightarrow \widetilde{h} > 0 \\ = & \forall \widetilde{h} \bullet \left(g = 0 \land \left(\widetilde{g} = 0 \lhd \widetilde{h} < 10 \vartriangleright \widetilde{g} = 1 \right) \right) \Rightarrow \widetilde{h} > 0 \\ = & g = 0 \Rightarrow \widetilde{g} = 1 \end{array}$$

R is a forwards simulation between $h := 0 \sqcap h := 1$ and the concrete action g := 0. This means the data refinement of B01 with respect to *R* is:

$$\langle g := 0 \rangle$$
; $\langle g = 0 \Rightarrow \tilde{g} = 1 \rangle = \langle g := 0 \rangle$; $\langle false \rangle$

This refinement is clearly too strong. One solution is to restructure the CA before applying data refinement:

$$B01 \quad = \quad \langle h := 0 \sqcap h := 1 \rangle ; \langle h = 0 \Rightarrow \tilde{h} = 1 \rangle$$

Data-refining this restructured specification with respect to *R* yields:

$$\langle g := 0 \rangle$$
; $\langle g = 0 \Rightarrow (\tilde{g} = 0 \lor \tilde{g} = 1) \rangle = \langle g := 0 \rangle$

Now, the CA is innocuous: the data refinement has resolved the nondeterministic choice in a secure way, so the CA plays no further part. \diamond

As Example 6.55 demonstrates, applying data refinement to processes featuring CAs has the potential to create inconsistencies between functionality and confidentiality. These inconsistencies can be detected by restoring a data-refined process to a fully justified form, just as with action refinement.

6.6.3 Process Refinement

In *Circus,* process refinement is defined as action refinement over the main actions of processes, with their state variables hidden:

$$P_1 \sqsubseteq_{\mathcal{P}} P_2 \triangleq (\exists v, v' \bullet P_1.Act) \sqsubseteq (\exists v, v' \bullet P_2.Act)$$

where P_1 .*Act* and P_2 .*Act* denote the main actions of P_1 and P_2 respectively. This definition can easily be extended to the lifted space, by hiding the fog state variables as well.

The *Circus* refinement strategy defines *splitting laws* for achieving process refinement (Cavalcanti et al., 2003). These laws enable a process P to be decomposed into multiple processes. These sub-processes can be joined together (using parallel composition, or another operator) to form a process that refines P. Hence, a monolithic process specification can be implemented by a coalition of communicating processes.

The splitting laws for *Circus* processes are equally applicable to lifted processes. Since these laws are defined in terms of action refinement, we can apply the *Circus* refinement calculus (extended with the laws of Section 6.4) to split a lifted process into multiple lifted processes. However, lifted processes need special attention for a couple of reasons:

- If processes are joined together in parallel composition on channel set *cs*, then synchronisations between the processes may constrain their behaviours. As described in Subsection 6.5.2, this effect can lead to inconsistencies in the parallel process which are not detectable by analysing its constituent processes individually. This problem can be avoided by enlarging Low's window on each process to include all channels in *cs*.
- An impediment to process refinement may arise if CAs reference state variables spanning across a partition of the process state. In such cases, it may be necessary to strengthen a CA with an equivalent (or stronger) CA before decomposing the process.

Example 6.57. Suppose a process incorporates a CA that references state variables *a* and *b*, such as $\langle \tilde{a} \neq a \lor \tilde{b} \neq b \rangle$. This CA could be refined to either $\langle \tilde{a} \neq a \rangle$ or $\langle \tilde{b} \neq b \rangle$, in order to partition the process into two components with states *a* and *b* respectively.

We do not consider the *Circus* process indexing mechanism, or the associated promotion techniques described by Cavalcanti et al. (2003).

6.7 Implementation

A program specification is ready to be implemented only if it exclusively consists of executable constructs (Morgan, 1994). Since CAs are partial constructs — i.e. they may exhibit miraculous behaviour — they are not executable and should therefore be eliminated from a process design to ensure the process is implementable.

6.7.1 Translating to Circus

The order embedding linking *Circus* constructs and lifted constructs can be applied to translate lifted processes to *Circus* processes. The purpose of making this translation is to enable existing techniques for implementing *Circus* to be applied.

Theorem 6.58. Provided *Q* is a fully justified lifted process with no miraculous CAs, D(Q) can be expressed as a pure *Circus* process.

The justification of Theorem 6.58 is instructive. Given a fully justified process *Q* such that no CA in *Q* is \langle **false** \rangle , then because *Q* is fully justified, the CAs in *Q* cannot induce miraculous behaviour. By Condition 5.11, the application of **D** to *Q* distributes through each the lifted operators of *Q*. Hence, by distributing **D** throughout the body of *Q*, **D**(*Q*) can be expressed as a process consisting of constructs of the form:

 $D(\langle A \rangle)$ and $D(\langle \theta \rangle)$

157

composed with the (unlifted) *Circus* operators. By Law 5.4 and Law 5.45, each of these constructs is itself a *Circus* action. It is important to stress this translation does *not* constitute a refinement of a lifted process, because the fog variables are absent from the alphabet of *Circus* processes.

A *Circus* process translated by applying **D** to a fully justified lifted process by this procedure does not leak secret information to Low. However, the *Circus* semantics does not guarantee that refinements of this *Circus* process are themselves secure. To guard against the danger of making the process insecure, the translation to *Circus* should be made only after all development work on the process design is finished. In particular, if the process is fully deterministic, then no refinement can reduce non-determinism further, thereby disposing of the danger described in Section 4.5).

6.7.2 Compilation to Code

Oliveira (2005) has described a scheme for translating *Circus* processes into Java code. Since the processes derived by applying **D** to lifted processes are just *Circus* processes, this scheme could be applied to implement them. However, the security of an implementation generated using this scheme is open to question, because it could leak secrets to the adversary in unexpected ways:

- At the software level, Oliveira's translation scheme implements the resolution of non-determinism according to the output of a software random number generator (RNG). It is well known that the output stream of a software RNG is not truly random, but exhibits patterns (Anderson, 2001). If an adversary can detect those patterns, it may exploit them to predict how non-determinism is resolved by the process implementation, with catastrophic consequences for confidentiality.
- A confederation of processes may be distributed across a network of nodes. Encrypting the information transmitted between nodes

would be necessary, but may not suffice to safeguard the confidentiality of that information. For instance, the encryption algorithm may be breakable with sufficient computational resources. Alternatively, the protocols used for authentication or key exchange may harbour subtle flaws in their design or implementation (Burrows et al., 1989; Carlsen, 1994), which an adversary may exploit to cause the protocol to deviate from the execution its designers intended.

• At the hardware level, power is consumed and electromagnetic radiation is generated in proportion to the computation being performed. Fluctuations in these *side channels* signal information about the computation in progress. If an adversary is able to monitor (or even tamper with) system hardware during operation, then it may exploit these side channels to extract information about the system's state or behaviour. These attacks are not merely hypothetical: Kocher et al. (1999) applied a form of power analysis to read cryptographic keys from smart cards, based on knowledge of the cryptographic algorithm being executed.

These characteristics of physical systems are not modelled in the semantics of *Circus* processes. In the words of Ryan and Schneider (1998):

"[A] formal model is necessarily an abstraction of the real system. As such it is always possible that some significant aspect of the real system is missed."

Since we have not formalised these characteristics, we can only justify a physical computing artefact that implements a process design as being secure if we assume the adversary cannot exploit these characteristics. In principle, these characteristics could be modelled by extending *Circus* (and its semantics), but that endeavour is beyond the scope of this thesis.

6.8 Conclusion

This chapter has detailed a confidentiality-sensitive extension of the *Circus* refinement strategy for developing abstract process designs into concrete processes. Since refinement in the lifted semantics is confidentiality-preserving but not feasibility-preserving, this strategy emphasises providing guidance to software engineers to maintain the feasibility of each process design in the refinement chain.

Backwards propagation is an effective device for executing the strategy, as it provides a platform for verifying the feasibility of a process and conducting refinement steps that preserve feasibility. Fully justifying a process is sufficient for verifying that the process is feasible. However, the rules given in Subsection 6.5.2 for propagating CAs through lifted constructs are not guaranteed to be complete, because they over-approximate Low's inferences in some circumstances. Therefore, in those cases where full justification fails to prove that a process is feasible, software engineers may be obliged to resort to first principles in order to establish the process is indeed feasible, or to adjust the process design in order to make full justification successful.

Strictly speaking, verification can be carried out at any stage during a process's development. It is arguably best left to software engineers to choose the stage of development where verification is most appropriate. However, delaying verification may be disadvantageous, because without the guidance of propagated CAs, a process design may be reached containing infeasible elements that cannot be refined away. It would then be necessary either:

- to backtrack to an earlier design in the development and proceed along a different refinement path;
- to break the refinement chain, by weakening the functionality or confidentiality attributes of the design.
- to weaken the notion of refinement; for instance, to co-operating or independent refinement (Subsection 3.6.2).

The first two options are not ideal: the first incurs wasted effort, the second may sacrifice the design's correctness. The third option has promise if the observational abilities of all users of the process — and the relationships between those users — are known. However, *Circus* does not provide facilities to specify this information, so extensions to the language would needed to support this option.

An important topic for future work would be to automate the propagation procedure. The connection between backwards propagation and the weakest precondition semantics suggests that verifying a process satisfies its CAs may be machine-checkable. Alternatively, it may be more efficient to automate propagation in the forwards direction.

The next chapter provides a case study, as empirical evidence to validate the comprehensiveness of the refinement strategy.

7 Case Study

7.1 Introduction

The previous chapters have expounded a formal platform for developing secure software which builds on the *Circus* formal method. This chapter exemplifies how the platform can be applied in practice, by executing a case study development. It explores the general issues that emerge when functionality and confidentiality concerns are united in a system development. It also aims to reinforce the reader's understanding of the techniques presented in the previous chapters.

Even though it is small in scope, this case study aims to validate the platform with respect to its facilities for incorporating confidentiality properties into system developments. Moreover, this case study clarifies the capabilities (and weaknesses) of the platform, enabling us to identify potential avenues for enhancing its scope.

Starting from a *Circus* specification of a system's functionality in Section 7.2, we apply backwards propagation to calculate the inferences of its users in Section 7.3. Thereafter, in Section 7.4, we consider a series of custom confidentiality properties which the system is required to uphold. For each property, we show how the specification can be adjusted to incorporate the property. We also apply backwards propagation to determine whether the resulting specifications are feasible. We aim to demonstrate that backwards propagation is a practicable technique for verifying confidentiality.

Having constructed a consistent process specification incorporating both functionality and confidentiality attributes, we consider a selection of possible design decisions that could be applied to that process in Sec-

7 Case Study

tion 7.5. Again, we verify the correctness of these design decisions with respect to confidentiality by applying backwards propagation. Finally, Section 7.6 summarises some lessons learnt from this case study.

7.2 A Sealed-Bid Auction

We study a *sealed-bid auction* contested by Alice and Bob. In a sealed-bid auction, each contestant independently submits a single bid (for our purposes, a positive integer) to the auctioneer. On receiving both bids, the auctioneer declares the winner of the auction as the contestant who bid the larger amount. Should the bids be equal, the auctioneer is free to choose the winner.

A (fictional) auction house commissions a *Circus* specification of the auctioneer, which is presented in Figure 7.1. This process first initialises its state, then accepts bids from Alice and Bob in an unspecified order, before it announces the winner of the auction and halts. For the purposes of this study, the process is able to choose which bid to accept first.

Notice that this process deadlocks if either contestant fails to submit their bid. A more realistic model would represent the contestants opting either to submit a bid or declining to bid (before a deadline expires), but this detail is irrelevant to our study.

The contestants and the auction house may expect a sealed-bid auction to uphold various kinds of confidentiality requirements. However, the *Auction* process specification tells us nothing about those requirements. These requirements need to be identified and formalised, in order to incorporate them within the process development.

7.3 Modelling Contestants' Inferences

In this section, we consider how backwards propagation can be applied to capture each user's maximal knowledge about the process state, based on its subsequent observation of the process's behaviour. To avoid com-

```
BDR ::= Alice \mid Bob
channel bidAlice, bidBob : \mathbb{N}_1
channel winner : BDR
process Auction \triangleq begin
state ASt \triangleq [a, b : \mathbb{N}]
Init \triangleq [ASt' \mid a' = 0 \land b' = 0]
BidA \triangleq bidAlice?n \rightarrow a := n?
BidB \triangleq bidBob?n \rightarrow b := n?
Submit \triangleq (BidA ; BidB) \sqcap (BidB ; BidA)
Declare \triangleq \begin{pmatrix} a > b \& winner!Alice \rightarrow Stop \\ \Box & a < b \& winner!Bob \rightarrow Stop \\ \Box & a = b \& \sqcap c : BDR \bullet winner!c \rightarrow Stop \end{pmatrix}
• Init ; Submit ; Declare
end
```

Figure 7.1: First model of auctioneer, capturing functionality alone.

plicating matters at this stage, we defer the introduction of confidentiality properties into the development until the next section.

We expect that Alice can observe her own bid (on the *bidAlice* channel), Bob can observe his own bid (on *bidBob*), and both contestants can observe the *winner* channel. In a sealed-bid auction, the contestants submit their bids in private, so we assume that Alice cannot observe *bidBob* and Bob cannot observe *bidAlice*. We specify Alice and Bob's windows as channel sets, as detailed in Subsection 5.4.1:

channelset $A \triangleq \{ | bidAlice, winner | \}$ **channelset** $B \triangleq \{ | bidBob, winner | \}$

Our next task is to lift the *Auction* process. We shall focus on the initialisation, submission and declaration phases of the auction. Hence, we elevate the actions *Init*, *Submit* and *Declare* to the status of blocks. Let:

```
Init_0 \triangleq \langle Init \rangle Submit_0 \triangleq \langle Submit \rangle Declare_0 \triangleq \langle Declare \rangle
```

7 Case Study

To model Alice and Bob's inferences about the behaviour of *Auction*, we apply full propagation to the lifted process. Since *Auction* is structured sequentially, it is expedient to propagate backwards from *Declare*, as explained in Remark 6.43.

7.3.1 Declaration

The backwards propagation laws given in Subsection 6.4.2 do not extend to blocks with the structure of $Declare_0$. Rather than refining $Declare_0$ into multiple blocks — which would over-approximate Alice and Bob's inferences about the winner of the auction (Subsection 6.2.2) — we instead apply the heuristics described in Subsection 6.4.1 to calculate the state information revealed to Alice and Bob by the behaviour of $Declare_0$.

Since both Alice and Bob can observe *winner* events, their interactions with *Declare*₀ are identical. We break *Declare*₀ apart as follows:

bw (*Declare*₀, θ)

$$= \begin{pmatrix} \mathsf{bw} (Declare_0 \land tr' - tr = \langle \rangle, \theta) \\ \land \mathsf{bw} (Declare_0 \land tr' - tr = \langle WA \rangle, \theta) \\ \land \mathsf{bw} (Declare_0 \land tr' - tr = \langle WB \rangle, \theta) \end{pmatrix}$$
[Lemma 6.21]
$$= \begin{pmatrix} \mathsf{bw} (Declare_0 \land tr' - tr = \langle \rangle \land WA \in ref', \theta) \\ \land \mathsf{bw} (Declare_0 \land tr' - tr = \langle \rangle \land WB \in ref', \theta) \\ \land \mathsf{bw} (Declare_0 \land tr' - tr = \langle WA \rangle, \theta) \\ \land \mathsf{bw} (Declare_0 \land tr' - tr = \langle WB \rangle, \theta) \end{pmatrix}$$

[Lemma 6.22]

where WA denotes (winner, Alice) and WB denotes (winner, Bob).

Since $Declare_0$ does not terminate successfully, the choice of θ is irrelevant, because **bw** only imposes θ on terminal states. We consider each part in isolation:

bw
$$(Declare_0 \land tr' - tr = \langle \rangle \land WA \in ref', \theta) = a \le b \Rightarrow \tilde{a} \le \tilde{b}$$

bw $(Declare_0 \land tr' - tr = \langle \rangle \land WB \in ref', \theta) = a \ge b \Rightarrow \tilde{a} \ge \tilde{b}$

7.3 Modelling Contestants' Inferences

bw
$$(Declare_0 \wedge tr' - tr = \langle WA \rangle, \theta) = a \ge b \Rightarrow \widetilde{a} \ge \widetilde{b}$$

bw $(Declare_0 \wedge tr' - tr = \langle WB \rangle, \theta) = a \le b \Rightarrow \widetilde{a} \le \widetilde{b}$

Hence, the weakest obligation which left-justifies *Declare*⁰ is:

$$bw0 = (a \ge b \Rightarrow \widetilde{a} \ge \widetilde{b}) \land (a \le b \Rightarrow \widetilde{a} \le \widetilde{b})$$

If a = b, the non-deterministic choice in $Declare_0$ reveals to the contestants either $a \le b$ or $a \ge b$. Since we cannot control demonic non-determinism, we must be prepared to accept either case. (The same problem manifests in Example 6.8.) By way of compensation, bw0 over-approximates the contestants' inferences about the process state if a = b:

$$a = b \Rightarrow bw0$$

= $a = b \Rightarrow \tilde{a} \ge \tilde{b} \land \tilde{a} \le \tilde{b}$ [prop calc]
= $a = b \Rightarrow \tilde{a} = \tilde{b}$ [property of \le]

and so Alice and Bob are modelled as being able to deduce if their bids are equal, even though a single run of *Declare*⁰ does not reveal that information to them.

7.3.2 Submission

To back-propagate an obligation through $Submit_0$, we can treat $Submit_0$ as a composite choice (in the sense of Subsection 6.2.2) between BidA; BidB and BidB; BidA. This approach is justified by instantiating Law 6.37:

$$\mathbf{bw}\left(Submit_{0},\theta\right) \subseteq \begin{pmatrix} \mathbf{bw}\left(\langle BidA; BidB \rangle, \theta\right) \\ \land \mathbf{bw}\left(\langle BidB; BidA \rangle, \theta\right) \end{pmatrix}$$

We consider Alice and Bob's inferences about this block separately. With Alice's window, assuming the scope of n? extends beyond the

7 Case Study

prefixing, the block $\langle A : BidA; BidB \rangle$ can be broken down as follows:

$$\langle A : bidAlice?n \to a := n?; bidBob?n \to b := n? \rangle$$

$$= \begin{cases} \langle A : bidAlice?n \to Skip \rangle; \\ \langle A : a := n?; bidBob?n \to b := n? \rangle \end{cases}$$

$$= \begin{cases} \langle A : bidAlice?n \to Skip \rangle; \langle A : a := n? \rangle; \\ \langle A : bidBob?n \to Skip \rangle; \langle A : b := n? \rangle \end{cases}$$
[Lemma 5.18, twice]

Hence, an obligation θ can be back-propagated through this specification fragment by appealing to the laws in Section 6.4. With Alice's window, the back-propagation of θ through *BidA* is:

$$\begin{aligned} & \mathsf{bw}\left(\langle A : bidAlice?n \to Skip \rangle; \langle A : a := n? \rangle, \theta\right) \\ &= & \mathsf{bw}\left(\langle A : bidAlice?n \to Skip \rangle, \mathsf{bw}\left(\langle A : a := n? \rangle, \theta\right)\right) \text{ [Law 6.32]} \\ &= & \mathsf{bw}\left(\langle A : bidAlice?n \to Skip \rangle, \theta[n?, \widetilde{n?}/a, \widetilde{a}]\right) \text{ [Law 6.31]} \\ &= & \forall n? : \mathbb{N}_1 \bullet \theta[n?, \widetilde{n?}/a, \widetilde{a}][n?/\widetilde{n?}] \text{ [Law 6.36]} \\ &= & \forall a : \mathbb{N}_1 \bullet \theta[a/\widetilde{a}] \text{ [renaming]} \end{aligned}$$

Here, the renaming of \tilde{a} to a is justified because $a = \tilde{a}$; that is, Alice knows the exact value of her own bid.

With Alice's window, the back-propagation of θ through *BidB* is:

$$\mathbf{bw} \left(\left\langle A : bidBob?n \to Skip \right\rangle; \left\langle A : b := n? \right\rangle, \theta \right) \\ = \mathbf{bw} \left(\left\langle A : bidBob?n \to Skip \right\rangle, \theta[n?, \tilde{n?}/b, \tilde{b}] \right) \qquad [as before] \\ = \forall n? : \mathbb{N}_1 \bullet \exists \tilde{n?} : \mathbb{N}_1 \bullet \theta[n?, \tilde{n?}/b, \tilde{b}] \qquad [Law 6.36] \\ = \forall b : \mathbb{N}_1 \bullet \exists \tilde{b} : \mathbb{N}_1 \bullet \theta \qquad [renaming] \end{cases}$$

In line with our intuition, this calculation indicates nothing is revealed to Alice about Bob's bid *b*, except that $b \in \mathbb{N}_1$. Therefore, we have:

bw
$$(\langle A : BidA; BidB \rangle, \theta) = \forall a, b : \mathbb{N}_1 \bullet \exists b : \mathbb{N}_1 \bullet \theta[a/\tilde{a}]$$

168

7.3 Modelling Contestants' Inferences

bw $(\langle A : BidB; BidA \rangle, \theta) = \forall a, b : \mathbb{N}_1 \bullet \exists \tilde{b} : \mathbb{N}_1 \bullet \theta[a/\tilde{a}]$

With Bob's window, the back-propagations of θ through *BidA* and *BidB* are symmetric:

bw $(\langle B : BidA; BidB \rangle, \theta) = \forall a, b : \mathbb{N}_1 \bullet \exists \tilde{a} : \mathbb{N}_1 \bullet \theta[b/\tilde{b}]$ **bw** $(\langle B : BidB; BidA \rangle, \theta) = \forall a, b : \mathbb{N}_1 \bullet \exists \tilde{a} : \mathbb{N}_1 \bullet \theta[b/\tilde{b}]$

Back-propagating *bw*0 through $\langle A : BidA; BidB \rangle$ yields:

$$\mathbf{bw} \left(\left\langle A : BidA; BidB \right\rangle, bw0 \right)$$

$$= \quad \forall a, b : \mathbb{N}_1 \bullet \exists \tilde{b} : \mathbb{N}_1 \bullet bw0[a/\tilde{a}] \qquad [as above]$$

$$= \quad \forall a, b : \mathbb{N}_1 \bullet \exists \tilde{b} : \mathbb{N}_1 \bullet (a \ge b \Rightarrow a \ge \tilde{b}) \land (a \le b \Rightarrow a \le \tilde{b})$$

$$= \quad \mathbf{true} \qquad [pred calc]$$

It should be no surprise that **bw** ($\langle B : BidA; BidB \rangle, bw0$) is also **true**.

By substituting these results into our split form of *Submit*₀:

$$\mathbf{bw} (Submit_0, \theta) \subseteq \begin{pmatrix} \mathbf{bw} (\langle BidA; BidB \rangle, \theta) \\ \land \mathbf{bw} (\langle BidB; BidA \rangle, \theta) \end{pmatrix}$$

we find that **bw** (*Submit*₀, *bw*0) \sqsubseteq **true**, so **bw** (*Submit*₀, *bw*0) = **true**.

7.3.3 Initialisation

The backwards propagation of **bw** (*Submit*₀, *bw*0) through *Init*₀ is trivial:

bw
$$(Init_0, true) = true$$

Figure 7.2 presents a fully-justified lifted process form of *Auction* — complete with CAs modelling Alice and Bob's inferences about the process state — based on the calculations we have performed.

 $BDR ::= Alice \mid Bob$ channel bidAlice, bidBob : \mathbb{N}_1 channel winner : BDR channelset $A \triangleq \{\mid bidAlice, winner \mid\}$ channelset $B \triangleq \{\mid bidBob, winner \mid\}$ process AuctionL \triangleq begin state $ASt \triangleq [a, b : \mathbb{N}]$ $Init_0 \triangleq \langle [ASt' \mid a' = 0 \land b' = 0] \rangle$ $BidA \triangleq bidAlice?n \rightarrow a := n?$ $BidB \triangleq bidBob?n \rightarrow b := n?$ $Submit_0 \triangleq \langle (BidA; BidB) \sqcap (BidB; BidA) \rangle$ $Declare_0 \triangleq \left\langle \begin{array}{c} a > b \& winner!Alice \rightarrow Stop \\ \Box & a > b \& winner!Bob \rightarrow Stop \\ \Box & a = b \& \sqcap c : BDR \bullet winner!c \rightarrow Stop \\ \Box & a = b \& \sqcap c : BDR \bullet winner!c \rightarrow Stop \\ bw0 \triangleq (a \ge b \Rightarrow \tilde{a} \ge \tilde{b}) \land (a \le b \Rightarrow \tilde{a} \le \tilde{b})$ $\bullet \langle true \rangle; Init_0; \langle true \rangle; Submit_0; \langle bw0 \rangle; Declare_0; \langle true \rangle$

Figure 7.2: Lifted and fully-justified model of auctioneer.

7.4 Specifying Confidentiality Properties

We now investigate how some example confidentiality properties can be expressed over *Auction*₀, and verify whether they are consistent with the pre-existing functionality of the process by applying backwards propagation.

In practice, it would be expected that backwards propagation should be performed *after* a specification covering functionality and confidentiality properties has been produced, to verify these properties are compatible before proceeding with the development. When we are reasonably confident that these properties are compatible, this approach would be most expedient. However, we take the opposite approach in this section for pedagogic reasons. Our goals here are to focus on how individual confidentiality requirements can be encoded; and to investigate how they may affect the feasibility of a process. These goals are best served by considering each requirement in isolation, as we proceed to do.

7.4.1 Comparing Bids

The auction house's first confidentiality requirement is:

CP1 The contestants cannot establish if Alice's bid is larger than Bob's.

We capture **CP1** by inserting a CA into the body of the process, between $Submit_0$ and $Declare_0$:

Init₀; Submit₀;
$$\langle a > b \Rightarrow \tilde{a} \leq b \rangle$$
; Declare₀

To determine whether this CA is consistent with the process, we compose it with the back-propagation *bw*0 of *Declare*₀:

$$\langle a > b \Rightarrow \widetilde{a} \le \widetilde{b} \rangle; \langle (a \ge b \Rightarrow \widetilde{a} \ge \widetilde{b}) \land (a \le b \Rightarrow \widetilde{a} \le \widetilde{b}) \rangle$$

$$= \langle a > b \Rightarrow \widetilde{a} \le \widetilde{b} \land (a \ge b \Rightarrow \widetilde{a} \ge \widetilde{b}) \land (a \le b \Rightarrow \widetilde{a} \le \widetilde{b}) \rangle [Law \ 5.41]$$

$$= \langle (a \ge b \Rightarrow \widetilde{a} = \widetilde{b}) \land (a < b \Rightarrow \widetilde{a} \le \widetilde{b}) \rangle [prop \ calc]$$

7 Case Study

The obligation tells us that $Declare_0$ does not leak information that violates **CP1**. Since this obligation is stronger than bw0, we still need to propagate it backwards through $Submit_0$, to ensure $Submit_0$ is left-justified. As in Subsection 7.3.2, we calculate:

$$\begin{aligned} & \mathsf{bw}\left(Submit_{0}, (a \geq b \Rightarrow \widetilde{a} = \widetilde{b}) \land (a < b \Rightarrow \widetilde{a} \leq \widetilde{b})\right) \\ &= \forall a, b : \mathbb{N}_{1} \bullet \exists \widetilde{b} : \mathbb{N}_{1} \bullet ((a \geq b \Rightarrow \widetilde{a} = \widetilde{b}) \land (a < b \Rightarrow \widetilde{a} \leq \widetilde{b}))[a/\widetilde{a}] \\ &= \mathbf{true} \end{aligned} \qquad [renaming, pred calc] \end{aligned}$$

This calculation indicates that **CP1** is upheld, because Alice cannot rule out the possibility that her own bid is equal to Bob's. Again, the calculation with Bob's window is symmetric.

Since this obligation matches the back-propagation of bw0 through $Submit_0$, we need not back-propagate it further. Hence, we can ascertain the process embedding the CA corresponding to **CP1** is self-consistent.

7.4.2 Value of Bids

The second confidentiality requirement we consider is:

CP2 Neither contestant can deduce if the other's bid exceeds £100.

We capture **CP2** by inserting two separate CAs over Alice and Bob's inferences about each other's bids. Let *CP2A* specify that Bob cannot deduce that Alice's bid exceeds £100:

 $CP2A \triangleq \langle B : a > 100 \Rightarrow \tilde{a} \le 100 \rangle$

Likewise, let CP2B specify the complementary condition for Alice:

 $CP2B \triangleq \langle A : b > 100 \Rightarrow \widetilde{b} \le 100 \rangle$

We cannot insert *CP2A* and *CP2B* into *BidA* and *BidB* without splitting *Submit*₀ into multiple blocks. Instead, we insert *CP2A* and *CP2B*

immediately following *Submit*⁰ in the process body:

$$Init_0$$
; $Submit_0$; $CP2A$; $CP2B$; $\langle bw0 \rangle$; $Declare_0$

To left-justify *Submit*⁰ with respect to *CP2A* and *CP2B*, we back-propagate each CA through *Submit*⁰ separately. Here, we focus on *CP2B*:

$$\begin{aligned} & \mathsf{bw} \left(Submit_0, bw0 \land CP2B \right) \\ & \forall a, b : \mathbb{N}_1 \bullet \exists \widetilde{b} : \mathbb{N}_1 \bullet (bw0 \land (b > 100 \Rightarrow \widetilde{b} \le 100)) [a/\widetilde{a}] \quad \text{[as above]} \\ & \forall a, b : \mathbb{N}_1 \bullet \exists \widetilde{b} : \mathbb{N}_1 \bullet \begin{pmatrix} bw0[a/\widetilde{a}] \land \widetilde{b} \le 100 \\ \lhd a > 100 \land a < b \rhd \\ bw0[a/\widetilde{a}] \land (b > 100 \Rightarrow \widetilde{b} \le 100) \end{pmatrix} \end{aligned}$$

[renaming; case split]

The case split is helpful for evaluating the truth of this predicate. Focusing on the case where Alice's bid exceeds £100 but is smaller than Bob's bid, we have:

$$\begin{array}{ll} a > 100 \land a < b \Rightarrow (a \ge b \Rightarrow a \ge \tilde{b}) \land (a \le b \Rightarrow a \le \tilde{b}) \land \tilde{b} \le 100 \\ = & a > 100 \land a < b \Rightarrow a \le \tilde{b} \land \tilde{b} \le 100 & \text{[pred calc]} \\ = & a > 100 \land a < b \Rightarrow a \le \tilde{b} \land \tilde{b} \le 100 \land a \le 100 & \text{[property of } \le] \\ = & a > 100 \land a < b \Rightarrow \text{false} & \text{[contradiction]} \end{array}$$

Because Bob is declared the winner in this case, Alice can infer that Bob's bid is over £100. Hence, **bw** (*Submit*₀; *CP2B*, *bw*0) reduces to:

 $\forall a, b : \mathbb{N}_1 \bullet \exists \tilde{b} : \mathbb{N}_1 \bullet$ $\neg (a > 100 \land a < b) \land bw0[a/\tilde{a}] \land (b > 100 \Rightarrow \tilde{b} \le 100)$ [as above] = false [pred calc]

This calculation tells us the specification fragment:

7 Case Study

may exhibit miraculous behaviour: should Alice's bid exceed £100 but be smaller than Bob's bid, the process would fail to announce the winner of the auction, to prevent Alice's inferences from violating **CP2**. Hence, the **CP2** property is inconsistent with the functionality of the process.

In this case study, we proceed with the process development in the absence of **CP2**. In practice, if the auction house's requirements for the system included **CP2**, then it would be the responsibility of the system developers to seek clarification from the auction house to resolve the inconsistency.

7.4.3 Order of Bids

The third confidentiality requirement we consider is:

CP3 Neither contestant can learn who placed the first bid.

We formalise **CP3** by refining $Submit_0$. Applying a tactic described in Subsection 5.6.2, we introduce a local variable *f* into $Submit_0$ to record which contestant bids first. Then, **CP3** can be expressed with a CA referencing *f*, as follows:

SubmitAux $\triangleq \langle f = Alice \& BidA ; BidB \Box f = Bob \& BidB ; BidA \rangle$ CP3 $\triangleq \langle \tilde{f} = opp(f) \rangle$ Submit₁ $\triangleq \mathbf{var} f \bullet \langle f := Alice \Box f := Bob \rangle ; CP3 ; SubmitAux$

where *opp* is a function that maps each contestant to its opponent:

$$opp: BDR \rightarrow BDR$$

 $opp(Alice) = Bob$
 $opp(Bob) = Alice$

Here, the CA *CP*3 specifies that Bob bidding first serves as a cover story for Alice bidding first ($f = Alice \Rightarrow \tilde{f} = Bob$) and vice versa. At first glance, we may be tempted to say this CA is upheld by the process, but we still need to apply back-propagation to be sure.

To back-propagate an obligation θ through *SubmitAux*, we again resort to heuristics. Reasoning from Alice's perspective, we have:

 $\mathbf{bw} (SubmitAux, \theta)$ = [Lemma 6.23; Lemma 6.21; Lemma 6.22] $\begin{pmatrix} \mathbf{bw} (SubmitAux \land wait' \land \neg TA \land RA, \theta) \\ \land \mathbf{bw} (SubmitAux \land wait' \land \neg TA \land \neg RA, \theta) \\ \land \mathbf{bw} (SubmitAux \land wait' \land TA, \theta) \\ \land \mathbf{bw} (SubmitAux \land \neg wait', \theta) \end{pmatrix}$

where A denotes the set of events in window A in the following:

$$TA = \exists n : \mathbb{N}_1 \bullet (tr' - tr) \upharpoonright \mathcal{A} = \langle (bidAlice, n) \rangle$$

$$RA = \exists n : \mathbb{N}_1 \bullet (bidAlice, n) \in ref' \cap \mathcal{A}$$

The first **bw** conjunct describes Alice's inferences when she attempts to submit her bid, but finds the process refuses to accept it. The only explanation for this circumstance is that the process is deadlocked, waiting for Bob to submit his bid. Hence:

bw
$$(SubmitAux \land wait' \land \neg TA \land RA, \theta) = (f = Bob \Rightarrow f = Bob)$$

(Obligation θ is irrelevant here, because we are considering only those behaviours where *SubmitAux* does not terminate.)

The second **bw** construct describes Alice's inferences when the process does not refuse her bid. It reveals no information to Alice about *f*:

bw (*SubmitAux*
$$\land$$
 wait' $\land \neg$ *TA* \land *RA*, θ) = **true**

The third **bw** conjunct describes Alice's inferences when the process is deadlocked, but she has submitted her bid. Again, this circumstance can only be explained by the process waiting for Bob to submit his bid. 7 Case Study

Hence:

bw
$$(SubmitAux \land wait' \land TA, \theta) = (f = Alice \Rightarrow \tilde{f} = Alice)$$

The fourth **bw** conjunct just back-propagates θ through *SubmitAux*:

bw
$$(SubmitAux \land \neg wait', \theta) = \forall a, b : \mathbb{N}_1 \bullet \exists \tilde{b} : \mathbb{N}_1 \bullet \theta[a/\tilde{a}]$$

Assembling these four conjuncts yields:

bw
$$(SubmitAux, bw0) = (f \in BDR \Rightarrow f = f)$$

We now back-propagate this obligation through *CP*3:

$$\begin{aligned} & \mathsf{bw}\left(\langle \widetilde{f} = opp(f) \rangle, (f \in BDR \Rightarrow f = \widetilde{f})\right) \\ &= \quad \widetilde{f} = opp(f) \land (f \in BDR \Rightarrow f = \widetilde{f}) \end{aligned} \tag{Law 6.18}$$

and then back-propagate through the remaining block of *Submit*₁:

$$\begin{aligned} & \mathsf{bw}\left(\langle f := Alice \sqcap f := Bob \rangle, \tilde{f} = opp(f) \land (f \in BDR \Rightarrow f = \tilde{f})\right) \\ &= \exists f, \tilde{f} \bullet f \in BDR \land \tilde{f} \in BDR \land \tilde{f} = opp(f) \land (f \in BDR \Rightarrow f = \tilde{f}) \\ &= \exists f, \tilde{f} \bullet f \in BDR \land \tilde{f} \in BDR \land \tilde{f} = opp(f) \land f = \tilde{f} \qquad [\text{prop calc}] \\ &= \exists f, \tilde{f} \bullet f \in BDR \land \tilde{f} \in BDR \land \tilde{f} \neq f \land f = \tilde{f} \qquad [\text{property of } opp] \\ &= \mathbf{false} \qquad [\text{pred calc}] \end{aligned}$$

This result informs us that our definition of $Submit_1$ is inconsistent with respect to its embedded CA. Intuitively, **CP3** is violated (by the unlifted *Auction* process) because if either contestant fails to submit a bid, then the other contestant can deduce the bidding order by observing:

- either the process's refusal to accept their own bid;
- or the process's refusal to supply the winner of the auction, after the contestant has submitted their own bid.

Were it the case that our model of Alice's and Bob's observations (as given by Definition 5.1) did not include a projection of the refusals of a process, then **CP3** would not induce an inconsistency with $Submit_1$.

If the auction house is prepared to relax the **CP3** requirement to:

CP4 *Under the circumstances that both bids are submitted,* neither contestant can learn who placed the first bid.

we can redefine $Submit_1$ as:

Submit₂
$$\triangleq$$
 var $f \bullet \langle f := Alice \sqcap f := Bob \rangle$; SubmitAux; CP3

Here, the CA *CP*3 is invoked only after Alice and Bob have submitted their bids. Even though *Submit*₂ uses the same CA as before, this specification is weaker than *Submit*₂.

Now, we have:

bw
$$\left(SubmitAux \land \neg wait', \tilde{f} = opp(f) \land bw0\right) = \tilde{f} = opp(f)$$

and back-propagating through the remaining block of *Submit*₂ gives:

bw
$$\left(\langle f := Alice \sqcap f := Bob \rangle, \tilde{f} = opp(f) \right) =$$
true

By the same argument as that given by Subsection 7.4.1, we can conclude the *AuctionLC* process (Figure 7.3) — featuring an explicit specification of **CP1** and **CP4** — is consistent with respect to our model of Alice and Bob's observational abilities.

7.5 Potential Refinement Steps

Following Section 6.6, we now consider briefly some refinement steps that could be applied to the consistent *AuctionLC* process.

 $BDR ::= Alice \mid Bob$ **channel** *bidAlice*, *bidBob* : \mathbb{N}_1 channel winner : BDR **channelset** $A \triangleq \{ | bidAlice, winner | \}$ **channelset** $B \triangleq \{ | bidBob, winner | \}$ **process** *AuctionLC* \triangleq **begin** state $ASt \triangleq [a, b : \mathbb{N}]$ $CP1 \triangleq \langle a > b \Rightarrow \widetilde{a} \leq \widetilde{b} \rangle$ $CP4 \triangleq \langle \widetilde{f} = opp(f) \rangle$ $Init_0 \triangleq \langle [ASt' | a' = 0 \land b' = 0] \rangle$ $BidA \triangleq bidAlice?n \rightarrow a := n?$ $BidB \triangleq bidBob?n \rightarrow b := n?$ SubmitAux $\triangleq \langle f = Alice \& BidA; BidB \Box f = Bob \& BidB; BidA \rangle$ $Submit_{2} \triangleq \operatorname{var} f \bullet \langle f := Alice \Box f := Bob \rangle ; SubmitAux ; CP4$ $Beclare_{0} \triangleq \left\langle \begin{array}{c} a > b \& winner!Alice \rightarrow Stop \\ \Box a < b \& winner!Bob \rightarrow Stop \\ \Box a = b \& \Box c : BDR \bullet winner!c \rightarrow Stop \end{array} \right\rangle$ • Init₀; Submit₂; CP1; Declare₀ end

Figure 7.3: Lifted model of auctioneer, capturing CP1 and CP4.

7.5.1 Bidding Order

We could refine *Submit*₂ so that Alice always bids first:

Submit₃ \triangleq **var** $f \bullet \langle f := Alice \rangle$; $\langle BidA; BidB \rangle$; CP4

This refinement is naïve: by fixing the order of bidding, it is clear that **CP4** is violated. Formally, this inconsistency would be detected by back-propagating *CP*4 through the preceding blocks.

An alternative refinement of $Submit_2$ would offer the choice between Alice and Bob's bids to the environment:

$$Submit_4 \triangleq \mathbf{var} f \bullet \left\langle \begin{array}{c} BidAlice ; BidBob ; f := Alice \\ \Box BidBob ; BidAlice ; f := Bob \end{array} \right\rangle ; CP4$$

Back-propagating *CP*4 through *Submit*₄ yields **true**, as did *Submit*₂. Hence, we conclude that *Submit*₄ is a feasibility-preserving refinement of *Submit*₂.

7.5.2 Deterministic Declaration

Suppose the auctioneer is instructed to declare Bob as the winner only if his bid is strictly greater than Alice's. This instruction clarifies the behaviour of *Declare*₀ by refinement:

 $Declare_1 \triangleq a \ge b \& winner! Alice \rightarrow Stop \Box a < b \& winner! Bob \rightarrow Stop$

By observing a *winner* event, the contestants can infer whether $a \ge b$ or a < b, as is reflected by:

bw (Declare₁, true) =
$$(a \ge b \Rightarrow \widetilde{a} \ge \widetilde{b}) \land (a < b \Rightarrow \widetilde{a} < \widetilde{b})$$

This obligation is neither weaker nor stronger than *bw*0. Back-propagating it through the CA *CP*1 in *AuctionLC* gives the obligation:

$$(a > b \Rightarrow \widetilde{a} = \widetilde{b}) \land (a = b \Rightarrow \widetilde{a} \ge \widetilde{b}) \land (a < b \Rightarrow \widetilde{a} < \widetilde{b})$$

7 Case Study

Now, back-propagating this obligation through $Submit_2$ yields **true**: the extra information that $Declare_1$ reveals does not violate **CP1** or **CP4**. Hence, the refinement of $Declare_0$ into $Declare_1$ is feasibility-preserving.

7.5.3 Data Refinement

One possible data refinement on *AuctionLC* would be to replace the state *ASt* with three new variables, linked by the simulation relation:

$$R = \begin{pmatrix} c = \max(a, b) \land (e = 1 \lhd a = b \triangleright e = 0) \\ \land (a > b \Rightarrow w = Alice) \land (a < b \Rightarrow w = Bob) \end{pmatrix}$$

Here, c stores the value of the larger bid, while w records the identity of the contestant who made that bid. The binary-valued variable e records whether the bids are equal. It serves no purpose with respect to the functionality of the specification.

By applying this data refinement to *AuctionLC*, we can derive the process presented in Figure 7.4. Notice how this data refinement entails the restructuring of *CP*1, by projecting its embedded obligation through the simulation:

$$\forall a, b, \widetilde{a}, \widetilde{b} \bullet \mathsf{U}(R) \Rightarrow (a > b \Rightarrow \widetilde{a} \le \widetilde{b})$$

= $e = 0 \land w = Alice \Rightarrow \widetilde{e} = 1 \lor \widetilde{w} = Bob$ [pred calc]

Here, the purpose of *e* becomes apparent. It is needed to capture the exact concrete fog states that *R* maps to $\tilde{a} \leq \tilde{b}$ in the abstract space, without strengthening the CA.

An implicit design decision captured by the $BidA_D$ and $BidB_D$ actions (see Figure 7.4) is that, if Alice and Bob's bids are equal, then the winner is the contestant who bid first. However, this decision harbours yet another subtle security flaw, with respect to **CP4**. Suppose that Alice bids the minimum amount £1 but is then declared the winner. Alice can infer that Bob's bid must also be £1 and, by extension, that she must have bid first, in violation of $CP4_D$. $BDR ::= Alice \mid Bob$ channel bidAlice, bidBob : \mathbb{N}_{1} channel winner : BDR channelset $A \triangleq \{\mid bidAlice, winner \mid\}$ process $AuctionLCD \triangleq$ begin state $CSt \triangleq [c : \mathbb{N}, e : \{0, 1\}, w : BDR]$ $CP1_{D} \triangleq \langle e = 0 \land w = Alice \Rightarrow \tilde{e} = 1 \lor \tilde{w} = Bob \rangle$ $CP4_{D} \triangleq \langle \tilde{f} = opp(f) \rangle$ $Init_{D} \triangleq \langle [CSt' \mid c' = 0 \land e' = 1 \land w' = Alice] \rangle$ $BidA_{D} \triangleq bidAlice?n \rightarrow \begin{bmatrix} \Delta CSt \\ n? : \mathbb{N}_{1} \end{vmatrix} \begin{vmatrix} c' = \max(c, n?) \\ e' = 1 \lhd c = n? \rhd e' = 0 \\ w' = Alice \lhd c > n? \rhd w' = w \end{bmatrix}$ $BidB_{D} \triangleq bidBob?n \rightarrow \begin{bmatrix} \Delta CSt \\ n? : \mathbb{N}_{1} \end{vmatrix} \begin{vmatrix} c' = \max(c, n?) \\ e' = 1 \lhd c = n? \rhd e' = 0 \\ w' = Bob \lhd c > n? \rhd w' = w \end{bmatrix}$ $SubmitAux_{D} \triangleq \langle f = Alice \& BidA_{D}; BidB_{D} \Box f = Bob \& BidB_{D}; BidA_{D} \rangle$ $Submit_{D} \triangleq varf \bullet \langle f := Alice \sqcap f := Bob \rangle; SubmitAux_{D}; CP4_{D}$ $Declare_{D} \triangleq \langle winner!w \rightarrow Stop \rangle$ end

Figure 7.4: Lifted and data-refined model of auctioneer.

7 Case Study

Ignoring $CP1_D$ for simplicity, we can confirm the presence of this security flaw in the usual fashion. Back-propagating the obligation:

bw
$$(CP4_D, \mathbf{bw} (Declare_D, \mathbf{true})) = w = \widetilde{w} \wedge \widetilde{f} = opp(f)$$

through $\langle A : Submit_D \rangle$ and simplifying yields the following:

$$\forall a, b : \mathbb{N}_1 \bullet \exists \, \widetilde{b} : \mathbb{N}_1 \bullet (a \ge b \land a > \widetilde{b}) \lor (a > b \land a \ge \widetilde{b})$$

When a = 1 and b = 1, no values of $\tilde{b} : \mathbb{N}_1$ can satisfy this condition. Hence, we have a contradiction and can therefore establish the process is insecure with respect to **CP4**.

7.6 Discussion

This case study has demonstrated that software development is made more complicated by the presence of confidentiality properties. Even a superficially simple process can harbour surprisingly deep subtleties that must be probed to acquire robust assurances of security.

We have used back-propagation as a tool to detect inconsistencies between the functionality and confidentiality requirements placed on the auctioneer. This analysis has revealed subtle security issues: in particular, the conflict between **CP3** and the *Auction* process (described in Subsection 7.4.3) may not be immediately obvious.

The verification procedure is tractable for a small case study, but the lack of tool support would obstruct its deployment in a larger case study. As we have seen, calculating back-propagations through *SubmitAux* and *Declare*₀ by hand is tricky and requires extra creativity in places. Moreover, the potential for introducing errors can only be magnified for more complex blocks.

8 Conclusions

8.1 Summary of the Thesis

This thesis presents a formal platform for specifying, verifying and developing software systems with respect to both functionality and confidentiality properties. This platform leads towards a methodology for constructing software that is "secure by design".

The platform's backbone is the Unifying Theories of Programming. We have created a generic approach for extending UTP theories to model users' observational abilities and, by extension, their inferences about a system's behaviour. This approach enables confidentiality properties to be specified alongside a functional model of a system. In this way, we reconcile the opposing concerns of functionality and confidentiality in a clean and appealing fashion.

We focus on making the approach practical by specialising it to *Circus*. We extend the *Circus* notation with facilities to embed confidentiality properties within a functional process specification. This work engenders a language that makes the specification of confidentiality properties accessible to formal methods practitioners who are non-specialists in the field of information security. These specifications can be subjected to rigorous consistency checks, to detect potential sources of insecurity.

Much of the work presented in this thesis draws inspiration from existing ideas on information flow security in the academic literature. The originality of the thesis lies in how we have distilled these ideas to their essence by characterising them in the UTP, enabling us to synthesise new connections between these ideas and the field of formal methods.

The work reported in this thesis could be ported to other formal

8 Conclusions

methods with a UTP semantics, or other languages in the *Circus* family. However, it is unclear how our encoding of confidentiality properties could be replicated for languages without a UTP-like semantics.

We now review the main contributions of the thesis, discuss their significance and highlight areas for improvement.

8.1.1 Modelling User Interactions

Chapter 3 proposes a framework for exploring specifications of systems from the perspective of their users. The key component of this framework is the view, which formalises a user's interface to a system. Views can be applied to translate a model of system behaviours into a model of user interactions, which can reveal defects in a system's functionality. Conversely, views can be used to calculate a user's inferences about a system's behaviour.

This framework is generic in the UTP theory framing a system model. We have applied the framework to the UTP theories of designs and reactive processes; and, in the latter case, we have demonstrated that the framework generalises Roscoe's notion of lazy abstraction.

8.1.2 Encoding Confidentiality

Chapter 4 extends the framework of Chapter 3 to formalise information flow from systems to adversarial users. Parameterised by an indistinguishability relation between system behaviours — constructed from a view — this encoding of information flow is generic across UTP theories.

We harness the lifted space of indistinguishable behaviours to capture notions of confidentiality. Obligations are closure conditions over the lifted space, which specify an upper bound on an adversary's inferences about a system's behaviour. We realise a specification of a secure system by conjoining obligations with a lifted model of the system's functionality, to eliminate any system behaviours that violate an obligation.

The fusion of functionality and confidentiality presented in this chapter is novel. The application of miracles to denote conflicting specifications of functionality and confidentiality is especially appealing, yet is only possible because miracles are valid specification artefacts in the UTP. Moreover, miracles solve the long-standing "refinement paradox" in a clean fashion: refinement can make system designs miraculous, but it cannot make them insecure.

Obligations capture a more liberal notion of confidentiality than noninterference. They are based on the premise that software engineers would find it advantageous to tailor the specification of confidentiality properties to the system domain, as opposed to aiming to satisfy rigid noninterference properties. Even though noninterference is arguably too strict for many classes of systems, our premise can only be vindicated with empirical evidence.

8.1.3 Integrating Confidentiality with Circus

Chapter 5 specialises the confidentiality encoding from Chapter 4 to cover *Circus* processes. Targeting *Circus* is a pragmatic choice, owing to its UTP semantics, its unified treatment of state and behavioural aspects of systems, and its comprehensive refinement strategy.

At the semantic level, the integration is accomplished by lifting *Circus* constructs to represent information flow to the adversary, using the techniques presented in Chapter 4. This approach complicates the semantics of these constructs, but in return, it affords reasoning about information flow in a compositional fashion.

Confidentiality annotations generalise *Circus* coercions to encode confidentiality properties over processes in terms of their intermediate state. The semantics of confidentiality annotations makes the insecure behaviours of a process miraculous, thereby preventing the process from leaking secret information to an adversary. A process specification can only be strengthened by inserting confidentiality annotations, so they can be introduced at any stage of a process's development.

Much of the expressive power of confidentiality annotations is exposed by composing them with lifted *Circus* constructs, to specify confidenti-

8 Conclusions

ality properties over the events a process performs. Taking this idea to an extreme, confidentiality annotations can be formulated as separate processes, which can be superposed with a functional process to create a secure specification of a system.

8.1.4 Verifying Security

Chapter 6 focuses on the development of processes with confidentiality annotations by stepwise refinement, from abstract specifications all the way to the implementation level.

While confidentiality annotations ensure confidentiality is preserved by refinement, ill-judged specifications or imprudent refinement steps can make a process infeasible. This effect motivates us to extend the *Circus* refinement strategy with special measures to ensure that each process design in the refinement chain is implementable.

This chapter presents backwards propagation, a technique for systematically and tractably over-approximating an adversary's inferences about the process state at each step of the process's execution. These derived inferences can be checked against the confidentiality annotations within the process body, to verify those annotations are respected throughout the process's execution.

Backwards propagation also enables software engineers to verify that refinement steps are feasibility-preserving. Rather than verifying each process design in isolation, we can record the adversary's inferences by inserting confidentiality annotations at each point of the process. Then, these annotations can be used to verify a localised refinement step, without incurring the expense of re-verifying the whole process.

8.2 Directions for Future Research

There is still much work to be done in scaling up the capabilities of the formal development platform beyond a proof of concept. We now review a selection of topics that merit further research.

8.2.1 Tool Support

Effective machine assistance is often cited as important for encouraging the industrial adoption of formal methods (Saiedian, 1996; Woodcock et al., 2009). The lack of any dedicated tool support for our platform is arguably the main impediment to its deployment.

Some possible goals of tool support for the platform would include the ability to (semi)-automatically verify that a specification of functionality and confidentiality properties is consistent; and to determine whether this consistency is maintained by refinement steps.

The design of the tools may be based on techniques in this thesis (such as backwards propagation) or may involve other techniques devised expressly with automation in mind. Work by Černý (2009) indicates that verifying a restricted space of terminating programs against a class of confidentiality properties (much like our own) can be reduced to a decision problem that can be tackled using a SMT solver. As with all automated verification techniques, restrictions on the program space are necessary to avoid the halting problem.

One pathway towards implementing effective tool support would be to adapt existing *Circus* tools, such as CRefine (Oliveira et al., 2008). Extensions to these tools would include facilities to specify confidentiality annotations; adaptations of the refinement rules to discriminate between unlifted and lifted constructs; and potentially an implementation of the backwards propagation calculus. Alternatively, if the platform is ported to a formalism such as CSP||B (Schneider and Treharne, 2005) or Event-B (Abrial, 2010), tool support for verifying confidentiality could potentially be implemented in the form of plug-ins for their respective toolsets.

8.2.2 Mechanising the Theory

All theorems, lemmas and laws presented in this thesis have been justified by hand proof. Nevertheless, it would be expedient to verify their correctness by encoding their proofs in a theorem prover. The *Saoithín* proof checker for UTP (Butterfield, 2010), or the *Isabelle/Circus* environment

8 Conclusions

(Feliachi et al., 2011), could be applied to this task.

8.2.3 Case Studies

The purpose of carrying out case studies is to evaluate the platform's suitability for the development of secure software. Owing to time constraints and the lack of tool support, the case study and examples presented in this thesis are small in scope. Larger case studies are needed to evaluate the platform comprehensively.

Two pertinent topics for larger case study developments with the platform are outlined thus:

- **Validating access controls** A specification of an access control mechanism could be augmented with confidentiality annotations, in order to specify explicitly the confidentiality policy that (it is claimed) is modelled by that mechanism. By combining this augmented specification with a system design, a specification could be derived that cannot leak information in contravention of the confidentiality policy, regardless of whether the mechanism captures the policy correctly. An insecure implementation along the lines of McLean's "System Z" (see Section 4.1) would then be unattainable.
- **Security protocol development** A security protocol could be formalised in lifted *Circus*, alongside a Dolev-Yao model of an adversary's abilities to tamper with the protocol (Dolev and Yao, 1983). The secrecy properties of the protocol would be specified by inserting confidentiality annotations into the protocol design. Once these secrecy properties are verified to hold, it would then be appropriate to verify that tactics for refining the protocol design — for instance, by fixing the order of tokens within messages — do not introduce insecurities into the concrete protocol design.

A complete delivery of these case studies would provide convincing evidence that, with suitable tool assistance, the platform could potentially be scaled up to accommodate real-world software engineering projects.

8.2.4 A Confidentiality Policy Language

Our encoding of confidentiality properties — as relations between two observational spaces of a system — is very general, but it is perhaps not readily comprehensible by software engineers. Moreover, interweaving confidentiality annotations with lifted *Circus* actions (as described in Section 5.6) is a somewhat esoteric means for indirectly expressing confidentiality properties over events. These tricks may confuse readers of a specification, especially if the meaning of the confidentiality annotations is not documented.

It would be worthwhile to define a dedicated language for expressing confidentiality policies in practice. Shaped by the experience grown through case studies, this language would feature idioms for expressing common confidentiality properties over the state and behaviour of processes in a clear and intuitive fashion. The language could be designed with automated analysis in mind: restricting the expressiveness of the idioms would simplify the task of building effective tool support to check those idioms.

Some degree of translation would be needed to integrate this policy specification language into our development platform (in its current form). This translation could be implemented as a tool, taking a confidentiality policy and a *Circus* process as input, and generating a process with embedded confidentiality annotations capturing the policy. The annotated process could then be developed using the platform.

8.2.5 Probabilistic Information Flow

The model of information flow considered in this thesis is *possibilistic*, rather than probabilistic: it deals only with what an adversary can infer *with certainty* about a system. The inherent limitations of possibilistic information flow models are discussed in Section 4.7.

Bresciani and Butterfield (2011) have developed a probabilistic variant of the UTP. If their work is extended to give *Circus* a probabilistic semantics, our platform could potentially be adapted to regulate in-

8 Conclusions

formation flow in a probabilistic setting. An analogue of confidentiality annotations would be embedded in specifications to prevent an adversary from deducing facts about the process state with high probability.

A further extension of this work would involve quantifying the flow of information from systems to users, in terms of Shannon's information theory. In this setting, confidentiality properties would then be defined as an upper bound on the *mutual information* between secret data and outputs to the adversary. There is potential for automatically testing (with statistical significance) whether a program satisfies these properties are satisfied: recent work by Chothia and Guha (2011) has shown how mutual information can be calculated from trial runs of a program alone.

8.2.6 Other Notions of Security

This thesis has focused on a narrow interpretation of security: specifically, limiting the flow of (qualitative) information to a specified adversary. However, information security also encompasses the *authenticity*, the *integrity* and the *availability* of information, as well as its confidentiality.

- **Integrity properties** are concerned with guarding against unauthorised modification of data. Jacob (1987) interprets integrity properties as a dual of confidentiality properties, in the sense that they prevent information from low-level users from flowing to (and corrupting) high-level users. Such properties can be formulated within our platform, simply by inverting the low-level and high-level users.
- **Availability properties** require that information is available to authorised users when it is needed. Techniques for ensuring availability include clustering (to provide fault tolerance), bandwidth management (to mitigate denial-of-service attacks) and failure recovery.

Confidentiality properties have received far more attention in the literature than integrity or availability (Foley, 2005). Indeed, in many classes of systems, safeguarding the integrity and availability of information has a higher priority than maintaining the confidentiality of that information (Clark and Wilson, 1987).

It is interesting to speculate about the potential for integrating specifications of desired security properties other than confidentiality into the software development process. Nevertheless, it may not be appropriate to characterise these properties in terms of information flow, so a new semantic foundation for them may be needed instead.

8.3 Coda

In their book, Hoare and He (1998) set out an agenda to unite the study of different programming paradigms, across multiple levels of abstraction. The work presented in this thesis hints at the potential of extending the UTP agenda to encompass topics in Computing Science that lie outside the standard models of software correctness. In so doing, it underscores the significance of the UTP project for building a firm foundation to address the contemporary challenges in software engineering.

Proofs

Proof of Lemma 3.4

$\mathbf{VH1}\left(V\right) = V$		
=	$[\operatorname{VH1}(V) \Rightarrow V] \land [V \Rightarrow \operatorname{VH1}(V)]$	[equivalence]
=	$\left[\left(\left(\exists x_V, x'_V \bullet V\right) \Rightarrow V\right) \Rightarrow V\right] \land \left[V \Rightarrow \left(\left(\forall x_V, x'_V \bullet V\right) \Rightarrow V\right) \Rightarrow V\right] \land \left[V \Rightarrow \left((\forall x_V, x'_V \bullet V) \Rightarrow V\right) \Rightarrow V\right] \land V \Rightarrow V$	$\exists x_V, x'_V \bullet V) \Rightarrow V)] [def VH1]$
=	$\left[\left((\exists x_V, x'_V \bullet V) \land \neg V\right) \lor V\right] \land \left[\left(V \land \exists x_V, x'_V \bullet V\right) \land \neg V\right] \land \left[\left(V \land \exists x_V, x'_V \bullet V\right) \land \neg V\right] \land \left[\left(V \land \exists x_V, x'_V \bullet V\right) \land \neg V\right] \land \left[\left(V \land \exists x_V, x'_V \bullet V\right) \land \neg V\right] \land \left[\left(V \land \exists x_V, x'_V \bullet V\right) \land \neg V\right] \land \left[\left(V \land \exists x_V, x'_V \bullet V\right) \land \neg V\right] \land \left[\left(V \land \exists x_V, x'_V \bullet V\right) \land \neg V\right] \land \left[\left(V \land \exists x_V, x'_V \bullet V\right) \land \neg V\right] \land \left[\left(V \land \exists x_V, x'_V \bullet V\right) \land \neg V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \neg V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \neg V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \neg V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \neg V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \forall x_V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \nabla V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \nabla V \land \nabla V\right) \land \nabla V\right] \land \left[\left(V \land \nabla V \land \nabla V\right) \land \nabla V\right] \land \nabla V \land $	$x_V, x'_V \bullet V) \Rightarrow V$ [prop calc]
=	$\left[\left((\exists x_V, x_V' \bullet V) \land \neg V\right) \lor V\right]$	[tautology]
=	$\left[\left((\exists x_V, x_V' \bullet V) \lor V\right) \land (V \lor \neg V)\right]$	[distributivity]
=	$\left[\left(\exists x_V, x_V' \bullet V\right) \lor V\right]$	[excluded middle]
=	$\left[\exists x_V, x'_V \bullet V\right]$	[lower bound]

Proof of Law 3.8

$VH1 \circ VH1(V)$		
=	$(\exists x_V, x'_V \bullet (\exists x_V, x'_V \bullet V) \Rightarrow V) \Rightarrow ((\exists x_V, x'_V \bullet V) \Rightarrow V)$	[def VH1]
=	$true \Rightarrow ((\exists x_V, x'_V \bullet V) \Rightarrow V)$	[pred calc]
=	$(\exists x_V, x'_V \bullet V) \Rightarrow V$	[prop calc]
=	VH1(V)	[def VH1]

Proof of Law 3.10

$$VH1 \circ VH2(V)$$

$$= (\exists x_V, x'_V \bullet (\exists x', x'_V \bullet V)) \Rightarrow (\exists x', x'_V \bullet V) \qquad [def VH1, VH2]$$

Proofs

$$= (\exists x', x'_{V} \bullet (\exists x_{V}, x'_{V} \bullet V)) \Rightarrow (\exists x', x'_{V} \bullet V)$$
 [pred calc]

$$= (\forall x', x'_{V} \bullet \neg (\exists x_{V}, x'_{V} \bullet V)) \lor (\exists x', x'_{V} \bullet V)$$
 [pred calc]

$$\Rightarrow (\exists x', x'_{V} \bullet \neg (\exists x_{V}, x'_{V} \bullet V)) \lor (\exists x', x'_{V} \bullet V)$$
 [pred calc]

$$= \exists x', x'_{V} \bullet \neg (\exists x_{V}, x'_{V} \bullet V) \lor V$$
 [pred calc]

$$= \forall H2 \circ \forall H1(V)$$
 [def $\forall H1, \forall H2$]

Proof of Lemma 3.13

Assume that $V_1 = \mathbf{VH}(V_1)$ and $V_2 = \mathbf{VH}(V_2)$. To prove $V_1 \wedge V_2$ is **VH**-healthy, we prove that $V_1 \wedge V_2$ is **VH1**-healthy and **VH2**-healthy separately.

V_1 /	$\setminus V_2$	
=	$VH1(V_1)\wedgeVH1(V_2)$	[assumption]
=	$(\exists x_1, x'_1 \bullet V_1) \Rightarrow V_1) \land (\exists x_2, x'_2 \bullet V_2) \Rightarrow V_2)$	[def VH1]
\Rightarrow	$((\exists x_1, x'_1 \bullet V_1) \land (\exists x_2, x'_2 \bullet V_2)) \Rightarrow V_1 \land V_2$	[prop calc]
=	$(\exists x_1, x_1', x_2, x_2' \bullet V_1 \land V_2) \Rightarrow V_1 \land V_2$	$[V_1, V_2 \text{ disjoint}]$
=	$\operatorname{VH1}(V_1 \wedge V_2)$	[def VH1]

V_1 /	$\setminus V_2$	
=	$VH2(V_1)\wedgeVH2(V_2)$	[assumption]
=	$(\exists x', x'_1 \bullet V_1) \land (\exists x', x'_2 \bullet V_2)$	[def VH2]
=	$\exists x', x_1', x_2' \bullet (\exists x', x_1' \bullet V_1) \land (\exists x', x_2' \bullet V_2)$	$[V_1, V_2 \text{ disjoint}]$
=	$VH2(VH2(V_1)\landVH2(V_2))$	[def VH2]
=	$\mathbf{VH2}(V_1 \wedge V_2)$	[assumption]

Proof of Law 3.19

 $L(V, P_1 \vee P_2)$

=	$\exists x, x' \bullet \Delta V \land (P_1 \lor P_2)$	[def L]
=	$(\exists x, x' \bullet \Delta V \land P_1) \lor (\exists x, x' \bullet \Delta V \land P_2)$	[pred calc]
=	$L(V,P_1) \vee L(V,P_2)$	[def L]

Proof of Law 3.20

P_1	$P_1 \sqsubseteq P_2$		
\Rightarrow	$(\Delta V \wedge P_1) \sqsubseteq (\Delta V \wedge P_2)$	[^-monotonicity]	
\Rightarrow	$(\exists x, x' \bullet \Delta V \land P_1) \sqsubseteq (\exists x, x' \bullet \Delta V \land P_2)$	[∃-monotonicity]	
=	$L(V,P_1) \sqsubseteq L(V,P_2)$	[def L]	

Proof of Law 3.21

L(($L\left(\left(\wedge i:1n\bullet V_{i}\right),P\right)$		
=	$\exists x, x' \bullet \Delta(\bigwedge i : 1n \bullet V_i) \land P$	[def L]	
\Rightarrow	$\wedge i: 1n \bullet \exists x, x' \bullet \Delta V_i \wedge P$	[pred calc]	
=	$\bigwedge i: 1n \bullet L(V_i, P)$	[def L]	

Proof of Lemma 3.24

$\mathbf{G}\left(V,U ight)$		
=	$\forall x_V, x_V' \bullet \Delta V \Rightarrow U$	[def G]
=	$\neg \exists x_V, x'_V \bullet \neg (\Delta V \Rightarrow U)$	[duality]
=	$\neg \exists x_V, x_V' \bullet \Delta V \land \neg U$	[prop calc]
=	$\exists x_V, x_V' \bullet \Delta V \wedge U$	[V functional (*)]
=	$\neg \forall x_V, x'_V \bullet \Delta V \Rightarrow \neg U$	[duality]
=	$\neg \mathbf{G}(V, \neg U)$	[def G]

Proofs

The step marked (*) is justified as follows. Since *V* is a total function from behaviours to interactions (i.e. **VH3**-healthy), each behaviour that *V* maps to a *U* interaction is *not* a behaviour that *V* maps to a \neg *U* interaction, and vice versa.

Proof of Theorem 3.25

$U \sqsubseteq L(V, P)$		
=	$[(\exists x, x' \bullet \Delta V \land P) \Rightarrow U]$	[def ⊑, L]
=	$[(\forall x, x' \bullet \neg \Delta V \lor \neg P) \lor U]$	[prop calc]
=	$[\neg \Delta V \lor \neg P \lor U]$	$[x, x' \text{ not in } \alpha U]$
=	$\left[\neg P \lor \left(\forall x_V, x'_V \bullet \neg \Delta V \lor U\right)\right]$	$[x_V, x'_V \text{ not in } \alpha P]$
=	$\left[P \Rightarrow (\forall x_V, x'_V \bullet \Delta V \Rightarrow U)\right]$	[prop calc]
=	$\mathbf{G}(V,U) \sqsubseteq P$	[def ⊑ , G]

Proof of Lemma 3.35

$L(V, Pre \vdash Post)$ $\exists x, x' \bullet \Delta V \land (ok \land Pre \Rightarrow ok' \land Post)$ [def **L**, ⊢] = $\exists x, x' \bullet \Delta V \land (\neg ok \lor \neg Pre \lor (ok' \land Post))$ [prop calc] = $(\exists x, x' \bullet \Delta V \land (\neg ok \lor \neg Pre)) \lor (\exists x, x' \bullet \Delta V \land ok' \land Post)$ = [distributivity] $(\neg ok_V \lor \exists x, x' \bullet \Delta V \land \neg Pre) \lor (ok'_V \land \exists x, x' \bullet \Delta V \land Post)$ = [V is VHD] $(\neg ok_V \lor \mathsf{L}(V, \neg Pre)) \lor (ok'_V \land \mathsf{L}(V, Post))$ [def L] = $\neg (ok_V \land \neg \mathsf{L}(V, \neg Pre)) \lor (ok'_V \land \mathsf{L}(V, Post))$ [de Morgan] = $(ok_V \land \neg \mathsf{L}(V, \neg Pre)) \Rightarrow (ok'_V \land \mathsf{L}(V, Post))$ [prop calc] = $\neg \mathbf{L}(V, \neg Pre) \vdash_{V} \mathbf{L}(V, Post)$ $[\operatorname{def} \vdash_V]$ =

Assuming that *Pre* is a condition over the *x* variables alone:

=	$(\neg \exists x, x' \bullet \Delta V \land \neg Pre) \vdash_V L(V, Post)$	[def L]

- $= (\forall x, x' \bullet \Delta V \Rightarrow Pre) \vdash_V \mathsf{L}(V, Post)$ [pred calc]
- $= (\forall x \bullet V \Rightarrow Pre) \vdash_V \mathsf{L}(V, Post)$ [assumption]

Proof of Lemma 3.47

Proof of Theorem 3.51

For a divergence-free CSP process *P*, Roscoe (1997) defines the failure set of Lazy (W, P) in the form:

$$\{(s \upharpoonright \mathcal{W}, r) \mid (s, r \cap \mathcal{W}) \in failures(P)\}$$

The failure set of a reactive design P_R (capturing the semantics of P) is defined by Cavalcanti and Woodcock (2006) as:

 $failures(P_R) =$

Proofs

$$\{(tr' - tr, ref') \mid ok \land \neg wait \land P_R \land ok'\}$$

$$\cup \{((tr' - tr) \land \langle \checkmark \rangle, ref') \mid ok \land \neg wait \land P_R \land ok' \land \neg wait'\}$$

$$\cup \{((tr' - tr) \land \langle \checkmark \rangle, ref' \cup \{\checkmark\}) \mid ok \land \neg wait \land P_R \land ok' \land \neg wait'\}$$

Since our reactive interface allows a user to monitor termination (as signalled by $wait'_V$), we need to extend W to include the \checkmark event. Let:

$$\mathcal{W}^\checkmark = \mathcal{W} \cup \{\checkmark\}$$

We now study the W^{\checkmark} -projection of the first set of *failures*(P_R):

$$\{ (tr'_{V} - tr_{V}, ref'_{V}) \mid \mathsf{LR} (\mathcal{W}^{\checkmark}, ok \land \neg wait \land P_{R} \land ok') \}$$

$$= \{ (tr'_{V} - tr_{V}, ref'_{V}) \mid \exists x, x' \bullet \mathcal{R} (\mathcal{W}^{\checkmark}) \land ok \land \neg wait \land P_{R} \land ok' \}$$

$$= \{ ((tr' - tr) \upharpoonright \mathcal{W}, ref' \cap \mathcal{W}) \mid \mathcal{R} (\mathcal{W}^{\checkmark}) \land ok \land \neg wait \land P_{R} \land ok' \}$$

The last proof step is justified because $\mathcal{R}(\mathcal{W}^{\checkmark})$ implies:

$$tr'_V - tr_V = (tr' - tr) \upharpoonright \mathcal{W}^{\checkmark}$$
 and $ref'_V \cap \mathcal{W}^{\checkmark} \subseteq ref'$

and refusal sets are downwards-closed.

The projections of the other two sets of $failures(P_R)$ are similar. Hence, the sets $failures(LR(W, P_R))$ and failures(Lazy(W, P)) are equivalent. We conclude that Lazy (W, P) yields a process isomorphic to $LR(W, P_R)$.

Proof of Theorem 4.4

P_1	$= P_2$	
=	$(P_1 \sqsubseteq P_2) \land (\widetilde{P_1} \sqsubseteq \widetilde{P_2})$	[property of \sqsubseteq]
=	$\left[(P_2 \Rightarrow P_1) \land (\widetilde{P_2} \Rightarrow \widetilde{P_1}) \right]$	$[def \sqsubseteq]$
=	$\left[\left(P_2 \wedge \widetilde{P_2}\right) \Rightarrow \left(P_1 \wedge \widetilde{P_1}\right)\right]$	[disjoint alphabets]
=	$(P_1 \wedge \widetilde{P_1}) \sqsubseteq (P_2 \wedge \widetilde{P_2})$	$[def \sqsubseteq]$
=	$\mathbf{U}\left(P_{1}\right)\sqsubseteq\mathbf{U}\left(P_{2}\right)$	[def U]

Proof of Lemma 4.6

D (U	$\mathbf{J}(P))$	
=	$\exists \widetilde{x}, \widetilde{x'} \bullet P \land \widetilde{P}$	[def U , D]
=	$P \wedge \exists \widetilde{x}, \widetilde{x'} \bullet \widetilde{P}$	$[\widetilde{x}, \widetilde{x'} \text{ free in } P]$
=	$P \wedge \mathbf{true}$	[pred calc]
=	Р	[unit of conjunction]

Proof of Law 4.9

$\mathbf{U}\left(P_{1}\vee P_{2}\right)$			
=	$(P_1 \lor P_2) \land (\widetilde{P_1} \lor \widetilde{P_2})$	[def U]	
\Leftarrow	$(P_1 \wedge \widetilde{P_1}) \vee (P_2 \wedge \widetilde{P_2})$	[prop calc]	
=	$\mathbf{U}\left(P_{1}\right)\vee\mathbf{U}\left(P_{2}\right)$	[def U]	

Proof of Law 4.10

$\mathbf{D}\left(Q_{1} ight)\wedge\mathbf{D}\left(Q_{2} ight)$		
=	$(\exists \widetilde{x}, \widetilde{x'} \bullet Q_1) \land (\exists \widetilde{x}, \widetilde{x'} \bullet Q_2)$	[def D]
\Leftarrow	$\exists \widetilde{x}, \widetilde{x'} \bullet Q_1 \land Q_2$	[prop calc]
=	$\mathbf{D}\left(Q_1 \wedge Q_2 ight)$	[def D]

Proof of Law 4.13

$\neg \mathbf{U}(P)$		
=	$\neg (P \land \widetilde{P})$	[def U]
=	$\neg P \lor \neg \widetilde{P}$	[de Morgan]
\Leftarrow	$\neg P \land \neg \widetilde{P}$	[prop calc]
=	$\mathbf{U}(\neg P)$	[def U]

199

Proofs

Proof of Law 4.14

$$D (\neg Q)$$

$$= \exists \tilde{x}, \tilde{x'} \bullet \neg Q \qquad [def D]$$

$$\Leftarrow \quad \forall \tilde{x}, \tilde{x'} \bullet \neg Q \qquad [pred calc]$$

$$= \neg \exists \tilde{x}, \tilde{x'} \bullet Q \qquad [pred calc]$$

$$= \neg D (Q) \qquad [def D]$$

Proof of Lemma 4.26

$$\begin{aligned} \theta_1 &\sqsubseteq \theta_2 \land P \propto_L \theta_2 \\ &= \theta_1 \sqsubseteq \theta_2 \land \left[P \Rightarrow \exists \widetilde{x}, \widetilde{x'} \bullet \mathsf{UI}(L, P) \land \theta_2 \right] & [\det \alpha_L] \\ &\Rightarrow \theta_1 \sqsubseteq \theta_2 \land \left[P \Rightarrow \exists \widetilde{x}, \widetilde{x'} \bullet \mathsf{UI}(L, P) \land \theta_1 \right] & [\theta_1 \sqsubseteq \theta_2] \\ &= \theta_1 \sqsubseteq \theta_2 \land P \propto_L \theta_1 & [\det \alpha_L] \end{aligned}$$

Proof of Lemma 4.40

$$\begin{array}{ll} Q_1 \sqsubseteq Q_2 \\ \Rightarrow & \mathsf{D}(Q_1) \sqsubseteq \mathsf{D}(Q_2) & [Law \ 4.12] \\ = & \overbrace{\mathsf{D}(Q_1)} \sqsubseteq \overbrace{\mathsf{D}(Q_2)} & [renaming] \\ \Rightarrow & Q_1 \land \overbrace{\mathsf{D}(Q_1)} \sqsubseteq Q_1 \land \overbrace{\mathsf{D}(Q_2)} & [monotonicity \ of \ conjunction] \\ \Rightarrow & Q_1 \land \overbrace{\mathsf{D}(Q_1)} \sqsubseteq Q_2 \land \overbrace{\mathsf{D}(Q_2)} & [because \ Q_1 \sqsubseteq Q_2] \\ = & \mathsf{C}(Q_1) \sqsubseteq \mathsf{C}(Q_2) & [def \ C] \end{array}$$

Proof of Lemma 4.43

$\mathbf{CC}\left(\mathbf{CC}\left(Q\right) ight)$				
=	$\mu X \bullet \mathbf{C} \left(\left(\mu X \bullet \mathbf{C} \left(Q \land X \right) \right) \land X \right)$	[def CC]		
=	$\mu X \bullet \mathbf{C} \left(Q \land X \right)$	[property of μX]		
=	$\mathbf{CC}(Q)$	[def CC]		

Proof of Theorem 4.44

Let $Q = UI(L, P) \land \theta$ in: $D(CC(Q)) \propto_L \theta$ $= [D(CC(Q)) \Rightarrow D(UI(L, D(CC(Q))) \land \theta)] \qquad [def \propto_L]$ $= [D(CC(Q)) \Rightarrow D(UI(L, D(CC(Q))))] \qquad [\theta \sqsubseteq CC(Q)]$ $= [D(CC(Q)) \Rightarrow D(CC(Q))] \qquad [Lemma 4.6]$ $= true \qquad [tautology]$

Proof of Lemma 5.3

This proof relies on two special properties of $\mathcal{I}(\mathcal{L})$. First, $\mathcal{I}(\mathcal{L})$ is tautologous if the fog space is unconstrained:

 $\exists \widetilde{x}, \widetilde{x'} \bullet \mathcal{I}(\mathcal{L}) = \text{true}$

Second, renaming the observable variables of $\mathcal{I}(\mathcal{L})$ to the fog variables results in a tautology:

$$\widetilde{\mathcal{I}(\mathcal{L})} = \widetilde{ok} = \widetilde{ok} \wedge \widetilde{ok'} = \widetilde{ok'} \wedge \ldots =$$
true

$$P_{1} \sqsubseteq P_{2}$$

$$= P_{1} \land \mathcal{I}(\mathcal{L}) \sqsubseteq P_{2} \land \mathcal{I}(\mathcal{L})$$
[property of $\mathcal{I}(\mathcal{L})$]
$$= \mathbf{U}(P_{1} \land \mathcal{I}(\mathcal{L})) \sqsubseteq \mathbf{U}(P_{2} \land \mathcal{I}(\mathcal{L}))$$
[Theorem 4.4]

Proofs

$$= P_1 \land \mathcal{I}(\mathcal{L}) \land \widetilde{P_1} \land \widetilde{\mathcal{I}(\mathcal{L})} \sqsubseteq P_2 \land \mathcal{I}(\mathcal{L}) \land \widetilde{P_2} \land \widetilde{\mathcal{I}(\mathcal{L})} \qquad [def U]$$

$$= P_1 \land \mathcal{I}(\mathcal{L}) \land \widetilde{P_1} \sqsubseteq P_2 \land \mathcal{I}(\mathcal{L}) \land \widetilde{P_2} \qquad [property of \mathcal{I}(\mathcal{L})]$$

$$= UC(\mathcal{L}, P_1) \sqsubseteq UC(\mathcal{L}, P_2) \qquad [def UC]$$

Proof of Law 5.4

A special property of $\mathcal{I}(\mathcal{L})$ is that it maps each behaviour ϕ to its fog counterpart $\phi[\tilde{x}, \tilde{x'}/x, x']$. Hence:

$D(UC(\mathcal{L},A))$			
=	$\exists \widetilde{x}, \widetilde{x'} \bullet A \land \widetilde{A} \land \mathcal{I}(\mathcal{L})$	[def D , UC]	
=	$A \wedge \exists \widetilde{x}, \widetilde{x'} \bullet \widetilde{A} \wedge \mathcal{I}(\mathcal{L})$	$[\widetilde{x}, \widetilde{x'} \text{ free in } A]$	
=	$A \wedge \exists \widetilde{x}, \widetilde{x'} \bullet \widetilde{A}$	[property of $\mathcal{I}(\mathcal{L})$]	
=	Α	[Lemma 4.6]	

Proof of Lemma 5.6

Proof of Theorem 5.7

$$= \widehat{\mathbf{R}} \left(\left(\neg A_{f}^{f} \vdash A_{f}^{t} \right) \land \left(\neg A_{f}^{f} \vdash A_{f}^{t} \right) [\widetilde{x}, \widetilde{x'} / x, x'] \right) \land \mathcal{I}(\mathcal{L}) \text{ [idempotence of } \widehat{\mathbf{R}} \text{]} \\ = \widehat{\mathbf{R}} \left(\left(ok \land \neg A_{f}^{f} \Rightarrow ok' \land A_{f}^{t} \right) \land \left(ok \land \neg \widetilde{A_{f}^{f}} \Rightarrow ok' \land \widetilde{A_{f}^{t}} \right) \right) \land \mathcal{I}(\mathcal{L}) \\ [\mathcal{I}(\mathcal{L}) \text{ implies } ok = \widetilde{ok} \land ok' = \widetilde{ok'} \text{]} \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land \left(\neg A_{f}^{f} \Rightarrow A_{f}^{t} \right) \land \left(\neg \widetilde{A_{f}^{f}} \Rightarrow \widetilde{A_{f}^{t}} \right) \right) \land \mathcal{I}(\mathcal{L}) \\ [\text{prop calc}] \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land A_{f}^{t} \land \widetilde{A_{f}^{t}} \right) \land \mathcal{I}(\mathcal{L}) \\ [\text{property of } A_{f}^{t}] \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land A_{f}^{t} \land \widetilde{A_{f}^{t}} \right) \land \mathcal{I}(\mathcal{L}) \\ [\text{property of } A_{f}^{t}] \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land A_{f}^{t} \land \widetilde{A_{f}^{t}} \right) \land \mathcal{I}(\mathcal{L}) \\ [\text{property of } A_{f}^{t}] \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land A_{f}^{t} \land \widetilde{A_{f}^{t}} \right) \land \mathcal{I}(\mathcal{L}) \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land A_{f}^{t} \land \widetilde{A_{f}^{t}} \right) \land \mathcal{I}(\mathcal{L}) \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land A_{f}^{t} \land \widetilde{A_{f}^{t}} \right) \land \widehat{\mathbf{R}}(\mathcal{L}) \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land A_{f}^{t} \land \widetilde{A_{f}^{t}} \right) \land \widehat{\mathbf{R}}(\mathcal{L}) \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land A_{f}^{t} \land \widetilde{A_{f}^{t}} \right) \land \widehat{\mathbf{R}}(\mathcal{L}) \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \Rightarrow ok' \land A_{f}^{t} \land \widetilde{A_{f}^{t}} \right) \land \widehat{\mathbf{R}}(\mathcal{L}) \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \land \widehat{\mathbf{R}}(\mathcal{L}) \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \land \widehat{\mathbf{R}}(\mathcal{L}) \right) \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \land \widehat{\mathbf{R}}(\mathcal{L}) \\ = \widehat{\mathbf{R}} \left(ok \land \neg \left(A_{f}^{f} \land \widetilde{A_{f}^{f}} \right) \land \widehat{\mathbf{R}}(\mathcal{L}) \right)$$

$$= \widehat{\mathbf{R}}\left(ok \wedge \neg \mathbf{U}\left(A_{f}^{t}\right) \Rightarrow ok' \wedge \mathbf{U}\left(A_{f}^{t}\right)\right) \wedge \mathcal{I}(\mathcal{L}) \qquad [\text{def }\mathbf{U}]$$

$$= \widehat{\mathbf{R}}\left(\neg \mathbf{U}\left(A_{f}^{f}\right) \vdash \mathbf{U}\left(A_{f}^{t}\right)\right) \land \mathcal{I}(\mathcal{L})$$
 [def $\widehat{\mathbf{R}}$]

Proof of Lemma 5.8

$$\begin{aligned} & \mathsf{UC} \left(\mathcal{L}, \mathsf{R} \left(Pre \vdash Post \right) \right) \\ &= \quad \widehat{\mathsf{R}} \left(\neg \mathsf{U} \left(\mathsf{R} \left(Pre \vdash Post \right) \right)_{ff}^{ff} \vdash \mathsf{U} \left(\mathsf{R} \left(Pre \vdash Post \right) \right)_{ff}^{tf} \right) \land \mathcal{I}(\mathcal{L}) \quad \text{[Theorem 5.7]} \\ &= \quad \widehat{\mathsf{R}} \left(\neg \mathsf{U} \left(\mathsf{R} \left(Pre \vdash Post \right)_{f}^{f} \right) \vdash \mathsf{U} \left(\mathsf{R} \left(Pre \vdash Post \right)_{f}^{t} \right) \right) \land \mathcal{I}(\mathcal{L}) \quad \text{[substitution]} \\ &= \quad \widehat{\mathsf{R}} \left(\neg \mathsf{U} \left(\neg Pre_{f}^{f} \right) \vdash \mathsf{U} \left(Pre_{f}^{t} \Rightarrow Post_{f}^{t} \right) \right) \land \mathcal{I}(\mathcal{L}) \quad \text{[property of reactive designs]} \\ &= \quad \widehat{\mathsf{R}} \left(\neg \mathsf{U} \left(\neg Pre_{f}^{f} \right) \vdash \mathsf{U} \left(Pre_{f} \Rightarrow Post_{f}^{t} \right) \right) \land \mathcal{I}(\mathcal{L}) \quad \text{[property of reactive designs]} \end{aligned}$$

$$= \widehat{\mathbf{R}} \left(\neg \mathbf{U} \left(\neg \operatorname{Pre} \right) \vdash \mathbf{U} \left(\operatorname{Pre} \Rightarrow \operatorname{Post} \right) \right) \land \mathcal{I}(\mathcal{L}) \qquad [\operatorname{def} \vdash, \mathbf{R3}; \operatorname{assumption}]$$

$$= \widehat{\mathbf{R}}\left(\operatorname{Pre} \lor \widetilde{\operatorname{Pre}} \vdash \mathbf{U}\left(\operatorname{Pre} \Rightarrow \operatorname{Post}\right)\right) \land \mathcal{I}(\mathcal{L}) \qquad [\operatorname{def} \mathbf{U}; \operatorname{prop calc}]$$

Proof of Lemma 5.9

$$UC (\mathcal{L}, w : [Pre, Post]) = UC (\mathcal{L}, \mathbf{R} (Pre \vdash Post \land \neg wait' \land tr' = tr \land u' = u)) \quad [def spec stmt] = \widehat{\mathbf{R}} \left(Pre \lor \widetilde{Pre} \vdash \mathbf{U} (Pre \Rightarrow (Post \land \neg wait' \land tr' = tr \land u' = u)) \right) \land \mathcal{I}(\mathcal{L})$$
[Lemma 5.8]

Proof of Theorem 5.14

This justification applies to the lifted operators defined in Section 5.3.

Condition 5.11

Condition 5.11 holds for sequential composition:

Condition 5.11 holds for internal choice:

$$D\left(UC\left(\mathcal{L},A_{1}\right) \cap UC\left(\mathcal{L},A_{2}\right)\right)$$

$$= D\left(UC\left(\mathcal{L},A_{1}\right)\right) \sqcap D\left(UC\left(\mathcal{L},A_{2}\right)\right) \qquad [Law \ 4.8]$$

$$= A_{1} \sqcap A_{2} \qquad [Law \ 5.4]$$

We appeal to Lemma 5.22 to treat the lifted prefixing operator in terms of lifted sequential composition:

$$\begin{split} \mathbf{D} \left(c.e \xrightarrow{\frown} \mathbf{UC} (\mathcal{L}, A) \right) \\ = \mathbf{D} \left(\mathbf{UC} (\mathcal{L}, c.e \rightarrow Skip) ; \mathbf{UC} (\mathcal{L}, A) \right) \end{split} \ [Lemma 5.22]$$

$$= c.e \rightarrow Skip; A \qquad [as above]$$
$$= c.e \rightarrow A$$

so Condition 5.11 holds for prefixing.

For each lifted operator defined as a lifted reactive design, we have:

$$\mathbf{D}\left(\widehat{\mathbf{R}}\left(Pre \vdash Post\right)\right) = \mathbf{R}\left(\neg \mathbf{D}\left(\neg Pre\right) \vdash \mathbf{D}\left(Post\right)\right)$$

By distributing **D** through the (lifted) *Pre* and *Post* terms, we derive the corresponding reactive design structure of the counterpart *Circus* operator, with each operand wrapped by **D**.

Taking the lifted guard as an example, we have:

$$\mathbf{D}\left(g \ \widehat{\&} \ \mathbf{UC}\left(\mathcal{L}, A\right)\right) = \mathbf{R}\left(\neg \mathbf{D}\left(\neg \left(\mathbf{U}\left(g\right) \Rightarrow \neg \mathbf{UC}\left(\mathcal{L}, A\right)_{ff}^{ff}\right)\right) \\ \vdash \mathbf{D}\left(\begin{matrix} \mathbf{U}\left(g\right) \land \mathbf{UC}\left(\mathcal{L}, A\right)_{ff}^{ff} \\ \lor \mathbf{U}\left(\neg g \land tr' = tr \land wait'\right) \end{matrix}\right) \right) \qquad [property above] \\ = \mathbf{R}\left(\neg \mathbf{D}\left(\mathbf{U}\left(g\right) \land \mathbf{UC}\left(\mathcal{L}, A\right)_{ff}^{ff}\right) \vdash \left(\begin{matrix} \mathbf{D}\left(\mathbf{U}\left(g\right) \land \mathbf{UC}\left(\mathcal{L}, A\right)_{ff}^{tf}\right) \\ \lor \mathbf{D}\left(\mathbf{U}\left(\neg g \land tr' = tr \land wait'\right)\right) \end{matrix}\right) \right) \qquad [prop calc; Law 4.8]$$

$$= \mathsf{R}\left(\neg \left(g \land \mathsf{D}\left(\mathsf{UC}\left(\mathcal{L},A\right)_{ff}^{ff}\right)\right) \vdash \left(\begin{array}{c}g \land \mathsf{D}\left(\mathsf{UC}\left(\mathcal{L},A\right)_{ff}^{tt}\right)\\ \lor \neg g \land tr' = tr \land wait'\end{array}\right)\right)$$

 $[\widetilde{x}, \widetilde{x'} \text{ free}]$

$$= \mathbf{R} \left(\left(g \Rightarrow \neg A_f^f \right) \vdash \left((g \land A_f^t) \lor (\neg g \land tr' = tr \land wait') \right) \right)$$
 [Law 5.4]
$$= g \& A$$
 [def &]

Condition 5.12

For each lifted operator $\widehat{\oplus}$, we provide a lemma to show that:

$$\mathsf{UC}\left(\mathcal{L},\bigoplus i\bullet A_i\right) \quad \sqsubseteq \quad \widehat{\bigoplus}\,i\bullet\mathsf{UC}\left(\mathcal{L},A_i\right)$$

or define $\widehat{\oplus}$ in terms of other lifted operators where that lemma holds. Hence, Condition 5.12 is satisfied by each lifted operator.

Condition 5.13

It is given that:

- conjunction and disjunction are monotonic;
- negation is antimonotonic;
- implication is antimonotonic in the antecedent and monotonic in the consequent.

Moreover, for each lifted operator $\widehat{\oplus}$, it is the case that:

- each operand of $\widehat{\oplus}$ within its precondition is negated; and
- each operand of $\widehat{\oplus}$ within its postcondition is not negated.

It follows from the structure of each lifted operator that refining each operand of a lifted construct yields a refinement of the whole construct. Hence, Condition 5.13 is satisfied by each lifted operator.

Proof of Lemma 5.15

$$\begin{aligned} & \mathsf{UC}\left(\mathcal{L}, A_{1}; A_{2}\right) \\ &= \begin{pmatrix} \exists x_{0} \bullet A_{1}[x_{0}/x'] \land A_{2}[x_{0}/x] \\ \land \quad (\exists x_{0} \bullet A_{1}[x_{0}/x'] \land A_{2}[x_{0}/x]) [\widetilde{x}, \widetilde{x'}/x, x'] \end{pmatrix} \land \mathcal{I}(\mathcal{L}) \quad [\mathsf{def} \, \mathsf{UC}, \, ;] \\ &= \begin{pmatrix} \exists x_{0} \bullet A_{1}[x_{0}/x'] \land A_{2}[x_{0}/x] \\ \land \quad \exists \widetilde{x_{0}} \bullet A_{1}[\widetilde{x}, \widetilde{x_{0}}/x, x'] \land A_{2}[\widetilde{x_{0}}, \widetilde{x'}/x, x'] \end{pmatrix} \land \mathcal{I}(\mathcal{L}) \quad [\mathsf{pred calc}] \end{aligned}$$

$$= \exists x_0, \widetilde{x_0} \bullet \mathsf{UC}(\mathcal{L}, A_1)[x_0, \widetilde{x_0}/x', \widetilde{x'}] \land \mathsf{UC}(\mathcal{L}, A_2)[x_0, \widetilde{x_0}/x, \widetilde{x}] \qquad [\mathsf{def} \ \mathsf{UC}]$$

Proof of Law 5.17

$$\begin{aligned} \mathbf{U}(A_{1};A_{2}) \\ &= \begin{pmatrix} \exists x_{0} \bullet A_{1}[x_{0}/x'] \land A_{2}[x_{0}/x] \\ \land \quad (\exists x_{0} \bullet A_{1}[x_{0}/x'] \land A_{2}[x_{0}/x])[\widetilde{x}, \widetilde{x'}/x, x'] \end{pmatrix} \quad [def \mathbf{U}, ;] \\ &= \begin{pmatrix} \exists x_{0} \bullet A_{1}[x_{0}/x'] \land A_{2}[x_{0}/x] \\ \land \quad \exists \widetilde{x_{0}} \bullet \widetilde{A_{1}}[\widetilde{x_{0}}/\widetilde{x'}] \land \widetilde{A_{2}}[\widetilde{x_{0}}/\widetilde{x}] \end{pmatrix} \quad [renaming] \\ &= \exists x_{0}, \widetilde{x_{0}} \bullet \begin{pmatrix} (A_{1} \land \widetilde{A_{1}})[x_{0}, \widetilde{x_{0}}/x', \widetilde{x'}] \\ \land \quad (A_{2} \land \widetilde{A_{2}})[x_{0}, \widetilde{x_{0}}/x, \widetilde{x}] \end{pmatrix} \quad [disjoint alphabets] \\ &= \mathbf{U}(A_{1}); \mathbf{U}(A_{2}) \qquad [def \mathbf{U}, ;] \end{aligned}$$

Proof of Lemma 5.18

Let A be a *Circus* action that terminates without changing the trace:

$$A = A \land (ok \land \neg wait \Rightarrow ok' \land \neg wait' \land tr' = tr)$$

Assuming this property holds, we have:

$$\begin{aligned} & \mathsf{U}(A \land ok \land \neg wait) \\ &= & \mathsf{U}(A) \land \mathsf{U}(ok \land \neg wait \land ok' \land \neg wait' \land tr' = tr) & \text{[assumption]} \\ &= & \mathsf{U}(A) \land ok \land \neg wait \land ok' \land \neg wait' \land tr' = tr \land \mathcal{I}(\mathcal{L}) & \text{[def } \mathcal{I}(\mathcal{L})] \\ &= & \mathsf{UC}(\mathcal{L}, A) \land ok \land \neg wait \land ok' \land \neg wait' \land tr' = tr & \text{[def } \mathsf{UC}] \end{aligned}$$

$$= UC(\mathcal{L}, A) \land ok \land \neg wait \qquad [assumption]$$

Now, consider the case where A_1 terminates immediately leaving the trace unchanged:

$$UC (\mathcal{L}, A_1; A_2) = U (A_1; A_2) \land \mathcal{I}(\mathcal{L}) \qquad [def UC] = \left(U (A_1); U (A_2) \right) \land \mathcal{I}(\mathcal{L}) \qquad [Law 5.17]$$

$$= \left(\mathsf{UC} \left(\mathcal{L}, A_1 \right) \, \widehat{;} \, \mathsf{U} \left(A_2 \right) \right) \wedge \mathcal{I}(\mathcal{L}) \qquad [A_1 \text{ terminates immediately}]$$

$$= \exists x_0, \widetilde{x_0} \bullet \mathsf{UC}(\mathcal{L}, A_1) [x_0, \widetilde{x_0}/x', \widetilde{x'}] \land \mathsf{U}(A_2) [x_0, \widetilde{x_0}/x, \widetilde{x}] \land \mathcal{I}(\mathcal{L}) \quad [\mathsf{def}];]$$

$$= \exists x_0, \widetilde{x_0} \bullet \begin{pmatrix} \mathsf{UC}(\mathcal{L}, A_1) [x_0, x_0/x^*, x^*] \\ \land \mathsf{U}(A_2) [x_0, \widetilde{x_0}/x, \widetilde{x}] & [A_1] \\ \land \mathcal{I}(\mathcal{L}) [x_0, \widetilde{x_0}/x, \widetilde{x}] \end{pmatrix} \text{ terminates immediately (*)]}$$

$$= \mathsf{UC}(\mathcal{L}, A_1) ; (\mathsf{U}(A_2) \land \mathcal{I}(\mathcal{L}))$$
 [def ;]

$$= UC(\mathcal{L}, A_1) ; UC(\mathcal{L}, A_2)$$
 [def UC]

The step marked (*) is justified because, since A_1 terminates immediately, it ensures ok' = ok', wait' = wait and tr' = tr, and the value of ref' is irrelevant upon termination of a *Circus* action.

The proof for the case where A_2 terminates immediately is similar.

Proof of Lemma 5.19

$UC\left(\mathcal{L}, A_1 \sqcap A_2\right)$			
=	$(A_1 \lor A_2) \land (\widetilde{A_1} \lor \widetilde{A_2}) \land \mathcal{I}(\mathcal{L})$	[def UC]	
\Leftarrow	$(A_1 \wedge \widetilde{A_1} \wedge \mathcal{I}(\mathcal{L})) \vee (A_2 \wedge \widetilde{A_2} \wedge \mathcal{I}(\mathcal{L}))$	[prop calc]	
=	$\operatorname{UC}\left(\mathcal{L},A_{1} ight)ee$ UC $\left(\mathcal{L},A_{2} ight)$	[def UC]	

Proof of Lemma 5.22

$$\begin{aligned} & \mathsf{UC} \left(\mathcal{L}, c.e \to Skip \right) \\ &= \quad \widehat{\mathsf{R}} \left(\mathsf{U} \left(c.e \to Skip \right)_{ff}^{ff} \vdash \mathsf{U} \left(c.e \to Skip \right)_{ff}^{tf} \right) \land \mathcal{I}(\mathcal{L}) \qquad [\text{Theorem 5.7}] \\ &= \quad \widehat{\mathsf{R}} \left(\mathsf{true} \vdash \mathsf{U} \left(v' = v \land \left(\begin{array}{c} tr' = tr \land c.e \notin ref' \\ \lhd wait' \vartriangleright \\ tr' = tr \frown \langle (c,e) \rangle \end{array} \right) \right) \right) \land \mathcal{I}(\mathcal{L}) \quad [\text{def} \to] \\ &= \quad \widehat{\mathsf{R}} \left(\mathsf{true} \vdash \mathsf{U} \left(v' = v \right) \land \left(\begin{array}{c} \mathsf{U} \left(tr' = tr \land c.e \notin ref' \right) \\ \lhd wait' \rhd \\ \mathsf{U} \left(tr' = tr \frown \langle (c,e) \rangle \right) \end{array} \right) \right) \land \mathcal{I}(\mathcal{L}) \\ &= \quad [\text{Law } 4.7; \mathcal{I}(\mathcal{L}) \text{ implies } wait = \widetilde{wait}] \end{aligned}$$

Proof of Lemma 5.24

$$\widehat{\mathbf{R}}\left(\left(\mathbf{U}\left(g\right)\Rightarrow\neg\mathbf{U}\left(A\right)_{ff}^{ff}\right)\vdash\mathbf{U}\left(A_{f}^{t}\lhd g\vartriangleright tr'=tr\land wait'\right)\right)\land\mathcal{I}(\mathcal{L})$$

[def **U**; prop calc]

$$\Leftarrow \quad \widehat{\mathsf{R}}\left(\left(\mathsf{U}\left(g\right) \Rightarrow \neg \:\mathsf{U}\left(A\right)_{ff}^{ff}\right) \vdash \left(\begin{array}{c} \mathsf{U}\left(g\right) \land \:\mathsf{U}\left(A\right)_{ff}^{tt} \\ \lor \:\mathsf{U}\left(\neg \:g \land tr' = tr \land wait'\right) \end{array}\right)\right) \land \mathcal{I}(\mathcal{L})$$

[strengthen postcondition (Law 4.11)]

Proof of Lemma 5.26

$$\mathbf{UC} (\mathcal{L}, A_1 \Box A_2) = \widehat{\mathbf{R}} \left(\neg \mathbf{U} (A_1 \Box A_2)_{ff}^{ff} \vdash \mathbf{U} (A_1 \Box A_2)_{ff}^{tt} \right) \land \mathcal{I}(\mathcal{L})$$
 [Theorem 5.7]

$$= \widehat{\mathbf{R}} \left(\neg \mathbf{U} \left(A_{1f}^{f} \lor A_{2f}^{f} \right) \vdash \mathbf{U} \left(\begin{array}{c} A_{1f}^{t} \land A_{2f}^{t} \\ \lhd tr' = tr \land wait' \rhd \\ A_{1f}^{t} \lor A_{2f}^{t} \end{array} \right) \right) \land \mathcal{I}(\mathcal{L}) \quad [\text{def } \Box]$$

$$\Leftrightarrow \widehat{\mathbf{R}} \left(\neg \mathbf{U} \left(A_{1f}^{f} \right) \land \neg \mathbf{U} \left(A_{2f}^{f} \right) \vdash \mathbf{U} \left(\begin{array}{c} A_{1f}^{t} \land A_{2f}^{t} \\ \lhd tr' = tr \land wait' \rhd \\ A_{1f}^{t} \lor A_{2f}^{t} \end{array} \right) \right) \land \mathcal{I}(\mathcal{L})$$

[weaken precondition (contrapositive of Law 4.9)]

$$\leftarrow \widehat{\mathbf{R}} \begin{pmatrix} \neg \mathbf{U} \left(A_{1f}^{f} \right) \land \neg \mathbf{U} \left(A_{2f}^{f} \right) \\ \vdash \begin{pmatrix} \mathbf{U} \left(tr' = tr \land wait' \right) \land \mathbf{U} \left(A_{1} \right)_{ff}^{tt} \land \mathbf{U} \left(A_{2} \right)_{ff}^{tt} \\ \lor \mathbf{U} \left(\neg \left(tr' = tr \land wait' \right) \right) \land \left(\mathbf{U} \left(A_{1} \right)_{ff}^{tt} \lor \mathbf{U} \left(A_{2} \right)_{ff}^{tt} \right) \end{pmatrix} \end{pmatrix} \\ \land \mathcal{I}(\mathcal{L}) \qquad [\text{strengthen postcondition (Law 4.11; Law 4.9)]}$$

Proof of Lemma 5.28

Let *EQ* denote $1.tr' \upharpoonright cs = 2.tr' \upharpoonright cs$ in:

$$(A_{2f}^{\prime}; 1.tr' = tr) \land (A_{2f}^{\prime}; 2.tr' = tr) \land EQ) \int dt$$

[Law 4.7, Law 4.13]

$$= \begin{pmatrix} \neg \exists 1.tr, 2.tr \bullet \begin{pmatrix} \mathsf{U}(A_1)_{ff}^{ff} ; \mathsf{U}(1.tr' = tr) \\ \land \mathsf{U}(A_2)_{ff} ; \mathsf{U}(2.tr' = tr) \end{pmatrix} \land \mathsf{U}(EQ) \\ \land \neg \exists 1.tr, 2.tr \bullet \begin{pmatrix} \mathsf{U}(A_1)_{ff} ; \mathsf{U}(1.tr' = tr) \\ \land \mathsf{U}(A_2)_{ff}^{ff} ; \mathsf{U}(2.tr' = tr) \end{pmatrix} \land \mathsf{U}(EQ) \end{pmatrix}$$

[Law 4.7, Law 5.17]

Proof of Lemma 5.30

Let HS denote
$$\begin{pmatrix} A[s, (cs \cup ref')/tr', ref'] \\ \land (tr' - tr) = (s - tr) \upharpoonright (EVENT - cs) \end{pmatrix}$$
 in:
UC $(\mathcal{L}, A \setminus cs)$
= UC $(\mathbb{R} (\exists s \bullet HS) ; Skip)$ [def \]
= UC $(\mathbb{R} (\exists s \bullet HS)) ;$ UC $(\mathcal{L}, Skip)$ [Lemma 5.18]
= $(\widehat{\mathbb{R}} (\mathbb{U} (\exists s \bullet HS)) \land \mathcal{I}(\mathcal{L})) ;$ UC $(\mathcal{L}, Skip)$ [Lemma 5.6]
= $(\widehat{\mathbb{R}} (\exists s, \tilde{s} \bullet \mathbb{U} (HS)) \land \mathcal{I}(\mathcal{L})) ;$ UC $(\mathcal{L}, Skip)$ [property of U]

Proof of Theorem 5.40

$$\begin{array}{ll} \theta_1 \sqsubseteq \theta_2 \\ = & (ok \land \neg \ wait \Rightarrow \theta_1) \sqsubseteq (ok \land \neg \ wait \Rightarrow \theta_2) & [pred \ calc \ (*)] \\ = & \operatorname{Conf}(\theta_1) \sqsubseteq \operatorname{Conf}(\theta_2) & [def \ Conf] \\ = & \operatorname{UC}(\mathcal{L}, Skip) \land \operatorname{Conf}(\theta_1) \sqsubseteq \operatorname{UC}(\mathcal{L}, Skip) \land \operatorname{Conf}(\theta_2) & [pred \ calc \ (*)] \\ = & \langle \theta_1 \rangle \sqsubseteq \langle \theta_2 \rangle & [def \ CA] \end{array}$$

It is assumed that obligations reference only the unprimed state variables v, \tilde{v} . The steps marked (*) are equivalences (rather than refinements), because they place no constraints on the unprimed state variables.

Proof of Law 5.41

We prove the result for two CAs. Since relational composition is associative, this result generalises to arbitrary sequences of CAs.

$$\begin{array}{l} \langle \theta_1 \rangle; \langle \theta_2 \rangle \\ = & (\mathbf{UC} \left(\mathcal{L}, Skip \right) \wedge \operatorname{Conf} \left(\theta_1 \right) \right) \hat{;} (\mathbf{UC} \left(\mathcal{L}, Skip \right) \wedge \operatorname{Conf} \left(\theta_2 \right)) & [\operatorname{def} \operatorname{CA}] \\ = & \exists x_0, \widetilde{x_0} \bullet \begin{pmatrix} & (\mathbf{UC} \left(\mathcal{L}, Skip \right) \wedge \operatorname{Conf} \left(\theta_1 \right)) [x_0, \widetilde{x_0} / x', \widetilde{x'}] \\ \wedge & (\mathbf{UC} \left(\mathcal{L}, Skip \right) \wedge \operatorname{Conf} \left(\theta_2 \right)) [x_0, \widetilde{x_0} / x, \widetilde{x}] \end{pmatrix} & [\operatorname{def} \hat{;}] \end{array}$$

$$= \operatorname{Conf}(\theta_{1}) \land \exists x_{0}, \widetilde{x_{0}} \bullet \begin{pmatrix} \operatorname{UC}(\mathcal{L}, Skip) [x_{0}, \widetilde{x_{0}}/x', \widetilde{x'}] \\ \land \quad (\operatorname{UC}(\mathcal{L}, Skip) \land \operatorname{Conf}(\theta_{2})) [x_{0}, \widetilde{x_{0}}/x, \widetilde{x}] \end{pmatrix}$$

$$= \operatorname{Conf}(\theta_{1}) \land \left(\operatorname{UC}(\mathcal{L}, Skip) \,\widehat{;} \left(\operatorname{UC}(\mathcal{L}, Skip) \land \operatorname{Conf}(\theta_{2})\right)\right) \quad [\operatorname{def} \,\widehat{;}]$$

$$= \operatorname{Conf}(\theta_{1}) \land \operatorname{UC}(\mathcal{L}, Skip) \land \operatorname{Conf}(\theta_{2}) \quad [Skip \ \operatorname{left} unit]$$

$$= \operatorname{UC}(\mathcal{L}, Skip) \land \operatorname{Conf}(\theta_{1} \land \theta_{2}) \quad [\operatorname{def} \ \operatorname{Conf}]$$

$$= \langle \bigsqcup \{\theta_{1}, \theta_{2}\} \rangle \quad [\operatorname{def} \ \operatorname{Cal}]$$

Proof of Law 5.43

$\begin{array}{ll} \langle \theta_1 \rangle \sqcap \cdots \sqcap \langle \theta_n \rangle \\ = & (\mathbf{UC} \left(\mathcal{L}, Skip \right) \land \mathsf{Conf} \left(\theta_1 \right) \right) \lor \cdots \lor (\mathbf{UC} \left(\mathcal{L}, Skip \right) \land \mathsf{Conf} \left(\theta_n \right)) [\mathsf{def} \mathsf{CA}, \sqcap] \\ = & \mathbf{UC} \left(\mathcal{L}, Skip \right) \land (\mathsf{Conf} \left(\theta_1 \right) \lor \cdots \lor \mathsf{Conf} \left(\theta_n \right)) & [\mathsf{prop} \mathsf{calc}] \\ = & \mathbf{UC} \left(\mathcal{L}, Skip \right) \land \mathsf{Conf} \left(\theta_1 \lor \cdots \lor \theta_n \right) & [\mathsf{def} \mathsf{Conf}] \\ = & \langle \sqcap \{ \theta_1, \dots, \theta_n \} \rangle & [\mathsf{def} \mathsf{CA}] \end{array}$

Proof of Law 5.44

$UC(\mathcal{L},[C])$				
=	$\widehat{R}\left(true \vdash U\left(C \land \neg \mathit{wait'} \land \mathit{tr'} = \mathit{tr} \land \mathit{v'} = \mathit{v}\right)\right) \land \mathcal{I}(\mathcal{L})$	C) [Lemma 5.9]		
=	$\widehat{R}\left(true \vdash U\left(\neg \textit{ wait'} \land tr' = tr \land v' = v\right)\right) \land \mathcal{I}(\mathcal{L}) \land Conf\left(U\left(C\right)\right)$			
	[prop calc; def Conf]		
=	$UC\left(\mathcal{L}, Skip ight) \wedge Conf\left(U\left(C ight) ight)$	[def Skip]		
=	$\langle \mathbf{U}(C) \rangle$	[def CA]		

Proof of Law 5.45

 $\mathbf{D}(\langle \theta \rangle)$

$$= \mathsf{D} (\mathsf{UC} (\mathcal{L}, Skip) \land (ok \land \neg wait \Rightarrow \theta)) \qquad [Law 5.44]$$

$$= \mathsf{D} (\widehat{\mathsf{R}} (true \vdash \theta \land \mathsf{U} (\neg wait' \land tr' = tr \land v' = v)) \land \mathcal{I}(\mathcal{L})) \qquad [Lemma 5.8, def Skip]$$

$$= \mathsf{R} (true \vdash \mathsf{D} (\theta) \land \neg wait' \land tr' = tr \land v' = v) \qquad [property of \mathsf{D}]$$

$$= [\mathsf{D} (\theta)] \qquad [def coercion]$$

Proof of Law 5.46

$$\begin{array}{ll} \langle \theta \rangle; \langle A \rangle \\ = & \exists x_0, \widetilde{x_0} \bullet \left(\begin{array}{c} (\mathsf{UC} \left(\mathcal{L}, Skip\right) \land (ok \land \neg wait \Rightarrow \theta))[x_0, \widetilde{x_0}/x', \widetilde{x'}] \\ \land & \mathsf{UC} \left(\mathcal{L}, A\right) [x_0, \widetilde{x_0}/x, \widetilde{x}] \end{array} \right) \\ & [def \ \widehat{;}, CA] \\ = & \exists x_0, \widetilde{x_0} \bullet \left(\begin{array}{c} (\mathsf{UC} \left(\mathcal{L}, Skip\right) \land (ok' \land \neg wait' \Rightarrow \theta'))[x_0, \widetilde{x_0}/x', \widetilde{x'}] \\ \land & \mathsf{UC} \left(\mathcal{L}, A\right) [x_0, \widetilde{x_0}/x, \widetilde{x}] \end{array} \right) \\ & [property of Skip] \\ = & \exists x_0, \widetilde{x_0} \bullet \left(\begin{array}{c} \mathsf{UC} \left(\mathcal{L}, Skip\right) \land (ok \land \neg wait' \Rightarrow \theta'))[x_0, \widetilde{x_0}/x', \widetilde{x'}] \\ \land & \mathsf{UC} \left(\mathcal{L}, A\right) [x_0, \widetilde{x_0}/x, \widetilde{x}] \end{array} \right) \end{array} \right) \end{array}$$

$$= \exists x_0, \widetilde{x_0} \bullet \begin{pmatrix} \mathsf{UC}(\mathcal{L}, Skip)[x_0, x_0/x', x'] \\ \land \quad (\mathsf{UC}(\mathcal{L}, A) \land (ok \land \neg wait \Rightarrow \theta))[x_0, \widetilde{x_0}/x, \widetilde{x}] \end{pmatrix}$$

[renaming]

$$= \mathsf{UC}(\mathcal{L}, Skip)\,\hat{;}\,(\mathsf{UC}(\mathcal{L}, A) \land (ok \land \neg wait \Rightarrow \theta)) \qquad [def\,\hat{;}]$$

$$= (UC(\mathcal{L}, A) \land (ok \land \neg wait \Rightarrow \theta))$$
[Skip left unit]
$$= \langle A \rangle \land Conf(\theta)$$
[def Conf]

Proof of Law 5.47

$$= \exists x_0, \tilde{x_0} \bullet \begin{pmatrix} (\mathsf{UC}(\mathcal{L}, A) \land (ok' \land \neg wait' \Rightarrow \theta'))[x_0, \tilde{x_0}/x', \tilde{x'}] \\ \land \quad \mathsf{UC}(\mathcal{L}, Skip)[x_0, \tilde{x_0}/x, \tilde{x}] \end{pmatrix}$$
[renaming]
$$= (\mathsf{UC}(\mathcal{L}, A) \land (ok' \land \neg wait' \Rightarrow \theta')) ; \mathsf{UC}(\mathcal{L}, Skip) \quad [\det \hat{f};]$$

$$= \mathsf{UC}(\mathcal{L}, A) \land (ok' \land \neg wait' \Rightarrow \theta') \quad [Skip right unit]$$

$$= \langle A \rangle \land \mathsf{Conf}'(\theta) \quad [\det \mathsf{Conf}']$$

Proof of Law 6.4

Let $C1 = c_1 \rightarrow A_1$ and $C2 = c_2 \rightarrow A_2$ in:

$$\begin{array}{ll} \langle L : C1 \sqcap C2 \rangle \\ = & \mathsf{UC} \left(\mathcal{L}, C1 \lor C2 \right) & [\det \langle \cdot \rangle, \sqcap] \\ = & (C1 \lor C2) \land (\widetilde{C1} \land \widetilde{C2}) \land \mathcal{I}(\mathcal{L}) & [\det \mathsf{UC}] \\ = & ((C1 \land \widetilde{C1}) \lor (C2 \land \widetilde{C2})) \land \mathcal{I}(\mathcal{L}) & [property \text{ of } \mathcal{I}(\mathcal{L})] \\ = & \mathsf{UC} \left(\mathcal{L}, C1 \right) \lor \mathsf{UC} \left(\mathcal{L}, C2 \right) & [\det \mathsf{UC}] \\ = & \langle L : C1 \rangle \sqcap \langle L : C2 \rangle & [\det \langle \cdot \rangle, \sqcap] \end{array}$$

Proof of Theorem 6.12

$$\begin{aligned} & \mathsf{bwQ}\left(\mathsf{UC}\left(\mathcal{L},\mathsf{R}\left(Pre\vdash Post\right)\right),\theta\right) \\ &= \quad \forall x' \bullet \left(\begin{array}{c} \mathsf{D}\left(\mathsf{UC}\left(\mathcal{L},\mathsf{R}\left(Pre\vdash Post\right)\right)\right) \Rightarrow \\ & \exists \widetilde{x'} \bullet \mathsf{UC}\left(\mathcal{L},\mathsf{R}\left(Pre\vdash Post\right)\right) \right) & \land \mathsf{Conf}'(\theta) \end{array}\right) \\ &= \quad \forall x' \bullet \left(\begin{array}{c} \mathsf{R}\left(Pre\vdash Post\right)! \Rightarrow \\ & \exists \widetilde{x'} \bullet \mathsf{UC}\left(\mathcal{L},\mathsf{R}\left(Pre\vdash Post\right)\right)! \land \mathsf{Conf}'(\theta) \end{array}\right) & [\mathsf{Law} \ 5.4] \\ &= \quad \forall x' \bullet \left(\begin{array}{c} \mathsf{R}\left(Pre\vdash Post\right)! \Rightarrow \\ & \exists \widetilde{x'} \bullet \left(\begin{array}{c} \mathsf{R}\left(Pre \lor \widetilde{Post} \vdash \mathsf{U}\left(Pre \Rightarrow Post\right)\right)! \\ & \land \ \mathcal{I}(\mathcal{L})! \land \mathsf{Conf}'(\theta) \end{array}\right) \end{array}\right) \end{aligned}\right) \end{aligned}$$

$$= \forall x' \bullet \begin{pmatrix} (Pre! \Rightarrow ok' \land Post!) \land tr \leq tr' \Rightarrow \\ \exists \tilde{x'} \bullet \begin{pmatrix} ((Pre! \lor \tilde{Pre!}) \Rightarrow ok' \land \mathbf{U} (Pre! \Rightarrow Post!)) \\ \land tr \leq tr' \land \tilde{tr} \leq \tilde{tr'} \land \mathcal{I}(\mathcal{L})! \land Conf'(\theta) \end{pmatrix} \end{pmatrix}$$
[!-apply]

 $= \qquad \forall x' \bullet Pre! \Rightarrow \begin{pmatrix} ok' \land Post! \land tr \leq tr' \Rightarrow \\ \exists \widetilde{x'} \bullet \begin{pmatrix} \widetilde{Pre!} \Rightarrow \widetilde{Post!} \\ \land \widetilde{tr} \leq t\widetilde{r'} \land \mathcal{I}(\mathcal{L})! \land \mathsf{Conf'}(\theta) \end{pmatrix} \end{pmatrix}$ $\land \forall x' \bullet \neg Pre! \Rightarrow \\ \begin{pmatrix} tr \leq tr' \Rightarrow \exists \widetilde{x'} \bullet \begin{pmatrix} \widetilde{Pre!} \Rightarrow ok' \land \widetilde{Post!} \\ \land \widetilde{tr} \leq t\widetilde{r'} \land \mathcal{I}(\mathcal{L})! \land \mathsf{Conf'}(\theta) \end{pmatrix} \end{pmatrix}$ [case split]

$$= \begin{pmatrix} \forall x' \bullet \begin{pmatrix} Pre! \land ok' \land Post! \land tr \leq tr' \Rightarrow \\ \exists \widetilde{x'} \bullet (\widetilde{Pre!} \Rightarrow \widetilde{Post!}) \land \widetilde{tr} \leq \widetilde{tr'} \land \mathcal{I}(\mathcal{L})! \land \operatorname{Conf'}(\theta) \end{pmatrix} \\ \land \forall x' \bullet \begin{pmatrix} \neg Pre! \land tr \leq tr' \Rightarrow \\ \exists \widetilde{x'} \bullet \begin{pmatrix} \widetilde{Pre!} \Rightarrow ok' \land \widetilde{Post!} \\ \land \ \widetilde{tr} \leq \widetilde{tr'} \land \mathcal{I}(\mathcal{L})! \land \operatorname{Conf'}(\theta) \end{pmatrix} \end{pmatrix}$$

[prop calc]

$$= \begin{pmatrix} \forall x' \bullet \begin{pmatrix} Pre! \land ok' \land Post! \land tr \leq tr' \Rightarrow \\ \exists \tilde{x'} \bullet (\tilde{Pre!} \Rightarrow \tilde{Post!}) \land \tilde{tr} \leq \tilde{tr'} \land \mathcal{I}(\mathcal{L})! \land Conf'(\theta) \end{pmatrix} \\ \land \forall x' \bullet \begin{pmatrix} \neg Pre! \land tr \leq tr' \Rightarrow \\ \exists \tilde{x'} \bullet \neg \tilde{Pre!} \land \tilde{tr} \leq \tilde{tr'} \land \mathcal{I}(\mathcal{L})! \land Conf'(\theta) \end{pmatrix} \end{pmatrix}$$

[*ok'* free in antecedent]

$$= \begin{array}{c} \forall x' \bullet \begin{pmatrix} Pre! \land ok' \land Post! \land tr \leq tr' \Rightarrow \\ \exists \widetilde{x'} \bullet (\widetilde{Pre!} \Rightarrow \widetilde{Post!}) \land \widetilde{tr} \leq \widetilde{tr'} \land \mathcal{I}(\mathcal{L})! \land \mathsf{Conf'}(\theta) \end{pmatrix} \\ \land \quad \forall x' \bullet (\neg Pre! \Rightarrow \exists \widetilde{x'} \bullet \neg \widetilde{Pre!} \land \theta') \end{array}$$

[x' free in *Pre*!; *ok'*, *wait'* free in antecendent]

The second conjunct is simplified by the proviso $\forall x \bullet \exists \tilde{x} \bullet \theta$, which is equivalent to $\forall x' \bullet \exists \tilde{x'} \bullet \theta'$:

$$\forall x' \bullet \left(\neg \operatorname{Pre!} \Rightarrow \exists \widetilde{x'} \bullet \neg \widetilde{\operatorname{Pre!}} \land \theta'\right)$$

= $\forall x' \bullet \left(\neg \operatorname{Pre!} \Rightarrow \exists \widetilde{x'} \bullet \neg \widetilde{\operatorname{Pre!}}\right)$ [proviso]

=	$\neg Pre! \Rightarrow \neg \widetilde{Pre!}$	[Pre a condition]
=	$\widetilde{Pre!} \Rightarrow Pre!$	[contrapositive]

Proof of Theorem 6.14

 $\mathsf{bwR}(\mathsf{UC}(\mathcal{L}, \mathsf{R}(\mathit{Pre} \vdash \mathit{Post})))$

$$= \exists x', \widetilde{x'} \bullet \mathsf{UC} (\mathcal{L}, \mathsf{R} (Pre \vdash Post))! \qquad [def \ \mathsf{bwR}]$$

$$= \exists x', \widetilde{x'} \bullet \widehat{\mathsf{R}} (Pre \lor \widetilde{Pre} \vdash \mathsf{U} (Pre \Rightarrow Post))! \land \mathcal{I}(\mathcal{L})! \qquad [Lemma 5.8]$$

$$= \exists x', \widetilde{x'} \bullet \mathsf{U} (tr \le tr') \land ((Pre! \lor \widetilde{Pre!}) \Rightarrow ok' \land \mathsf{U} (Pre! \Rightarrow Post!)) \land \mathcal{I}(\mathcal{L})! \qquad [!-apply]$$

Proof of Lemma 6.16

To prove **bw** (B, θ) is monotonic in θ , it suffices to prove that **bwQ** (B, θ) is monotonic in θ . Let $\theta_1 \sqsubseteq \theta_2$. Then:

$$\begin{aligned} & \mathbf{bwQ} \left(B, \theta_2 \right) \\ &= & \forall x' \bullet \mathbf{D} \left(B! \right) \Rightarrow \exists \widetilde{x'} \bullet B! \land \left(ok' \land \neg wait' \Rightarrow \theta_2' \right) & \text{[def bwQ]} \\ &\Rightarrow & \forall x' \bullet \mathbf{D} \left(B! \right) \Rightarrow \exists \widetilde{x'} \bullet B! \land \left(ok' \land \neg wait' \Rightarrow \theta_1' \right) & \text{[}\theta_1 \sqsubseteq \theta_2 \text{]} \\ &= & \mathbf{bwQ} \left(B, \theta_1 \right) & \text{[def bwQ]} \end{aligned}$$

Proof of Law 6.18

First, we calculate **bwR** ($\langle \theta_1 \rangle$):

$$\begin{aligned} & \mathbf{bwR}\left(\langle \theta_{1} \rangle\right) \\ &= \quad \exists x', \widetilde{x'} \bullet \left(\mathbf{UC}\left(\mathcal{L}, Skip\right) \land \left(ok' \land \neg wait' \Rightarrow \theta_{1}'\right)\right)! & \text{[def bwR, def CA]} \\ &= \quad \exists x', \widetilde{x'} \bullet \mathbf{UC}\left(\mathcal{L}, Skip\right)! \land \theta_{1}' & \text{[def !]} \\ &= \quad \theta_{1} & \text{[def Skip]} \end{aligned}$$

Second, we calculate **bwQ** $(\langle \theta_1 \rangle, \theta_2)$:

$$\begin{aligned} & \mathsf{bwQ}\left(\langle \theta_1 \rangle, \theta_2\right) \\ &= & \forall x' \bullet \mathsf{D}\left(\langle \theta_1 \rangle!\right) \Rightarrow \exists \tilde{x'} \bullet \langle \theta_1 \rangle! \land (ok' \land \neg wait' \Rightarrow \theta'_2) \qquad [\text{def } \mathsf{bw}] \\ &= & \forall x' \bullet \mathsf{D}\left(\langle \theta_1 \rangle!\right) \Rightarrow \exists \tilde{x'} \bullet \langle \theta_1 \rangle! \land \theta_2 \qquad [\text{property of } Skip] \\ &= & \forall x' \bullet \mathsf{D}\left(\theta_1\right) \Rightarrow \exists \tilde{x'} \bullet \theta_1 \land \theta_2 \qquad [\text{as above}] \\ &= & \mathsf{D}\left(\theta_1\right) \Rightarrow \theta_1 \land \theta_2 \qquad [x', \tilde{x'} \text{ free in } \theta_1, \theta_2] \end{aligned}$$

The result follows by conjoining the **bwR** and **bwQ** terms:

bwR $(\langle heta_1 angle) \wedge $ bwQ $(\langle heta_1 angle, heta_2)$		
=	$ heta_{1} \wedge \left(\mathbf{D}\left(heta_{1} ight) \Rightarrow heta_{1} \wedge heta_{2} ight)$	[as above]
=	$ heta_1 \wedge heta_2$	$\left[D\left(\theta_{1}\right) \sqsubseteq\theta_{1}\right]$

Proof of Lemma 6.19

By the definition of **bwR**, we have:

bwR (B)			
=	$\exists x', \widetilde{x'} \bullet B!$	[def bwR]	
=	$\exists x', \widetilde{x'} \bullet (B \land (C \lor \neg C))!$	[excluded middle]	
=	$\exists x', \widetilde{x'} \bullet (B \land C)! \lor \exists x', \widetilde{x'} \bullet (B \land \neg C)!$	[prop calc]	
=	$bwR(B\wedge C)\lorbwR(B\wedge\neg C)$	[def bwR]	
\Leftarrow	bwR $(B \land C) \land$ bwR $(B \land \neg C)$	[prop calc]	

The proof for $\boldsymbol{\mathsf{bwQ}}$ has a similar structure:

$$\mathbf{bwQ}(B,\theta) = \forall x' \bullet \left(\mathbf{D}(B!) \Rightarrow \exists \tilde{x'} \bullet B! \land \operatorname{Conf}'(\theta) \right) \qquad [\operatorname{def} \mathbf{bwQ}]$$

= $\begin{pmatrix} \forall x' \bullet \mathbf{D}((B \land (C \lor \neg C))!) \Rightarrow \\ \exists \tilde{x'} \bullet (B \land (C \lor \neg C))! \land \operatorname{Conf}'(\theta) \end{pmatrix} \qquad [\operatorname{excluded} \operatorname{middle}]$

$$\leftarrow \begin{pmatrix} \forall x' \bullet \left(\mathsf{D} \left((B \land C)! \right) \Rightarrow \exists \tilde{x'} \bullet (B \land C)! \land \mathsf{Conf'}(\theta) \right) \\ \land \forall x' \bullet \left(\mathsf{D} \left((B \land \neg C)! \right) \Rightarrow \exists \tilde{x'} \bullet (B \land \neg C)! \land \mathsf{Conf'}(\theta) \right) \end{pmatrix}$$
 [pred calc]
= $\mathsf{bwQ} \left(B \land C, \theta \right) \land \mathsf{bwQ} \left(B \land \neg C, \theta \right)$ [def bwQ]

The result follows by conjoining the \mathbf{bwR} and \mathbf{bwQ} terms together.

Proof of Theorem 6.20

The proof for **bwQ** depends on side-conditions 3 and 4:

$$\begin{aligned} \mathbf{bwQ} \left(B, \theta \right) &= \forall x' \bullet \left(\mathsf{D} \left(B! \right) \Rightarrow \exists \tilde{x'} \bullet B! \land \mathsf{Conf'} \left(\theta \right) \right) & [\det \mathsf{bwQ}] \\ &= \begin{pmatrix} \forall x' \bullet \mathsf{D} \left(B! \right) \land C \Rightarrow \exists \tilde{x'} \bullet B! \land \mathsf{Conf'} \left(\theta \right) \\ \land \forall x' \bullet \mathsf{D} \left(B! \right) \land \neg C \Rightarrow \exists \tilde{x'} \bullet B! \land \mathsf{Conf'} \left(\theta \right) \end{pmatrix} & [\mathsf{case split}] \\ &= \begin{pmatrix} \forall x' \bullet \mathsf{D} \left(B! \right) \land C \Rightarrow \exists \tilde{x'} \bullet B! \land \mathsf{U} \left(C \right) \land \mathsf{Conf'} \left(\theta \right) \\ \land \forall x' \bullet \mathsf{D} \left(B! \right) \land \neg C \Rightarrow \exists \tilde{x'} \bullet B! \land \mathsf{U} \left(\neg C \right) \land \mathsf{Conf'} \left(\theta \right) \end{pmatrix} & [\mathsf{SC 4}] \\ &= \begin{pmatrix} \forall x' \bullet \mathsf{D} \left((B \land \mathsf{U} \left(C \right))! \right) \Rightarrow \exists \tilde{x'} \bullet \left(B \land \mathsf{U} \left(C \right) \right)! \land \mathsf{Conf'} \left(\theta \right) \\ \land \forall x' \bullet \mathsf{D} \left((B \land \mathsf{U} \left(\neg C \right))! \right) \Rightarrow \exists \tilde{x'} \bullet \left(B \land \mathsf{U} \left(\neg C \right) \right)! \land \mathsf{Conf'} \left(\theta \right) \end{pmatrix} & [\mathsf{SC 4}] \\ &= \mathsf{bwQ} \left(B \land \mathsf{U} \left(C \right), \theta \right) \land \mathsf{bwQ} \left(B \land \mathsf{U} \left(\neg C \right), \theta \right) & [\mathsf{def} \mathsf{bwQ}] \\ &= \mathsf{bwQ} \left(B \land \mathsf{C}, \theta \right) \land \mathsf{bwQ} \left(B \land \mathsf{V} \left(\neg C \right), \theta \right) & [\mathsf{SC 4}] \end{aligned}$$

Hence:

The property of **bwR** in question is that side-conditions 1 and 2 each imply **bwR** (B) = **true**.

Proof of Lemma 6.21

It suffices to establish *HT* respects side-condition 4 of Theorem 6.20:

Proofs for Lemma 6.22 and Lemma 6.23 are similar.

Proof of Law 6.24

Skip is not miraculous, so **bw** ($\langle Skip \rangle$, θ) = **bwQ** ($\langle Skip \rangle$, θ).

$$\begin{aligned} \mathsf{bwQ}\left(\langle Skip \rangle, \theta\right) &= \mathsf{bwQ}\left(\mathsf{UC}\left(\mathcal{L}, \mathsf{R}\left(\mathsf{true} \vdash tr' = tr \land \neg wait' \land v' = v\right)\right), \theta\right) & [\det Skip] \\ &= \forall x' \bullet \begin{pmatrix} ok' \land tr' = tr \land \neg wait' \land v' = v \\ \Rightarrow \exists \widetilde{x'} \bullet \widetilde{tr'} = \widetilde{tr} \land \neg \widetilde{wait'} \land \widetilde{v'} = \widetilde{v} \land \mathcal{I}(\mathcal{L})! \land \mathsf{Conf'}(\theta) \end{pmatrix} & [\mathsf{Theorem 6.12}] \\ &= \forall v' \bullet (v' = v \Rightarrow \exists \widetilde{v'} \bullet \widetilde{v'} = \widetilde{v} \land \theta') & [\mathsf{pred calc}] \\ &= \theta & [\mathsf{renaming}] \end{aligned}$$

Proof of Law 6.25

Stop is not miraculous, so **bw** ($\langle Stop \rangle, \theta$) = **bwQ** ($\langle Stop \rangle, \theta$).

$$\begin{aligned} \mathbf{bwQ}\left(\langle Stop \rangle, \theta\right) &= \mathbf{bwQ}\left(\mathbf{UC}\left(\mathcal{L}, \mathbf{R}\left(\mathbf{true} \vdash tr' = tr \land wait'\right)\right), \theta\right) & [\det Stop] \\ &= \forall x' \bullet \begin{pmatrix} ok' \land tr' = tr \land wait' \\ \Rightarrow \exists \widetilde{x'} \bullet \widetilde{tr'} = \widetilde{tr} \land \widetilde{wait'} \land \mathcal{I}(\mathcal{L})! \land \mathsf{Conf'}(\theta) \end{pmatrix} \\ & [Theorem 6.12] \\ &= \forall x' \bullet \begin{pmatrix} ok' \land tr' = tr \land wait' \\ \Rightarrow \exists \widetilde{x'} \bullet \widetilde{tr'} = \widetilde{tr} \land \widetilde{wait'} \land \mathcal{I}(\mathcal{L})! \end{pmatrix} & [\det \mathsf{Conf'}] \end{aligned}$$

$$= true$$
 [pred calc]

Proof of Law 6.26

Chaos is not miraculous, so **bw** ($\langle Chaos \rangle, \theta$) = **bwQ** ($\langle Chaos \rangle, \theta$).

Proof of Law 6.27

The scope constructs are not miraculous, so we have **bw** ($\langle \mathbf{var} \ a : T \rangle, \theta$) = **bwQ** ($\langle \mathbf{var} \ a : T \rangle, \theta$) and **bw** ($\langle \mathbf{end} \ a : T \rangle, \theta$) = **bwQ** ($\langle \mathbf{end} \ a : T \rangle, \theta$).

$$\begin{aligned} & \mathsf{bw}\left(\langle \operatorname{var} a:T \rangle, \theta\right) \\ &= \operatorname{bwQ}\left(\langle \operatorname{var} a:T \rangle, \theta\right) & [\operatorname{var} \operatorname{non-miraculous}] \\ &= \forall x' \bullet \mathsf{D}\left(\mathsf{UC}\left(\mathcal{L}, \operatorname{var} a:T\right)\right)! \Rightarrow \exists \widetilde{x'} \bullet \mathsf{UC}\left(\mathcal{L}, \operatorname{var} a:T\right)! \land \operatorname{Conf'}(\theta) \end{aligned}$$

[def **bwQ**] $\forall x' \bullet (\exists a : T \bullet \pi)! \Rightarrow \exists \widetilde{x'} \bullet \mathsf{UC} (\mathcal{L}, \exists a : T \bullet \pi)! \land (ok' \land \neg wait' \Rightarrow \theta')$ = [def var, Conf'] $\forall x' \bullet (\exists a : T \bullet \pi)! \Rightarrow \exists \widetilde{x'} \bullet \mathsf{UC} (\mathcal{L}, \exists a : T \bullet \pi)! \land \theta' \quad [\text{property of } \pi]$ = $\forall v' \bullet (\exists a : T \bullet \pi!) \Rightarrow \exists \widetilde{v'} \bullet (\exists \widetilde{a} : T \bullet \widetilde{\pi}!) \land \theta'$ $[w', \widetilde{w'} \text{ free in } \theta']$ = $\forall v' \bullet (\exists a : T \bullet \Pi!) \land a' \in T \Rightarrow \exists \widetilde{v'} \bullet (\exists \widetilde{a} : T \bullet \widetilde{\Pi}!) \land \widetilde{a'} \in T \land \theta'$ = [property of II] $\forall v' \bullet (\exists a : T \bullet \Pi!) \land a' \in T \Rightarrow (\exists \widetilde{a'} : T \bullet \theta')[\widetilde{x}/\widetilde{x'}]$ [one point rule] = $(\forall a': T \bullet (\exists \widetilde{a'}: T \bullet \theta')[\widetilde{x}/\widetilde{x'}])[x/x']$ [one point rule] = $\forall a: T \bullet \exists \widetilde{a}: T \bullet \theta$ [renaming] _

The proof that **bw** (\langle **end** *a* : *T* \rangle , θ) = $\forall a \bullet \exists \tilde{a} : T \bullet \theta$ is similar. Hence:

$$bw (var a : T \bullet B, \theta)$$

$$= bw (\langle var a : T \rangle; B; \langle end a : T \rangle, \theta) \qquad [def var]$$

$$= bw (\langle var a : T \rangle, bw (B, bw (\langle end a : T \rangle, \theta))) \qquad [Law 6.32, twice]$$

$$= \forall a : T \bullet \exists \tilde{a} : T \bullet bw (B, \forall a : T \bullet \exists \tilde{a} : T \bullet \theta)$$

Proof of Law 6.28

Let *X* abbreviate \neg *wait'* \wedge *tr'* = *tr* \wedge *u'* = *u* in the following:

$$\begin{aligned} \mathbf{bwQ} \left(\mathbf{UC} \left(\mathcal{L}, w : [Pre, Post] \right), \theta \right) &= \mathbf{bwQ} \left(\mathbf{UC} \left(\mathcal{L}, \mathbf{R} \left(Pre \vdash Post \land X \right) \right), \theta \right) & [\text{def } \mathbf{bwQ}, \text{spec stmt}] \\ &= & \forall x' \bullet \begin{pmatrix} Pre! \land ok' \land Post! \land X \\ \Rightarrow \exists \tilde{x'} \bullet \begin{pmatrix} \widetilde{Pre!} \Rightarrow (\widetilde{Post!} \land \tilde{X}) \\ \land \quad \widetilde{tr} \leq \widetilde{tr'} \land \mathcal{I}(\mathcal{L})! \land \text{Conf'}(\theta) \end{pmatrix} \end{pmatrix} \\ &\wedge \quad \widetilde{Pre!} \Rightarrow Pre! \end{aligned}$$

[Theorem 6.12]

$$= \forall v' \bullet \left(\begin{array}{c} (Pre \land Post \land u' = u) \Rightarrow \\ \exists \, \widetilde{v'} \bullet (\widetilde{Pre} \Rightarrow \widetilde{Post} \land \widetilde{u'} = \widetilde{u}) \land \theta' \end{array} \right) \land (\widetilde{Pre} \Rightarrow Pre)$$

[*Pre*, *Post*, θ over state variables]

Proof of Law 6.29

Assumptions are not miraculous, so **bw** $(\langle \{C\} \rangle, \theta) =$ **bwQ** $(\langle \{C\} \rangle, \theta)$.

 $\begin{aligned} & \mathsf{bwQ}\left(\mathsf{UC}\left(\mathcal{L}, \{C\}\right), \theta\right) \\ &= \mathsf{bwQ}\left(\mathsf{UC}\left(\mathcal{L}, \left[C, \mathsf{true}\right]\right), \theta\right) & [\text{def assumption}] \\ &= \forall v' \bullet \left(\left(C \land v' = v\right) \Rightarrow \exists \widetilde{v'} \bullet \left(\widetilde{C} \Rightarrow \widetilde{v'} = \widetilde{v}\right) \land \theta'\right) \land \left(\widetilde{C} \Rightarrow C\right) [\text{Law 6.28}] \\ &= (C \Rightarrow \exists \widetilde{v'} \bullet \left(\widetilde{C} \Rightarrow \widetilde{v'} = \widetilde{v}\right) \land \theta' [v/v']) \land \left(\widetilde{C} \Rightarrow C\right) & [\text{one point rule}] \end{aligned}$

Proof of Law 6.31

Assignment is not miraculous, so **bw** $(\langle a := E \rangle, \theta) =$ **bwQ** $(\langle a := E \rangle, \theta)$.

 $\mathsf{bwQ}\left(\mathsf{UC}\left(\mathcal{L},\left\{\,C\,\right\}\right),\theta\right)$

$$= \mathbf{bwQ} \left(\mathbf{UC} \left(\mathcal{L}, a : [\mathbf{true}, a' = E] \right), \theta \right)$$
 [def assignment]

$$= \forall v' \bullet \left((a' = E \land u' = u) \Rightarrow \exists \widetilde{v'} \bullet (\widetilde{a'} = \widetilde{E} \land \widetilde{u'} = \widetilde{u}) \land \theta' \right) \quad [Law \ 6.28]$$

$$= \theta'[E, \tilde{E}/a', \tilde{a'}][u/u']$$
 [one point rule]

$$= \theta[E, \widetilde{E}/a, \widetilde{a}]$$

Proof of Law 6.32

$$\begin{aligned} & \mathsf{bwR}\left(B_{1}\,\widehat{;}\,B_{2}\right) \\ &= \quad \exists x', \widetilde{x'} \bullet \left(B_{1}\,\widehat{;}\,B_{2}\right)! & [def\,\mathsf{bwR}] \\ &= \quad \exists x', \widetilde{x'} \bullet B_{1}!\,\widehat{;}\,\mathsf{Conf}\left(B_{2}!\right) & [property of sequential composition] \\ &= \quad \exists x', \widetilde{x'} \bullet \exists x_{0}, \widetilde{x_{0}} \bullet B_{1}![x_{0}, \widetilde{x_{0}}/x', \widetilde{x'}] \land \mathsf{Conf}\left(B_{2}!\right)[x_{0}, \widetilde{x_{0}}/x, \widetilde{x}] & [def\,\widehat{;}] \\ &= \quad \exists x_{0}, \widetilde{x_{0}} \bullet B_{1}![x_{0}, \widetilde{x_{0}}/x', \widetilde{x'}] \land \mathsf{Conf}\left(\exists x', \widetilde{x'} \bullet B_{2}!\right)[x_{0}, \widetilde{x_{0}}/x, \widetilde{x}] \\ &\quad [x', \widetilde{x'} \text{ free in } B_{1}[x_{0}, \widetilde{x_{0}}/x', \widetilde{x'}]] \\ &= \quad \exists x', \widetilde{x'} \bullet B_{1}! \land \mathsf{Conf}\left(\exists x', \widetilde{x'} \bullet B_{2}!\right)[x', \widetilde{x'}/x, \widetilde{x}] & [renaming] \\ &= \quad \mathsf{bwR}\left(B_{1} \land \mathsf{Conf}\left(\mathsf{bwR}\left(B_{2}\right)\right)[x', \widetilde{x'}/x, \widetilde{x}]\right) & [def\,\mathsf{bwR}] \\ &= \quad \mathsf{bwR}\left(B_{1} \land \mathsf{Conf'}\left(\mathsf{bwR}\left(B_{2}[x', \widetilde{x'}/x, \widetilde{x}]\right)\right)\right) & [renaming] \end{aligned}$$

$$\begin{aligned} \mathbf{bwQ} \left(B_{1} \ \hat{\mathbf{j}} \ B_{2}, \theta \right) \\ &= \forall x' \bullet \left(\mathsf{D} \left(B_{1} \ \hat{\mathbf{j}} \ B_{2} \right) ! \Rightarrow \exists \tilde{x'} \bullet \left(B_{1} \ \hat{\mathbf{j}} \ B_{2} \right) ! \land \mathsf{Conf'} \left(\theta \right) \right) \qquad [\text{def } \mathbf{bwQ}] \\ &= \forall x' \bullet \left(\left(\mathsf{D} \left(B_{1} \right) \ \mathbf{j} \ \mathsf{D} \left(B_{2} \right) \right) ! \Rightarrow \exists \tilde{x'} \bullet \left(B_{1} \ \hat{\mathbf{j}} \ B_{2} \right) ! \land \mathsf{Conf'} \left(\theta \right) \right) \qquad [\text{Condition } 5.11] \\ &= \forall x' \bullet \left(\begin{array}{c} (\exists x_{0} \bullet \left(\mathsf{D} \left(B_{1} \right) \left[x_{0} / x' \right] \land \mathsf{D} \left(B_{2} \right) \left[x_{0} / x \right] \right] ! \Rightarrow \\ &\exists \tilde{x'} \bullet \left(\exists x_{0}, \tilde{x_{0}} \bullet \left(\begin{array}{c} B_{1} \left[x_{0}, \tilde{x_{0}} / x' , \tilde{x'} \right] \\ \land & B_{2} \left[x_{0}, \tilde{x_{0}} / x , \tilde{x} \right] \end{array} \right) ! \land \mathsf{Conf'} \left(\theta \right) \right) \\ &= [\text{def } \mathbf{j}, \mathbf{j}] \end{aligned}$$

$$= \forall x' \bullet \begin{pmatrix} (\exists x_0 \bullet \mathbf{D}(B_1)![x_0/x'] \land \operatorname{Conf}(\mathbf{D}(B_2))[x_0/x]) \Rightarrow \\ \exists \widetilde{x'} \bullet \begin{pmatrix} \exists x_0, \widetilde{x_0} \bullet \begin{pmatrix} B_1![x_0, \widetilde{x_0}/x', \widetilde{x'}] \\ \land & \operatorname{Conf}(B_2)[x_0, \widetilde{x_0}/x, \widetilde{x}] \end{pmatrix} \land \operatorname{Conf}'(\theta) \end{pmatrix} \end{pmatrix}$$

[property of sequential composition]

$$= \forall x', x_0 \bullet \begin{pmatrix} \mathsf{D}(B_1)![x_0/x'] \Rightarrow \\ (\operatorname{Conf}(\mathsf{D}(B_2))[x_0/x] \Rightarrow \\ (\exists \tilde{x}_0 \bullet B_1![x_0, \tilde{x}_0/x', \tilde{x'}] \Rightarrow \\ \exists \tilde{x'} \bullet \operatorname{Conf}(B_2)[x_0, \tilde{x}_0/x, \tilde{x}] \wedge \operatorname{Conf'}(\theta) \end{pmatrix} \end{pmatrix} \end{pmatrix}$$

$$[pred calc]$$

$$= \forall x_0 \bullet \begin{pmatrix} \mathsf{D}(B_1)![x_0/x'] \Rightarrow \\ (\exists \tilde{x}_0 \bullet B_1![x_0, \tilde{x}_0/x', \tilde{x'}] \\ \wedge \forall x' \bullet \begin{pmatrix} \operatorname{Conf}(\mathsf{D}(B_2))[x_0/x] \Rightarrow \\ \exists \tilde{x'} \bullet \operatorname{Conf}(B_2)[x_0, \tilde{x}_0/x, \tilde{x}] \wedge \operatorname{Conf'}(\theta) \end{pmatrix} \end{pmatrix} \end{pmatrix}$$

$$[pred calc]$$

$$= \forall x_0 \bullet \begin{pmatrix} \mathsf{D}(B_1)![x_0/x'] \Rightarrow \\ \exists \tilde{x}_0 \bullet B_1![x_0, \tilde{x}_0/x', \tilde{x'}] \wedge \operatorname{Conf}(\mathsf{bwQ}(B_2, \theta))[x_0, \tilde{x}_0/x, \tilde{x}] \end{pmatrix}$$

$$[def \ \mathsf{bwQ}]$$

$$= \forall x' \bullet (\mathsf{D}(B_1)! \Rightarrow \exists \tilde{x'} \bullet B_1! \wedge \operatorname{Conf'}(\mathsf{bwQ}(B_2, \theta))) \quad [renaming]$$

$$= \mathsf{bwQ}(B_1, \mathsf{bwQ}(B_2, \theta)) \quad [def \ \mathsf{bwQ}]$$

Proof of Law 6.33

$$\begin{aligned} & \mathsf{bwR}\left(g\ \widehat{\&}\ B\right) \\ &= \left(\begin{array}{c} & \mathsf{U}\left(g\right)\wedge\mathsf{bwR}\left(g\ \widehat{\&}\ B\right) \\ & \lor & \mathsf{U}\left(\neg g\right)\wedge\mathsf{bwR}\left(g\ \widehat{\&}\ B\right) \\ & \lor & (g\neq \widetilde{g})\wedge\mathsf{bwR}\left(g\ \widehat{\&}\ B\right) \end{array}\right) & \text{[case split]} \end{aligned}$$

$$&= \left(\begin{array}{c} & \mathsf{U}\left(g\right)\wedge\mathsf{bwR}\left(B\right) \\ & \lor & \mathsf{U}\left(\neg g\right)\wedge\mathsf{bwR}\left(\widehat{R}\left(\mathsf{true}\vdash\mathsf{U}\left(tr'=tr\wedge\mathit{wait'}\right)\right)\right) \\ & \lor & (g\neq \widetilde{g})\wedge\mathsf{bwR}\left(\widehat{R}\left(\mathsf{true}\vdash\mathsf{false}\right)\right) \end{array}\right) & \text{[def}\ \widehat{\&}\] \end{aligned}$$

$$&= \left(\begin{array}{c} & \mathsf{U}\left(g\right)\wedge\mathsf{bwR}\left(B\right) \\ & \lor & \mathsf{U}\left(\neg g\right)\wedge\mathsf{true} \\ & \lor & (g\neq \widetilde{g})\wedge\mathsf{false} \end{array}\right) & \text{[def}\ \mathsf{bwR}\] \end{aligned}$$

This result simplifies the calculation of **bwQ**:

Proof of Law 6.34

The prefixing operator is not miraculous, so **bw** $(\langle c.E \rightarrow Skip \rangle, \theta) =$ **bwQ** ($\langle c.E \rightarrow Skip \rangle$, θ). Let *X* abbreviate:

$$(tr' = tr \land (c, E) \notin ref' \lhd wait' \rhd tr' = tr \land \langle (c, E) \rangle) \land v' = v$$

in the following:

$$\begin{aligned} & \mathsf{bwQ}\left(\mathsf{UC}\left(\mathcal{L}, c.E \to Skip\right), \theta\right) \\ &= \mathsf{bwQ}\left(\mathsf{UC}\left(\mathcal{L}, \mathsf{R}\left(\mathsf{true} \vdash X\right)\right), \theta\right) & [\text{def prefix}] \\ &= \forall x' \bullet \left(ok' \land X! \Rightarrow \exists \widetilde{x'} \bullet \widetilde{X}! \land \mathcal{I}(\mathcal{L})! \land \mathsf{Conf}'(\theta)\right) & [\text{Theorem 6.12}] \\ &\quad \forall x' \bullet \left(ok' \land wait' \land X! \Rightarrow \exists \widetilde{x'} \bullet \widetilde{X}! \land \mathcal{I}(\mathcal{L})!\right) \\ &= \land c \in \mathcal{L} \Rightarrow \forall x' \bullet \left(ok' \land \neg wait' \land X! \Rightarrow \exists \widetilde{x'} \bullet \widetilde{X}! \land \mathcal{I}(\mathcal{L})! \land \theta'\right) \\ &\quad \land c \notin \mathcal{L} \Rightarrow \forall x' \bullet \left(ok' \land \neg wait' \land X! \Rightarrow \exists \widetilde{x'} \bullet \widetilde{X}! \land \mathcal{I}(\mathcal{L})! \land \theta'\right) \\ &\quad (cse split) \end{aligned}$$

$$= c \in \mathcal{L} \Rightarrow \forall x' \bullet \begin{pmatrix} tr' = tr \land \langle (c, E) \rangle \land v' = v \Rightarrow \\ t\widetilde{r'} = \widetilde{tr} \land \langle (c, \widetilde{E}) \rangle \\ \land \widetilde{v'} = \widetilde{tr} \land \langle (c, \widetilde{E}) \rangle \\ \land \widetilde{v'} = \widetilde{v} \land \theta' \end{pmatrix} \end{pmatrix}$$

$$\land c \notin \mathcal{L} \Rightarrow \forall x' \bullet (v' = v \Rightarrow \exists \widetilde{x'} \bullet \widetilde{v'} = \widetilde{v} \land \theta')$$

$$= (c \in \mathcal{L} \Rightarrow \theta' [v, \widetilde{v}/v', \widetilde{v'}] \land E = \widetilde{E}) \land (c \notin \mathcal{L} \Rightarrow \theta' [v, \widetilde{v}/v', \widetilde{v'}])$$

$$[simplify]$$

$$= (c \in \mathcal{L} \Rightarrow \theta \land E = \widetilde{E}) \land (c \notin \mathcal{L} \Rightarrow \theta)$$

$$[renaming]$$

Proof of Law 6.37

$$\mathbf{bwR} \begin{pmatrix} B_1 \ \cap B_2 \end{pmatrix}$$

$$= \exists x', \widetilde{x'} \bullet B_1! \lor B_2! \qquad [def \mathbf{bwR}]$$

$$\Leftrightarrow \exists x', \widetilde{x'} \bullet B_1! \land \exists x', \widetilde{x'} \bullet B_2! \qquad [pred calc]$$

$$= \mathbf{bwR} (B_1) \land \mathbf{bwR} (B_2) \qquad [def \mathbf{bwR}]$$

$$\begin{aligned} & \mathbf{bwQ}\left(B_{1}\ \widehat{\sqcap}\ B_{2},\theta\right) \\ &= \forall x' \bullet \left(\mathbf{D}\left((B_{1} \lor B_{2})!\right) \Rightarrow \exists \widetilde{x'} \bullet (B_{1} \lor B_{2})! \land \mathsf{Conf'}(\theta)\right) & [\mathsf{def}\ \mathsf{bwQ}] \\ &= \forall x' \bullet \left(\mathbf{D}\left(B_{1}\right)! \lor \mathbf{D}\left(B_{2}\right)! \Rightarrow \exists \widetilde{x'} \bullet (B_{1}! \lor B_{2}!) \land \mathsf{Conf'}(\theta)\right) & [\mathsf{Law}\ 4.8] \\ &\leftarrow \left(\begin{array}{c} \forall x' \bullet \left(\mathbf{D}\left(B_{1}\right)! \Rightarrow \exists \widetilde{x'} \bullet B_{1}! \land \mathsf{Conf'}(\theta)\right) \\ \land \forall x' \bullet \left(\mathbf{D}\left(B_{2}\right)! \Rightarrow \exists \widetilde{x'} \bullet B_{2}! \land \mathsf{Conf'}(\theta)\right) \end{array} \right) & [\mathsf{pred}\ \mathsf{calc}] \\ &= \mathbf{bwQ}\left(B_{1},\theta\right) \land \mathsf{bwQ}\left(B_{2},\theta\right) & [\mathsf{def}\ \mathsf{bwQ}] \end{aligned}$$

A final refinement step glues the pieces together.

Proof of Law 6.38

We prove the result for an external choice between two alternatives. It is straightforward to generalise this proof to arbitary numbers of alternatives.

First, we apply the approach described in Subsection 6.4.1 to derive:

$$\mathbf{bw} \begin{pmatrix} B_1 \ \Box B_2, \theta \end{pmatrix}$$

$$\Leftrightarrow \begin{pmatrix} \mathbf{bw} ((B_1 \ \Box B_2) \land \mathbf{U} (tr' = tr \land wait'), \theta) \\ \land \mathbf{bw} ((B_1 \ \Box B_2) \land \neg \mathbf{U} (tr' = tr \land wait'), \theta) \end{pmatrix} [Lemma 6.19]$$

$$\Leftrightarrow \begin{pmatrix} \mathbf{bw} ((B_1 \ \Box B_2) \land \mathbf{U} (tr' = tr \land wait'), \theta) \\ \land \mathbf{bw} ((B_1 \ \Box B_2) \land \mathbf{U} (\neg (tr' = tr \land wait')), \theta) \end{pmatrix} [Law 4.13]$$

$$= \begin{pmatrix} \mathbf{bw} ((B_1 \ \Box B_2) \land \mathbf{U} (\neg (tr' = tr \land wait'), \theta) \\ \land \mathbf{bw} ((B_1 \ \Box B_2) \land \mathbf{U} (\neg (tr' = tr \land wait'), \theta) \end{pmatrix}$$
[property of \Box]

The last step is a consequence of Definition 5.27: when \neg ($tr' = tr \land wait'$) holds, (lifted) external choice is equivalent to the (lifted) reactive design form of (lifted) internal choice.

Assuming at least one of the guards holds in each state, we sketch how this result can be specialised to an external choice of guarded prefixes:

$$\mathbf{bw} \begin{pmatrix} g_1 \widehat{\&} c_1 \xrightarrow{\frown} B_1 \widehat{\Box} g_2 \widehat{\&} c_2 \xrightarrow{\frown} B_2, \theta \end{pmatrix}$$

$$\Leftarrow \begin{pmatrix} \mathbf{bw} \begin{pmatrix} g_1 \widehat{\&} c_1 \xrightarrow{\frown} B_1 \widehat{\Box} g_2 \widehat{\&} c_2 \xrightarrow{\frown} B_2 \\ \wedge \mathbf{U} (tr' = tr \land wait') \\ g_1 \widehat{\&} c_1 \xrightarrow{\frown} B_1 \widehat{\Box} g_2 \widehat{\&} c_2 \xrightarrow{\frown} B_2 \\ \wedge \mathbf{bw} \begin{pmatrix} d & \mathbf{U} (r' = tr \land wait') \\ & \mathbf{U} (\neg (tr' = tr \land wait')) \end{pmatrix}, \theta \end{pmatrix} \end{pmatrix}$$
[as above]

$$= \begin{pmatrix} \mathbf{bw} \left(\left(g_1 \ \widehat{\&} \ Stop \ \widehat{\Box} \ g_2 \ \widehat{\&} \ Stop \right) \land \mathbf{U} \left(tr' = tr \land wait' \right), \theta \right) \\ \land \mathbf{bw} \left(\left(g_1 \ \widehat{\&} \ B_1 \ \widehat{\Box} \ g_2 \ \widehat{\&} \ B_2 \right) \land \mathbf{U} \left(\neg \left(tr' = tr \land wait' \right) \right), \theta \right) \end{pmatrix}$$

[property of **bw**]

$$= \begin{pmatrix} \mathbf{bw} \left(g_1 \ \widehat{\&} \ Stop \ \widehat{\sqcap} g_2 \ \widehat{\&} \ Stop, \theta\right) \\ \wedge \ \mathbf{bw} \left(g_1 \ \widehat{\&} \ B_1 \ \widehat{\sqcap} g_2 \ \widehat{\&} \ B_2, \theta\right) \end{pmatrix}$$
[assumption]
$$\Leftrightarrow \begin{pmatrix} \mathbf{bw} \left(g_1 \ \widehat{\&} \ Stop, \theta\right) \land \mathbf{bw} \left(g_2 \ \widehat{\&} \ Stop, \theta\right) \\ \wedge \ \mathbf{bw} \left(g_1 \ \widehat{\&} \ B_1, \theta\right) \land \mathbf{bw} \left(g_2 \ \widehat{\&} \ B_2, \theta\right) \end{pmatrix}$$
[Law 6.37]
$$= \begin{pmatrix} (\mathbf{U} (g_1) \land \mathbf{true} \lor \mathbf{U} (\neg g_1)) \\ \wedge \ (\mathbf{U} (g_2) \land \mathbf{true} \lor \mathbf{U} (\neg g_2)) \\ \wedge \ (\mathbf{U} (g_1) \land \mathbf{bw} (B_1, \theta) \lor \mathbf{U} (\neg g_1)) \\ \wedge \ (\mathbf{U} (g_2) \land \mathbf{bw} (B_2, \theta) \lor \mathbf{U} (\neg g_2)) \end{pmatrix}$$
[Law 6.25, Law 6.33]
$$= \begin{pmatrix} (\mathbf{U} (g_1) \land \mathbf{bw} (B_1, \theta) \lor \mathbf{U} (\neg g_1)) \\ \wedge \ (\mathbf{U} (g_2) \land \mathbf{bw} (B_2, \theta) \lor \mathbf{U} (\neg g_2)) \end{pmatrix}$$
[prop calc]

List of References

- Jean-Raymond Abrial. *The B-Book: Assigning Programs to Meanings*. Cambridge University Press, New York, NY, USA, 1996. ISBN 0-521-49619-5.
- Jean-Raymond Abrial. *Modeling in Event-B: System and Software Engineering*. Cambridge University Press, first edition, June 2010. ISBN 0-521-89556-1.
- Rajeev Alur, Pavol Černý, and Steve Zdancewic. Preserving secrecy under refinement. In Automata, Languages and Programming, volume 4052 of Lecture Notes in Computer Science, pages 107–118. Springer Berlin / Heidelberg, 2006. doi: 10.1007/11787006_10.
- James M. Anderson. Why we need a new definition of information security. *Computers & Security*, 22(4):308–313, May 2003. doi: 10.1016/S0167-4048(03)00407-3.
- Ross Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems.* Wiley Computer Publishing, first edition, 2001.
- R. J. R. Back and K. Sere. Superposition refinement of reactive systems. *Formal Aspects of Computing*, 8(3):324–346, May 1996. ISSN 0934-5043. doi: 10.1007/BF01214918.
- Michael J. Banks and Jeremy L. Jacob. Unifying theories of confidentiality. In Shengchao Qin, editor, *3rd International Symposium on Unifying Theories of Programming (UTP 2010)*, volume 6445 of *Lecture Notes in Computer Science*, pages 120–136. Springer Berlin / Heidelberg, 2010a. doi: 10.1007/978-3-642-16690-7_5.

List of References

- Michael J. Banks and Jeremy L. Jacob. On modelling user observations in the UTP. In Shengchao Qin, editor, 3rd International Symposium on Unifying Theories of Programming (UTP 2010), volume 6445 of Lecture Notes in Computer Science, pages 101–119. Springer Berlin / Heidelberg, 2010b. doi: 10.1007/978-3-642-16690-7_4.
- Michael J. Banks and Jeremy L. Jacob. Specifying confidentiality in Circus. In FM 2011: Formal Methods, volume 6664 of Lecture Notes in Computer Science, pages 215–230. Springer Berlin / Heidelberg, 2011a. doi: 10.1007/978-3-642-21437-0_18.
- Michael J. Banks and Jeremy L. Jacob. Confidentiality annotations for Circus. In Christopher M. Poskitt, editor, *Proceedings of the Fourth York Doctoral Symposium on Computer Science (YDS 2011)*, pages 15–24.
 Department of Computer Science, University of York, UK, October 2011b.
- David E. Bell and Leonard J. La Padula. Secure computer systems: Volume I – mathematical foundations, Volume II – A mathematical model, Volume III – A refinement of the mathematical model. Technical Report MTR-2547, The MITRE Corporation, Bedford, Massachusetts, 1973.
- David E. Bell and Leonard J. La Padula. Secure computer system: Unified exposition and MULTICS interpretation. Technical Report ESD-TR-75-306, The MITRE Corporation, 1976.
- Riccardo Bresciani and Andrew Butterfield. Towards a UTP-style framework to deal with probabilities. Technical report, Foundations and Methods Group, Trinity College Dublin, Dublin, Ireland, August 2011.
- Michael Burrows, Martín Abadi, and Roger Needham. A logic of authentication. Technical Report 39, Digital Systems Research Center, February 1989.

- Andrew Butterfield. Saoíthin: A theorem prover for UTP. In *3rd International Symposium on Unifying Theories of Programming (UTP 2010),* Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2010.
- Ulf Carlsen. Cryptographic protocol flaws: know your enemy. In *7th IEEE Computer Security Foundations Workshop (CSFW '94)*, pages 192– 200. IEEE Computer Society Press, 1994. doi: 10.1109/CSFW.1994. 315934.
- A. L. C. Cavalcanti and J. C. P. Woodcock. Predicate transformers in the semantics of Circus. *IEE Proceedings Software*, 150(2):85–94, April 2003. ISSN 1462-5970. doi: 10.1049/ip-sen:20030131.
- Ana Cavalcanti and Jim Woodcock. ZRC A refinement calculus for Z. *Formal Aspects of Computing*, 10(3):267–289–289, March 1998. ISSN 0934-5043. doi: 10.1007/s001650050016.
- Ana Cavalcanti and Jim Woodcock. A tutorial introduction to CSP in Unifying Theories of Programming. In *Refinement Techniques in Software Engineering*, volume 3167 of *Lecture Notes in Computer Science*, pages 220–268. Springer Berlin / Heidelberg, 2006. doi: 10.1007/11889229_6.
- Ana Cavalcanti, Augusto Sampaio, and Jim Woodcock. A refinement strategy for Circus. *Formal Aspects of Computing*, 15(2-3):146–181, November 2003. doi: 10.1007/s00165-003-0006-5.
- Ana Cavalcanti, Augusto Sampaio, and Jim Woodcock. Unifying classes and processes. *Software and Systems Modeling*, 4(3):277–296–296, July 2005. ISSN 1619-1366. doi: 10.1007/s10270-005-0085-2.
- K. Mani Chandy and Jayadev Misra. Parallel Program Design: A Foundation. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1988. ISBN 0-201-05866-9.
- Tom Chothia and Apratim Guha. A statistical test for information leaks using continuous mutual information. In 2011 IEEE 24th Computer Security Foundations Symposium (CSF), pages 177–190, Los Alamitos,

CA, USA, June 2011. IEEE. ISBN 9781-612-8464-4-6. doi: 10.1109/CSF. 2011.19.

- David D. Clark and David R. Wilson. A comparison of commercial and military computer security policies. In *Proceedings of the 1987 IEEE Symposium on Security and Privacy*, pages 184–194. IEEE Computer Society Press, 1987. doi: 10.1109/SP.1987.10001.
- Patrick Cousot and Radhia Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In POPL '77: Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of Programming Languages, pages 238– 252, New York, NY, USA, 1977. ACM. doi: 10.1145/512950.512973.
- Patrick Cousot and Radhia Cousot. Abstract interpretation and application to logic programs. *The Journal of Logic Programming*, 13(2-3): 103–179, July 1992. ISSN 07431066. doi: 10.1016/0743-1066(92)90030-7.
- Dorothy E. Denning. *Cryptography and Data Security*. Addison-Wesley Longman Publishing Company, Inc., Boston, MA, USA, 1982. ISBN 0-201-10150-5.
- Edsger W. Dijkstra. *A Discipline of Programming*. Prentice-Hall Series in Automatic Computation. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, October 1976. ISBN 0-13-215871-X.
- Edsger W. Dijkstra and Carel S. Scholten. *Predicate calculus and program semantics*. Texts and Monographs in Computer Science. Springer-Verlag New York, Inc., New York, NY, USA, 1990. ISBN 0-387-96957-8.
- Danny Dolev and Andrew C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, March 1983. doi: 10.1109/TIT.1983.1056650.
- Abderrahmane Feliachi, Marie-Claude Gaudel, and Burkhart Wolff. Isabelle/Circus: a process specification and verification environment. Technical Report 1547, LRI, Université Paris-Sud XI, November 2011.

- Riccardo Focardi and Roberto Gorrieri. A taxonomy of security properties for process algebras. *Journal of Computer Security*, 3(1):5–34, 1995.
- Simon N. Foley. Believing the integrity of a system (invited talk). *Electronic Notes in Theoretical Computer Science*, 125(1):3–12, 2005. doi: 10.1016/j.entcs.2004.09.037.
- Paweł Gancarski and Andrew Butterfield. The denotational semantics of slotted-Circus. In Ana Cavalcanti and Dennis R. Dams, editors, FM 2009: Formal Methods, volume 5850 of Lecture Notes in Computer Science, pages 451–466, Berlin, Heidelberg, 2009. Springer Berlin / Heidelberg. ISBN 978-3-642-05088-6. doi: 10.1007/978-3-642-05089-3_29.
- Joseph A. Goguen and José Meseguer. Security policies and security models. In *Proceedings of the 1982 IEEE Symposium on Security and Privacy*, pages 11–20. IEEE Computer Society, April 1982.
- Joseph A. Goguen and José Meseguer. Unwinding and inference control. In *Proceedings of the 1984 IEEE Symposium on Security and Privacy*, pages 75–86. IEEE Computer Society, 1984. doi: 10.1109/SP.1984.10019.
- David Gries. *The Science of Programming*. Texts and Monographs in Computer Science. Springer, 1981. ISBN 978-0-387-96480-5.
- Jifeng He, C. A. R. Hoare, and Jeff W. Sanders. Data refinement refined (resumé). In Bernard Robinet and Reinhard Wilhelm, editors, ESOP'86: Proceedings of the European Symposium on Programming, volume 213 of Lecture Notes in Computer Science, pages 187–196. Springer Berlin / Heidelberg, 1986. doi: 10.1007/3-540-16442-1_14.
- Eric C. R. Hehner. Predicative programming part I. *Communications of the ACM*, 27(2):134–143, 1984a. ISSN 0001-0782. doi: 10.1145/69610.357988.
- Eric C. R. Hehner. Predicative programming part II. *Communications of the ACM*, 27(2):144–151, 1984b. ISSN 0001-0782. doi: 10.1145/69610.357990.

List of References

- C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, October 1969. ISSN 0001-0782. doi: 10.1145/363235.363259.
- C. A. R. Hoare. Programming: Sorcery or science? *IEEE Software*, 1(2): 5–16, April 1984a. ISSN 0740-7459. doi: 10.1109/MS.1984.234042.
- C. A. R. Hoare. Programs are predicates. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 312 (1522), 1984b. ISSN 00804614. doi: 10.2307/37446.
- C. A. R. Hoare. *Communicating Sequential Processes*. Prentice Hall International Series in Computer Science. Prentice Hall, Inc., 1985a.
- C. A. R. Hoare. A couple of novelties in the propositional calculus. Zeitschrift für Mathematische Logik und Grundlagen der Mathematik, 31 (9-12):173–178, 1985b. ISSN 1521-3870. doi: 10.1002/malq.19850310905.
- C. A. R. Hoare and Jifeng He. *Unifying Theories of Programming*. Prentice Hall International Series in Computer Science. Prentice Hall Inc., 1998.
- Graham Hutton. A tutorial on the universality and expressiveness of fold. *Journal of Functional Programming*, 9(4):355–372, July 1999. ISSN 0956-7968. doi: 10.1017/S0956796899003500.
- Jeremy L. Jacob. On Shared Systems. PhD thesis, Oxford University, 1987.
- Jeremy L. Jacob. Security specifications. In *Proceedings of the 1988 IEEE* Symposium on Security and Privacy, pages 14–23, 1988. doi: 10.1109/ SECPRI.1988.8094.
- Jeremy L. Jacob. On the derivation of secure components. In *Proceedings of the 1989 IEEE Symposium on Security and Privacy*, pages 242–247. IEEE Computer Society, 1989a. doi: 10.1109/SECPRI.1989.36298.
- Jeremy L. Jacob. Refinement of shared systems. In John A. McDermid, editor, *The Theory and Practice of Refinement: Approaches to the Develop-*

ment of Large-Scale Software Systems, pages 27–36. Butterworths, 1989b. ISBN 0-408-03981-7.

- Jeremy L. Jacob. Separability and the detection of hidden channels. *Information Processing Letters*, 34(1):27–29, 1990. doi: 10.1016/ 0020-0190(90)90225-M.
- Jeremy L. Jacob. A uniform presentation of confidentiality properties. *IEEE Transactions on Software Engineering*, 17(11):1186–1194, 1991. doi: 10.1109/32.106973.
- Jeremy L. Jacob. Basic theorems about security. *Journal of Computer Security*, 1(4):385–411, 1992.
- Dale M. Johnson and F. Javier Thayer. Security and the composition of machines. In *First IEEE Computer Security Foundations Workshop* -*CSFW'88*, pages 72–89. IEEE Computer Society, May 1988.
- Cliff B. Jones. *Systematic software development using VDM*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, second edition, 1990. ISBN 0-13-880733-7.
- Nigel A. Jones. Building in... information security, privacy and assurance
 A high-level roadmap. Technical report, Cyber Security Knowledge Transfer Network (CSKTN), 2009.
- Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In *Advances in Cryptology – CRYPTO '99*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397. Springer-Verlag Berlin / Heidelberg, 1999. doi: 10.1007/3-540-48405-1_25.
- Butler W. Lampson. A note on the confinement problem. *Communications* of the ACM, 16(10):613–615, October 1973. doi: 10.1145/362375.362389.
- Heiko Mantel. Unwinding possibilistic security properties. In *Computer Security ESORICS 2000*, volume 1895 of *Lecture Notes in Computer Science*, pages 238–254. Springer Berlin / Heidelberg, 2000a. doi: 10. 1007/10722599_15.

List of References

- Heiko Mantel. Possibilistic definitions of security an assembly kit. In 13th IEEE Computer Security Foundations Workshop (CSFW '00), pages 185–199, 2000b. doi: 10.1109/CSFW.2000.856936.
- Heiko Mantel. Preserving information flow properties under refinement. In *Proceedings of the 2001 IEEE Symposium on Security and Privacy*, pages 78–91. IEEE Computer Society Press, 2001. doi: 10.1109/SECPRI.2001. 924289.
- Heiko Mantel. A Uniform Framework for the Formal Specification and Verification of Information Flow Security. PhD thesis, Universität Saarbrücken, July 2003.
- Daryl McCullough. Specifications for multi-level security and a hook-up property. In *Proceedings of the 1987 IEEE Symposium on Security and Privacy*, pages 161–166, Los Alamitos, CA, USA, 1987. IEEE Computer Society. doi: 10.1109/SP.1987.10009.
- Annabelle McIver. The secret art of computer programming. In Martin Leucker and Carroll Morgan, editors, *Theoretical Aspects of Computing ICTAC 2009*, volume 5684 of *Lecture Notes in Computer Science*, chapter 3, pages 61–78–78. Springer Berlin / Heidelberg, Berlin, Heidelberg, 2009. ISBN 978-3-642-03465-7. doi: 10.1007/978-3-642-03466-4_3.
- Annabelle McIver and Carroll Morgan. Compositional refinement in agent-based security protocols. *Formal Aspects of Computing*, pages 1–27, November 2010. ISSN 0934-5043. doi: 10.1007/s00165-010-0164-1.
- John McLean. Reasoning about security models. In *Proceedings of the 1987 IEEE Symposium on Security and Privacy*, pages 123–131. IEEE Computer Society Press, 1987. doi: 10.1109/SP.1987.10020.
- John McLean. A general theory of composition for trace sets closed under selective interleaving functions. In *Proceedings of the 1994 IEEE Symposium on Security and Privacy*, pages 79–93, 1994a. doi: 10.1109/ RISP.1994.296590.

- John McLean. Security models. In John Marciniak, editor, *Encyclopedia of Software Engineering*, volume 2, pages 1136–1145. John Wiley & Sons, Inc., 1994b.
- Carroll Morgan. *Programming from Specifications*. Prentice Hall International Series in Computer Science. Prentice Hall Inc., Hertfordshire, UK, second edition, 1994.
- Carroll Morgan. The shadow knows: Refinement and security in sequential programs. *Science of Computer Programming*, 74(8):629–653, June 2009. ISSN 0167-6423. doi: 10.1016/j.scico.2007.09.003.
- Carroll Morgan. Compositional noninterference from first principles. *Formal Aspects of Computing*, 24(1):3–26, January 2012. ISSN 1433-299X. doi: 10.1007/s00165-010-0167-y.
- Chunyan Mu. Quantitative information flow for security: a survey. Technical Report TR-08-06, Department of Computer Science, King's College London, September 2008.
- Greg Nelson. A generalization of Dijkstra's calculus. *ACM Transactions on Programming Languages and Systems*, 11(4):517–561, October 1989. ISSN 0164-0925. doi: 10.1145/69558.69559.
- NIST. Minimum security requirements for federal information and information systems. Technical Report 200, National Institute of Standards and Technology, Gaithersburg, MD, 2006.
- Colin O'Halloran. A calculus of information flow. In *ESORICS 90 First European Symposium on Research in Computer Security*, pages 147–159. AFCET, 1990.
- M. V. M. Oliveira, A. C. Gurgel, and C. G. Castro. CRefine: Support for the Circus refinement calculus. In *Sixth IEEE International Conference on Software Engineering and Formal Methods*, pages 281–290. IEEE, November 2008. ISBN 978-0-7695-3437-4. doi: 10.1109/SEFM.2008.9.

List of References

- Marcel Oliveira, Ana Cavalcanti, and Jim Woodcock. A UTP semantics for Circus. *Formal Aspects of Computing*, 21(1):3–32, February 2009. doi: 10.1007/s00165-007-0052-5.
- Marcel V. Oliveira. *Formal Derivation of State-Rich Reactive Programs using Circus*. PhD thesis, Department of Computer Science, University of York, 2005.
- A. W. Roscoe. Model-checking CSP. In A. W. Roscoe, editor, A Classical Mind, Essays in Honour of C.A.R. Hoare, chapter Model-checking CSP, pages 353–378. Prentice Hall International (UK) Ltd., Hertfordshire, UK, 1994. ISBN 0-13-294844-3.
- A. W. Roscoe. CSP and determinism in security modelling. In *Proceedings* of the 1995 IEEE Symposium on Security and Privacy, pages 114–127. IEEE Computer Society, 1995. doi: 10.1109/SECPRI.1995.398927.
- A. W. Roscoe. *The Theory and Practice of Concurrency*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1997. ISBN 0-13-674409-5.
- A. W. Roscoe, J. C. P. Woodcock, and L. Wulf. Non-Interference through determinism. In ESORICS '94: Proceedings of the Third European Symposium on Research in Computer Security, volume 875 of Lecture Notes in Computer Science, pages 33–53. Springer Berlin / Heidelberg, 1994. doi: 10.1007/3-540-58618-0_55.
- John Rushby. The design and verification of secure systems. In *Eighth ACM Symposium on Operating System Principles (SOSP '81)*, pages 12–21. ACM, December 1981. doi: 10.1145/800216.806586.
- Peter Ryan. Mathematical models of computer security. In *Foundations of Security Analysis and Design*, volume 2171 of *Lecture Notes in Computer Science*, pages 1–62. Springer Berlin / Heidelberg, 2001. doi: 10.1007/3-540-45608-2_1.

Peter Ryan and Steve Schneider. An attack on a recursive authentication

protocol. A cautionary tale. *Information Processing Letters*, 65(1):7–10, 1998. ISSN 0020-0190. doi: 10.1016/S0020-0190(97)00180-4.

- Peter Ryan and Steve A. Schneider. Process algebra and non-interference. In *12th IEEE Computer Security Foundations Workshop (CSFW '99)*, pages 214–227. IEEE Computer Society, 1999. doi: 10.1109/CSFW.1999.779775.
- Peter Ryan, John McLean, Jon Millen, and Vergil Gligor. Non-interference: who needs it? In 14th IEEE Computer Security Foundations Workshop (CSFW '01), pages 237–238. IEEE Computer Society, 2001. doi: 10.1109/ CSFW.2001.930149.
- Hossein Saiedian. An invitation to formal methods. *Computer*, 29(4): 16–17, April 1996. ISSN 0018-9162. doi: 10.1109/MC.1996.488298.
- Thomas Santen. A formal framework for confidentiality-preserving refinement. In *Computer Security – ESORICS 2006*, volume 4189 of *Lecture Notes in Computer Science*, pages 225–242. Springer Berlin / Heidelberg, 2006. doi: 10.1007/11863908_15.
- Thomas Santen. Preservation of probabilistic information flow under refinement. *Information and Computation*, 206(2-4):213–249, February 2008. doi: 10.1016/j.ic.2007.07.008.
- Thomas Santen, Maritta Heisel, and Andreas Pfitzmann. Confidentiality-Preserving refinement is compositional – sometimes. In Dieter Gollmann, Günther Karjoth, and Michael Waidner, editors, *Computer Security – ESORICS 2002*, volume 2502 of *Lecture Notes in Computer Science*, pages 194–211. Springer Berlin / Heidelberg, 2002. ISBN 978-3-540-44345-2. doi: 10.1007/3-540-45853-0_12.
- Steve Schneider. Concurrent and Real Time Systems: The CSP Approach (Worldwide Series in Computer Science). John Wiley & Sons, September 1999. ISBN 0-471-62373-3.
- Steve Schneider. *The B-Method: An Introduction*. Cornerstones of Computing. Palgrave Macmillan, October 2001. ISBN 0-333-79284-X.

List of References

- Steve Schneider and Helen Treharne. CSP theorems for communicating B machines. *Formal Aspects of Computing*, 17(4):390–422, December 2005. ISSN 0934-5043. doi: 10.1007/s00165-005-0076-7.
- Bruce Schneier. *Secrets and Lies: Digital Security in a Networked World*. Wiley, 2004.
- Dana Scott and Christopher Strachey. Toward A mathematical semantics for computer languages. In Jerome Fox, editor, *Proceedings of the Symposium on Computers and Automata*, volume XXI, pages 19–46, Brooklyn, N.Y., April 1971. Polytechnic Press.
- Fredrik Seehusen and Ketil Stølen. Maintaining information flow security under refinement and transformation. In *Formal Aspects in Security and Trust*, volume 4691 of *Lecture Notes in Computer Science*, pages 143–157.
 Springer Berlin / Heidelberg, 2007. doi: 10.1007/978-3-540-75227-1_10.
- Fredrik Seehusen and Ketil Stølen. Information flow security, abstraction and composition. *IET Information Security*, 3(1):9–33, 2009. doi: 10. 1049/iet-ifs:20080069.
- Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, July 1948.
- Claude E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, 1949.
- J. M. Spivey. *The Z Notation: a reference manual*. Prentice Hall, second edition, 1992.
- Clifford Stoll. Stalking the wily hacker. *Communications of the ACM*, 31 (5):484–497, May 1988. ISSN 0001-0782. doi: 10.1145/42411.42412.
- Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955.

- Pavol Černý. *Software Model Checking for Confidentiality*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, 2009.
- Kun Wei, Jim Woodcock, and Alan Burns. Formalising the timebands model in timed Circus. Technical report, Department of Computer Science, University of York, York, UK, August 2010.
- Jeanette M. Wing. A symbiotic relationship between formal methods and security. In *Conference on Computer Security, Dependability, and Assurance: From Needs to Solutions,* pages 26–38. IEEE Computer Society, 1998. doi: 10.1109/CSDA.1998.798355.
- Jeannette M. Wing. A specifier's introduction to formal methods. *Computer*, 23(9):8–24, 1990. doi: 10.1109/2.58215.
- Niklaus Wirth. Program development by stepwise refinement. *Communications of the ACM*, 14(4):221–227, April 1971. ISSN 0001-0782. doi: 10.1145/362575.362577.
- Jim Woodcock. The miracle of reactive programming. In Andrew Butterfield, editor, *Unifying Theories of Programming*, volume 5713 of *Lecture Notes in Computer Science*, chapter 12, pages 202–217. Springer Berlin / Heidelberg, Berlin, Heidelberg, 2010. ISBN 978-3-642-14520-9. doi: 10.1007/978-3-642-14521-6_12.
- Jim Woodcock and Ana Cavalcanti. A concurrent language for refinement. In Andrew Butterfield, Glenn Strong, and Claus Pahl, editors, *5th Irish Workshop on Formal Methods*, Workshops in Computing. BCS, 2001.
- Jim Woodcock and Ana Cavalcanti. The semantics of Circus. In ZB 2002: Formal Specification and Development in Z and B, volume 2272 of Lecture Notes in Computer Science, chapter 10, pages 184–203. Springer Berlin / Heidelberg, 2002. doi: 10.1007/3-540-45648-1_10.
- Jim Woodcock and Ana Cavalcanti. A tutorial introduction to designs in Unifying Theories of Programming. In *Integrated Formal Methods*,

List of References

volume 2999 of *Lecture Notes in Computer Science*, pages 40–66. Springer Berlin / Heidelberg, 2004. ISBN 978-3-540-21377-2. doi: 10.1007/ b96106.

- Jim Woodcock and Jim Davies. *Using Z: Specification, Refinement, and Proof.* Prentice Hall International Series in Computer Science. Prentice Hall, 1996.
- Jim Woodcock, Peter G. Larsen, Juan Bicarregui, and John Fitzgerald.
 Formal methods: Practice and experience. ACM Computing Surveys, 41
 (4):1–36, October 2009. ISSN 0360-0300. doi: 10.1145/1592434.1592436.
- Michael Workman and John Gathegi. Punishment and ethics deterrents: A study of insider security contravention. *Journal of the American Society for Information Science and Technology*, 58(2):212–222, January 2007. ISSN 1532-2890. doi: 10.1002/asi.20474.
- Aris Zakinthinos and E. S. Lee. A general theory of security properties. In *Proceedings of the 1997 IEEE Symposium on Research in Security and Privacy*, pages 94–102, 1997. doi: 10.1109/SECPRI.1997.601322.
- Frank Zeyda, Bill Stoddart, and Steve Dunne. The refinement of reversible computations. In *2nd International Workshop on Refinement of Critical Systems*, 2003.