

**South African
Computer
Journal
Number 19
February 1997**

**Suid-Afrikaanse
Rekenaar-
tydskrif
Nommer 19
Februarie 1997**

**Computer Science
and
Information Systems**

Proceedings of WOFACS'96

**Rekenaarwetenskap
en
Inligtingstelsels**

**The South African
Computer Journal**

*An official publication of the Computer Society
of South Africa and the South African Institute of
Computer Scientists*

**Die Suid-Afrikaanse
Rekenaartydskrif**

*'n Amptelike publikasie van die Rekenaarvereniging
van Suid-Afrika en die Suid-Afrikaanse Instituut
vir Rekenaarwetenskaplikes*

Editor

Professor Derrick G Kourie
Department of Computer Science
University of Pretoria
Hatfield 0083
dkourie@dos-lan.cs.up.ac.za

Subeditor: Information Systems

Prof Lucas Introna
Department of Informatics
University of Pretoria
Hatfield 0083
lintrona@econ.up.ac.za

Production Editor

Dr Riël Smit
Mosaic Software (Pty) Ltd
P.O.Box 23906
Claremont 7735
gds@mosaic.co.za

World-Wide Web: <http://www.mosaic.co.za/sacj/>

Editorial Board

Professor Judy M Bishop
University of Pretoria, South Africa
jbishop@cs.up.ac.za

Professor R Nigel Horspool
University of Victoria, Canada
nigelh@csr.csc.uvic.ca

Professor Richard J Boland
Case Western Reserve University, USA
boland@spider.cwrw.edu

Professor Fred H Lochovsky
University of Science and Technology, Hong Kong
fred@cs.ust.hk

Professor Ian Cloete
University of Stellenbosch, South Africa
ian@cs.sun.ac.za

Professor Kalle Lyytinen
University of Jyvaskyla, Finland
kalle@cs.jyu.fi

Professor Trevor D Crossman
University of Natal, South Africa
crossman@bis.und.ac.za

Doctor Jonathan Miller
University of Cape Town, South Africa
jmiller@gsb2.uct.ac.za

Professor Donald D Cowan
University of Waterloo, Canada
dcowan@csg.uwaterloo.ca

Professor Mary L Soffa
University of Pittsburgh, USA
soffa@cs.pitt.edu

Professor Jürg Gutknecht
ETH, Zürich, Switzerland
gutknecht@inf.ethz.ch

Professor Basie H von Solms
Rand Afrikaanse Universiteit, South Africa
basie@rkw.rau.ac.za

Subscriptions

	Annual	Single copy
Southern Africa:	R50,00	R25,00
Elsewhere:	\$30,00	\$15,00

An additional \$15 per year is charged for airmail outside Southern Africa

to be sent to:

*Computer Society of South Africa
Box 1714 Halfway House 1685*

Introduction

WOFACS '96: Workshop on Formal and Applied Computer Science

Chris Brink

Laboratory for Formal Aspects and Complexity in Computer Science, Department of Mathematics and Applied Mathematics, University of Cape Town
cbrink@maths.uct.ac.za

“What I tell you three times is true”, said the Bellman in Lewis Carroll’s *The Hunting of the Snark*. By somewhat the same principle, it is often held that three events of the same kind serve to make a series. And so, after WOFACS '92, '94 and now '96, we may claim to have a well-established biennial Southern African series of Workshops in the area of Formal Aspects of Computer Science.

This issue of the SACJ is devoted to the *Proceedings* of WOFACS '96. In other words, it contains survey articles written especially for this volume by the invited speakers (and their collaborators), on the topics they lectured on during the Workshop. These were:

- **Dr Maarten de Rijke** (University of Warwick): *Reasoning with Incomplete and Changing Information*.
- **Dr Holger Schlingloff** (Technische Universität München): *Verification of Finite-state Systems with Temporal Model Checking*.
- **Prof Jan Peleska** (Universität Bremen): *Test Automation for Safety-Critical Reactive Systems*.
- **Dr Jeff Sanders** (Oxford University): *Application-oriented Program Semantics*.

The format of WOFACS '96 followed the same pattern as before. Each speaker gave a course of 10 lectures, at a rate of one lecture per day. Material was pitched at about Honours level, and students had the opportunity of offering WOFACS courses for credit in their degree programmes at their respective home universities. Those who took up this option did some exercises and assignments and were evaluated by the speaker(s) concerned, thus gaining valuable

insight into material, methods and expectations at an international level. WOFACS '96 was organised by FACCs-Lab (the Laboratory for Formal Aspects and Complexity in Computer Science), and was co-hosted by the Department of Mathematics and Applied Mathematics and the Department of Computer Science at the University of Cape Town. Accommodation was available in a University residence, and we were able to make available some financial support for travel and accommodation to participants (especially students) who could not obtain funding from their home institutions. Attendance stood at about 50 participants. Cape Town is a pleasant place to visit, even in winter, and we took care to have suitable outings and social events for our visitors. I am pleased to be able to mention that the WOFACS series has now attracted international attention, and that WOFACS 98 is being planned under the auspices of IFIP (the International Federation of Information Processing), specifically Working Group 2.3 on Programming Methodology.

No event of this nature can succeed without the hard work of a number of people. I would like to express my grateful thanks to:

- the invited speakers, for the care they took and the quality of their presentations;
- the Foundation for Research Development and the UCT Research Committee, for sponsorship;
- Diana Dixon, Jeanne Weir, Peter Jipsen and other FACCs-Lab staff members, for their hard work, and
- all participants, for attending.

SACJ is produced with kind support from
Mosaic Software (Pty) Ltd.

Application-Oriented Program Semantics

A K McIver Carroll Morgan J W Sanders
Programming Research Group, Oxford University, U K

Abstract

This paper abridges lecture notes from WOFACS 96. It provides semantic models for a variety of programming and development formalisms, showing how different models for the same formalism are related by Galois connections. The formalisms are: the language of guarded commands and its refinement calculus; communicating sequential processes; and probabilistic imperative code. The semantic models include: binary relations; predicate transformers; traces; failures; and distributions. In each case denotations of code are characterised, in the more general specification space, by the healthiness conditions they satisfy. Models are evaluated for their elegance, expressive power and calculational facility.

Expanded notes containing proofs, an extra chapter on data refinement and exercises for each chapter, are available from the authors.

Keywords: *Semantics, formal methods, refinement, data refinement, process algebra, probability, Galois connection, binary relations, predicate transformers.*

Computing Review Categories: *D.1, D.2, F.3, G.3.*

1 Introduction

It is usual for a course in program semantics to be either general (comparing the different styles of semantics: operational, denotational and algebraic) or specific, studying in one style the semantics of a programming language. Each type of course has its advantages and disadvantages. The former conveys an appreciation of the conditions under which each type of semantics is appropriate, by being broad though shallow. The latter is usually more pragmatic, concentrating on the details of a (typically obsolete) language at the expense of theory.

This course in *Application-Oriented Program Semantics* has been designed to combine the advantages of both those types of course, in a contemporary setting. It emphasises the uses of the semantics of imperative programming languages without concentrating exclusively on theory. For each language studied, an appropriate domain is introduced and its salient features characterised mathematically; Galois connections are used to relate the various semantic models. A feature is our insistence that semantic domains be broad enough to embrace not just programs but also the constructs which arise in program development (like miracles and data refinement).

Outline

After settling upon notation and studying what we need of Galois connections between partial orders, we begin the course proper by studying the intuitively compelling relational model of imperative programs. It provides denotations of code, specification constructs, and development rules including two rules for data refinement which are together complete. However use of weakest pre and post-specification for program development is not simple, so the model has limited applicability.

On investigating the lattice structure of the relational

model we are led to study predicate transformers: the *premier* semantic model of imperative programs. The two models are related by a Galois connection, which imparts healthiness conditions to the predicate-transformer denotations of code. Predicate-transformers are shown to provide a model of the refinement calculus in which there is a single complete refinement rule. In the guise of the refinement calculus they support a viable development method.

Next, the notation CSP (for Communicating Sequential Processes) is summarised, and its three basic semantic models related by Galois connections. An important feature of the study is its use of both the maximal elements and the compact elements in its semantic domains. It provides further evidence for our underlying claims about the utility of Galois connections and the particular form of their constituent functions.

Finally a probabilistic lifting is investigated on semantic domains and applied to the study of some random algorithms! Properties of the probabilistic semantic domain are considered and probabilistic refinement rules demonstrated. This example is included to demonstrate the utility of the techniques, exposed in these lectures, in a novel situation.

The semantic domains which arise in this course are studied both in their own right, and as vehicles for conveying the aims of the course.

Aims

The general aims of this course are:

1. to further familiarity with the use of mathematical techniques (including algebra, category theory and topology) in Computer Science;
2. to observe how domains arise in different semantic contexts and to learn how to recognise them and their salient features (e.g. compact elements, maximal elements);

3. to understand how Galois connections are used to relate different semantic models;
4. to appreciate the four major semantic models studied in the course (relational, predicate-transformer, communicating-process and probabilistic), their relationship and individual characteristics;
5. to understand the manner in which the laws of a refinement calculus are proved sound in a semantic model and be able to translate that understanding to specific domains;
6. to be able to choose an appropriate semantic model for a given application, to identify an appropriate 'healthy' subdomain of it, and to relate it to alternatives.

2 Notation

The purpose of this chapter is to provide basic notation for what follows. Since most of the underlying concepts are expected to be familiar, the treatment is brief.

Sets and logic

The symbol $\hat{=}$ means *equal by definition*. The *natural numbers*, *real numbers* and *Booleans* are denoted respectively by \mathbf{N} , \mathbf{R} and \mathbf{B} . Intervals of natural numbers are expressed $a..b \hat{=} \{n : \mathbf{N} \mid a \leq n < b\}$.

When the "universe of discourse" \mathbf{X} is clear, if $\mathbf{E} \subseteq \mathbf{X}$ then $\bar{\mathbf{E}}$ means the complement of \mathbf{E} in \mathbf{X} ; otherwise we write it explicitly as $\mathbf{X} \setminus \mathbf{E}$.

If \mathbf{X} is a finite set, $\#\mathbf{X}$ denotes its cardinality.

The predicate "*property q holds for all x of type X*" is written $\forall x : \mathbf{X} \bullet q$. The set of functions from \mathbf{X} to \mathbf{Y} is denoted $\mathbf{X} \rightarrow \mathbf{Y}$ with one exception: we regard a predicate as a Boolean-valued function, and write $\mathbf{P}\mathbf{X}$ for the set of all predicates on \mathbf{X} $\mathbf{P}\mathbf{X} \hat{=} \mathbf{X} \rightarrow \mathbf{B}$. The characteristic function of a set $\mathbf{E} \subseteq \mathbf{X}$ is denoted $\xi_{\mathbf{E}} : \mathbf{P}\mathbf{X}$.

When clarity requires it, the conjunction of two predicates, traditionally written *why* \wedge *wherefore*, is written

$$\left(\begin{array}{l} \text{why} \\ \text{wherefore} \end{array} \right).$$

We use ternary infix notation for the conditional: for predicates p , q and r ,

$$p \triangleleft r \triangleright q \hat{=} (p \wedge r) \vee (q \wedge \neg r)$$

which is read "p if r else q".

The remaining logical symbols have their usual meanings. However we write $p \equiv q$ iff $p \Leftrightarrow q$ is a theorem; similarly for $p \Rightarrow q$ iff $p \Rightarrow q$ is a theorem. One advantage of that notation is that by identifying \equiv -equivalent formulae, \Rightarrow becomes antisymmetric (see the example at the end of this section) although implication is not. Another is that we can unambiguously "chain" predicates to reason. Following Dijkstra and Scholten we include a justification for each link in the chain of such a proof. For example:

$$p \triangleleft r \triangleright q$$

\equiv

$$\left(\begin{array}{l} r \Rightarrow p \\ \neg r \Rightarrow q \end{array} \right).$$

propositional calculus

If \mathbf{X} is a set, $\text{seq.}\mathbf{X}$ denotes the set of all finite sequences of elements of \mathbf{X} . Sequence comprehension is expressed using angle brackets so that, for example, the empty sequence is written $\langle \rangle$.

Relations and functions

A *relation* from \mathbf{X} to \mathbf{Y} is a predicate on the Cartesian product of \mathbf{X} with \mathbf{Y} . Thus the set of all relations from \mathbf{X} to \mathbf{Y} is defined $\mathbf{X} \leftrightarrow \mathbf{Y} \hat{=} \mathbf{P}(\mathbf{X} \times \mathbf{Y})$. Relations are usually written in *infix*, so that for $r : \mathbf{X} \leftrightarrow \mathbf{Y}$, $x : \mathbf{X}$ and $y : \mathbf{Y}$, r relates x to y is written $x r y$ (traditionally expressed $(x, y) \in r$ with r viewed as a *subset* of $\mathbf{X} \times \mathbf{Y}$).

If $r : \mathbf{X} \leftrightarrow \mathbf{Y}$ and $s : \mathbf{Y} \leftrightarrow \mathbf{Z}$, their *forward relational composition* $r ; s : \mathbf{X} \leftrightarrow \mathbf{Z}$ is defined

$$x(r ; s)z \hat{=} \exists y : \mathbf{Y} \bullet \left(\begin{array}{l} x r y \\ y s z \end{array} \right).$$

The *identity*, $\iota_{\mathbf{X}}$, and *universal relation*, $\omega_{\mathbf{X}}$, on \mathbf{X} have their standard meanings: $\iota_{\mathbf{X}}$ relates every element of \mathbf{X} to itself whilst $\omega_{\mathbf{X}}$ relates all elements of \mathbf{X}

$$\forall x : \mathbf{X} \bullet x \iota_{\mathbf{X}} x$$

$$\forall x, y : \mathbf{X} \bullet x \omega_{\mathbf{X}} y.$$

The former gains its name by being the identity for relational composition. Further laws are investigated when we discuss relational semantics.

If $r : \mathbf{X} \leftrightarrow \mathbf{Y}$ then $r^{\sim} : \mathbf{Y} \leftrightarrow \mathbf{X}$ denotes the *converse* of r , defined $y r^{\sim} x \hat{=} x r y$.

A function can be viewed as a total relation f which at each point is single-valued:

$$\left(\begin{array}{l} x f y \\ x f y' \end{array} \right) \Rightarrow y = y'.$$

If $f : \mathbf{X} \rightarrow \mathbf{Y}$ then we write $f.x = y$ instead of $x f y$. Thus we avoid parentheses by writing $f.x$ for the application of f to argument $x : \mathbf{X}$. Functional application associates to the left. We shall also use notation f for the lifting of $f : \mathbf{X} \rightarrow \mathbf{Y}$ to its action on subsets by defining, for $\mathbf{E} \subseteq \mathbf{X}$, $f.\mathbf{E} \hat{=} \{f.e \mid e \in \mathbf{E}\}$. Ambiguity about the type of f will be avoided by context.

The functions *dom* and *ran* denote the domain and range, respectively, of a relation. Thus for $r : \mathbf{X} \leftrightarrow \mathbf{Y}$,

$$\text{dom.}r \hat{=} \{x : \mathbf{X} \mid \exists y : \mathbf{Y} \bullet x r y\}$$

$$\text{ran.}r \hat{=} \{y : \mathbf{Y} \mid \exists x : \mathbf{X} \bullet x r y\}.$$

The set of *partial functions* from \mathbf{X} to \mathbf{Y} is denoted $\mathbf{X} \rightarrow \mathbf{Y}$. For (partial) functions it is often convenient to exploit single-valuedness to rewrite relational composition in the standard way: if $f : \mathbf{X} \rightarrow \mathbf{Y}$ and $g : \mathbf{Y} \rightarrow \mathbf{Z}$, $g \circ f$ denotes

the functional composition f followed by g , defined

$$(g \circ f).x \hat{=} g.(f.x) \quad (= (f ; g).x).$$

The converse f^{\sim} of a partial function f is usually called its *inverse*, although the result is again a partial function iff f is injective.

Order

(X, \leq) is *partially ordered* means that X is a nonempty set on which \leq is a reflexive, antisymmetric and transitive relation:

$$\iota x \subseteq \leq,$$

$$(\leq \cap \geq) \subseteq \iota x,$$

where \geq denotes the relational converse of \leq

$$(\leq ; \leq) \subseteq \leq.$$

If (X, \leq) is partially ordered then the relation $<$ denotes the *strict order* $< \hat{=} (\leq \cap \neq)$. If $x \leq y$ in a partial order we say that y *dominates* x ; and if $x < y$ then we say that y *dominates* x *strictly*.

Basic examples of partial orders are provided by the standard number lines (\mathbf{N}, \leq) and (\mathbf{R}, \leq) , and their subsets with the induced orderings. The most important example of the latter type is the two-element set \mathbf{B} consisting of 0 (or *false*) and 1 (or *true*), with $0 < 1$; it is called the *Boolean order*.

Derived examples are provided by the product of a partial order with itself under the product order, and its extension: the set of all functions from a set Y into a set X where (X, \leq) is partially ordered, under the pointwise order

$$f \leq g \hat{=} \forall y : Y \bullet f.y \leq g.y.$$

An important example of such a partial order is the set PX of all predicates on X under implication

$$p \leq q \hat{=} p \Rightarrow q.$$

That is thought of as expressing p is *stronger than* q or, equivalently, q is *weaker than* p .

The space (PX, \leq) is isomorphic to the space of all subsets of X with the inclusion ordering, under the isotope bijection which assigns to predicate p the set of all elements satisfying p .

Extrema

Assume that (X, \leq) is partially ordered.

If $x : X$ and $E \subseteq X$ then x is a *lower bound* of E means that it is dominated by each member of E

$$\forall e : E \bullet x \leq e.$$

A *least element* of E is a lower bound which belongs to E ; if it exists, it is unique (why?).

A *greatest lower bound*, or glb, or *infimum*, of E is a lower bound of E which dominates any lower bound of E . Equivalently, but preferable for calculation, x is a greatest lower bound of E iff

$$y \leq x \equiv \forall e : E \bullet y \leq e.$$

If it exists a glb of E is unique; it is written $\sqcap E$. Thus

$$y \leq \sqcap E \equiv \forall e : E \bullet y \leq e.$$

Similarly are defined *upper bound*, *greatest element*, *least upper bound* (or lub, or supremum):

$$\sqcup E \leq y \equiv \forall e : E \bullet e \leq y.$$

Observe that in order to reduce reasoning about \sqcap to reasoning about \sqcup ,

$$\sqcap E = \sqcup \{x : X \mid \forall e : E \bullet x \leq e\}$$

(and analogously for \sqcup in terms of \sqcap).

Completeness

Partially ordered (X, \leq) is a *lattice* means that each nonempty finite subset of X has a glb and a lub. Equivalently, the binary versions of \sqcap and \sqcup are idempotent, commutative, associative and absorptive.

A lattice is *complete* means that every subset has a glb and lub. In particular a complete lattice has a least element, $\sqcap X$, and a greatest element, $\sqcup X$. The former is denoted \perp and the latter \top .

A *directed* subset of partially ordered (X, \leq) means a subset E of X of which every finite subset has an upper bound in E . (X, \leq) is a *complete partial order*, or cpo, means every directed subset of X has a lub.

An element $x : X$ of a cpo is *compact* means that for any directed set $E \subseteq X$, if the lub of E dominates x then some member of E dominates x

$$x \leq \sqcup E \Rightarrow \exists e : E \bullet x \leq e.$$

The set of all compact elements of (X, \leq) is denoted $K(X)$.

A cpo is a *domain*, or *algebraic*, means each element is the lub of the compact elements it dominates

$$\forall x : X \bullet x = \sqcup \{y : K(X) \mid y \leq x\}.$$

Morphisms

Suppose that (X, \leq) and (Y, \leq) are partially ordered; their orders, glb's, lub's, etc. will be distinguished by context (rather than by say subscripts).

A function $f : X \rightarrow Y$ is *monotone* means that if x' dominates x in X then $f.x'$ dominates $f.x$ in Y

$$x \leq x' \Rightarrow f.x \leq f.x'.$$

f is *isotone* means $x \leq x' \equiv f.x \leq f.x'$.

Now assume that (X, \leq) is a complete lattice and con-

sider

$$f. \sqcup E = \sqcup f.E. \quad (1)$$

Then f is

- *universally* \sqcup -*junctive* means (1) holds for all $E \subseteq X$;
- *positively* \sqcup -*junctive* means (1) holds for all nonempty $E \subseteq X$; and
- \sqcup -*junctive* means (1) holds for all nonempty finite $E \subseteq X$.

Similarly for \sqcap -*junctivity*.

If (X, \leq) is a cpo then $f : X \rightarrow Y$ is *continuous* means f is monotone and (1) holds for each directed $E \subseteq X$.

Fixed points

If $f : X \rightarrow X$ then $x : X$ is a *fixed point of f* means

$$f.x = x.$$

If (X, \leq) is a cpo with least element \perp and $f : X \rightarrow X$ is monotone then f has a least fixed point $\mu f : X$

$$\left(\begin{array}{l} f. \mu f = \mu f \\ \forall x : X \bullet x = f.x \Rightarrow \mu f \leq x \end{array} \right).$$

(Proof is by transfinite induction.) If moreover f is continuous then $\mu f = \sqcup \{f^n. \perp \mid n : \mathbb{N}\}$; that is, the least fixed point is attained over the natural numbers.

Complements and conjugation

A lattice (X, \leq) is *complemented* means firstly that it has greatest and least elements, and secondly that each element of the lattice has a complement

$$\forall x : X \bullet \exists y : X \bullet \left(\begin{array}{l} x \sqcup y = \top \\ x \sqcap y = \perp \end{array} \right).$$

Complement y is unique, for each x , if the lattice is *distributive*

$$\left(\begin{array}{l} x \sqcup (y \sqcap z) = (x \sqcup y) \sqcap (x \sqcup z) \\ x \sqcap (y \sqcup z) = (x \sqcap y) \sqcup (x \sqcap z) \end{array} \right),$$

in which case the complement of x is written $\neg x$. A complemented distributive lattice is called a *Boolean algebra*.

In a Boolean algebra each function $f : X \rightarrow X$ has *conjugate* $f^* : X \rightarrow X$ defined $f^*.x \hat{=} \neg f. \neg x$. Evidently the least element \perp of X is a fixed point of f iff the greatest element \top is a fixed point of f^* .

Example

Predicates, since they express constraints between observables, are fundamental for describing (discrete) systems. The obvious order for our purposes is that which corresponds to the "implements" relation on descriptions. Thus if p and q are predicates whose free variables represent observations of a system then $p \leq q$ is the proof obligation for showing that description p meets specification q (or that p is a refinement of q). Since it holds iff every

observation described by p is an observation specified by q , we have

$$p \leq q \hat{=} p \Rightarrow q.$$

Of course that ordering coincides with the pointwise ordering on predicates as elements of $X \rightarrow B$.

We consider the structure of the space of predicates (PX, \leq) . It is partially ordered, as already observed, since if (Y, \leq) is partially ordered and X is a set then $X \rightarrow Y$ is partially ordered with the pointwise-lifting order. Also it is a complete Boolean algebra with: least element the predicate which is always false; greatest element the predicate which is always true; negation as complement; disjunction as least upper bound; and conjunction as greatest lower bound.

It is a cpo since a set of predicates $Q \subseteq PX$ has as lub the predicate p for which

$$p.x \hat{=} \vee \{q.x \mid q \in Q\};$$

thus Q need not be directed.

A predicate is compact iff it is the characteristic function of a *finite* subset of X . Indeed for $E \subseteq X$ and set Q of predicates,

$$\xi_E \leq \sqcup Q$$

$$\hat{=}$$

definitions of order and sup

$$\forall x : X \bullet \xi_E.x \leq \vee \{q.x \mid q \in Q\}$$

$$\hat{=}$$

definitions of characteristic function and \vee

$$\forall x : X \bullet x \in E \Rightarrow \exists q : Q \bullet q.x.$$

Now if E is finite that implies, if Q is directed,

$$\exists q : Q \bullet \forall x : X \bullet x \in E \Rightarrow q.x$$

$$\hat{=}$$

definitions of characteristic function and order

$$\exists q : Q \bullet \xi_E \leq q,$$

as required for compactness of ξ_E . But conversely if E is infinite, letting \mathcal{F} denote a family of increasing finite subsets whose union equals X (why does \mathcal{F} exist?), the previous reasoning applies to the increasing (and hence directed) set $Q \hat{=} \{\xi_{E \cap F} \mid F \in \mathcal{F}\}$ to show

$$\left(\begin{array}{l} \xi_E \leq \sqcup Q \\ \neg \exists q : Q \bullet \xi_E \leq q \end{array} \right);$$

thus ξ_E is not compact.

Finally (PX, \leq) is a domain. For if $F : \mathcal{F}$ and $p : PX$, letting $p_F \hat{=} p \wedge \xi_F$, then for each F ,

$$\left(\begin{array}{l} p_F \leq p \\ p_F \text{ compact} \\ p = \sqcup \{p_F \mid F \in \mathcal{F}\} \end{array} \right).$$

Thus p is the lub of the compact elements it dominates. Since p was arbitrary, (PX, \leq) is a domain.

We return to this example in section 3. In the meantime we have established:

Theorem(predicates). If X is a set then the

space $(\mathbf{P}\mathbf{X}, \leq)$ of predicates on \mathbf{X} with the point-wise order forms a complete Boolean algebra which is a domain whose compact elements are the characteristic functions of finite sets.

3 Galois connections

Introduction

We wish to compare different semantic models of the same formal method. In this chapter we study our main technique: Galois connections.

We consider the basic properties and examples of such connections, concentrating on their dual advantages:

- the theoretical properties Galois connections establish;
- the support Galois connections provide for calculation.

But we begin by motivating their definition.

Motivation

Two partially ordered spaces (\mathbf{X}, \leq) and (\mathbf{Y}, \leq) are *isomorphic* means there is an isotone bijection $\phi : \mathbf{X} \rightarrow \mathbf{Y}$ between them; that is, a bijection for which

$$\mathbf{x} \leq \mathbf{x}' \equiv \phi.\mathbf{x} \leq \phi.\mathbf{x}'.$$

Equivalently, since ϕ has a well-defined inverse ϕ^{\sim} , choosing the unique $\mathbf{y} : \mathbf{Y}$ for which $\mathbf{x}' = \phi^{\sim}.\mathbf{y}$,

$$\begin{aligned} \mathbf{x} &\leq \phi^{\sim}.\mathbf{y} \\ \equiv & && \text{by choice of } \mathbf{y} \\ \mathbf{x} &\leq \mathbf{x}' \\ \equiv & && \text{equivalence above} \\ \phi.\mathbf{x} &\leq \phi.\mathbf{x}' \\ \equiv & && \text{in terms of } \mathbf{y} \text{ again} \\ \phi.\mathbf{x} &\leq \phi.\phi^{\sim}.\mathbf{y} \\ \equiv & && \phi \circ \phi^{\sim} = \iota \\ \phi.\mathbf{x} &\leq \mathbf{y}. \end{aligned}$$

Thus $\mathbf{x} \leq \phi^{\sim}.\mathbf{y} \equiv \phi.\mathbf{x} \leq \mathbf{y}$.

We are interested in pairs of semantic spaces of which one is "finer" or "contains more information" than the other. In that case ϕ is merely a monotone injection and ϕ^{\sim} is replaced by a monotone projection π , with the same equivalence

$$\mathbf{x} \leq \pi.\mathbf{y} \equiv \phi.\mathbf{x} \leq \mathbf{y}.$$

Here are two examples, one from arithmetic and the other from program semantics, which support a definition of that form. Captured in these simple examples is much of the theory to follow.

Floor

The partially-ordered space (\mathbf{R}, \leq) of real numbers with its standard order is finer than the partially-ordered space (\mathbf{Z}, \leq) of integers. The latter is embedded in the former by

the natural monotone injection

$$\begin{aligned} \varepsilon : \mathbf{Z} &\rightarrow \mathbf{R} \\ \varepsilon.\mathbf{n} &= \mathbf{n}, \end{aligned}$$

in which \mathbf{n} on the right-hand side is regarded as a real number. But there are many monotone injections of \mathbf{Z} into \mathbf{R} ; the natural one is characterised by having as a kind of inverse the *floor*, or integer-part-of, function $\lfloor \cdot \rfloor : \mathbf{R} \rightarrow \mathbf{Z}$, traditionally defined

$$\lfloor \mathbf{r} \rfloor \hat{=} \text{the (unique) integer } \mathbf{n} \text{ for which } \mathbf{n} \leq \mathbf{r} < \mathbf{n} + 1.$$

Equivalently,

$$\lfloor \mathbf{r} \rfloor = \sqcup \{ \mathbf{n} : \mathbf{Z} \mid \varepsilon.\mathbf{n} \leq \mathbf{r} \}. \tag{2}$$

Those definitions have a third equivalent form (why is it equivalent?) which is superior for calculations:

$$\mathbf{n} \leq \lfloor \mathbf{r} \rfloor \equiv \varepsilon.\mathbf{n} \leq \mathbf{r}. \tag{3}$$

All of those equations express that $\lfloor \mathbf{r} \rfloor$ is the greatest integer which is at most \mathbf{r} . Although it is clear that the definition of π characterises ε amongst monotone embeddings of \mathbf{Z} in \mathbf{R} , it is surprising (at first sight) that it is in fact determined by ε and that equation (2) follows from equation (3); see the treatment at the end of this section.

Now $\lfloor \cdot \rfloor$ is monotone and inverse to ε in this sense

$$\forall \mathbf{n} : \mathbf{Z} \bullet \lfloor \varepsilon.\mathbf{n} \rfloor = \mathbf{n}.$$

Of course the dual identity fails, since $\varepsilon.\lfloor \mathbf{r} \rfloor \neq \mathbf{r}$ unless $\mathbf{r} \in \text{ran } \varepsilon$. Fortunately equivalence (3) demonstrates the sense in which ε and $\lfloor \cdot \rfloor$ are inverse; indeed in spite of the previous inequality it follows from (3) that $\varepsilon.\lfloor \mathbf{r} \rfloor \leq \mathbf{r}$.

In summary,

$$\left(\begin{array}{l} \varepsilon : \mathbf{Z} \rightarrow \mathbf{R} \text{ monotone} \\ \lfloor \cdot \rfloor : \mathbf{R} \rightarrow \mathbf{Z} \text{ monotone} \\ \lfloor \cdot \rfloor \circ \varepsilon = \iota_{\mathbf{Z}} \\ \varepsilon \circ \lfloor \cdot \rfloor \leq \iota_{\mathbf{R}} \end{array} \right). \tag{4}$$

We shall see (section 3) that most of those conditions follow simply from equivalence (3); and the one that does not is of independent interest.

We now demonstrate the calculational facility of equivalence (3) (following Aarts) by proving

$$\forall \mathbf{r} : \mathbf{R} \bullet \mathbf{r} \geq 0 \Rightarrow \lfloor \sqrt{\lfloor \mathbf{r} \rfloor} \rfloor = \lfloor \sqrt{\mathbf{r}} \rfloor.$$

The layout of the following proof is to be compared with that, essentially equivalent but expressed without a Galois connection, in Graham, Knuth and Patashnik, *Concrete Mathematics: A Foundation for Computer Science*, Addison Wesley, 1989, result 3.9.

Proof. Recall that for $\mathbf{r} : \mathbf{R}$, if $\mathbf{r} \geq 0$ then

$$\sqrt{\mathbf{r}} = \mathbf{x} \equiv \mathbf{r} = \mathbf{x}^2.$$

Of course x may be negative, so there may be two solutions x . For such r we reason as follows, being slightly more careful about extracting the root of only non-negative *real* numbers:

$$\begin{aligned}
n &\leq \lfloor \sqrt{\varepsilon} \cdot \lfloor r \rfloor \rfloor \\
&\equiv \text{gc}(\varepsilon, \lfloor \cdot \rfloor; \mathbf{Z}, \mathbf{R}) \\
\varepsilon \cdot n &\leq \sqrt{\varepsilon} \cdot \lfloor r \rfloor \\
&\equiv r \leq \sqrt{s} \equiv (r \geq 0 \Rightarrow r^2 \leq s) \\
\varepsilon \cdot n \geq 0 &\Rightarrow (\varepsilon \cdot n)^2 \leq \varepsilon \cdot \lfloor r \rfloor \\
&\equiv (\varepsilon \cdot n)^2 \in \mathbf{Z} \\
\varepsilon \cdot n \geq 0 &\Rightarrow \varepsilon \cdot (\varepsilon \cdot n)^2 \leq \varepsilon \cdot \lfloor r \rfloor \\
&\equiv \varepsilon \cdot m \leq \varepsilon \cdot n \equiv m \leq n \\
\varepsilon \cdot n \geq 0 &\Rightarrow (\varepsilon \cdot n)^2 \leq \lfloor r \rfloor \\
&\equiv \text{gc}(\varepsilon, \lfloor \cdot \rfloor; \mathbf{Z}, \mathbf{R}) \text{ again} \\
\varepsilon \cdot n \geq 0 &\Rightarrow \varepsilon \cdot (\varepsilon \cdot n)^2 \leq r \\
&\equiv \varepsilon \cdot (\varepsilon \cdot n)^2 = (\varepsilon \cdot n)^2 \\
\varepsilon \cdot n \geq 0 &\Rightarrow (\varepsilon \cdot n)^2 \leq r \\
&\equiv r \leq \sqrt{s} \equiv (r \geq 0 \Rightarrow r^2 \leq s) \text{ again} \\
\varepsilon \cdot n &\leq \sqrt{r} \\
&\equiv \text{gc}(\varepsilon, \lfloor \cdot \rfloor; \mathbf{Z}, \mathbf{R}) \text{ one last time} \\
n &\leq \lfloor \sqrt{r} \rfloor.
\end{aligned}$$

Though protracted, that proof requires little motivation and so is readily followed and equally readily constructed. Its use of Galois-connection *clichés* (like switching between the sides of the Galois equivalence) distinguishes it from the proof of Graham *et al.* in which mathematical creativity is seen as playing an important role. It is for such reasons that we exploit the calculational facility of Galois connections.

Pre- and post-conditions

If \mathbf{X} is a set representing the state space (including input) on which a program, say r , operates and \mathbf{Y} is a set representing the state space in which the result (including output) lies, then r may be viewed as a relation from \mathbf{X} to \mathbf{Y} . As a predicate transformer there are (at least) two interpretations of r

- $\text{wp}.r : \mathbf{PY} \rightarrow \mathbf{PX}$ gives the *weakest precondition* of its argument. Thus for *postcondition*, (or predicate on final states) $q : \mathbf{PY}$, it holds at just those states of \mathbf{X} from which execution of r terminates in a state satisfying q

$$\text{wp}.r.q.x \hat{=} \forall y : \mathbf{Y} \bullet x r y \Rightarrow q.y$$

- $\text{sp}.r : \mathbf{PX} \rightarrow \mathbf{PY}$ gives the *strongest postcondition* of its argument. Thus for *precondition*, (or predicate on initial states) $p : \mathbf{PX}$, it holds at just those states of \mathbf{Y} which can be reached if execution of r starts in a

state satisfying p

$$\text{sp}.r.p.y \hat{=} \exists x : \mathbf{X} \bullet \begin{pmatrix} x r y \\ p.x \end{pmatrix}.$$

Now equivalence of those two representations of r is expressed: $p \Rightarrow \text{wp}.r.q \equiv \text{sp}.r.p \Rightarrow q$.

In summary, for $r : \mathbf{X} \leftrightarrow \mathbf{Y}$,

$$\begin{pmatrix} \text{sp}.r : \mathbf{PX} \rightarrow \mathbf{PY} \\ \text{wp}.r : \mathbf{PY} \rightarrow \mathbf{PX} \\ p \leq \text{wp}.r.q \equiv \text{sp}.r.p \leq q \end{pmatrix}.$$

As we shall see in section 7, important properties of $\text{wp}.r$ and $\text{sp}.r$ follow directly from that equivalence.

These considerations lead us to the following.

Definition

Suppose that (\mathbf{X}, \leq) and (\mathbf{Y}, \leq) are partially ordered with $\varepsilon : \mathbf{X} \rightarrow \mathbf{Y}$ and $\pi : \mathbf{Y} \rightarrow \mathbf{X}$. Then the pair ε, π forms a *Galois connection*, written $\text{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$, means

$$\forall x : \mathbf{X} \bullet \forall y : \mathbf{Y} \bullet x \leq \pi.y \Leftrightarrow \varepsilon.x \leq y.$$

In that case ε is called the *embedding* and π the *projection*; each is described (loosely) as being adjoint to the other.

The next two results explore basic properties of the set $\text{gc}(\mathbf{X}, \mathbf{Y})$ of all Galois connections between partially-ordered spaces \mathbf{X} and \mathbf{Y} .

Galois connections act as morphisms between partially-ordered spaces:

Theorem(combining connections). The identity (and its converse) form a Galois connection on any partially-ordered space:

$$\text{gc}(\iota_{\mathbf{X}}, \iota_{\mathbf{X}}; \mathbf{X}, \mathbf{X}).$$

Galois connections compose as follows:

$$\begin{aligned}
&\begin{pmatrix} \text{gc}(\varepsilon_0, \pi_0; \mathbf{X}, \mathbf{Y}) \\ \text{gc}(\varepsilon_1, \pi_1; \mathbf{Y}, \mathbf{Z}) \end{pmatrix} \\
&\Rightarrow \\
&\text{gc}(\varepsilon_1 \circ \varepsilon_0, \pi_0 \circ \pi_1; \mathbf{X}, \mathbf{Z}).
\end{aligned}$$

The space of Galois connections inherits order properties:

Theorem(ordering connections). If \mathbf{X} and \mathbf{Y} are partially ordered so too is $\text{gc}(\mathbf{X}, \mathbf{Y})$, under

$$(\varepsilon_0, \pi_0) \leq (\varepsilon_1, \pi_1) \hat{=} \begin{pmatrix} \varepsilon_0 \leq \varepsilon_1 \\ \pi_1 \leq \pi_0 \end{pmatrix}.$$

If \mathbf{X} and \mathbf{Y} are lattices then so too is $\text{gc}(\mathbf{X}, \mathbf{Y})$ with: minimum (\perp, \top) ; maximum (\top, \perp) ; and with (ε_0, π_0) and (ε_1, π_1) having $\text{glb}(\varepsilon_0 \sqcap \varepsilon_1, \pi_0 \sqcup \pi_1)$ and $\text{lub}(\varepsilon_0 \sqcup \varepsilon_1, \pi_0 \sqcap \pi_1)$.

In the remaining sections of this chapter we investigate the more detailed structure exhibited by a single Galois connection.

Duality

The following result will enable us to expend half as much effort in proving subsequent results. Indeed it establishes the two steps in the previous proof that were justified as "similarly".

Theorem(duality). Writing \mathbf{X}^{\sim} for the partial order (\mathbf{X}, \geq) converse to (\mathbf{X}, \leq) ,

$$\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y}) \equiv \mathbf{gc}(\pi, \varepsilon; \mathbf{Y}^{\sim}, \mathbf{X}^{\sim}).$$

Equivalence

We now head towards characterising a Galois connection in terms that have already arisen in the first motivating example.

Lemma(weak inversion). If $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then, in terms of the pointwise ordering on functions (see section 2),

$$\left(\begin{array}{l} \iota_{\mathbf{X}} \leq \pi \circ \varepsilon \\ \varepsilon \circ \pi \leq \iota_{\mathbf{Y}} \end{array} \right).$$

Lemma(monotonicity). If $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then ε and π are monotone.

The result we have been aiming at is

Theorem(equivalence).

$$\begin{aligned} &\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y}) \\ &\equiv \\ &\left(\begin{array}{l} \varepsilon, \pi \text{ monotone} \\ \iota_{\mathbf{X}} \leq \pi \circ \varepsilon \\ \varepsilon \circ \pi \leq \iota_{\mathbf{Y}} \end{array} \right). \end{aligned}$$

Isomorphism

We now work towards establishing that in a Galois connection the ranges of the two functions are isomorphic.

Lemma(trading). If $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then these are equivalent

$$\begin{aligned} &\mathbf{x} \leq \pi.\mathbf{y} \\ &\varepsilon.\mathbf{x} \leq \varepsilon.\pi.\mathbf{y} \\ &\varepsilon.\mathbf{x} \leq \mathbf{y} \\ &\pi.\varepsilon.\mathbf{x} \leq \pi.\mathbf{y}. \end{aligned}$$

Corollary (of trading). If $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then these are equivalent

$$\begin{aligned} &\pi.\mathbf{y}' \leq \pi.\mathbf{y} \\ &\varepsilon.\pi.\mathbf{y}' \leq \varepsilon.\pi.\mathbf{y} \\ &\varepsilon.\pi.\mathbf{y}' \leq \mathbf{y}. \end{aligned}$$

In particular $\pi.\mathbf{y}' = \pi.\mathbf{y} \equiv \varepsilon.\pi.\mathbf{y}' = \varepsilon.\pi.\mathbf{y}$. Dually, these are also equivalent

$$\begin{aligned} &\varepsilon.\mathbf{x} \leq \varepsilon.\mathbf{x}' \\ &\pi.\varepsilon.\mathbf{x} \leq \pi.\varepsilon.\mathbf{x}' \\ &\mathbf{x} \leq \pi.\varepsilon.\mathbf{x}'. \end{aligned}$$

In particular $\varepsilon.\mathbf{x} = \varepsilon.\mathbf{x}' \equiv \pi.\varepsilon.\mathbf{x} = \pi.\varepsilon.\mathbf{x}'$.

Lemma(sandwiching). If $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then

$$\left(\begin{array}{l} \varepsilon = \varepsilon \circ \pi \circ \varepsilon \\ \pi = \pi \circ \varepsilon \circ \pi \end{array} \right).$$

The aim of the last few results has been the next theorem. It is to be compared with the case of an isotone bijection ϕ between partial orders \mathbf{X} and \mathbf{Y} : the range of ϕ is isomorphic to its domain. Hence the range of ϕ is isomorphic to the range of its converse ϕ^{\sim} , and ϕ itself provides an isotone bijection. That rather trivial observation is extended, to Galois connections, in the next result.

Theorem(isomorphism). If $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then $\mathbf{ran}.\varepsilon$ and $\mathbf{ran}.\pi$ are isomorphic and the restrictions of ε and π are isotone bijections.

The isomorphism theorem has an interesting and simple restatement in terms of fixed points.

Definition

If $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then $\mathbf{x} : \mathbf{X}$ is *closed* means it is a fixed point of $\pi \circ \varepsilon$. \mathbf{X}_0 denotes the set of closed elements of \mathbf{X} , defined $\mathbf{X}_0 \hat{=} \{\mathbf{x} : \mathbf{X} \mid \mathbf{x} = \pi.\varepsilon.\mathbf{x}\}$. Analogously $\mathbf{y} : \mathbf{Y}$ is *open* means it is a fixed point of $\varepsilon \circ \pi$, and $\mathbf{Y}_0 \hat{=} \{\mathbf{y} : \mathbf{Y} \mid \varepsilon.\pi.\mathbf{y} = \mathbf{y}\}$.

Theorem(fixed-point isomorphism).

If $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then

$$\left(\begin{array}{l} \mathbf{X}_0 = \mathbf{ran}.\pi \\ \mathbf{Y}_0 = \mathbf{ran}.\varepsilon \end{array} \right),$$

so that \mathbf{X}_0 and \mathbf{Y}_0 are isomorphic.

Adjunction

Next we show that each function in a Galois connection determines the other, and that it does so via an explicit formula (provided the underlying spaces are complete lattices) which we call adjunction. That characterises, for such underlying lattices, when a single function belongs to a Galois connection, and the form of its adjoint.

Theorem(determinism). If $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then each function determines the other.

The next result motivates adjunction as well as establishing further important properties of a Galois connection.

Theorem(junctivity). If (\mathbf{X}, \leq) and (\mathbf{Y}, \leq) are complete lattices and $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ then ε is universally \sqcup -junctive and π is universally \sqcap -junctive. In particular $\varepsilon.\perp = \perp$ and $\pi.\top = \top$.

There is an important converse to the previous result: if the underlying partial order \mathbf{X} is a complete lattice and ε is universally \sqcup -junctive then a simple formula defines a projection which together with ε forms a Galois connection.

Theorem(adjunction). If (\mathbf{X}, \leq) is a complete lattice and ε is universally \sqcup -junctive then the function $\pi : \mathbf{Y} \rightarrow \mathbf{X}$ given by

$$\pi.\mathbf{y} = \sqcup\{\mathbf{x} : \mathbf{X} \mid \varepsilon.\mathbf{x} \leq \mathbf{y}\}$$

is well defined and $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$.

Dually, for universally \sqcap -junctive π , the function $\varepsilon : \mathbf{X} \rightarrow \mathbf{Y}$ defined

$$\varepsilon.x = \sqcap\{y : \mathbf{Y} \mid x \leq \pi.y\}.$$

Galois embeddings

So far we have accounted, in a Galois connection, for all of the properties encountered in the motivating examples, except for one property of floor: $\iota_{\mathbf{X}} = \pi \circ \varepsilon$. Galois connections which satisfy that stronger property occur when one semantic domain is stronger than the other, which accounts for the name we give them.

Definition

A Galois embedding, $\mathbf{ge}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$, is a Galois connection $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ in which $\pi \circ \varepsilon = \iota_{\mathbf{X}}$.

The following result is readily proved using the same techniques (mainly sandwiching) as in previous results. It means that in a Galois embedding \mathbf{X} can be considered a subspace of \mathbf{Y} .

Theorem (equivalence for embeddings). These are equivalent

- $\mathbf{ge}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$ (i.e. $\pi \circ \varepsilon = \iota_{\mathbf{X}}$)
- ε is injective
- π is surjective.

Defining a Galois connection

Suppose we wish to define a Galois connection with \mathbf{X} and \mathbf{Y} complete lattices, $\mathbf{gc}(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$, and we already have a definition of an embedding function $\varepsilon : \mathbf{X} \rightarrow \mathbf{Y}$. To complete the Galois connection there are three possibilities.

- (a) Define $\pi : \mathbf{Y} \rightarrow \mathbf{X}$ and prove the Galois equivalence (or an equivalent form). This approach is recommended if an obvious simple π suggests itself. However it might be difficult to think of the appropriate π ; use of adjunction to do that provides the second approach.
- (b) Show ε is universally \sqcup -junctive and define π by adjunction. The advantage of this approach is that the Galois equivalence need not be checked and that no inspiration is required in defining π . Its disadvantage is primarily that universal \sqcup -junctivity of ε needs to be checked (otherwise the defining formula may not be well defined); also, to a lesser extent, that some manipulation may be necessary to simplify the form of π . The third approach obviates the former disadvantage.
- (c) Substitute the definition of ε into the Galois equivalence and, by manipulation, obtain a definition of π . That approach shares with the previous one the advantage that the resulting π automatically completes a Galois connection. Its disadvantage is mainly that the manipulations leading to a definition may be difficult to achieve; and to a lesser extent that care must be exercised to avoid undefined cases (which might not be evident from the manipulation, comparable to

use of approach (b) without confirming universal \sqcup -junctivity of ε).

In these notes we shall free to use whichever approach seems easiest.

Examples

Now that we have completed the basic theory of Galois connections, we consider three examples: first that of predicates, considered in section 2, then the motivating examples from this section. We use them to compare the three approaches above and to exemplify the isomorphism theorem.

Predicates

The most useful predicate combinator for constructing designs is conjunction: $\mathbf{p} \wedge \mathbf{r}$ is a design satisfying the constraints \mathbf{p} and \mathbf{r} . If \mathbf{q} is a specification then $\mathbf{p} \wedge \mathbf{r} \leq \mathbf{q}$ is the proof obligation for establishing that the design meets specification \mathbf{q} . Let us consider the adjoint of that obligation.

For any set \mathbf{X} and predicate $\mathbf{r} : \mathbf{PX}$, define

$$\begin{aligned} \varepsilon : \mathbf{PX} &\rightarrow \mathbf{PX} \\ \varepsilon.p &\hat{=} \mathbf{p} \wedge \mathbf{r}. \end{aligned}$$

We wish to make embedding ε into a Galois connection on \mathbf{PX} . Following (b) above, we observe that ε is universally \sqcup -junctive (why?) so its adjoint equals the function satisfying, for each $\mathbf{q} : \mathbf{PX}$,

$$\begin{aligned} \pi.q & & & \text{adjunction} \\ \equiv & & & \\ \sqcup\{\mathbf{p} : \mathbf{PX} \mid \varepsilon.p \leq \mathbf{q}\} & & & \\ \equiv & & & \text{definition of } \varepsilon \text{ and of order} \\ \sqcup\{\mathbf{p} : \mathbf{PX} \mid \mathbf{p} \wedge \mathbf{r} \Rightarrow \mathbf{q}\} & & & \\ \equiv & & & \text{propositional calculus} \\ \sqcup\{\mathbf{p} : \mathbf{PX} \mid \mathbf{p} \Rightarrow (\mathbf{r} \Rightarrow \mathbf{q})\} & & & \\ \equiv & & & \text{lub attained} \\ (\mathbf{r} \Rightarrow \mathbf{q}). & & & \end{aligned}$$

Thus $\pi.q \hat{=} \mathbf{r} \Rightarrow \mathbf{q}$ completes a Galois connection with ε .

Were (c) above to be used instead, we would reason

$$\begin{aligned} \varepsilon.p \leq \mathbf{q} & & & \text{definition of } \varepsilon \text{ and of order} \\ \equiv & & & \\ (\mathbf{p} \wedge \mathbf{r}) \Rightarrow \mathbf{q} & & & \\ \equiv & & & \\ \mathbf{p} \Rightarrow (\mathbf{r} \Rightarrow \mathbf{q}) & & & \\ \equiv & & & \\ \mathbf{p} \leq (\mathbf{r} \Rightarrow \mathbf{q}); & & & \end{aligned}$$

which leads us to define $\pi.q \hat{=} \mathbf{r} \Rightarrow \mathbf{q}$.

Now to identify isomorphism, we must identify the ranges of the two functions. First

$$\mathbf{ran}.\varepsilon$$

$$\begin{aligned}
 &= && \text{definition of } \mathbf{ran} \\
 &\{\varepsilon.p \mid p \in \mathbf{PX}\} \\
 &= && \text{definition of } \varepsilon \\
 &\{p \wedge r \mid p \in \mathbf{PX}\} \\
 &= && \text{logic} \\
 &\{q : \mathbf{PX} \mid q \Rightarrow r\}.
 \end{aligned}$$

Thus the range of the embedding ε consists of all predicates at least as strong as r . Next

$$\begin{aligned}
 &\mathbf{ran}.\pi \\
 &= && \text{definition of } \mathbf{ran} \\
 &\{\pi.q \mid q \in \mathbf{PX}\} \\
 &= && \text{definition of } \varepsilon \\
 &\{r \Rightarrow q \mid q \in \mathbf{PX}\} \\
 &= && \text{logic} \\
 &\{p : \mathbf{PX} \mid \neg r \Rightarrow p\}.
 \end{aligned}$$

Thus the range of projection π consists of all predicates at least as weak as $\neg r$.

Isomorphism of $\mathbf{ran}.\varepsilon$ with $\mathbf{ran}.\pi$ is thus the isomorphism of the predicates at least as strong as r with those at least as weak as $\neg r$. It is achieved by negation, from the contrapositive law.

The floor

Equivalence (3) establishes that the embedding from the integers to the reals (both with their standard order), together with the floor function as projection in the reverse direction, forms a Galois connection.

The equivalence theorem shows that properties (4) follow straight from the definition of a Galois connection, except for $\forall n : \mathbf{Z} \bullet [\varepsilon.n] = n$, which states that the functions form a Galois embedding.

In the orders on \mathbf{Z} and \mathbf{R} , $m \sqcup n$ is the maximum of m and n . Thus the embedding is disjunctive and so adjunction constructs the floor function (in the guise of equation (2)) straight from definition (3), a possibility which was probably not clear originally, but which corresponds to either (b) or (c) above.

The range of the embedding is the integral real numbers, and the range of the floor function is the integers. Thus the isomorphism theorem is no more than the order isomorphism of the integers with the integral reals. Of course every integer is a fixed point of the embedding followed by the floor, so $\mathbf{Z}_0 = \mathbf{Z}$. The fixed points of the floor followed by the embedding are the integral real numbers, so \mathbf{R}_0 consists of the integral reals.

Predicate transformers

Earlier in section 3, approach (a) was used to show that, for any $r : \mathbf{X} \leftrightarrow \mathbf{Y}$, $\mathbf{gc}(\mathbf{sp}.\mathbf{r}, \mathbf{wp}.\mathbf{r}; \mathbf{PX}, \mathbf{PY})$. Recall that $\mathbf{sp}.\mathbf{r}$ maps precondition p to its image through relation r , whilst $\mathbf{wp}.\mathbf{r}$ maps postcondition q to the weakest precondition whose image through r is at least as strong as q .

How does use of approach (c) compare? It generates the definition of $\mathbf{wp}.\mathbf{r}$ from that of $\mathbf{sp}.\mathbf{r}$ as follows

$$\begin{aligned}
 &\mathbf{sp}.\mathbf{r}.p \leq q \\
 &\equiv && \text{pointwise order} \\
 &\forall y : \mathbf{Y} \bullet \mathbf{sp}.\mathbf{r}.p.y \leq q.y \\
 &\equiv && \text{definitions of } \mathbf{sp} \text{ and order} \\
 &\forall y : \mathbf{Y} \bullet \exists x : \mathbf{X} \bullet \begin{pmatrix} x r y \\ p.x \end{pmatrix} \Rightarrow q.y \\
 &\equiv && \text{logic} \\
 &\forall x : \mathbf{X} \bullet p.x \Rightarrow \forall y : \mathbf{Y} \bullet (x r y \Rightarrow q.y) \\
 &\equiv && \text{defining } \mathbf{wp}.\mathbf{r}.q.x \hat{=} \forall y : \mathbf{Y} \bullet (x r y \Rightarrow q.y) \\
 &\forall x : \mathbf{X} \bullet p.x \leq \mathbf{wp}.\mathbf{r}.q.x \\
 &\equiv && \text{pointwise order} \\
 &p \leq \mathbf{wp}.\mathbf{r}.q.
 \end{aligned}$$

Approach (b) is almost identical, though universal \sqcup -junctivity of $\mathbf{sp}.\mathbf{r}$ must be checked too.

To identify the claim of the isomorphism theorem, we observe that

$$\begin{aligned}
 &q \in \mathbf{ran}.\mathbf{sp}.\mathbf{r} \\
 &\equiv && \text{definition of } \mathbf{ran} \\
 &\exists p : \mathbf{PX} \bullet q = \mathbf{sp}.\mathbf{r}.p \\
 &\equiv && \text{pointwise equality} \\
 &\exists p : \mathbf{PX} \bullet \forall y : \mathbf{Y} \bullet q.y = \mathbf{sp}.\mathbf{r}.p.y \\
 &\equiv && \text{definition of } \mathbf{sp}.\mathbf{r} \\
 &\exists p : \mathbf{PX} \bullet \forall y : \mathbf{Y} \bullet q.y = \exists x : \mathbf{X} \bullet \begin{pmatrix} x r y \\ p.x \end{pmatrix}.
 \end{aligned}$$

Thus q is the image through r of a predicate on r 's domain.

Analogously

$$\begin{aligned}
 &p \in \mathbf{ran}.\mathbf{wp}.\mathbf{r} \\
 &\equiv && \text{pointwise equality} \\
 &\exists q : \mathbf{PX} \bullet \forall x : \mathbf{X} \bullet p.x = \mathbf{wp}.\mathbf{r}.q.x \\
 &\equiv && \text{definition of } \mathbf{wp}.\mathbf{r} \\
 &\exists q : \mathbf{PX} \bullet \forall x : \mathbf{X} \bullet p.x = \forall y : \mathbf{Y} \bullet (x r y \Rightarrow q.y) \\
 &\equiv && \text{definition of characteristic function} \\
 &\exists q : \mathbf{PX} \bullet p \Rightarrow \xi_{\mathbf{dom}.\mathbf{r}}.
 \end{aligned}$$

Thus p is at least as strong as $\xi_{\mathbf{dom}.\mathbf{r}}$.

To identify the isomorphism of $\mathbf{ran}.\mathbf{sp}.\mathbf{r}$ with $\mathbf{ran}.\mathbf{wp}.\mathbf{r}$ achieved by the restriction of $\mathbf{wp}.\mathbf{r}$, we observe that for $\mathbf{wp}.\mathbf{r}$ evaluated at an arbitrary member of $\mathbf{ran}.\mathbf{sp}.\mathbf{r}$,

$$\begin{aligned}
 &\mathbf{wp}.\mathbf{r}.\mathbf{sp}.\mathbf{r}.p.x \\
 &\equiv && \text{definition of } \mathbf{wp}.\mathbf{r} \\
 &\forall y : \mathbf{Y} \bullet x r y \Rightarrow \mathbf{sp}.\mathbf{r}.p.y \\
 &\equiv && \text{definition of } \mathbf{sp}.\mathbf{r} \\
 &\forall y : \mathbf{Y} \bullet x r y \Rightarrow \exists z : \mathbf{X} \bullet \begin{pmatrix} z r y \\ p.z \end{pmatrix}.
 \end{aligned}$$

Informally that affects the “ r -closure” of p by assigning to p the strongest predicate weaker than it and invariant under “having the same r image”. In those terms p is a fixed point of $(wp.r) \circ (sp.r)$ iff p equals its own r -closure.

Similar elementary reasoning demonstrates that conjugation transforms the strongest postcondition of r to the weakest precondition of the converse of r (and dually) $(sp.r)^* = wp.(r^{\sim})$.

4 Relational semantics

Introduction

We now begin the study of program semantics by considering a relational model of the guarded-command language, and its extension to a relational model of specifications. We shall find that in spite of its intuitive appeal, the model is neither as straightforward nor as useful for calculation as might have been hoped.

The guarded-command language

A computation in the guarded-command language consists of a sequence of assignments, configured by the control structures of sequential composition, conditional choice and repetition (although in the next section we also allow invocation of procedures). The vector of variables used by a computation is called its *state*, whose type consists of the cartesian product, called the *state space*, of the type of each variable. Each *terminating* assignment transforms state, from the *state before* the assignment to a *state after* it.

Without loss of generality we shall assume a homogeneous state space. For firstly we can ignore input and output by incorporating the former in the state before and the latter in the state after, by suitable extensions of those states. And secondly if those extended kinds of state have different types we can form their (discriminated) union. The result is a homogeneous state space which we denote X .

As the computation proceeds a succession of states is traced out in state space X . This is the view of a process or computing agent as a *dynamical system*, more commonly referred to in Computation as that of a *transition system*.

We are interested in semantic support for the development of implementations from their specifications, and so in the description of computations at all levels of abstraction: from abstract, simple choice of X through successively more complicated choices, to the final concrete choice in the implementation. Typically, abstract behaviour is described in a state-based specification language, intermediate-level behaviour is expressed in a refinement calculus, and implementations are written in code.

But whatever the level of abstraction, we must decide which behaviours we deem equivalent. Here we take the view of *total correctness*, the simplest outlook for ensuring correct, termination of designs. (Alternatives include *partial correctness*, which fails to distinguish nontermina-

tion from termination and so is too weak for our purposes, and *generalised correctness*, which incorporates both total and partial correctness and so is unnecessarily extensive for us here.) However we shall settle on that view only after a brief excursion into partial correctness. Let us begin by considering the behaviour of implementations.

E. W. Dijkstra’s *guarded-command language* provides a convenient notation for describing algorithms imperatively. We shall refer to it as *code*, even though it deliberately abstracts many features of imperative programming languages and incorporates nondeterminism. However it is close enough to several popular programming languages so that algorithms designed in it can be “downcoded” routinely without introducing design errors. Its atomic statements are

- **skip**, which terminates immediately, changing no variables
- **abort**, which models the (undesirable) behaviour of nontermination
- **assignment**, $x := e$ (where e is an expression whose evaluation terminates with a single value) which terminates with x having the value of e .

The constructors of the guarded-command language are

- **sequential composition**, $P ; Q$, which first executes statement P and, if P terminates, then executes statement Q
- **conditional**, $\text{if } \square_i g_i \rightarrow P_i \text{ fi}$, which executes one of the statements P_i whose *guard* g_i , a predicate on state space, is true; if none of the guards is true it aborts; note that evaluation of a predicate is presumed to terminate at each state
- **repetition**, $\text{do } \square_i g_i \rightarrow P_i \text{ od}$, which—as long as at least one of the guards holds—executes one of the statements P_i whose guard g_i is true; otherwise it terminates.

In order to give a semantics to the guarded command language it is convenient to consider instead the language \mathcal{L} in which conditional and repetition are replaced by

- **binary conditional**, $P \triangleleft b \triangleright Q$, read P if b else Q , which executes P if predicate b holds and executes Q otherwise
- **choice**, $P \square Q$, which executes P or Q ; the choice between them is nondeterministic
- **recursion**, μF , which for monotone (indeed, in practice, continuous) function F on programs denotes the least fixed point of F .

In executable programming languages the expressions which appear in assignments are defined by syntax over variables and ‘built-in’ constants, functions and operations. It is assumed that their evaluation terminates with a single result. Accordingly, in the guarded-command language, an expression is a (total) function on state space X . Computability (in the sense of Recursion Theory) is ignored. In an executable implementation computability is not an issue. In an implementation expressed in the guarded-command language it is assured by the particular form chosen for expressions.

Such freedom from a specific syntax for expressions allows the guarded-command language to be used at various levels of abstraction. A program can be specified as an 'abstract' assignment over 'not necessarily executable' state space (for example doubling a natural number: $x := 2x$), but implemented using only executable expressions (for example x may be represented as an array of bits and the assignment achieved by a loop moving bits one place to the left and appending 0 in the least-significant place). Indeed the first program developments were performed in the guarded-command language.

The guarded-command language thus has deterministic assignments, and introduces nondeterminism through overlapping guards. In language \mathcal{L} an alternative approach is taken: an expression is permitted not to terminate and to be nondeterministic. As a result \mathcal{L} is even more useful in supporting program developments (though we shall extend it further).

Languages \mathcal{L} and the guarded command language are endowed with the pre-order \sqsubseteq of program refinement, defined by the law $P \sqsubseteq Q \Leftrightarrow P \sqcap Q = P$. Though \sqsubseteq is not antisymmetric we shall consider as equivalent computations which are "equally deterministic". Firstly that decision implies in particular that the pre-order is made into a partial order on equivalence classes of programs, just as we did for the partially-ordered space of predicates under \Rightarrow modulo \equiv . But it also means that syntactically distinct programs deemed to have the same behaviour will be considered equivalent. For example the three programs **abort**, **if false** \rightarrow **skip fi**, and **do true** \rightarrow **skip od** are deemed equivalent; evidently each aborts. Such equivalences reflect our commitment to total correctness. The implications of that will be discussed later; in summary, possible and certain nontermination are identified.

It is to language \mathcal{L} with that partial order, $(\mathcal{L}, \sqsubseteq)$, that we shall give a relational semantics in this chapter and a transformer semantics later. In each semantic model considered we shall write $[P]$ for the denotation of construct P in \mathcal{L} .

Partial functions

We start our semantic quest (ignoring, for the moment, total correctness and) by recalling the form of semantics used in Recursion Theory. It provides the traditional method for modelling deterministic computations which do not necessarily terminate. Thus it is too limited for our (nondeterministic) needs. However it provides a simple starting point by separating, as far as possible, the two issues of termination and nondeterminism. Our work will not be wasted since we shall find that its useful features can be lifted, by a Galois connection, to a more comprehensive model.

The model

Each program is assumed to be deterministic and so is represented by a partial function on homogeneous state space \mathbf{X} . In terms of the semantic denotation $[P]$ of program P , $[P] : \mathbf{X} \rightarrow \mathbf{X}$ with the domain of $[P]$ equal to the set of

states from which program P terminates. Semantic denotations thus lie in $\mathcal{P} \cong \mathbf{X} \rightarrow \mathbf{X}$.

Program P is refined by program Q iff partial function $[Q]$ is an extension of partial function $[P]$. For then Q gives the same results as P when the latter terminates, though Q may terminate from states in which P does not. We thus define the order on \mathcal{P} to be inclusion: $f \sqsubseteq g \Leftrightarrow f \subseteq g$.

In $(\mathcal{P}, \sqsubseteq)$ the least, or most unrefined, element is the empty partial function, $\{\}$, which denotes the nonterminating program **abort**. An element is maximal iff it can be extended no further: it is a function. And an element is compact iff it is a finite extension of the empty partial function: it is finite. Thus:

Theorem(\mathcal{P} a domain). $(\mathcal{P}, \sqsubseteq)$ is a domain with least element $\{\}$, whose maximal elements are the total functions and whose compact elements are the finite partial functions: $\mathbf{K}(\mathcal{P}) = \{f : \mathcal{P} \mid \#f < \infty\}$.

Semantics

The approach of Recursion Theory provides, for the deterministic subset of language \mathcal{L} , the following semantics. Each denotation is motivated by the operational explanation of its corresponding construct given model \mathcal{P} .

$$\begin{aligned} [\text{skip}] &\cong \iota_{\mathbf{X}} \\ [\text{abort}] &\cong \{\} \\ [x := e] &\cong \{(x, e.x) \mid x \in \mathbf{X}\} \\ &\quad e \text{ total deterministic} \\ [P ; Q] &\cong [P] ; [Q] \\ [P \triangleleft b \triangleright Q] &\cong \{(x, y) \mid x[P]y \triangleleft b.x \triangleright \\ &\quad x[Q]y\} \\ [P \square Q] &\cong [P] \cup [Q] \\ &\quad [P] \upharpoonright \text{dom}.Q = [Q] \upharpoonright \text{dom}.P \\ [\mu F] &\cong \bigcup_{\kappa} F^{\kappa} . \{\} \\ &\quad F \text{ monotone on } \mathcal{P} . \end{aligned}$$

The last union is over all ordinals κ . Fortunately if F is continuous then that reduces to a countable union (attained from a countable iteration of F applied to the least element, **abort**). The next result ensures that is the case whenever F is a composition of the three program combinators: sequential composition, binary conditional and nondeterministic choice.

Theorem(semantics of \mathcal{P}). All denotations are well defined and the three combinators are continuous (hence monotone).

The restriction to determinism, imposed simply because we have chosen to start from the traditional approach of Recursion Theory, will be removed in the next section since it conflicts with one of the strengths of the guarded-

command language: its treatment of nondeterminism.

Nonetheless the semantic denotations of programs in \mathcal{L} (with assignments having terminating deterministic expressions) exhaust model \mathcal{P} ; that is, $[\mathcal{L}] = \mathcal{P}$. For if $f \in \mathcal{P}$ then choosing

$$\mathbf{P} \hat{=} (\mathbf{x} := \mathbf{f}.\mathbf{x}) \triangleleft \mathbf{x} \in \text{dom}.\mathbf{f} \triangleright \mathbf{abort},$$

yields $\mathbf{P} \in \mathcal{L}$ and $[\mathbf{P}] = \mathbf{f}$.

Laws

The utility of that (or any) semantics must be confirmed by proving the soundness of the *laws of programming*, which reflect our view of program behaviour. For example the law $\mathbf{abort}; \mathbf{P} = \mathbf{abort}$ expresses the inability of any program to follow nontermination; it holds because the composition of the empty partial function with any function is empty. The law $\mathbf{skip}; \mathbf{P} = \mathbf{P}$ expresses the fact that \mathbf{skip} terminates without effect, and thus that any program can follow termination; it holds because function $\iota_{\mathbf{X}}$ is the identity for composition. The law $\mathbf{abort} \square \mathbf{P} = \mathbf{abort}$ expresses the fact that a program which might not terminate is as unreliable as one which is certain not to terminate; thus the behaviour of the two programs should be identified to avoid errors of nontermination. However it *fails* in this model because the union of the empty partial function with any function yields that function. Thus in the model \mathcal{P} , $\mathbf{abort} \square \mathbf{P} = \mathbf{P}$. A result of that fallacious law would therefore be that program development based on equivalence between programs (or, even worse, merely on refinement) would ignore termination. In terms introduced earlier, $(\mathcal{P}, \sqsubseteq)$ is appropriate to partial correctness and not total correctness. We return to this when studying predicate transformers.

There is a second problem. Were inclusion to be used for program development, a (partial) specification could have inconsistent refinements!

Thus even for deterministic programs, model \mathcal{P} is unacceptable for reasoning about programs. It also does not contain development features like miracles. There is therefore little point in attempting a simple modification to incorporate nondeterminism. We must find an alternative model, which we begin to do by distinguishing termination in the usual way. That will leave us free to model nondeterminism as many-valuedness, and to exploit the relational calculus.

Total relations

In this section we provide a model for \mathcal{L} which avoids the weaknesses of model \mathcal{P} .

Suppose that \perp denotes an element not in (homogeneous) state space \mathbf{X} and let \mathbf{X}_{\perp} denote \mathbf{X} suitably augmented $\mathbf{X}_{\perp} \hat{=} \mathbf{X} \cup \{\perp\}$.

In the new model program \mathbf{P} is represented by a relation $[\mathbf{P}]$ on \mathbf{X}_{\perp} . The condition "program \mathbf{P} may not terminate from initial state \mathbf{x} " is now expressed $\mathbf{x}[\mathbf{P}] \perp$. The phrase "may not" is required since a nondeterministic \mathbf{P} may both terminate and not terminate from a given

initial state.

Introducing the model

The idea of the new model is that the refinement ordering on programs be represented by containment of the relations that represent them. Thus, for example, the least element in the model, $[\mathbf{abort}]$, must be the universal relation ω . However that means that for a program to be able to refine program \mathbf{P} that "terminates less", \mathbf{P} must behave like ω at those states from which it may not terminate; for only then is refinement represented as containment.

Thus if the ordering on our new model is to be containment of relations, the relations must satisfy a healthiness condition: if relation \mathbf{d} represents a computation which may not terminate from state \mathbf{x} , then for each $\mathbf{y} : \mathbf{X}_{\perp}$, $\mathbf{x} \mathbf{d} \mathbf{y}$:

$$\mathbf{x} \mathbf{d} \perp \Rightarrow \forall \mathbf{y} : \mathbf{X}_{\perp} \bullet \mathbf{x} \mathbf{d} \mathbf{y}. \quad (5)$$

In other words, for each $\mathbf{x} : \mathbf{X}_{\perp}$ the set $\mathbf{d}[\mathbf{x}] \hat{=} \{\mathbf{y} : \mathbf{X}_{\perp} \mid \mathbf{x} \mathbf{d} \mathbf{y}\}$ is up-closed in the flat order on \mathbf{X}_{\perp} (in which the elements of \mathbf{X} are incomparable but each dominates \perp).

Now recursion is represented as the intersection of a *decreasing* family of relations, so some condition is required to ensure that it is nonempty. At this point it is convenient to consider the degree to which the guarded-command language, or indeed language \mathcal{L} , exhibits nondeterminism. In both it is explicitly assumed, for semantic purposes, that only bounded nondeterminism occurs. Thus the expression in an assignment assigns to each state at most a finite number of alternative values (if it terminates). In other words, the relational denotation of each program is finitary. The semantic reason behind that choice is precisely that it ensures that a recursion has a nonempty solution—and that it is again finitary. So finiteness is our final healthiness condition.

Those healthiness conditions ensure that the refinement ordering is as simple as possible. It is possible to simplify the healthiness conditions at the expense of complicating the ordering.

Those considerations are conveniently described using further relational notation. For $\mathbf{r} : \mathbf{X}_{\perp} \leftrightarrow \mathbf{X}_{\perp}$, \mathbf{r} is *finitary* means that for each $\mathbf{x} : \mathbf{X}_{\perp}$, the set $\mathbf{r}[\mathbf{x}]$ is either nonempty and finite, or equal to \mathbf{X}_{\perp} . Thus in particular a finitary relation is total.

Defining the model

\mathcal{D} denotes the space of relations on \mathbf{X}_{\perp} which are finitary and relate \perp to every element of \mathbf{X}_{\perp}

$$\mathcal{D} \hat{=} \{\mathbf{d} : \mathbf{X}_{\perp} \leftrightarrow \mathbf{X}_{\perp} \mid \left(\begin{array}{l} \mathbf{d} \text{ is finitary} \\ \mathbf{d}[\perp] = \mathbf{X}_{\perp} \end{array} \right)\}.$$

The space (\mathcal{D}, \supseteq) is partially ordered; note that the ordering is containment rather than inclusion. Its least, or most unrefined, element is the universal relation ω on \mathbf{X}_{\perp} ; its maximal elements are those which can be refined no fur-

ther: they are functions. Moreover:

Theorem(\mathcal{D} a domain). $(\mathcal{D}, \sqsupseteq)$ is a domain with least element ω , whose maximal elements are (total) functions and whose compact elements are those relations relating all but finitely many members of \mathbf{X}_\perp to each member of \mathbf{X}_\perp .

Galois connection

We wish to have a Galois connection from $(\mathcal{P}, \sqsubseteq)$ to $(\mathcal{D}, \sqsupseteq)$ which converts abortion in \mathcal{P} (where it is represented as partiality) to abortion in \mathcal{D} (where it is represented as arbitrary behaviour). Thus we must define $\varepsilon : \mathcal{P} \rightarrow \mathcal{D}$ by (recalling that \bar{r} denotes the complement of relation r), $\varepsilon.f \hat{=} f \cup \bar{f} ; \omega$. On the domain of f , $\varepsilon.f$ equals f ; but outside it $\varepsilon.f$ behaves like the universal relation. That captures the intuition behind our model, expressed in condition (5). We refer informally to $\varepsilon.f$ as the *fluffing up* of f !

Approach (c) yields the following definition of $\pi : \mathcal{D} \rightarrow \mathcal{P}$ forming a Galois connection:

$$\pi.d \hat{=} \{(x, y) : d \mid \left(\begin{array}{l} y \neq \perp \\ x d y' \Rightarrow y' = y \end{array} \right)\},$$

$\pi.d$ denotes the largest partial function in d which ‘‘accounts for all of d ’s results at its arguments’’. It may be thought of as the largest partial function which approximates, in \mathcal{P} , total relation d ; indeed that is the form for π given by adjunction.

Theorem(connecting \mathcal{P} to \mathcal{D}). Functions ε and π form a Galois connection from $(\mathcal{P}, \sqsubseteq)$ to $(\mathcal{D}, \sqsupseteq)$: $\mathbf{gc}(\varepsilon, \pi; \mathcal{P}, \mathcal{D})$.

Semantics

The Galois connection can now be used to make easier the task of describing the semantics of \mathcal{L} in \mathcal{D} . Indeed it enables us simply to lift, via ε , the semantics from \mathcal{P} .

Most of the denotations follow from those for \mathcal{P} . For the semantics of assignment we could either assume expression e is deterministic (as in the \mathcal{P} case), using non-deterministic choice to generate nondeterministic assignments; or we could permit e to be nondeterministic from the start. We adopt the latter approach, which requires e to be finitary as a relation in order for $\varepsilon.e$ to be well defined.

$$[\mathbf{skip}] \hat{=} \varepsilon.\iota_x \quad (= \iota_x \cup \omega \uparrow \perp)$$

$$[\mathbf{abort}] \hat{=} \varepsilon.\{\} \quad (= \omega)$$

$$[x := e] \hat{=} \varepsilon.\{(x, e.x) \mid e.x \text{ terminates}\}$$

$$e \text{ finitary}$$

$$[P ; Q] \hat{=} [P] ; [Q]$$

$$[P \triangleleft b \triangleright Q] \hat{=} \{(x, y) \mid x[P]y \triangleleft b.x \triangleright x[Q]y\}$$

$$[P \square Q] \hat{=} [P] \cup [Q]$$

$$[\mu F] \hat{=} \bigcap_{\kappa} F^{\kappa}.\omega$$

$$F \text{ monotone on } \mathcal{D}.$$

We wish to confirm, as we did for \mathcal{P} , that each denotation is well defined and that recursion can be given a countable unfolding because each combinator is continuous.

Theorem(semantics of \mathcal{D}). Each denotation in $[\mathcal{L}]$ is well defined, and the three combinators are continuous.

By virtue of the more general nature of its expressions, language \mathcal{L} is more expressive than the guarded-command language. We have seen that, with total functions for expressions, $[\mathcal{L}] = \mathcal{P}$. Evidently with arbitrary finitely-non-deterministic expressions, $[\mathcal{L}] = \mathcal{D}$.

Specification space

We now consider how the relational model can be used to support program development. Model \mathcal{D} captures exactly the semantics of code, as we have just seen. So first we must extend it to incorporate development constructs.

The ‘‘specification space’’ we consider consists of partial relations. Miracles are represented by non-total relations; and examples of combinators not in \mathcal{D} include complement and intersection (with some given relation). Though not code, they are extremely useful specification constructs.

Definition

The *specification space of relations* means the space $(\mathcal{R}, \sqsupseteq)$ where $\mathcal{R} \hat{=} \mathbf{X}_\perp \leftrightarrow \mathbf{X}_\perp$.

Since the specification space $(\mathcal{R}, \sqsupseteq)$ actually *contains* the space $(\mathcal{D}, \sqsupseteq)$ denoting ‘‘code’’, no Galois connection is required (in fact none exists).

Development

A development consists of a sequence of designs in \mathcal{R} , starting from the specification, each design refining its predecessor, and ending with the desired implementation (in \mathcal{D}). Laws are thus required, amongst others, for introducing the various constructs of language \mathcal{L} . Perhaps the most important are laws for introducing sequential composition. We consider those now.

In developing sequential programs we must be able to solve for S in $Q \sqsubseteq S ; P$. For only then can we implement

specification Q by calculation, given the final component P in a sequential composition. Similarly we must be able to solve for S in $Q \sqsubseteq P \circ S$, since that corresponds to being given the first component in the sequential composition.

Weakest pre- and postspecification

The weakest such S for which $Q \sqsubseteq S \circ P$, is called the *weakest prespecification of P in Q* , and is written Q/P . Thus $Q/P \sqsubseteq S \hat{=} Q \sqsubseteq S \circ P$. In that form, as a Galois connection, we see that Q/P exists by disjunctivity of \circ .

Similarly the weakest S for which $Q \sqsubseteq P \circ S$, is called the *weakest postspecification of P in Q* , defined $P \setminus Q \sqsubseteq S \hat{=} Q \sqsubseteq P \circ S$.

We shall need both forms of weakest specification in these notes. Here are the properties of weakest prespecification; exploration of analogues for weakest postspecification is left as an exercise. Proofs are mostly routine application of the Galois equivalence.

Theorem(basic properties).

1. $(Q/P) \circ P \sqsubseteq Q$
2. $Q \sqsubseteq (Q \circ P)/P$
3. $\iota \sqsubseteq Q/Q$
4. $Q = Q/\iota$
5. $Q/(P \circ R) = (Q/R)/P$
6. $(Q \cap P)/R = (Q/R) \cap (P/R)$
7. $Q/(P \cup R) = (Q/P) \cap (Q/R)$.

Theorem(further properties).

1. $Q/P = \bigcup \{R : \mathcal{R} \mid R \circ P \sqsubseteq Q\}$
2. $x_0(Q/P)x \hat{=} \forall y : X \bullet (xPy \Rightarrow x_0Qy)$
3. $Q/P = \overline{Q \circ P}$
4. For each total $F : \mathcal{P}$, $Q/F = Q \circ (\iota/F) = Q \circ F$
5. For each $f : \mathcal{P}$, $Q/\varepsilon.f = Q/f \cap \overline{\{\}/f \cup Q/\omega}$
6. $Q/(\mu F) = \bigcup \{Q/(F^n.\omega) \mid n : \mathbb{N}\}$.

The results for weakest pre and post-specification may be used for program development. But it is to be observed that they are not particularly simple to apply.

5 Relational data refinement

Introduction

So far we have considered only computations (programs or specifications) over the same state space. However a vital technique in program development permits a computation to be specified (clearly but at the expense of being unimplementable) using abstract state variables but implemented (efficiently but at the expense of being complicated) using so called *concrete* state variables.

Such use of local variables is controlled through use of *datatypes*, and the technique for developing correct concrete computations from their more abstract specifications is called *data refinement*. In this chapter we introduce those concepts and analyse them in the relational model. We consider two simulation rules which are sound methods of data refinement. By itself neither rule is complete,

yet together they are complete. This reflects essential properties of the relational model, and is to be compared with the completeness of just one of those rules in the predicate-transformer model (established in section 7).

Datatypes

A *datatype*, or module, controls access to its local variables through invocation of its named operations and by communication using their input and output parameters.

Example

The type *Mean* contains a bag of real numbers, which is initially empty. Numbers can be added to the bag, by invoking operation **add** with input parameter the number to be added; and the mean of a nonempty bag can be obtained by invoking operation **mean** whose output parameter gives that mean. We write bag comprehension like set comprehension but with $[\dots]$ replacing $\{ \dots \}$.

Type Mean

state	$b : \text{bag } R$
initially	$b := []$
add ($x? : R$)	$b := b + [x]$
mean ($y! : R$)	$(y := (\sum b)/\#b$ $\triangleleft b \neq [] \triangleright$ abort)

end

Definition

A *datatype* for use by programs having global state space G , consists of a state space L (for its local variables), an initialisation operation *in* from G to L , a finalisation operation *fin* from L to G (which is **skip** for type *Mean*, and is hence omitted), and a family \mathcal{O} of (named) operations on L having input and output. It is represented $T = (\text{in}, \mathcal{O}, \text{fin})$, with L left implicit to emphasise that it cannot be directly accessed by the programs on G which use T .

Use of datatype T by a program P on G consists of a sequence $\text{in} \circ P \circ \text{fin}$ in which P invokes the operations in \mathcal{O} .

Refinement

The two operations of *Mean* can be executed efficiently, without altering observable behaviour, by replacing the local variable, bag b , by its sum and cardinality. Although that loses information (like which numbers have been added) that information is not discernible from the operations of the type.

Type Mean'

state	$s : R; n : N$
initially	$s, n := 0, 0$
add' ($x? : R$)	$s, n := s + x, n + 1$
mean' ($y! : R$)	$(y := s/n$ $\triangleleft n \neq 0 \triangleright$ abort)

end

Each use of *Mean'* by a program of language \mathcal{L} (achieved by invoking operations of the type) could have been achieved by using type *Mean* instead. It is that which we wish to formalise in defining refinement between datatypes.

Conformality

Consider two datatypes which are *conformal* in the sense of having the same global state space and with their operations paired, so that for each $O : \mathcal{O}$, operation O of one type has input and output parameters of the same type as operation O' of the other:

$$\begin{aligned} \mathbf{T} &\cong (\mathbf{in}, \mathcal{O}, \mathbf{fin}) \\ \mathbf{T}' &\cong (\mathbf{in}', \mathcal{O}', \mathbf{fin}') \end{aligned}$$

Extend language \mathcal{L} to include named-procedure calls. Thus invocations of operations O and O' now become program statements. Suppose that $P \in \mathcal{L}$ includes procedure calls from just type \mathbf{T} and let P' denote the program which is like P except that it uses O' where P uses O . P' thus includes procedure calls from just type \mathbf{T}' .

Definition

Type \mathbf{T}' *refines* type \mathbf{T} , written $\mathbf{T} \sqsubseteq \mathbf{T}'$, means

$$\forall P : \mathcal{L} \bullet \mathbf{in} ; P ; \mathbf{fin} \sqsubseteq \mathbf{in}' ; P' ; \mathbf{fin}'.$$

In terms of the relational model,

$$\forall P : \mathcal{L} \bullet [\mathbf{in} ; P ; \mathbf{fin}] \supseteq [\mathbf{in}' ; P' ; \mathbf{fin}'].$$

Observe that proof of refinement between data types thus requires proof of that program refinement for all $P \in \mathcal{L}$. Our next task is to find a sufficient condition for that which is simple to establish but which has wide applicability.

Simulations

We require a condition sufficient for data refinement but which does not require reasoning over all $P \in \mathcal{L}$. Such a condition can be found by exploiting the states of the datatypes.

Example

Consider the conformal datatypes *Mean* and *Mean'*. The intuition that the former is refined by the latter can be based on the observation that the state $\mathbf{b} : \mathbf{bag} \mathbf{R}$ of the former is represented by the state $\mathbf{s} : \mathbf{R}; \mathbf{n} : \mathbf{N}$ of the latter, iff

$$\left(\begin{array}{l} \mathbf{s} = \sum \mathbf{b} \\ \mathbf{n} = \#\mathbf{b} \end{array} \right).$$

Now the intuition is completed by using that correspondence to pair uses of the type.

Definition

Given conformal datatypes \mathbf{T} and \mathbf{T}' as above with local states $\mathbf{l} : \mathbf{L}$ and $\mathbf{l}' : \mathbf{L}'$ respectively, relation $\mathbf{d} : \mathbf{L} \leftrightarrow \mathbf{L}'$ is a *downwards simulation*, or forwards simulation, or *dsim*

for short, from \mathbf{T} to \mathbf{T}' means

$$\begin{aligned} \mathbf{in} &\quad \mathbf{in} ; \mathbf{d} \sqsubseteq \mathbf{in}' \\ \mathbf{op} &\quad \forall O : \mathcal{O} \bullet O ; \mathbf{d} \sqsubseteq \mathbf{d} ; O' \\ \mathbf{fin} &\quad \mathbf{fin} \sqsubseteq \mathbf{d} ; \mathbf{fin}'. \end{aligned}$$

Definition

With \mathbf{T} and \mathbf{T}' as above, relation $\mathbf{u} : \mathbf{L}' \leftrightarrow \mathbf{L}$ is an *upwards simulation*, or backwards simulation, or *usim* for short, from \mathbf{T}' to \mathbf{T} means

$$\begin{aligned} &\quad \mathbf{u} \text{ is finitary} \\ \mathbf{in} &\quad \mathbf{in} \sqsubseteq \mathbf{in}' ; \mathbf{u} \\ \mathbf{op} &\quad \forall O : \mathcal{O} \bullet \mathbf{u} ; O \sqsubseteq O' ; \mathbf{u} \\ \mathbf{fin} &\quad \mathbf{u} ; \mathbf{fin} \sqsubseteq \mathbf{fin}'. \end{aligned}$$

Of course it is vital that those rules, intended as sufficient conditions for data refinement, are correct in the following sense.

Theorem(Soundness of simulations). If there is a *dsim* from \mathbf{T} to \mathbf{T}' or a *usim* from \mathbf{T}' to \mathbf{T} then \mathbf{T} is refined by \mathbf{T}' : $\mathbf{T} \sqsubseteq \mathbf{T}'$.

Utility

Simulations are used in documenting developments since, like loop invariants, they clarify otherwise complex code; in this case, they clarify the relationship that is supposed to exist between an abstract datatype and its implementation.

Perhaps more importantly simulations can be used in deriving a datatype from its more abstract specification. In terms of weakest postspecification (see section 4), the concrete operations \mathbf{in}' , O' and \mathbf{fin}' can be *specified*:

$$\begin{aligned} \mathbf{in}' &\cong \mathbf{in} ; \mathbf{d} \\ O' &\cong \mathbf{d} \setminus (O ; \mathbf{d}) \\ \mathbf{fin}' &\cong \mathbf{d} \setminus \mathbf{fin}. \end{aligned}$$

Completeness of simulations

We now consider the breadth of applicability of the simulation rules for ensuring data refinement.

First we observe that by itself neither rule is complete for establishing refinement.

Incompleteness of dsim

The abstract datatype, \mathbf{T} , has five states and a single operation, O , which after initialisation outputs 1 in one state and 0 in the other. It is defined using \mathbf{x}_0 for the initial state,

\mathbf{x} for the final state and \mathbf{y} for output:

Type T	
state	$\mathbf{x} : \{0, 1, 2, 3, 4\}$
initially	$\mathbf{x} := 0$
$O(\mathbf{y}! : \{0, 1\})$	$\left(\begin{array}{l} \mathbf{x}_0 = 0 \wedge \mathbf{x} = 1 \wedge \mathbf{y} = 0 \\ \vee \\ \mathbf{x}_0 = 0 \wedge \mathbf{x} = 2 \wedge \mathbf{y} = 0 \\ \vee \\ \mathbf{x}_0 = 1 \wedge \mathbf{x} = 3 \wedge \mathbf{y} = 1 \\ \vee \\ \mathbf{x}_0 = 2 \wedge \mathbf{x} = 4 \wedge \mathbf{y} = 0 \end{array} \right)$
finally	skip
end.	

In the concrete type, T' , there are four states and operation O' is nondeterministic.

Type T'	
state'	$\mathbf{x} : \{0, 1, 2, 3\}$
initially'	$\mathbf{x} := 0$
$O'(\mathbf{y}! : \{0, 1\})$	$\left(\begin{array}{l} \mathbf{x}_0 = 0 \wedge \mathbf{x} = 1 \wedge \mathbf{y} = 0 \\ \vee \\ \mathbf{x}_0 = 1 \wedge \mathbf{x} = 2 \wedge \mathbf{y} = 1 \\ \vee \\ \mathbf{x}_0 = 1 \wedge \mathbf{x} = 3 \wedge \mathbf{y} = 0 \end{array} \right)$
finally'	skip
end.	

Evidently there is no *dsim* from T to T' (although there is a *usim* from T' to T).

Dually, by an analogous example, *usim* is incomplete.

Completeness

However jointly the two simulation rules are complete for data refinement.

Theorem(Completeness theorem). If $T \sqsubseteq T'$ then there is a *usim* \mathbf{u} and a *dsim* \mathbf{d} for which $\mathbf{d} \circ \mathbf{u}$ is a relation from L to L' .

Corollary(of the method). If the operations of T are functions and if $T \sqsubseteq T'$ then there is a *dsim* from T to T' .

Thus the simulation rule *dsim* is complete for data types whose operations are total and deterministic.

6 Predicate-transformer semantics

Introduction

We now consider a more expressive model for computation which contains not only \mathcal{R} but also other important computational entities. Its lattice structure ensures that a single simulation rule is complete for data refinement, and it provides reasonable calculational facility.

Predicate transformers

Definition

For state space \mathbf{X} , we continue to identify \mathbf{PX} with the space of Boolean-valued functions (that is, predicates) on \mathbf{X} . The set, \mathcal{T} , of *predicate transformers* on \mathbf{X} is defined $\mathcal{T} \hat{=} \mathbf{PX} \rightarrow \mathbf{PX}$.

A predicate, \mathbf{q} , on \mathbf{X} which is thought of as being a predicate on final states (that is, after a computation) is called a *postcondition*. One, \mathbf{p} , which is thought of as being a predicate on initial states (that is, before a computation) is called a *precondition*. Predicate transformer $\mathbf{t} : \mathcal{T}$ represents a computation by mapping postconditions to preconditions as follows. For each postcondition $\mathbf{q} : \mathbf{PX}$, precondition $\mathbf{t}.\mathbf{q}$ holds at just those initial states from which computation \mathbf{t} is guaranteed to terminate in a state satisfying \mathbf{q} . (The phrase "is guaranteed to terminate" is included to emphasise application of the definition to nondeterministic computations.)

For each $\mathbf{t} : \mathcal{T}$ and $\mathbf{q} : \mathbf{PX}$, the predicate $\mathbf{t}.\mathbf{q}$ is called the *weakest precondition* of computation \mathbf{t} at postcondition \mathbf{q} .

Predicate transformers are ordered with the partial order lifted pointwise from implication on \mathbf{PX}

$$\mathbf{t} \sqsubseteq \mathbf{u} \hat{=} \forall \mathbf{q} : \mathbf{PX} \bullet \mathbf{t}.\mathbf{q} \leq \mathbf{u}.\mathbf{q}.$$

Thus \mathbf{t} is refined by \mathbf{u} iff for each postcondition \mathbf{q} , the weakest precondition of \mathbf{t} at \mathbf{q} is at least as strong as the weakest precondition of \mathbf{u} at \mathbf{q} . Thus whenever \mathbf{t} achieves \mathbf{q} so too does \mathbf{u} .

The least transformer is the function returning *false* everywhere (representing **abort**); the greatest transformer is the function returning *true* everywhere (representing **magic**).

The Galois connection with relations

We now define two functions

$$\mathbf{wp} : \mathcal{R} \rightarrow \mathcal{T} \quad \text{the embedding } \varepsilon$$

$$\mathbf{rp} : \mathcal{T} \rightarrow \mathcal{R} \quad \text{the projection } \pi$$

between the lattices (\mathcal{R}, \supseteq) and $(\mathcal{T}, \sqsubseteq)$.

The embedding \mathbf{wp}

For a relational computation $\mathbf{r} : \mathcal{R}$ and postcondition $\mathbf{q} : \mathbf{PX}$, define

$$\mathbf{wp}.\mathbf{r}.\mathbf{q}.\mathbf{x} \hat{=} \forall \mathbf{y} : \mathbf{X}_\perp \bullet \mathbf{x} \mathbf{r} \mathbf{y} \Rightarrow \mathbf{q}.\mathbf{y}.$$

The projection \mathbf{rp}

For a predicate transformer $\mathbf{t} : \mathcal{T}$, an initial state $\mathbf{x} : \mathbf{X}$ and a possibly improper final state $\mathbf{y} : \mathbf{X}_\perp$, define

$$\mathbf{x}(\mathbf{rp}.\mathbf{t})\mathbf{y} \hat{=} \forall \mathbf{q} : \mathbf{PX} \bullet \mathbf{t}.\mathbf{q}.\mathbf{x} \Rightarrow \mathbf{q}.\mathbf{y}.$$

Galois properties

For comparison with the definition of $gc(\varepsilon, \pi; \mathbf{X}, \mathbf{Y})$, note the correspondence:

$$\begin{array}{ccc} \mathbf{x} \leq \pi.\mathbf{y} & \equiv & \varepsilon.\mathbf{x} \leq \mathbf{y} \\ \updownarrow & & \updownarrow \\ \mathbf{r} \subseteq \mathbf{rp.t} & \equiv & \mathbf{wp.r} \sqsubseteq \mathbf{t}. \end{array}$$

Of course both formalise "the transformer is refined by the relation".

Consequences

The first two properties follow directly from the Galois connection; proof of the last is an exercise.

1. For all $\mathbf{r} : \mathcal{R}$, $\mathbf{r} \subseteq \mathbf{rp}(\mathbf{wp.r})$.
2. For all $\mathbf{t} : \mathcal{T}$, $\mathbf{t} \sqsubseteq \mathbf{wp}(\mathbf{rp.r})$.
3. The embedding is injective: $\mathbf{rp}(\mathbf{wp.r}) = \mathbf{r}$. Thus the pair of functions forms a Galois embedding

$$gc(\mathbf{wp}, \mathbf{rp}; \mathcal{R}, \mathcal{T}).$$

Healthiness conditions

We now recall Dijkstra's healthiness conditions for language \mathcal{L} , in the predicate-transformer model. For $\mathbf{r} : \mathcal{D}$ we have

1. *The excluded miracle*: $\mathbf{wp.r.false} = \mathbf{false}$.
2. *Monotonicity*: if $\mathbf{q} \Rightarrow \mathbf{q}'$ then $\mathbf{wp.r.q} \Rightarrow \mathbf{wp.r.q}'$.
3. *Conjunctivity*: $\mathbf{wp.r}(\mathbf{q} \wedge \mathbf{q}') = \mathbf{wp.r.q} \wedge \mathbf{wp.r.q}'$.
4. *Disjunctivity*: if \mathbf{r} is predetermined (deterministic unless aborting) then

$$\mathbf{wp.r}(\mathbf{q} \vee \mathbf{q}') = \mathbf{wp.r.q} \vee \mathbf{wp.r.q}'.$$

5. *Continuity*: if \mathbf{Q} is a directed set of predicates in the space (\mathbf{PX}, \leq) then

$$\mathbf{wp.r}.\sqcup \mathbf{Q} = \sqcup \mathbf{wp.r.Q}.$$

The \mathbf{wp} -image of \mathcal{R} is characterised by *positive* \sqcup -*junction*: for any nonempty set \mathbf{Q} of postconditions,

$$\mathbf{wp.r}.\sqcup \mathbf{Q} = \sqcup \{\mathbf{wp.r.q} \mid \mathbf{q} : \mathbf{Q}\}.$$

Note that (positive) \sqcup -junction implies monotonicity.

Other healthiness-style facts include

1. If \mathbf{r} is finitary then $\mathbf{wp.r}$ is \vee -continuous.
2. If \mathbf{r} is total then $\mathbf{wp.r}$ is strict (no miracles).
3. If \mathbf{t} is \vee -continuous then positive conjunctivity is equivalent to binary conjunctivity.
4. $\mathbf{wp.r}$ is disjunctive iff \mathbf{r} is predetermined.

 \mathbf{wp} is a functor

Since both \mathcal{R} and \mathcal{T} contain identities and are closed under composition, it is meaningful to seek their preservation under \mathbf{wp} . We have

1. $\mathbf{wp}.\iota_{\mathcal{R}} = \iota_{\mathcal{T}}$ and
2. $\mathbf{wp}(\mathbf{r} \circ \mathbf{s}) = \mathbf{wp.r} \circ \mathbf{wp.s}$,

which together mean that \mathbf{wp} is a (covariant?) functor.

The structure on transformers**Self-duality: conjugates**

For transformer $\mathbf{t} : \mathcal{T}$, recall that its *conjugate* \mathbf{t}^* is defined (see section 2) $\mathbf{t}^*.\mathbf{q} \hat{=} \neg \mathbf{t}.\neg \mathbf{q}$.

The dual lattice $(\mathcal{T}^*, \sqsupseteq)$ is isomorphic to the original $(\mathcal{T}, \sqsubseteq)$ with, for example,

1. $\mathbf{false}^* = \mathbf{true}$, and $\mathbf{true}^* = \mathbf{false}$;
2. $\mathbf{t}^{**} = \mathbf{t}$;
3. \mathbf{t} monotonic iff \mathbf{t}^* monotonic;
4. $(\mathbf{t} \circ \mathbf{u})^* = \mathbf{t}^* \circ \mathbf{u}^*$;
5. $\mathbf{t} \sqsubseteq \mathbf{u}$ iff $\mathbf{t}^* \sqsupseteq \mathbf{u}^*$; and
6. $(\mathbf{t} \sqcap \mathbf{u})^* = \mathbf{t}^* \sqcup \mathbf{u}^*$.

The operator \sqcup , created in the last item above by duality, is *angelic choice*: $(\mathbf{t} \sqcup \mathbf{u}).\mathbf{q} \hat{=} \mathbf{t}.\mathbf{q} \vee \mathbf{u}.\mathbf{q}$.

Non-relational transformers

Transformers \mathbf{t} for which $\mathbf{t}.\mathbf{false} \neq \mathbf{false}$ are *infeasible*, or *miraculous*, and correspond (if conjunctive) to relations that are not total: for some initial states there is "no" final state, and the computation "cannot be started".

Transformers \mathbf{t} for which $\mathbf{t}(\sqcup \mathbf{Q}) \neq \sqcup \{\mathbf{t}.\mathbf{q} \mid \mathbf{q} : \mathbf{Q}\}$ are *not continuous*, and correspond (if conjunctive) to relations that are not image finite; for some initial states there are infinitely many possible final states.

Transformers \mathbf{t} for which $\mathbf{t}(\mathbf{q} \wedge \mathbf{q}') \neq \mathbf{t}.\mathbf{q} \wedge \mathbf{t}.\mathbf{q}'$ are *not conjunctive*; they do not correspond to relations at all. Angelic choice does not preserve conjunctivity.

We continue to concentrate on transformers that are monotonic.

Angelic nondeterminism

Let \mathbf{X} be the integers, and let $\mathbf{t}^+, \mathbf{t}^-$ be the predicate transformers corresponding to programs that set the final state demonically to some non-negative, non-positive value.

The program $\mathbf{t}^0 \hat{=} \mathbf{t}^+ \sqcup \mathbf{t}^-$ is not conjunctive, and so does not correspond to any relation; in particular it does not represent the (deterministic) program "set the final state to 0" (which is the intersection of the two relations $\mathbf{rp.t}^+$ and $\mathbf{rp.t}^-$).

Execution can however be arranged by backtracking; but that requires a more elaborate model than the relational.

Characterisation of the monotonic transformers

The following theorem shows that there are no more unusual elements lurking in the monotonic transformer space than those already discovered:

Theorem(transformer factorisation). Any monotonic predicate transformer can be written as the angelic choice of conjunctive (that is, terminating relational) transformers.

The theorem means that execution of a (general) monotonic predicate transformer can be interpreted as a game of moves alternating between a player (who moves angelically) and an adversary (who moves demonically).

Strongest postconditions

A second embedding

Define two functions

$$\mathbf{sp} : \mathcal{R} \rightarrow \mathcal{T} \quad \text{the embedding } \varepsilon$$

$$\mathbf{qp} : \mathcal{T} \rightarrow \mathcal{R} \quad \text{the projection } \pi$$

between the lattices (\mathcal{R}, \supseteq) and $(\mathcal{T}, \sqsubseteq)$, with

$$\mathbf{sp.r.q} \supseteq \mathbf{x} \hat{=} \exists \mathbf{y} : \mathbf{X} \bullet \begin{pmatrix} \mathbf{q.y} \\ \mathbf{y.r.x} \end{pmatrix}.$$

The name **sp** stands for *strongest postcondition* (and “ \supseteq ” is pronounced “epsilonoff”).

It is possible to define (indeed, to calculate) a definition for **qp** that makes $(\mathbf{sp}, \mathbf{qp})$ a Galois connection.

wp versus sp

For each relation $\mathbf{r} : \mathcal{R}$ the two transformers $\mathbf{wp.r}$ and $\mathbf{sp.r}$ are adjoint (but see section 7).

Weakest liberal preconditions

Purpose of wlp

The weakest *liberal* precondition makes behavioural distinctions that **wp** cannot detect, and *vice versa*. For example

$$\mathbf{skip} =_{\mathbf{wlp}} \mathbf{skip} \sqcap \mathbf{abort} =_{\mathbf{wp}} \mathbf{abort}.$$

The distinction is based on *not* assuming that **abort** “can do anything” (as well as fail to terminate). For example, in that interpretation the program

$$\mathbf{x} := 0 \sqcap \mathbf{abort}$$

might fail to set variable \mathbf{x} to 0, but it can't set \mathbf{x} to 1.

The relational model is as before, except that we no longer require $\mathbf{x.r.y}$ whenever $\mathbf{x.r} \perp$.

Definition of wlp

We define two functions

$$\mathbf{wlp} : \mathcal{R} \rightarrow \mathcal{T} \quad \text{the embedding } \varepsilon$$

$$\mathbf{rlp} : \mathcal{R} \rightarrow \mathcal{T} \quad \text{the projection } \pi$$

between the lattices (\mathcal{R}, \supseteq) and $(\mathcal{T}, \sqsubseteq)$, with

$$\mathbf{wlp.r.q.x} \hat{=} \forall \mathbf{y} : \mathbf{X}_{\perp} \bullet \mathbf{x.r.y} \Rightarrow (\mathbf{y} = \perp \vee \mathbf{q.x}).$$

Thus $\mathbf{wlp.r.q}$ holds at just those initial states from which execution of \mathbf{r} either diverges or terminates in a state satisfying \mathbf{q} .

The Galois property can be used to calculate a corresponding definition for **rlp**.

The wp and wlp orders

On the (new) relational space, the **wp** order is derived from the *Smyth* preorder on subsets of \mathbf{X} :

$$\mathbf{L} \sqsubseteq_{\mathbf{S}} \mathbf{H} \hat{=} \forall \mathbf{h} : \mathbf{H} \bullet \exists \mathbf{l} : \mathbf{L} \bullet \mathbf{l} \sqsubseteq \mathbf{h}.$$

We have

$$\mathbf{r} \sqsubseteq_{\mathbf{wp}} \mathbf{s} \hat{=} \forall \mathbf{x} : \mathbf{X} \bullet \mathbf{x}.\langle \mathbf{r} \rangle \sqsubseteq_{\mathbf{S}} \mathbf{x}.\langle \mathbf{s} \rangle,$$

where as before $\mathbf{x}.\langle \mathbf{r} \rangle$ denotes the image of $\{\mathbf{x}\}$ through relation \mathbf{r} .

“Fluffing up” converts the preorder to an ordinary partial order, and $\sqsubseteq_{\mathbf{S}}$ becomes simple \supseteq .

The **wlp** order is derived from the *Hoare* pre-order on subsets of \mathbf{X} :

$$\mathbf{L} \sqsubseteq_{\mathbf{H}} \mathbf{H} \hat{=} \forall \mathbf{l} : \mathbf{L} \bullet \exists \mathbf{h} : \mathbf{H} \bullet \mathbf{l} \sqsubseteq \mathbf{h}.$$

We have $\mathbf{r} \sqsubseteq_{\mathbf{wlp}} \mathbf{s} \hat{=} \forall \mathbf{x} : \mathbf{X} \bullet \mathbf{x}.\langle \mathbf{r} \rangle \sqsupseteq_{\mathbf{H}} \mathbf{x}.\langle \mathbf{s} \rangle$; note the reversal of the order. “Fluffing down” converts the preorder to an ordinary partial order, and $\sqsubseteq_{\mathbf{H}}$ becomes simple \sqsubseteq .

7 Predicate-transformer data refinement

In this chapter we return to *data refinement*—last considered relationally—in the setting of predicate transformers. The result is more than a simple reformulation; with the extra structure of predicate transformers we have a single complete rule for data refinement.

Data refinement

Recall that a *datatype* is defined by its initialisation **in**, its family of operations \mathcal{O} (an indexed set) and its finalisation **fin**.

A use of a datatype is a composition $\mathbf{in}; \mathbf{P}(\mathcal{O}); \mathbf{fin}$ for some program scheme \mathbf{P} referring by index to individual operations in \mathcal{O} .

One datatype $(\mathbf{in}, \mathcal{O}, \mathbf{fin})$ is *refined* by another datatype $(\mathbf{in}', \mathcal{O}', \mathbf{fin}')$ if for all uses \mathbf{P} we have

$$\mathbf{in}; \mathbf{P}(\mathcal{O}); \mathbf{fin} \sqsubseteq \mathbf{in}'; \mathbf{P}'(\mathcal{O}'); \mathbf{fin}',$$

with \sqsubseteq denoting ordinary refinement between programs.

Informally, $(\mathbf{in}, \mathcal{O}, \mathbf{fin})$ is refined by $(\mathbf{in}', \mathcal{O}', \mathbf{fin}')$ if no program \mathbf{P} can show that $(\mathbf{in}', \mathcal{O}', \mathbf{fin}')$ behaves in a way that $(\mathbf{in}, \mathcal{O}, \mathbf{fin})$ could not.

Simulation

Simulation is a relationship established component-wise between two datatypes, as below, in order to show data refinement between the datatypes. It avoids having to consider all their possible uses \mathbf{P} .

Recall that there are two kinds of simulation. A *dsim* is an operation **rep** satisfying

$$\mathbf{in} : \mathbf{in}; \mathbf{rep} \sqsubseteq \mathbf{in}',$$

op: for all corresponding O, O' in $\mathcal{O}, \mathcal{O}'$ respectively

$$O \text{ ; rep } \sqsubseteq \text{ rep ; } O' \text{ and}$$

fin: $\text{fin} \sqsubseteq \text{rep ; fin}'$.

A *usim* is an operation **rep** (which in the presence of recursion **rep** must also be continuous and strict) satisfying

in: $\text{in} \sqsubseteq \text{rep ; in}'$,

op: for all corresponding O, O' in $\mathcal{O}, \mathcal{O}'$ respectively

$$\text{rep ; } O \sqsubseteq O' \text{ ; rep, and}$$

fin: $\text{fin ; rep} \sqsubseteq \text{fin}'$.

We have seen that in the relational setting, those simulations are sound and jointly complete.

Weak inverses

Any conjunctive and terminating predicate transformer \mathbf{P} corresponds to some relation and so can be written wp.r for that relation r .¹ It therefore has an adjoint sp.r , which we write $\widehat{\mathbf{P}}$, satisfying

$$\begin{array}{l} \iota \sqsubseteq \mathbf{P ; } \widehat{\mathbf{P}} \\ \widehat{\mathbf{P}} \text{ ; } \mathbf{P} \sqsubseteq \iota. \end{array}$$

The requirement that \mathbf{P} be terminating is necessary, corresponding in lattice terms to $\mathbf{P.true} = \text{true}$. To that extent section 6 was not telling the whole truth.

Construction of the data refinement

Given two datatypes with

$$(\text{in}, \mathcal{O}, \text{fin}) \sqsubseteq (\text{in}', \mathcal{O}', \text{fin}')$$

in proving the completeness theorem in section 5 we constructed a simulation whose form we mimick now. Set

$$\text{rep} \hat{=} \sqcup \{ (\widehat{\text{in}'} \text{ ; } \mathbf{S}') \text{ ; } \text{in ; } \mathbf{S} \mid \mathbf{S} \in \mathcal{O}' \},$$

where \mathcal{O}' denotes the closure of \mathcal{O} under sequential composition.

The proofs of the three properties are then direct, if not completely obvious.

Nontermination

When \mathbf{P} does not necessarily terminate, it may be that the adjoint $\widehat{\mathbf{P}}$ does not exist. Luckily it is possible to generalise that definition, by using assumptions, to

$$\begin{array}{l} \{\text{wp.P.true}\} \sqsubseteq \mathbf{P ; } \widehat{\mathbf{P}} \\ \widehat{\mathbf{P}} \text{ ; } \mathbf{P} \sqsubseteq \iota, \end{array}$$

where $\{\text{wp.P.true}\}$ is a program that aborts when \mathbf{P} would abort, and behaves like **skip** otherwise (when \mathbf{P} is guaranteed to terminate).

¹We write \mathbf{P} rather than \mathbf{t} for transformers here, to be consistent with the convention for datatypes.

The proof of completeness then goes through with the modified definition, but for *soundness* in the presence of possible failure to terminate we require that **rep** be strict and \sqcup -continuous. And the **rep** generated by the completeness construction is guaranteed to be so only if the abstract program itself is continuous—it must contain no unbounded nondeterminism.

8 Process semantics

Introduction

We now turn to another programming paradigm: that of process algebra. Settling upon the notation CSP we study its hierarchy of three progressively finer semantic models which are usually linked by the natural embeddings up the hierarchy. In order to obtain Galois connections we replace those embeddings by alternatives, based on a version of “fluffing up” appropriate to interactive processes, and identify their adjoints.

The result is that the standard hierarchy of models for CSP now becomes founded on embeddings that are capable of straightforward extension when it is desired to incorporate the results of further observations in the semantics. Such extensions are to be found, for example, in timed models and probabilistic models.

Processes

A process, \mathbf{P} , has a finite *alphabet*, $\alpha\mathbf{P}$, of *events* in which it may engage. It is expressed algebraically, according to which events it may next perform. We consider a slightly restricted subset of standard CSP, ignoring termination, sequential composition and interleaving using only the following combinators.

- *General choice*, $(\mathbf{x} : \mathbf{X} \rightarrow \mathbf{P}(\mathbf{x}))$, offers its environment any of the events $\mathbf{x} : \mathbf{X}$; after performing \mathbf{x} it behaves like $\mathbf{P}(\mathbf{x})$. Special cases are
 - $\mathbf{X} = \{ \}$, called *deadlock* and written **stop**
 - $\mathbf{X} = \{ \mathbf{x} \}$, called *prefix* and written $\mathbf{x} \rightarrow \mathbf{P}$
 - $\mathbf{X} = \{ \mathbf{x}, \mathbf{y} \}$, called *choice* and written

$$(\mathbf{x} \rightarrow \mathbf{P}(\mathbf{x}) \mid \mathbf{y} \rightarrow \mathbf{P}(\mathbf{y})).$$

- *Recursion*, $\mu \mathbf{Z} : \mathbf{A} \bullet \mathbf{F.Z}$, is the least fixed point of continuous function \mathbf{F} on processes having alphabet \mathbf{A} . For example

$$\mu \mathbf{Z} : \{0, 1\} \bullet 0 \rightarrow 1 \rightarrow \mathbf{Z}$$

alternately performs 0 then 1 forever.

- *Parallel composition*, $\mathbf{P} \parallel \mathbf{Q}$, synchronises on events common to \mathbf{P} and \mathbf{Q} and interleaves others.
- *External choice*, $\mathbf{P} \square \mathbf{Q}$, offers its environment the choice between \mathbf{P} and \mathbf{Q} on the first interaction.
- *Abstraction*, $\mathbf{P} \text{ hide } \mathbf{E}$, for $\mathbf{E} \subseteq \alpha\mathbf{P}$, behaves like \mathbf{P} but with the events in \mathbf{E} concealed, or internalised.
- *Nondeterministic choice*, $\mathbf{P} \sqcap \mathbf{Q}$, arises from abstraction. It behaves like either \mathbf{P} or \mathbf{Q} , but the environment

has no influence over which.

- *Abort, chaos*, exhibits arbitrary behaviour.

Refinement between processes, $P \sqsubseteq Q$, holds iff P and Q have the same alphabet, $\alpha P = \alpha Q$, and in every environment P can be replaced by Q . As usual equality between processes, $P = Q$, means refinement of each side by the other. Laws to formalise the concept of refinement include

$$x \rightarrow (P \parallel Q) = (x \rightarrow P) \parallel (x \rightarrow Q)$$

$$\begin{aligned} & (x \rightarrow P \square x \rightarrow Q) \text{ hide } \{x\} \\ & = \\ & (P \text{ hide } \{x\}) \sqcap (Q \text{ hide } \{x\}) \end{aligned}$$

$$P \sqcap Q \sqsubseteq P.$$

Models

There are three standard models of CSP: traces; failures; and divergences. They capture increasingly more detailed behaviour of a process: safety; liveness; and divergence. We now summarise them, reminding the reader of the healthiness conditions for each.

Traces

The *traces model* of CSP captures safety properties and is complete for deterministic processes. In it, each process is denoted by the set of finite sequences of events in which it can engage, called its *traces* (defined by structural induction). One process refines another if it can perform all the traces of that other.

Thus for finite universe A of events, the *traces model* over A is defined to be the space $(\mathcal{T}, \sqsubseteq_{\mathcal{T}})$ where

$$\begin{aligned} & \mathcal{T} \\ & \cong \\ & \{\mathbf{T} \subseteq \text{seq } A \mid \mathbf{T} \text{ is nonempty and prefix closed}\} \end{aligned}$$

and where $\mathbf{T} \sqsubseteq_{\mathcal{T}} \mathbf{U} \hat{=} \mathbf{T} \subseteq \mathbf{U}$.

The least process is **stop** which performs no events at all. So recursion is given, for continuous F , by

$$[\mu Z \bullet F.Z] = \bigcup_{n:N} [F^n.\text{stop}].$$

However traces refinement is unable to distinguish $P \sqcap \text{stop}$ from P , since their traces are the same. Thus development based on it could not ensure absence of deadlock! It is also unable to distinguish internal and external choice $[x \rightarrow P \square y \rightarrow Q] = [x \rightarrow P \sqcap y \rightarrow Q]$. That distinction lies with the failures model!

Failures

The *failures model* of CSP captures liveness properties and is complete for nondivergent processes. In it, each process is denoted by a relation between the set of traces of the process and the set of all subsets of its alphabet, relating each trace to the sets of events which may be refused (that is, may lead to deadlock if its environment offers them) by the process immediately after it has engaged in that trace. Such failures information is sufficient to express nondeter-

minism, with the result that a process which is less deterministic than another contains the failures of that other.

The failures of CSP are defined by structural induction. In the failures model one process refines another if its behaviour is at least as deterministic.

Thus for universe A of events, the *failures model* of CSP is defined to be the space $(\mathcal{F}, \sqsubseteq_{\mathcal{F}})$ where \mathcal{F} denotes the set of relations F from $\text{seq } A$ to subsets of A satisfying:

$$\left(\begin{array}{l} \text{dom } F \in \mathcal{T} \\ \left(\begin{array}{l} t F E \\ D \subseteq E \end{array} \right) \Rightarrow t F D \\ \left(\begin{array}{l} t F E \\ x \in A \end{array} \right) \Rightarrow \begin{array}{l} t F (E \cup \{x\}) \\ \vee \\ t \langle x \rangle F \{ \} \end{array} \end{array} \right),$$

and where $F \sqsubseteq_{\mathcal{F}} G \hat{=} F \supseteq G$.

The least process is **havoc** which may perform or refuse any element in A at any interaction (but never diverges). So recursion is given, for continuous F , by

$$[\mu Z \bullet F.Z] = \bigcap_{n:N} [F^n.\text{havoc}].$$

In the failures model nondeterminism can be distinguished; in particular P is distinguished from $P \sqcap \text{stop}$ (which it cannot in the traces model). Indeed a process is *nondeterministic* means at some interaction it can both perform and refuse an event. Thus $P : \mathcal{F}$ is *deterministic* means

$$t \langle a \rangle \in \text{dom } P \Rightarrow \neg t P \{a\},$$

and otherwise P is nondeterministic. For example, $x \rightarrow P \square y \rightarrow Q$ is deterministic but $x \rightarrow P \sqcap y \rightarrow Q$ is not.

However the failures model is unable to distinguish arbitrary refusal behaviour from *livelock* (or divergence): $[\text{havoc}] = [\text{chaos}]$. For that we must turn to the divergences model.

Divergences

The *divergences model* of CSP captures divergence information and is complete for CSP. In it, each process is denoted by a pair (F, D) where F is the failures, as above, and D is the set of traces after which the process may diverge. For consistency with \mathcal{F} , if P diverges at trace t then after having engaged in t , P can also subsequently perform or refuse any event. One process refines another if it is at least as deterministic and at least as nondivergent.

Thus for universe A , the *divergences model* of CSP is defined to be the space $(\mathcal{C}, \sqsubseteq_{\mathcal{C}})$ where \mathcal{C} denotes the set of pairs (F, D) with $F \in \mathcal{F}$ and $D \subseteq \text{dom}.F$ satisfying

$$\left(\begin{array}{l} \left(\begin{array}{l} t \in D \\ u \in \text{seq } A \end{array} \right) \Rightarrow tu \in D \\ \left(\begin{array}{l} t \in D \\ E \subseteq A \end{array} \right) \Rightarrow t F E \end{array} \right),$$

and where

$$(\mathbf{F}, \mathbf{D}) \sqsubseteq_{\mathcal{C}} (\mathbf{G}, \mathbf{E}) \hat{=} \left(\begin{array}{l} \mathbf{F} \supseteq \mathbf{G} \\ \mathbf{D} \supseteq \mathbf{E} \end{array} \right).$$

The least process is **chaos** which may diverge, or perform, or refuse any element in **A** at any interaction.

In the divergences model divergence is distinguished from arbitrary refusal behaviour; in particular **havoc** is distinct from **chaos** since the former has no divergences whilst the latter has every possible sequence as a divergence.

Summary

As summary of those models we have:

Theorem(CSP domains). Each of the spaces $(\mathcal{T}, \sqsubseteq_{\mathcal{T}})$, $(\mathcal{F}, \sqsubseteq_{\mathcal{F}})$ and $(\mathcal{C}, \sqsubseteq_{\mathcal{C}})$ is a domain and the operations of CSP provide continuous functions on each of them.

Embedding \mathcal{T} in \mathcal{F}

In this section we replace the standard embedding of \mathcal{T} in \mathcal{F} by one which forms a Galois embedding.

The standard embedding

The *deterministic embedding* $\hat{\cdot} : \mathcal{T} \rightarrow \mathcal{F}$ assigns to traces process **D** the failures process $\hat{\mathbf{D}}$ which performs all events that **D** performs and refuses all others:

$$t \hat{\mathbf{D}} \hat{\mathbf{E}} \hat{=} \left(\begin{array}{l} t \in \mathbf{D} \\ \forall a : \mathbf{A} \bullet t \langle a \rangle \in \mathbf{D} \Rightarrow a \notin \mathbf{E} \end{array} \right).$$

Since it is unable to refuse any event it performs, $\hat{\mathbf{D}}$ is deterministic. Each trace it performs is a trace of **D** and conversely. Thus:

Lemma(Deterministic embedding). $\hat{\mathbf{D}}$ is deterministic with the same traces as **D**.

However in order to be the embedding of a Galois connection, $\hat{\cdot}$ would have to be monotone (by the equivalence theorem of section 3). Yet it is not; for in \mathcal{T} , $\mathbf{stop} \sqsubseteq_{\mathcal{T}} \mathbf{a} \rightarrow \mathbf{stop}$ since the latter can perform every trace the former can. But in \mathcal{F} that refinement fails since

$$\left(\begin{array}{l} \neg \langle a \rangle \mathbf{stop} \{ \} \\ \langle a \rangle (\mathbf{a} \rightarrow \mathbf{stop}) \{ \} \end{array} \right).$$

We must look a little further.

Recall that in section 4 we embedded the space \mathcal{P} of partial functions in the space \mathcal{D} of total relations by an embedding ε which "fluffed up" its argument: outside the domain of partial function **f**, $\varepsilon.f$ was as unconstrained as possible, whilst on $\mathbf{dom.f}$ it equalled **f**. Being as unconstrained as possible meant behaving like the least element, **abort**, the universal relation.

Now we are in an identical situation. In order to find a Galois connection from \mathcal{T} to \mathcal{F} we must "fluff up" a traces process **D** so that it behaves in as unconstrained a manner as possible unless it behaves like **D**. Thus $\varepsilon.\mathbf{D}$ behaves like the least element **chaos** unless **D** performs an event in

which case $\varepsilon.\mathbf{D}$ cannot refuse that event. This leads us to the following definition.

The liberal embedding

The *liberal embedding* $^+ : \mathcal{T} \rightarrow \mathcal{F}$ assigns to $\mathbf{D} : \mathcal{T}$ the process $\mathbf{D}^+ : \mathcal{F}$ whose only constraint is that it does not refuse an event performed by **D**:

$$t \mathbf{D}^+ \hat{\mathbf{E}} \hat{=} \forall a : \mathbf{A} \bullet t \langle a \rangle \in \mathbf{D} \Rightarrow a \notin \mathbf{E}.$$

Since \mathbf{D}^+ has arbitrary traces it is, in general, nondeterministic (unlike $\hat{\mathbf{D}}$). The following result establishes that claim together with two results about refinement of $\hat{\mathbf{D}}$ in \mathcal{F} .

Lemma(Liberal embedding). \mathbf{D}^+ is refined in \mathcal{F} by $\hat{\mathbf{D}}$. Every sequence from its alphabet is a trace of \mathbf{D}^+ , and if \mathbf{D}^+ is refined in \mathcal{F} by **P** then every trace of **P** is a trace of **D**.

Calculating the adjoint

To identify the adjoint of $^+$ we have (approach (c) works well) $^- : \mathcal{F} \rightarrow \mathcal{T}$ defined

$$\mathbf{P}^- \hat{=} \{ r : \mathbf{dom.P} \mid u \langle a \rangle \leq r \Rightarrow \neg u \mathbf{P} \{ a \} \}.$$

Function $^-$ thus retains deterministic behaviour by discarding events that are both performed and refused.

A Galois connection

The manipulation above provides us with the Galois connection we seek between \mathcal{T} and \mathcal{F} . In fact the embedding and projection form a Galois embedding (and moreover the projection is also \sqsubseteq -junctive:

Theorem(Liberal connection). The liberal embedding $^+$ and its adjoint $^-$ form a Galois embedding: $\mathbf{ge}(^+, -; \mathcal{T}, \mathcal{F})$.

Since the Galois connection is a Galois embedding, the isomorphism theorem of section 3 simply yields the isomorphism between \mathcal{T} and the image of the embedding in \mathcal{F} . Every element of \mathcal{T} is a fixed point of the embedding followed by the projection, whilst the fixed points of the projection followed by the embedding are those processes which at each interaction and for each event either perform the event or behave chaotically by both performing and refusing it.

Embedding \mathcal{F} in \mathcal{C}

In this section we replace the standard embedding of \mathcal{F} in \mathcal{C} by one which forms the basis of a Galois embedding. Combining this with the result of the previous section we will obtain a Galois embedding of \mathcal{T} in \mathcal{C} .

The standard embedding

The *divergence-free embedding* $^\circ : \mathcal{R} \rightarrow \mathcal{C}$ assigns to failures process **F** the divergences process having the same failures as **F** but having no divergences:

$$\mathbf{F}^\circ \hat{=} (\mathbf{F}, \{ \}).$$

Now \circ is monotone and continuous, but no process is mapped by it to **chaos**. Hence (by the junctivity theorem of section 3) it cannot be the embedding in a Galois connection.

The divergent embedding

Were we to apply the reasoning from earlier in the section, we would expect to embed \mathcal{F} in \mathcal{C} by mapping $\mathbf{F} : \mathcal{F}$ to a process in \mathcal{C} which is as unrefined as possible; it certainly cannot in general have empty divergences.

A natural first attempt is to map \mathbf{F} to $(\mathbf{F}, \text{seq.A})$, which has all possible divergences. Unfortunately that fails the (first and third) healthiness conditions for \mathcal{C} . But it *does* satisfy them when \mathbf{F} behaves like **havoc**, which suggests that instead \mathbf{F} be mapped to a process which diverges only when \mathbf{F} behaves like **havoc**; for then the result does lie in \mathcal{C} . That leads us to the following definition.

The *divergent embedding* $\oplus : \mathcal{F} \rightarrow \mathcal{C}$ ensures that \mathbf{F}^\oplus diverges at t iff \mathbf{F} after $t = \text{havoc}$:

$$\begin{aligned} & \mathbf{F}^\oplus \\ & \cong \\ & (\mathbf{F}, \{t : \text{seq.A} \mid \forall s : \text{seq.A} \bullet \forall E \subseteq A \bullet \text{st } \mathbf{F} E\}). \end{aligned}$$

Calculating the adjoint

The adjoint of \oplus is readily calculated by either approach (b) or (c) to be \ominus where

$$(\mathbf{G}, \mathbf{D})^\ominus \cong \mathbf{G}.$$

A Galois connection

Theorem(Divergent embedding). The divergent embedding and its adjoint \ominus form a Galois embedding:

$$\text{ge}(\oplus, \ominus; \mathcal{F}, \mathcal{C}).$$

Furthermore \ominus is \sqcup -junctive.

Embedding \mathcal{T} in \mathcal{C}

Combining the results of the previous sections we obtain a Galois embedding from \mathcal{T} to \mathcal{C} .

Theorem(CSP embedding). The liberal embedding followed by the divergent embedding and their inverses composed in the reverse order, form a Galois embedding

$$\text{ge}(\oplus \circ \hat{}, \ominus \circ \hat{}; \mathcal{T}, \mathcal{C}).$$

Furthermore the composition of the projections is \sqcup -junctive.

9 Probabilistic semantics

Introduction

We have considered two standard programming paradigms (the imperative language \mathcal{L} and the communicating process language CSP) and studied several semantic models

for each. In this chapter we turn to a more novel topic, that of probabilistic imperative programs, and show how the same techniques can be used to provide a semantics for it.

The language \mathcal{L}_\oplus consists of \mathcal{L} extended with a combinator $\mathbf{p}\oplus$, for each $\mathbf{p} : [0, 1]$, which chooses its left-hand argument with probability \mathbf{p} and its right-hand argument with probability $1 - \mathbf{p}$. Such a language is sufficient to express the probabilistic algorithms championed by Rabin and others (for a survey of which, see R. Motwani and P. Raghavan, *Randomized Algorithms*, CUP, 1995).

For language \mathcal{L}_\oplus we consider: both a relational semantics and an extended-predicate-transformer semantics; a Galois connection between them; healthiness conditions; and some programming techniques for probabilistic algorithms (namely for reasoning about loops). Subtleties of the theory are due to the interplay between probability and nondeterminism.

Programming-language syntax

Our probabilistic programming language \mathcal{L}_\oplus is the language \mathcal{L} extended with probabilistic choice $\mathbf{p}\oplus$, for $\mathbf{p} : [0, 1]$.

Let \mathbf{P} range over programs, \mathbf{b} over Boolean expressions, and \mathbf{p} over real number expressions between 0 and 1 inclusive; assume that \mathbf{x} stands for a list of distinct variables, and \mathbf{E} for a list of expressions; and let the program *scheme* \mathcal{C} be a program in which the program *name* \mathbf{Z} can appear.

$$\begin{aligned} \mathbf{P} \cong & \text{abort} \mid \text{Skip} \mid \mathbf{x} := \mathbf{E} \mid \mathbf{P} ; \mathbf{P} \mid \\ & \mathbf{P} \triangleleft \mathbf{b} \triangleright \mathbf{P} \mid \mathbf{P} \mathbf{p}\oplus \mathbf{P} \mid \mathbf{P} \sqcap \mathbf{P} \mid \\ & (\mu \mathbf{Z} \bullet \mathcal{C}) \end{aligned}$$

Relational semantics

The result of a *deterministic* probabilistic program is a probability distribution over its state space. For a *nondeterministic* probabilistic program the result is a set of distributions, one for each resolution of the nondeterminism. A probabilistic program is therefore modelled relationally by a function from its initial state to a set of distributions over final state.

Distributions

We formalise that by introducing the following notation for distributions. As usual, \mathbf{X} denotes the underlying state space. $\Pi.\mathbf{X}$ denotes the set of all subsets of \mathbf{X} .

A (discrete probability) *distribution* means a function $\mathbf{F} : \Pi.\mathbf{X} \rightarrow [0, 1]$ satisfying, for all $\mathbf{U}, \mathbf{V} \in \Pi.\mathbf{X}$,

$$\begin{aligned} \mathbf{U} \subseteq \mathbf{V} & \Rightarrow \mathbf{F}.\mathbf{U} \leq \mathbf{F}.\mathbf{V} \\ \mathbf{F}.\mathbf{U} \cup \mathbf{V} & \geq \mathbf{F}.\mathbf{U} + \mathbf{F}.\mathbf{V} - \mathbf{F}.\mathbf{U} \cap \mathbf{V} \\ \mathbf{F}.\{\} & = 0. \end{aligned}$$

Note that distributions do not necessarily satisfy $\mathbf{F}.\mathbf{X} = 1$. Often we shall write $\mathbf{F}.\mathbf{x}$ for $\mathbf{F}.\{\mathbf{x}\}$.

We write $(\tilde{\mathbf{X}}, \sqsubseteq)$ for the space of distributions over \mathbf{X} ,

with order defined

$$\mathbf{F} \sqsubseteq \mathbf{G} \hat{=} \forall \mathbf{x} : \mathbf{X} \bullet \mathbf{F}.\mathbf{x} \leq \mathbf{G}.\mathbf{x}.$$

Note that \mathbf{X} is embedded in $\tilde{\mathbf{X}}$ via the point-mass function: for $\mathbf{x} : \mathbf{X}$ its point mass $\tilde{\mathbf{x}} : \tilde{\mathbf{X}}$ assigns 1 to any set containing \mathbf{x} and 0 to any other set

$$\tilde{\mathbf{x}}.U = (\mathbf{x} \in U).$$

The relational semantic space for \mathcal{L}_Θ is thus a subset of $\mathbf{X} \rightarrow \Pi.\tilde{\mathbf{X}}$. That leads us to consider the partially-ordered space $(\mathcal{H}, \sqsubseteq)$, where

$$\mathcal{H} \hat{=} \mathbf{X} \rightarrow \Pi.\tilde{\mathbf{X}},$$

and \sqsubseteq is containment; for $\mathbf{P}, \mathbf{P}' : \mathbf{X} \rightarrow \Pi.\tilde{\mathbf{X}}$, it is defined

$$\mathbf{P} \sqsubseteq \mathbf{P}' \hat{=} \forall \mathbf{x} \in \mathbf{X} \bullet \mathbf{P}.\mathbf{x} \supseteq \mathbf{P}'.\mathbf{x}.$$

Relational semantics

The semantics of language \mathcal{L}_Θ is given by the function

$$[] : \mathcal{L}_\Theta \rightarrow \mathcal{H}$$

defined as follows.

$$[\mathbf{Skip}].\mathbf{x} \hat{=} \{\tilde{\mathbf{x}}\}$$

$$[\mathbf{abort}].\mathbf{x} \hat{=} \tilde{\mathbf{X}}$$

$$[\mathbf{x} := \mathbf{E}].\mathbf{s} \hat{=} \mathbf{s}[\tilde{\mathbf{x}} := \mathbf{E}]$$

$$[\mathbf{P} \sqcap \mathbf{P}'].\mathbf{x} \hat{=} \cup\{(\mathbf{P} \mathbf{p} \oplus \mathbf{P}').\mathbf{x} \mid \mathbf{p} \in [0, 1]\}$$

$$[\mathbf{P} \triangleleft \mathbf{b} \triangleright \mathbf{P}'].\mathbf{x} \hat{=} \mathbf{P}.\mathbf{x} \triangleleft \mathbf{b}.\mathbf{x} \triangleright \mathbf{P}'.\mathbf{x}$$

$$[\mathbf{P} \mathbf{p} \oplus \mathbf{P}'].\mathbf{x} \hat{=} \{\mathbf{p}\mathbf{E} + (1 - \mathbf{p})\mathbf{F} \mid \left(\begin{array}{l} \mathbf{E} \in \mathbf{P}.\mathbf{x} \\ \mathbf{F} \in \mathbf{P}'.\mathbf{x} \end{array} \right)\}$$

$$[\mathbf{P} ; \mathbf{P}'].\mathbf{x} \hat{=} \sum\{\mathbf{F}.\mathbf{x}' \times \mathbf{G}_{\mathbf{x}'} \mid \left(\begin{array}{l} \mathbf{x}' \in \mathbf{X} \\ \mathbf{F} \in \mathbf{P}.\mathbf{x} \\ \mathbf{G}_{\mathbf{x}'} \in \mathbf{P}'.\mathbf{x}' \end{array} \right)\}$$

$$[\mu \mathbf{Z} \bullet \mathcal{C}] \hat{=} \text{least fixed point of} \\ \text{cntx: } \mathcal{H} \rightarrow \mathcal{H}, \text{ defined} \\ \text{wp}.\mathcal{C} = \text{cntx}.\text{(wp}.\mathbf{Z}\text{)}.$$

Healthiness

The set of distributions of a probabilistic program satisfy three conditions: they are convex-closed, Cauchy-closed and up-closed.

1. A set \mathcal{F} of distributions is **convex** means for every $\mathbf{E}, \mathbf{F} \in \mathcal{F}$ and $\mathbf{p} \in [0, 1]$, $\mathbf{p}\mathbf{E} + (1 - \mathbf{p})\mathbf{F} \in \mathcal{F}$ also.
2. A set \mathcal{F} of distributions is **Cauchy-closed** means it is a closed subset of $\mathbf{R}^{\#\mathbf{X}}$ in the usual Euclidean sense.

3. A set \mathcal{F} of distributions is **up-closed** means it is closed under refinement of its elements: for all $\mathbf{E}, \mathbf{F} \in \tilde{\mathbf{X}}$,

$$\left(\begin{array}{l} \mathbf{E} \in \mathcal{F} \\ \mathbf{E} \sqsubseteq \mathbf{F} \end{array} \right) \Rightarrow \mathbf{F} \in \mathcal{F}.$$

Up-closure means that program refinement corresponds to reverse inclusion of resulting-distribution sets. Convex-closure means that a nondeterministic choice between two programs is refined by any probabilistic choice. Cauchy-closure makes the space of programs into a complete partial order.

Predicate-transformer semantics

We now consider an alternative semantics, based on predicate transformers, for language \mathcal{L}_Θ . Again, we require a little notation first.

Probabilistic predicates

A *probabilistic predicate* is a non-negative-real-valued predicate on state space; the set of all such is defined $\mathcal{Q} \hat{=} \mathbf{X} \rightarrow \mathbf{R}^{\geq}$, where \mathbf{R}^{\geq} denotes the set of non-negative reals.

Standard predicates are embedded in \mathcal{Q} by identifying *true* with 1 and *false* with 0. For standard predicate \mathbf{q} we write $[\mathbf{q}]$ for its embedding.

The logical operations, and their arithmetic equivalents, on probabilistic predicates are summarised as follows.

logic	arithmetic
$[\neg \mathbf{q}]$	$1 - [\mathbf{q}]$
$[\mathbf{q} \vee \mathbf{q}']$	$[\mathbf{q}] \sqcup [\mathbf{q}']$, or $[\mathbf{q}] + [\mathbf{q}']$ if \mathbf{q}, \mathbf{q}' disjoint
$[\mathbf{q} \wedge \mathbf{q}']$	$[\mathbf{q}] \sqcap [\mathbf{q}']$, or $[\mathbf{q}] \times [\mathbf{q}']$
$\mathbf{q} \Rightarrow \mathbf{q}'$	$[\mathbf{q}] \leq [\mathbf{q}']$ on \mathbf{X}
$\mathbf{q} \equiv \mathbf{q}'$	$[\mathbf{q}] = [\mathbf{q}']$ on \mathbf{X}
$\mathbf{q} \Leftarrow \mathbf{q}'$	$[\mathbf{q}] \geq [\mathbf{q}']$ on \mathbf{X}

The *entailment* relation \Rightarrow between standard predicates becomes pointwise \leq over their probabilistic counterparts.

For distribution $\mathbf{F} : \tilde{\mathbf{X}}$ and probabilistic predicate $\mathbf{q} : \mathcal{Q}$, the integral over \mathbf{X} of \mathbf{q} against \mathbf{F} with the discrete, or counting, measure is denoted

$$\int_{\mathbf{F}} \mathbf{q} \hat{=} \sum_{\mathbf{x} \in \mathbf{X}} \mathbf{q}.\mathbf{x} \times \mathbf{F}.\mathbf{x}.$$

Semantics

Let $(\mathcal{J}, \sqsubseteq)$ denote the space of probabilistic-predicate transformers ordered

$$\mathbf{j} \sqsubseteq \mathbf{k} \hat{=} \forall \mathbf{q} : \mathcal{Q} \bullet \mathbf{j}.\mathbf{q} \Rightarrow \mathbf{k}.\mathbf{q}.$$

We can now interpret a program in \mathcal{L}_Θ to be a transformer from probabilistic postconditions to probabilistic precon-

ditions as follows.

$$\mathbf{wp}.\mathbf{abort}.q \hat{=} 0$$

$$\mathbf{wp}.\mathbf{Skip}.q \hat{=} q$$

$$\mathbf{wp}.(x := E).q \hat{=} q[x := E]$$

$$\mathbf{wp}.(P \text{ ; } P').q \hat{=} \mathbf{wp}.P.(\mathbf{wp}.P'.q)$$

$$\mathbf{wp}.(P \triangleleft b \triangleright P').q \hat{=} [b] \times \mathbf{wp}.P.q + [\neg b] \times \mathbf{wp}.P'.q$$

$$\mathbf{wp}.(P \sqcap P').q \hat{=} \mathbf{wp}.P.q \sqcap \mathbf{wp}.P'.q$$

$$\mathbf{wp}.(P \text{ p } \oplus P').q \hat{=} p \times \mathbf{wp}.P.q + (1 - p) \times \mathbf{wp}.P'.q$$

$$\mathbf{wp}.(μ Z \bullet C) \hat{=} \text{least fixed point of } \mathbf{cntx}: \mathcal{J} \rightarrow \mathcal{J}, \text{ defined } \mathbf{wp}.C = \mathbf{cntx}.\mathbf{wp}.Z.$$

There, where appropriate, \sqcap means pointwise minimum between functions.

Galois connection

For program P and probabilistic postcondition q ,

$$\mathbf{wp}.P.q.x = \min\{\int_{\mathbf{F}} q \mid \mathbf{F} \in P.x\}.$$

We shall define two functions $\mathbf{wp} : \mathcal{H} \rightarrow \mathcal{J}$ and $\mathbf{rp} : \mathcal{J} \rightarrow \mathcal{H}$ between the lattices (\mathcal{H}, \supseteq) and $(\mathcal{J}, \sqsubseteq)$.

The embedding \mathbf{wp}

For probabilistic relation $h : \mathcal{H}$, probabilistic predicate $q : \mathcal{Q}$, and state $x : \mathbf{X}$,

$$\mathbf{wp}.h.q.x \hat{=} \sqcap\{\int_{\mathbf{F}} q \mid \mathbf{F} \in h.x\}.$$

The projection \mathbf{rp}

For probabilistic predicate transformer $j : \mathcal{J}$ and state $x : \mathbf{X}$,

$$\mathbf{rp}.j.x \hat{=} \{\mathbf{F} : \tilde{\mathbf{X}} \mid \forall q : \mathcal{Q} \bullet j.q.x \leq \int_{\mathbf{F}} q\}.$$

Note that \mathbf{rp} is not always defined: recall that relations in \mathcal{H} are non-empty. However the application of \mathbf{rp} to a \mathbf{wp} image of a relation is always defined.

Partial weak inverse

When \mathbf{rp} is defined then \mathbf{wp} is its weak inverse. For $h : \mathcal{H}$ and $j : \mathcal{J}$,

$$\mathbf{rp}.\mathbf{wp}.h = h \\ j \sqsubseteq \mathbf{wp}.\mathbf{rp}.j.$$

Healthiness conditions

The \mathbf{wp} images of relations within \mathcal{J} no longer satisfy distribution of \sqcap . Instead they are characterised by continuity

and sub-linearity.

Transformer $j : \mathcal{J}$ is *sub-linear* if for all $q, q' : \mathcal{Q}$ and all $a, b, c : \mathbf{R}^{\geq}$

$$aj.q + bj.q' - c \Rightarrow j.(aq + bq' - c).$$

Sub-linearity implies monotonicity, feasibility and scaling. In fact when the state space is finite it implies continuity as well.

Operator $\&$ is defined on probabilistic predicates:

$$(A \& B).s \hat{=} (A.s + B.s - 1) \sqcup 0.$$

Sub-linearity implies $\&$ -distribution and when specialised to standard programs $\&$ -distribution is equivalent to positive conjunctivity.

Loop rules

One of the strengths of the standard predicate-transformer calculus is its facility for proof—in particular variant/invariant-based rules for correctness of iterations. This section deals with the probabilistic counterparts to the standard loop rules where, here too, the predicate-transformer calculus pays off.

We shall use laws to prove termination and correctness about loops of the form

$$\mathbf{loop} \hat{=} \mathbf{do} G \rightarrow \mathbf{body} \mathbf{od}$$

where the loop guard G is a standard (non-probabilistic) predicate and \mathbf{body} is a relational (thus sub-linear) predicate transformer. A program, h , terminates from an initial state x if $\mathbf{wp}.h.[\mathbf{true}].x = 1$.

Let $\mathbf{T} = \mathbf{wp}.\mathbf{loop}.[\mathbf{true}]$ be the termination condition for \mathbf{loop} . We will assume (although it is not necessary to do so) that \mathbf{body} always terminates. (The rules are actually derived from a probabilistic version of \mathbf{wlp} which agrees with \mathbf{wp} if termination is assumed.)

Standard rules for reasoning about probabilistic loops separate correctness from termination. An invariant-based argument is usually sufficient to establish partial correctness, whereas a variant argument is used to show termination.

A predicate, I , is an invariant for the loop if

$$(G \sqcap I) \Rightarrow \mathbf{wp}.\mathbf{body}.I.$$

Combining partial and total correctness

Let I be an invariant for \mathbf{loop} . Then

$$\mathbf{Rule 1}: I \& \mathbf{T} \Rightarrow \mathbf{wp}.\mathbf{loop}.(I \sqcap \tilde{G}).$$

Notice that reduces to the standard rule when specialised to standard programs.

In many cases $\&$ gives too pessimistic a bound. We can do better, though. If $I \Rightarrow \mathbf{T}$ then

$$\mathbf{Rule 2}: I \Rightarrow \mathbf{wp}.\mathbf{loop}.(I \sqcap \tilde{G}).$$

Variant-based arguments*The 0-1 law*

The probabilistic variant arguments (over a finite state space) we use are based on a so-called 0-1 law.

Informally, the idea is as follows. Let Y be a fixed subset of state space X and suppose that from every state in Y there is a non-zero probability of the loop escaping from Y . Then in fact we can deduce that the probability of eventual escape from Y is 1.

The formal rule is

$$\mathbf{0-1\ Law} : (\exists p > 0 \bullet pI \Rightarrow T) \Rightarrow I \Rightarrow T.$$

From that we can derive Rule 3.

Rule 3 : Let v be an integer-valued expression on X defined at least over some subset I of the state space X . Suppose further that for iteration **loop**

1. there are fixed integer constants L (low) and H (high) such that

$$G \cap I \Rightarrow [L \leq V < H],$$

2. the subset I , as a (standard) predicate, is a **wp**-invariant for **loop**
3. for some fixed probability $p \neq 0$ and for all $M : \mathbf{N}$, $p(G \cap I \cap [V = M]) \Rightarrow \mathbf{wp.body}.[V < M]$.

Then termination is certain from any state in which I holds: we have $I \Rightarrow T$.

Finitary completeness of variants

Rule 3 is sound and complete for termination over a finite state space.

Self-stabilisation

We now give an example of a probabilistic algorithm. It is based on an algorithm due to Ted Hermann. The problem is how to achieve, efficiently and symmetrically, a leader from a ring of processes.

Consider M identical processors connected clockwise in a ring. A single processor—a leader—is chosen from them as follows.

Initially each processor is given a single token: the leader will be the first processor to obtain all M tokens. Fix some probability $p : [0, 1]$. On each step all processors synchronously perform the following actions:

1. Make a local probabilistic decision either to **pass** (with probability p) or to **keep** (with probability $1 - p$) all its tokens.
2. If **pass**, then send **all** its tokens to its clockwise successor; if **keep**, skip.
3. Receive any tokens passed from its clockwise predecessor.

Termination of this algorithm follows directly from Rule 3 taking as variant the largest arc of the ring containing all tokens.

**The South African
Computer Journal**

*An official publication of the Computer Society
of South Africa and the South African Institute of
Computer Scientists*

**Die Suid-Afrikaanse
Rekenaartydskrif**

*'n Amptelike publikasie van die Rekenaarvereniging
van Suid-Afrika en die Suid-Afrikaanse Instituut
vir Rekenaarwetenskaplikes*

Editor

Professor Derrick G Kourie
Department of Computer Science
University of Pretoria
Hatfield 0083
dkourie@dos-lan.cs.up.ac.za

Subeditor: Information Systems

Prof Lucas Introna
Department of Informatics
University of Pretoria
Hatfield 0083
lintrona@econ.up.ac.za

Production Editor

Dr Riël Smit
Mosaic Software (Pty) Ltd
P.O.Box 23906
Claremont 7735
gds@mosaic.co.za

World-Wide Web: <http://www.mosaic.co.za/sacj/>

Editorial Board

Professor Judy M Bishop
University of Pretoria, South Africa
jbishop@cs.up.ac.za

Professor R Nigel Horspool
University of Victoria, Canada
nigelh@csr.csc.uvic.ca

Professor Richard J Boland
Case Western Reserve University, USA
boland@spider.cwrw.edu

Professor Fred H Lochovsky
University of Science and Technology, Hong Kong
fred@cs.ust.hk

Professor Ian Cloete
University of Stellenbosch, South Africa
ian@cs.sun.ac.za

Professor Kalle Lyytinen
University of Jyvaskyla, Finland
kalle@cs.jyu.fi

Professor Trevor D Crossman
University of Natal, South Africa
crossman@bis.und.ac.za

Doctor Jonathan Miller
University of Cape Town, South Africa
jmiller@gsb2.uct.ac.za

Professor Donald D Cowan
University of Waterloo, Canada
dcowan@csg.uwaterloo.ca

Professor Mary L Soffa
University of Pittsburgh, USA
soffa@cs.pitt.edu

Professor Jürg Gutknecht
ETH, Zürich, Switzerland
gutknecht@inf.ethz.ch

Professor Basie H von Solms
Rand Afrikaanse Universiteit, South Africa
basie@rkw.rau.ac.za

Subscriptions

	Annual	Single copy
Southern Africa:	R50,00	R25,00
Elsewhere:	\$30,00	\$15,00

An additional \$15 per year is charged for airmail outside Southern Africa

to be sent to:

*Computer Society of South Africa
Box 1714 Halfway House 1685*

Contents

INTRODUCTION

WOFACS '96: Workshop on Formal and Applied Computer Science C Brink	1
--	---

PROCEEDINGS

Reasoning about Changing Information P Blackburn, J Jaspars and M de Rijke	2
Verification of Finite State Systems with Temporal Logic Model Checking B-H Schlingloff	27
Test Automation of Safety-Critical Reactive Systems J Peleska and M Siegel	53
Application-Oriented Program Semantics AK McIver, C Morgan and JW Sanders	78
