The United Nations
University

UNU-IIST

**International Institute for
Software Technology**

Reasoning about a distributed probabilistic system

Ukachukwu Ndukwu and J. W. Sanders

August 2008

$\mathcal{R}$

UNU-IIST and UNU-IIST Reports

UNU-IIST (United Nations University International Institute for Software Technology) is a Research and Training Centre of the United Nations University (UNU). It is based in Macao, and was founded in 1991. It started operations in July 1992. UNU-IIST is jointly funded by the government of Macao and the governments of the People's Republic of China and Portugal through a contribution to the UNU Endowment Fund. As well as providing two-thirds of the endowment fund, the Macao authorities also supply UNU-IIST with its office premises and furniture and subsidise fellow accommodation.

The mission of UNU-IIST is to assist developing countries in the application and development of software technology.

UNU-IIST contributes through its programmatic activities:

1. **Advanced development projects**, in which software techniques supported by tools are applied,

2. **Research projects**, in which new techniques for software development are investigated,

3. **Curriculum development projects**, in which courses of software technology for universities in developing countries are developed,

4. **University development projects**, which complement the curriculum development projects by aiming to strengthen all aspects of computer science teaching in universities in developing countries,

5. **Schools and Courses**, which typically teach advanced software development techniques,

6. **Events**, in which conferences and workshops are organised or supported by UNU-IIST, and

7. **Dissemination**, in which UNU-IIST regularly distributes to developing countries information on international progress of software technology.

Fellows, who are young scientists and engineers from developing countries, are invited to actively participate in all these projects. By doing the projects they are trained.

At present, the technical focus of UNU-IIST is on **formal methods** for software development. UNU-IIST is an internationally recognised center in the area of formal methods. However, no software technique is universally applicable. We are prepared to choose complementary techniques for our projects, if necessary.

UNU-IIST produces a report series. Reports are either Research $\boxed{\text{R}}$, Technical $\boxed{\text{T}}$, Compendia $\boxed{\text{C}}$ or Administrative $\boxed{\text{A}}$. They are records of UNU-IIST activities and research and development achievements. Many of the reports are also published in conference proceedings and journals.

Please write to UNU-IIST at P.O. Box 3058, Macao or visit UNU-IIST's home page: http://www.iist.unu.edu, if you would like to know more about UNU-IIST and its report series.

G. M. Reed, Director

The United Nations University

UNU-IIST

International Institute for Software Technology

P.O. Box 3058
Macao

Reasoning about a distributed probabilistic system

Ukachukwu Ndukwu and J. W. Sanders

Abstract

Reasoning about a distributed system that exhibits a combination of probabilistic and temporal behaviour does not seem to be easy with current techniques. The reason is the interaction between probability and abstraction, made worse by remote synchronisation. In this paper the recently proposed language PTSC (for probability, time and shared-variable concurrency) is extended by constructs for interleaving and local block. Both enhance a designer's ability to modularise a design; the latter also permits a design to be compared with its more abstract specification, by concealing appropriately chosen design variables. Laws of the extended language are studied and applied in a case study consisting of a faulty register-transfer-level design.

Ukachukwu Ndukwu obtained a *B.Tech.* in Mathematics and Computer Science (First Class Honours) at the Federal University of Technology Owerri, Nigeria in 2004. From June 2007 to February 2008 he was a UNU-IIST fellow in Macao under the supervision of Jeff Sanders. Currently he is pursuing a *Ph.D.* in the Computer Science department of Macquarie University, Sydney, NSW, Australia under the supervision of Prof. Annabelle McIver. His interests span probabilistic, concurrent and real-time systems and his thesis concerns formal models for safety-critical systems.

Jeff Sanders is Principal Research Fellow at UNU-IIST with interests largely in Formal Methods.

Contents

1   Introduction

Probabilistic programs afford one means of both efficiency and elegance. In a situation where error can be tolerated (and in many probabilistic programs the error can be decreased at the expense of longer execution time — like further iterations of a loop) then there may exist a probabilistic program that is far more efficient than the best non-probabilistic programs. Primality testing was the original example, but there are now very many more [20]. For elegance: in some situations there may be a symmetrical efficient probabilistic program where no non-probabilistic symmetrical program exists. Arbitration is an example. But perhaps most important of all: probabilistic programs provide a way of modelling, reasoning about and simulating, faulty systems.

There is a highly satisfactory theory of sequential probabilistic programs in the Dijkstra style, due largely to McIver and Morgan [13] but built on work of Clare Jones [10] and Kozen [12]. However for distributed systems it seems that little satisfactory progress has been made. The interaction between (demonic) nondeterminism and probabilism seems more intricate there because of the scope for demonic behaviour.

Current approaches to probabilistic distributed systems reflect the twin approaches to the theory of general distributed systems theory: namely process algebra and shared-variable concurrency. The former theory has been shaped by CSP [7, 25], CCS [15] and ACP [2]. The latter forms the basis for the memory models of programming languages (for example the Java memory model [14]) and for reasoning about hardware (for example Verilog [1, 6]).

Both are well developed but in slightly different directions. Process algebra offers hierarchies of semantic models and complete families of sound laws. However the incorporation of probability there exhibits unsatisfactory features [18]; the incorporation of time is intricate (see the survey by Ouaknine and Schneider [21]); and the two together is problematic. For shared-variable concurrency, on the other hand, much effort has been spent on formalisms for reasoning (for a thorough summary see the text by de Roever *et al.* [24]) and less on semantic models. Indeed, since each program has access to the variables of all the other programs, there is scope for mayhem (as unprincipled as unstructured low-level programming). Important methods have been proposed by Owicki and Gries (non-interference [22]), Cliff Jones (rely and guarantee [11]), Misra and Chandy (assumption and commitment [16]), Brookes [3] and de Roever *et al.* (temporal orderings [24]). The variety of methods attests to the subtlety of the problem.

In this paper we consider shared-variable concurrency, and build on the recent model of Zhu, Qin, He and Bowen [26]. They propose a language, PTSC, containing probability, time and shared-variable concurrency. Since concurrency is achieved by the interleaving of atomic events, rather than by synchronisation or message passing as in process algebra, it is important to decide in any such model how various threads interact: at which points a thread can be interspersed by another. In [26] that is achieved by a structured operational semantics [23] with four kinds of transition (corresponding to the progress of an atomic action, time, resolution of nondeterminism, and the triggering of a guarded action). Bisimulation is then used to establish

laws. However there the authors did not consider any probability-free operator for combining programs in parallel, of the kind one would expect to use in modularising a system using parallel components. Nor did they consider local block, of use for controlling the scope of variables in interleaving programs and for enabling a low-level program to be compared, after concealing variables appropriately, with a higher-level specification.

In this paper we add interleaving and local block to PTSC. We choose to present laws without providing an operational semantics (though in fact some of the laws have been confirmed by bisimulation from a structured operational semantics). That is no deficiency because laws also document the manner in which one program can intersperse another. In other words, the computational intuition so often used to justify an operational semantics is equally evident in the laws. And an operational model anyway provides no guarantee of consistency of the laws it posits, in the way that a denotational model does.

We attempt to avoid the difficult problem of reasoning about shared-variable concurrency by using the laws to reason algebraically. A register-transfer-level program is offered as an example, first to test the expressive power of our extension to PTSC and secondly to explore how laws might be used to establish that an implementation with probabilistic behaviour conforms to its probabilistic specification.

In the next section, Section 2, we give an outline of PTSC. We use slightly different syntax from the original [26] in order to conform to what might be called the Dutch style, that is beneficial for calculation; for example the expression $(+ \, n \in \mathbb{N} : even(n) : x_n)$ represents the summation $\sum_{n \in 2\mathbb{N}} x_n$. Section 3 introduces interleaving, Section 4 introduces local block and Section 5 presents the case study. In concluding we propose further work.

2   Language PTSC

The language PTSC is based on Dikstra's guarded-command language [5], as reflected in the first six constructs of Figure 1. Command **abort** is of course present as **do** $true \; \rightarrow \;$ **skip od**, but we follow the authors of [26] and leave it implicit since our concerns here do not address it. On the other hand **skip**, which could equally well have been left implicit as $x := x$, is required here.

A *guarded assignment*, $(@i \in I : b_i : x_i := e_i \; \mathring{,} \; P_i)$, assumes an index set I, predicates b_i on state space (that need be neither exhaustive nor pairwise disjoint) and a program with a leading assignment of expression e_i to variable x_i and body P_i (which may of course be **skip**). If a guard holds the subsequent program may be scheduled for execution (so nondeterminism arises from overlapping guards); but otherwise the guarded assignment lets time advance. Execution cannot be interspersed between evaluation of the guard and execution of the leading assignment, but of course can be after execution of that assignment. We allow several guarded assignments to be combined with the @ combinator. For example this program waits until either the value of

inaction	**skip**
assignment	$x := e$
conditional	$P \lhd b \rhd Q$
iteration	**do** $b \rightarrow P$ **od**
sequential composition	$P \mathbin{\overset{\circ}{,}} Q$
demonic nondeterminism	$P \sqcap Q$
guarded assignment	$(@i \in I : b_i : x_i := e_i \mathbin{\overset{\circ}{,}} P_i)$
delay	$\#n$
binary probabilistic choice	$P \,_r\oplus Q$
probabilistic parallel	$P \,_r\| Q$
interleaving	$P \,\|\| Q$
local block	$_x[\![\, x\!:\!X \circ P \,]\!]_x$

Figure 1: The language PTSC augmented by constructs for interleaving and local block (with mild notational variations on the syntax of [26]). We also exploit recursion and mutual recursion when convenient in defining programs.

a is 2 or that of b is 1 and then performs the corresponding assignment.

$$(@ : a = 2 : x := 0) \; @ \; (@ : b = 1 : y := 3)$$

A *delay*, $\#n$, (for $n : \mathbb{N}$) lets time advance n steps before it terminates. By using sequential composition a *delay-guarded program* is defined

$$\#n \, P \;\; := \;\; \#n \mathbin{\overset{\circ}{,}} P$$

Evidently it can also be constructed from $\#1$ and P.

A *binary probabilistic choice*, $P \,_r\oplus Q$, selects program P with probability r, and program Q with probability $1 - r$. The expression r is assumed to be a function of state with values in the real unit interval $[0, 1]$. For convenience, we write $1 - r$ as $\overline{r}$. The program cannot be interrupted between evaluation of r and execution of the first atomic event of the selected component.

A *probabilistic parallel composition*, $P \,_r\| Q$, interleaves (any) initial atomic actions of P and Q with probabilities r and $\overline{r}$ respectively. If the initial atomic events are offered as a probabilistic choice, the resulting probabilities of their being offered for scheduling represent their conditional probabilities.

To deal with multi-way probabilistic choices we adopt the notation $(\oplus i \in I : p_i : P_i)$ in which the probabilities sum to 1: $(+i \in I :: p_i) = 1$. The binary probabilistic choice $P_0 \,_r\oplus P_1$ may be expressed in that notation with index set $I := \{0, 1\}$ and probabilities $p_0 := r$, $p_1 := \overline{r}$.

The language PTSC is enriched with a *guarded choice* as follows. Components with leading assignments may be expressed as guarded assignments and combined with (multi-way) probabilistic

choice. In the program

$$[\![]\!]_{i \in I}\{[p_i] \ choice_{j \in J_i}(b_{ij} \& (x_{ij} := e_{ij}) \ P_{ij})\}$$

from [26] it is assumed that for each $i : I$ the predicates b_{ij} partition unity (*i.e.* are exhaustive $(\forall i : I :: \exists j : J_i \cdot b_{ij})$ and pairwise disjoint $(\forall i : I \cdot \forall j, j' : J_i : j \neq j' : \neg(b_{ij} \wedge b_{ij'}))$ and the probabilities sum to 1. Thus in any state, for each $i \in I$, exactly one of the b_{ij} holds (for $j \in I_i$); between the $i \in I$, program $x_{ij} := e_{ij} \ \mathbin{\mathring{,}} \ P_{ij}$ is chosen for (potential) scheduling with probability p_i. In our notation that guarded choice is expressed as a combination of probabilism and guarded assignment

$$(\oplus i \in I : p_i : (@j \in J_i : b_{ij} : x_{ij} := e_{ij} \ \mathbin{\mathring{,}} \ P_{ij})) \,.$$

No action can intersperse between evaluation of the probability, the guard and completion of the leading assignment.

Guarded choices are not combined indiscriminately. For example since the guards in a guarded choice are exhaustive, the passage of time (initially) plays no part in it, and there is no point in combining it with a delay. However if a guarded choice is combined with a guarded assignment (the disjunction of whose guards, it will be recalled, need not be exhaustive), the latter has priority; so only outside the disjunction of the guards of the guarded assignment does the guarded choice apply. Similarly a choice between a delay and a guarded assignment attaches priority to the guarded assignment.

In the presence of probabilistic behaviour, programs are regarded as the same if their behaviours are the same to within unit probability (rather than 'absolutely'). For example a random walk on the integers with equal probability of incrementing and decrementing its value n has probability 1 of returning to the origin, although there is an execution that increments forever. So a loop with guard $n \neq 0$ and body that increments and decrements n with equal probability, is for us the same as the assignment $n := 0$.

As a result, a conditional is a special case of a probabilistic choice. Indeed for predicate b on state space,

$$P \triangleleft b \triangleright Q \quad = \quad P \ {}_{[b]}\!\oplus Q$$

where $[b]$ is the function on state space with value 1 if b and 0 if $\neg b$. So laws for conditional may be inferred from those for probabilism.

In a shared-variable program, each component program can in principle read and write-to all variables. However in practice that is seldom necessary. Practical designs exploit severe modularisation in which component programs access variables selectively. We write $vars(P)$ for the set of variables that program P accesses (either reads or writes-to).

For example here is a PTSC program C (similar to one that will arise in Section 5). $vars(C) = \{a, b, c, d, e\}$ where all variables are of type $\mathbb{N}$ except e which is Boolean. After updating variable

a, the program enters a loop in which it waits for variables b and c to be 0. When they are, it reads the value of variable Boolean e; if high, it updates b and c with equal likelihood and iterates; if low, it updates d and returns to the start of C.

$$
\begin{aligned}
C \ := \ & a := 2 \ \fatsemi \\
& \mathbf{do} \ \mathit{true} \ \rightarrow \\
& \quad (@ : (b = 0 \wedge c = 0) : ((b := 2 \ _{\frac{1}{2}} \oplus c := 1) \lhd e \rhd (d := 1 \ \fatsemi \ C))) \\
& \mathbf{od}
\end{aligned}
$$

The last two lines of Figure 1 introduce the constructs added here to PTSC. They are introduced as follows.

3   Interleaving

Programs that communicate *via* shared variables execute in parallel, interleaving their atomic actions. The laws for interleaving, $|||$, make precise the points in a program at which another may intersperse its own atomic actions. That idea goes back, as with so much, to a paper of Dikjstra, in this case [4] in which he exploited concurrency with **parbegin**, **parend** blocks.

If $x := e$ and $y := f$ are atomic assignments to distinct variables then their concurrent execution must achieve one before the other, yet the order is not determined

$$
x := e \ ||| \ y := f \ = \ (x := e \ \fatsemi \ y := f) \sqcap (x := e \ \fatsemi \ y := f). \tag{1}
$$

Evidently the result depends, in general, on the order (since an updated value of one variable may be used in the other expression). This situation is thus usually to be avoided (even if $x = y$). By comparison the *multiple assignment*,

$$
x, y := e, f \,,
$$

evaluates the two assignments atomically, concurrently, and without interference between the expressions (by taking a copy of each of x and y, for use in evaluating the expressions, just before the assignments begin it can be ensured that the initial and final values are related thus: $x = e[x_0, y_0/x, y] \wedge y = f[x_0, y_0/x, y]$).

The appropriate law for interleaving is

$$
\begin{aligned}
(x := e \ \fatsemi \ P) \ ||| \ (y := f \ \fatsemi \ Q) \ = \ & x := e \ \fatsemi \ (P \ ||| \ (y := f \ \fatsemi \ Q)) \\
& \sqcap \\
& y := f \ \fatsemi \ ((x := e \ \fatsemi \ P) \ ||| \ Q).
\end{aligned}
\tag{2}
$$

The interleaving of a program P with inaction consists just of the actions of P.

$$\textbf{skip} \;|\!|\!|\; P \;\; = \;\; P \tag{3}$$

Interleaving attributes no importance to the order of its component programs (it is commutative) nor to the order in which more than two are combined (it is associative).

$$P \;|\!|\!|\; Q \;\; = \;\; Q \;|\!|\!|\; P \tag{4}$$
$$P \;|\!|\!|\; (Q \;|\!|\!|\; R) \;\; = \;\; (P \;|\!|\!|\; Q) \;|\!|\!|\; R \tag{5}$$

The interleaving of a program P with a nondeterministic choice between programs Q and R consists of the interleaving of P with either Q or R, the choice being again nondeterministic.

$$P \;|\!|\!|\; (Q \sqcap R) \;\; = \;\; (P \;|\!|\!|\; Q) \sqcap (P \;|\!|\!|\; R) \tag{6}$$

Similar reasoning leads to laws for probabilistic choice and probabilistic parallel, provided that execution of R does not change the probabilistic expression r which is evaluated first on the right-hand side.

$$(P \;{}_r\!\oplus Q) \;|\!|\!|\; R \;\; = \;\; (P \;|\!|\!|\; R) \;{}_r\!\oplus (Q \;|\!|\!|\; R) \tag{7}$$
$$(P \;{}_r\!\| Q) \;|\!|\!|\; R \;\; = \;\; (P \;|\!|\!|\; R) \;{}_r\!\| (Q \;|\!|\!|\; R) \tag{8}$$

The laws for guarded assignment ensure that no action can be interspersed between the test of a guard which is true and the assignment it guards.

$$
\begin{aligned}
&(@i \in I : b_i : x_i := e_i \mathbin{\mathring{,}} P_i) \;|\!|\!|\; (@j \in J : c_j : y_j := f_j \mathbin{\mathring{,}} Q_j) \\
&= \\
&(@i \in I, j \in J : b_i \wedge c_j : x_i := e_i \mathbin{\mathring{,}} (P_i \;|\!|\!|\; (y_j := f_j \mathbin{\mathring{,}} Q_j)) \\
&\qquad\qquad\qquad\qquad \sqcap \\
&\qquad\qquad\qquad y_j := f_j \mathbin{\mathring{,}} ((x_i := e_i \mathbin{\mathring{,}} P_i) \;|\!|\!|\; Q_j)) \\
&@ \\
&(@i \in I, j \in J : b_i \wedge \neg c_j : x_i := e_i \mathbin{\mathring{,}} (P_i \;|\!|\!|\; (y_j := f_j \mathbin{\mathring{,}} Q_j))) \\
&@ \\
&(@i \in I, j \in J : \neg b_i \wedge c_j : y_j := f_j \mathbin{\mathring{,}} ((x_i := e_i \mathbin{\mathring{,}} P_i) \;|\!|\!|\; Q_j)))
\end{aligned}
\tag{9}
$$

Priority is given to a guarded assignment: an unguarded assignment in parallel has the chance to be scheduled only if no guard holds.

$$(@i \in I : b_i : x_i := e_i \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} P_i) \mathbin{|\!|\!|} (y := f \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} Q) \tag{10}$$
$$=$$
$$((@i \in I : b_i : x_i := e_i \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} (P_i \mathbin{|\!|\!|} (y := f \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} Q))$$
$$\mathbin{\lhd} \exists i : I \cdot b_i \mathbin{\rhd}$$
$$y := f \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} (@i \in I : b_i : (x_i := e_i \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} P_i) \mathbin{|\!|\!|} Q))$$

The passage of time is not interleaved, but experienced equally by both component programs.

$$(\#1\, P) \mathbin{|\!|\!|} (\#1\, Q) \quad = \quad \#1\,(P \mathbin{|\!|\!|} Q) \tag{11}$$

The interleaving of a delay with a guarded assignment gives priority to the guarded assignment and only outside the disjunction of its guards is time allowed to advance.

$$(\#1\, P) \mathbin{|\!|\!|} (y := f \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} Q) \quad = \quad (y := f) \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} (\#1\, P \mathbin{|\!|\!|} Q) \tag{12}$$
$$(\#1\, P) \mathbin{|\!|\!|} (@i \in I : b_i : x_i := e_i \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} P_i) \quad = \quad ((@i \in I : b_i : x_i := e_i \mathbin{\raise0.3ex\hbox{$\mathchar"203B$}} (\#1\, P \mathbin{|\!|\!|} P_i)) \tag{13}$$
$$\mathbin{\lhd} \exists i : I \cdot b_i \mathbin{\rhd}$$
$$x_i := e_i)$$

4 Local block

Local blocks are used to delimit the scope of variables. Our notation for a block in which the variable x, of type X, is local in program Q is

$$_x[\![\, x \,{:}\, X \circ Q \,]\!]_x \,.$$

We treat a local block as a whole, matching the beginning of the block with its end. An exposition in the setting of wp and the refinement calculus is Morgan's textbook [17]. For an alternative treatment, separating the two halves of a block, see Hoare and He's UTP text [8] (noting that the separation is independent of their UTP treatment of it).

Many of the laws are standard, even if they are not well advertised. For example, a localised variable is invisible outside its block.

$$_x[\![\, x \,{:}\, X \circ x := e \,]\!]_x \quad = \quad \mathbf{skip} \tag{14}$$

As a result, if x does not appear in program P then since it plays no part in P, even if it is initialised, its localisation leaves P unchanged.

$$_x[\![\, x\!:\!X \circ P \,]\!]_x \;\; = \;\; P \tag{15}$$

$$_x[\![\, x\!:\!X \circ x := e \mathbin{\overset{\circ}{,}} P \,]\!]_x \;\; = \;\; P \tag{16}$$

Localising a nondeterministic choice between two programs gives a localised version of one or the other, the choice being again nondeterministic.

$$_x[\![\, x\!:\!X \circ P \sqcap Q \,]\!]_x \;\; = \;\; {}_x[\![\, x\!:\!X \circ P \,]\!]_x \;\sqcap\; {}_x[\![\, x\!:\!X \circ Q \,]\!]_x \tag{17}$$

We assume that law also for arbitrary (nonvoid) nondeterminism.

Our treatment assumes that a block begins with a type declaration for the new local variable, but not necessarily with its initialisation. An uninitialised local variable assumes a nondeterministically-chosen value of its type. Thus,

$$x\!:\!X \circ P \;\; = \;\; x\!:\!X \circ \sqcap_{v \in X}\ x := v \mathbin{\overset{\circ}{,}} P . \tag{18}$$

That involves finite nondeterminism iff the type X is finite. Of course typically the type declaration is followed immediately by initialisation in P, which 'supersedes' that nondeterministic choice. For example if x is a Boolean variable and P has a leading assignment initialising x to *false*, then the nondeterministic default is indeed superseded:

$$(\sqcap_{v:\mathbb{B}}\ x := v \mathbin{\overset{\circ}{,}} x := \textit{false}) \;\; = \;\; (x := \textit{false}) .$$

The companion law to (16) can now be stated. Assume $y \neq x$. If the local variable x is not free in expression e and y does not appear in program P then

$$_x[\![\, x\!:\!X \circ y := e \mathbin{\overset{\circ}{,}} P \,]\!]_x \;\; = \;\; y := e \mathbin{\overset{\circ}{,}} {}_x[\![\, x\!:\!X \circ P \,]\!]_x . \tag{19}$$

To deal with the case x free in e, if x is initialised (and $x \neq y$ and y does not appear in program P) then

$$_x[\![\, x\!:\!X \circ x := x_0 \mathbin{\overset{\circ}{,}} y := e \mathbin{\overset{\circ}{,}} P \,]\!]_x \;\; = \;\; y := e[x_0/x] \mathbin{\overset{\circ}{,}} {}_x[\![\, x\!:\!X \circ x := x_0 \mathbin{\overset{\circ}{,}} P \,]\!]_x . \tag{20}$$

If local variable x is not initialised then Law (18) applies. Assuming the extension of Law (17) to arbitrary nondeterminism (and $x \neq y$ and y not in program P) yields

$$_x[\![\, x\!:\!X \circ y := e \mathbin{\fatsemi} P \,]\!]_x \quad = \quad \sqcap_{v \in X} \; y := e[v/x] \mathbin{\fatsemi} {}_x[\![\, x\!:\!X \circ P \,]\!]_x. \tag{21}$$

Because a localised variable is invisible outside its block, localising it immediately a second time has no further effect.

$$_x[\![\, x\!:\!X \circ {}_x[\![\, x\!:\!X \circ P \,]\!]_x \,]\!]_x \quad = \quad {}_x[\![\, x\!:\!X \circ P \,]\!]_x \tag{22}$$

The localisation of (distinct) variables can be achieved in either order.

$$_x[\![\, x\!:\!X \circ {}_y[\![\, y\!:\!Y \circ P \,]\!]_y \,]\!]_x \quad = \quad {}_y[\![\, y\!:\!Y \circ {}_x[\![\, x\!:\!X \circ P \,]\!]_x \,]\!]_y \tag{23}$$

That means the definition of block can be extended from one local variable to any finite set of local variables

$$
\begin{aligned}
_{\{\}}[\![\, P \,]\!]_{\{\}} \quad &:= \quad P \\
_{x,y}[\![\, x\!:\!X \mathbin{\fatsemi} y\!:\!Y \circ P \,]\!]_{x,y} \quad &:= \quad {}_x[\![\, x\!:\!X \circ {}_y[\![\, y\!:\!Y \circ P \,]\!]_y \,]\!]_x.
\end{aligned}
$$

Indeed Laws (22) and (23) ensure the consistency of that definition: blocks do indeed behave as a 'set-like' function of their local variables. They demonstrate, for example, that the block with local variables $\{x, x, y\}$ is the same, just as it ought to be, as that with local variables $\{y, x\}$.

Recall (from Section 2) that since a conditional is a special case of a probabilistic choice, the laws for localising a conditional are subsumed in those for localising probabilism.

If x is not free in the probabilistic expression r then localising a probabilistic choice between P and Q results in a probabilistic choice between the localised versions of P and Q.

$$_x[\![\, x\!:\!X \circ P \mathbin{_r\oplus} Q \,]\!]_x \quad = \quad {}_x[\![\, x\!:\!X \circ P \,]\!]_x \mathbin{_r\oplus} {}_x[\![\, x\!:\!X \circ Q \,]\!]_x \tag{24}$$

If the probability r depends on an initialised local variable x then the probabilism is resolved by substitution (as in Law (20)).

$$_x[\![\, x\!:\!X \circ x := x_0 \mathbin{\fatsemi} P \mathbin{_r\oplus} Q \,]\!]_x \quad = \quad {}_x[\![\, x\!:\!X \circ x := x_0 \mathbin{\fatsemi} P \,]\!]_x \mathbin{_{r[x_0/x]}\oplus} {}_x[\![\, x\!:\!X \circ x := x_0 \mathbin{\fatsemi} P \,]\!]_x \tag{25}$$

But if the local variable is uninitialised then the default is a nondeterministic choice over all possible initialisations (Law (21))

$$_x[\![\, x\!:\!X \circ P \,_r\!\oplus Q \,]\!]_x \quad = \quad \sqcap_{v \in X} \,_x[\![\, x\!:\!X \circ P \,]\!]_x \,_{r[v/x]}\!\oplus \,_x[\![\, x\!:\!X \circ Q \,]\!]_x. \tag{26}$$

For example we can infer that a conditional with an unitialised Boolean test is nondeterministic:

$_x[\![\, x\!:\!\mathbb{B} \circ P \lhd x \rhd Q \,]\!]_x$
=   Law (18)

$_x[\![\, x\!:\!\mathbb{B} \circ (\sqcap_{v \in \mathbb{B}} \; x := v) \,\mathbin{\substack{\circ \\ 9}}\, (P \lhd x \rhd Q) \,]\!]_x$
=   standard

$_x[\![\, x\!:\!\mathbb{B} \circ \sqcap_{v \in \mathbb{B}} \; (x := v \,\mathbin{\substack{\circ \\ 9}}\, (P \lhd x \rhd Q)) \,]\!]_x$
=   standard

$_x[\![\, x\!:\!\mathbb{B} \circ \sqcap_{v \in \mathbb{B}} \; (P \lhd x[v/x] \rhd Q) \,]\!]_x$
=   calculus

$_x[\![\, x\!:\!\mathbb{B} \circ (P \lhd \mathit{true} \rhd Q) \sqcap (P \lhd \mathit{false} \rhd Q) \,]\!]_x$
=   definition of conditional

$_x[\![\, x\!:\!\mathbb{B} \circ P \sqcap Q \,]\!]_x$
=   Law (17)

$_x[\![\, x\!:\!\mathbb{B} \circ P \,]\!]_x \; \sqcap \; _x[\![\, x\!:\!\mathbb{B} \circ Q \,]\!]_x.$

Probabilistic parallel composition is similar to probabilistic choice. Its analogue of (24), for example, is (provided x is not free in r),

$$_x[\![\, x\!:\!X \circ P \,_r\!\| Q \,]\!]_x \quad = \quad _x[\![\, x\!:\!X \circ P \,]\!]_x \,_r\!\| \,_x[\![\, x\!:\!X \circ Q \,]\!]_x. \tag{27}$$

Again, guarded choice follows a similar pattern. The simple case, in which the guards do not depend on the local variable, is

$$_x[\![\, x\!:\!X \circ (@i \in I : b_i : P_i) \,]\!]_x \quad = \quad (@i \in I : b_i : \,_x[\![\, x\!:\!X \circ P_i \,]\!]_x). \tag{28}$$

Time is not localisable.

$$_x[\![\, x\!:\!X \circ \#n \, P \,]\!]_x \quad = \quad \#n \,_x[\![\, x\!:\!X \circ P \,]\!]_x \tag{29}$$

The localisation of an interleaving of P and Q interleaves the localisations of P and Q.

$$_x[\![\,x\!:\!X \circ P \,|\!|\!|\, Q\,]\!]_x \quad = \quad _x[\![\,x\!:\!X \circ P\,]\!]_x \,|\!|\!|\, _x[\![\,x\!:\!X \circ Q\,]\!]_x \tag{30}$$

5   Case study

In this section we consider a circuit, at the register-transfer level, that adds its two integer inputs inx and iny and outputs their sum to variable out. However the register containing inx is faulty. Our first task is to model the circuit accurately in PTSC and our second is to reason about its behaviour.

5.1   Informal description

We assume that the environment of the circuit supplies the two inputs when the circuit is ready for them, and receives the single output when the circuit is ready to provide it. The circuit is abstract in the sense that the registers that hold the two inputs are of type natural number (rather than being bit strings, and rather than assuming the numbers to be *a priori* bounded). One register performs natural-number increment and the other the decrement of a positive natural number. The circuit attempts to achieve the required addition by successive concurrent faulty-increments and decrements until the decrement register contains zero, at which point its initial value ought (in the absence of faults) to have been 'transferred' to the other register.

The specification thus has three variables, all integers. When the circuit behaves correctly, it achieves the sum:
$$\begin{aligned} Spec \;:=\;\; & inx, iny, out : \mathbb{N} \circ \\ & out := inx + iny\,. \end{aligned}$$

The implementation contains five local variables plus the two register variables. It consists of an initialisation of the local variables, followed by the parallel composition of the increment register X, the decrement register Y and the control unit C.

$$\begin{aligned} Imp \;:=\;\; & x, y : \mathbb{N} \,\fatsemi \\ & ie, op : 0 \mid 1 \mid 2 \,\fatsemi \\ & oe, f : 0 \mid 1 \,\fatsemi \\ & s : \mathbb{B} \circ \\ & ie, op, oe, f := 0, 0, 0, 1 \,\fatsemi \\ & X \,|\!|\!|\, Y \,|\!|\!|\, C \end{aligned}$$

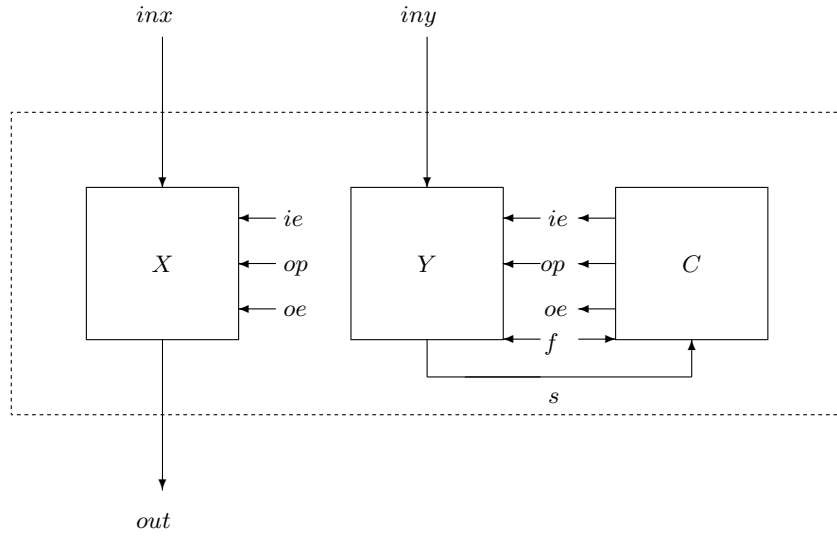The three component programs share variables as follows. Recalling (Section 2) the notation

Figure 2: Register-transfer circuit interpreted as a data-flow design. The enables ie, op and oe are shared variables essentially written-to by C and read by X and Y; but to avoid races they also act as semaphores and so are written also by the two registers. The status signal s, however, is written-to by Y and read by C. The semaphore f is written-to and read by both Y and C.

for the variables accessed by a program:

$$
\begin{aligned}
vars(X) &:= \{inx, x, ie, op, oe, out\} \\
vars(Y) &:= \{iny, y, ie, op, oe, s, f\} \\
vars(C) &:= \{ie, op, oe, s, f\} \,.
\end{aligned}
$$

Thus ie, op and oe are common to all three but f and s to just Y and C, and oe to just X and C. Variables xin and out are common to X and the environment and variable yin is common to Y and the environment. However x and y are unique to their programs. See Figure 2.

However the increment register X increments x with probability r but overlooks the increment with probability $\overline{r}$.

We must first express X, Y and C in PTSC. Then we must establish the behaviour of $_E[\![\,Imp\,]\!]_E$, where $E := \{x, y, ie, op, oe, s, f\}$, and relate it to $Spec$.

5.2 Definition in PTSC

The increment register X consists of a ternary guarded choice, determined by the (pairwise) disjoint conditions $ie = 2$, $op = 2$ and $oe = 1$. If the first guard holds, it inputs inx, decrements ie (to avoid repeated inputs) and iterates. If the second guard holds, it tries to increment its x register but does so successfully with probability r whilst with probability $\overline{r}$ it skips; then

it decrements *op* (again, to avoid unwanted—this time critical—repetition) and iterates. We assume that r is a positive constant, so the faulty register always has the same positive chance of behaving correctly. If the third guard holds it outputs the value of x and terminates.

$$
\begin{aligned}
X \;:=\; & (@ : ie = 2 : x := inx \; \mathring{,} \; ie := ie - 1 \; \mathring{,} \; X) \; @ \\
& (@ : op = 2 : (x := x + 1 \; {}_r\oplus \mathbf{skip}) \; \mathring{,} \; op := op - 1 \; \mathring{,} \; X) \; @ \\
& (@ : oe = 1 : out := x)
\end{aligned}
$$

The decrement register Y consists of the same guarded choice. In the case $ie = 2$ it begins like X by inputting to y and decrementing ie; but then it writes the Boolean $(y \neq 0)$ to s, lowers the flag f and iterates. In the case $op = 2$ it decrements its y register and op then awaits the semaphore f; when $f = 1$ it updates s to the usual Boolean, lowers f and iterates. In the third case, $oe = 1$, it terminates.

$$
\begin{aligned}
Y \;:=\; & (@ : ie = 2 : y := iny \; \mathring{,} \; ie := ie - 1 \; \mathring{,} \; s := (y \neq 0) \; \mathring{,} \; f := 0 \; \mathring{,} \; Y) \; @ \\
& (@ : op = 2 : y := y - 1 \; \mathring{,} \; op := op - 1 \; \mathring{,} \; (@ : f = 1 : s := (y \neq 0) \; \mathring{,} \; f := 0) \; \mathring{,} \; Y) \; @ \\
& (@ : oe = 1 : \mathbf{skip})
\end{aligned}
$$

The control unit C begins by enabling register inputs by setting $ie = 2$. It then enters a loop in which it waits for the semaphore f to be high and any previous *op* action to be complete ($op = 0$) before reading the Boolean s and, if $y \neq 0$ enabling *op* and setting the semaphore $f = 1$; otherwise it enables output and terminates.[1] It seems slightly more natural (and facilitates comparison in the next section) to use (tail) recursion rather than a loop with guard *true*, as follows.

$$
\begin{aligned}
C \;:=\; & ie := 2 \; \mathring{,} \; D \\
D \;:=\; & (@ : (f = 0 \wedge op = 0) : (op := 2 \; \mathring{,} \; f := 1 \; \mathring{,} \; D) \lhd s \rhd oe := 1)
\end{aligned}
$$

We conclude that whilst the design can be expressed in PTSC, the use of various values (*ie*, *op*, *oe* and f) as clocking variables, or semaphores, betrays the existence of a simpler more abstract representation. For the purposes of informing the PTSC description, we consider it now.

5.3 Comparison with process algebra

In process algebra, channels can be used to communicate the status value s from Y to the control unit C and the values of the enables *ie*, *op* and *oe* from C to the registers. However the occurrence of probability requires us to use probabilistic process algebra. We choose the process algebra CSP [25] and its probabilistic extension pCSP [18], with contains the binary combinator ${}_r\oplus$ for probabilistic choice between its two arguments.

We recall that in CSP an assignment is a process, and so is combined with its successor process using $\mathring{,}$, whilst an event is combined with its successor process using $\rightarrow$. In CSP we continue to write the conditional construct 'A if b else C' as $A \lhd b \rhd C$.

---

[1]The simple modification which returns to the start, ready for another cycle, requires '$\mathring{,} C$' after the assignment to *oe* in C and similarly in X and Y. However it adds unnecessarily to the example.

The increment register X offers its environment a choice of three enables and in each case performs the expected actions

$$X := (ie \to x := inx \,\mathbin{\text{\scriptsize 9}}\, X) \;[\!]\; (op \to (x := x + 1 \,_r\!\oplus\, \mathbf{skip}) \,\mathbin{\text{\scriptsize 9}}\, X) \;[\!]\; (oe \to out := x).$$

The decrement register Y offers its environment a choice of the three enables, but also communication along channel s on which it outputs the Boolean $(y \neq 0)$

$$Y := (ie \to y := iny \,\mathbin{\text{\scriptsize 9}}\, Y) \;[\!]\; (op \to y := y - 1 \,\mathbin{\text{\scriptsize 9}}\, Y) \;[\!]\; (s!(y \neq 0) \to Y) \;[\!]\; (oe \to \mathbf{skip}).$$

The control unit C starts by enabling ie and then iteratively inputs the value from channel s and, if it is true enables op and reiterates but, if it is false then enables oe and terminates

$$C := ie \to \mu D \cdot (s?t \to ((op \to D) \lhd t \rhd (oe \to \mathbf{skip}))).$$

The implementation is as before

$$Imp := (X\|Y\|C) \setminus E$$

where we have been slightly cavalier with CSP notation in combining the hiding of the events with those of the internal variables.

The synchronisation between sender and receiver that is part of CSP obviates the need for the semaphores that cluttered our PTSC description. In particular Y and C appear simpler, as does the design as a whole.

5.4 Behaviour identified

In this section we outline an algebraic approach to identifying the behaviour of the design Imp. Consider first just the interleaving of C and Y.

$\quad ie, op, oe, f := 0, 0, 0, 1 \,\mathbin{\text{\scriptsize 9}}$
$\quad C \;\|\|\; Y$
$\quad =$ $\hfill$ definitions of C and Y

$\quad ie, op, oe, f := 0, 0, 0, 1 \,\mathbin{\text{\scriptsize 9}}$
$\quad (ie := 2 \,\mathbin{\text{\scriptsize 9}}\, \mu D \;(@ : (f = 0 \wedge op = 0) : (op := 2 \,\mathbin{\text{\scriptsize 9}}\, f := 1 \,\mathbin{\text{\scriptsize 9}}\, D) \lhd s \rhd (oe \,\mathbin{\text{\scriptsize 9}}\, C))$
$\quad \|\|$
$\quad (@ : ie = 2 : y := iny \,\mathbin{\text{\scriptsize 9}}\, ie := ie - 1 \,\mathbin{\text{\scriptsize 9}}\, s := (y \neq 0) \,\mathbin{\text{\scriptsize 9}}\, f := 0 \,\mathbin{\text{\scriptsize 9}}\, Y) \;@$
$\quad (@ : op = 2 : y := y - 1 \,\mathbin{\text{\scriptsize 9}}\, op := op - 1 \,\mathbin{\text{\scriptsize 9}}\, (@ : f = 1 : s := (y \neq 0) \,\mathbin{\text{\scriptsize 9}}\, f := 0) \,\mathbin{\text{\scriptsize 9}}\, Y))$
$\quad =$ $\hfill$ standard reasoning and Laws (9) and (10)

$\quad ie, op, oe, f := 0, 0, 0, 1 \,\mathbin{\text{\scriptsize 9}}$
$\quad ie := 2 \,\mathbin{\text{\scriptsize 9}}\, y := iny \,\mathbin{\text{\scriptsize 9}}\, ie := ie - 1 \,\mathbin{\text{\scriptsize 9}}\, s := (y \neq 0) \,\mathbin{\text{\scriptsize 9}}\, f := 0 \,\mathbin{\text{\scriptsize 9}}$
$\quad \mu Z \;(op := 2 \,\mathbin{\text{\scriptsize 9}}\, (f := 1 \;\|\|\; (y := y - 1 \,\mathbin{\text{\scriptsize 9}}\, op := op - 1)) \,\mathbin{\text{\scriptsize 9}}\, s := (y \neq 0) \,\mathbin{\text{\scriptsize 9}}\, f := 0 \,\mathbin{\text{\scriptsize 9}}\, Z$
$\qquad\qquad \lhd s \rhd$
$\qquad\qquad oe := 1)$

Next, including also X, we find

$$ie, op, oe, f := 0, 0, 0, 1 \,\fatsemi$$
$$(C \parallel\!\mid Y \parallel\!\mid X)$$
$$= \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\text{definition of } X \text{ and Law (5)}$$

$$ie, op, oe, f := 0, 0, 0, 1 \,\fatsemi$$
$$\quad (ie := 2 \,\fatsemi\, y := iny \,\fatsemi\, ie := ie - 1 \,\fatsemi\, s := (y \neq 0) \,\fatsemi\, f := 0 \,\fatsemi$$
$$\quad \mu Z \ (op := 2 \,\fatsemi\, (f := 1 \parallel\!\mid (y := y - 1 \,\fatsemi\, op := op - 1)) \,\fatsemi\, s := (y \neq 0) \,\fatsemi\, f := 0 \,\fatsemi\, Z$$
$$\qquad\qquad \vartriangleleft s \vartriangleright$$
$$\qquad\quad oe := 1)$$
$$\quad \parallel\!\mid$$
$$\quad (@ : ie = 2 : x := inx \,\fatsemi\, ie := ie - 1 \,\fatsemi\, X) \ @$$
$$\quad (@ : op = 2 : (x := x + 1 \,{}_r\!\oplus\, \mathbf{skip}) \,\fatsemi\, op := op - 1 \,\fatsemi\, X) \ @$$
$$\quad (@ : oe = 1 : out := x))$$
$$= \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\text{standard reasoning and Laws (9) and (10)}$$

$$op, oe, f := 0, 0, 1 \,\fatsemi$$
$$ie := 2 \,\fatsemi\, ((x := xin \,\fatsemi\, ie := ie - 1) \parallel\!\mid (y := yin \,\fatsemi\, ie := ie - 1 \,\fatsemi\, s := (y \neq 0) \,\fatsemi\, f := 0)) \,\fatsemi$$
$$\mu Z \ (op := 2 \,\fatsemi\, (((x := x + 1 \,{}_r\!\oplus\, \mathbf{skip}) \,\fatsemi\, op := op - 1 \,\fatsemi\, X) \parallel\!\mid f := 1 \parallel\!\mid (y := y - 1 \,\fatsemi\, op := op - 1)) \,\fatsemi$$
$$\qquad\quad s := (y \neq 0) \,\fatsemi\, f := 0 \,\fatsemi\, Z$$
$$\qquad \vartriangleleft s \vartriangleright$$
$$\qquad oe := 1 \,\fatsemi\, out := x) \,.$$

Thus localising the variables in the set $F := \{ie, op, oe, f\}$,

$${}_F[\![\, ie, op, oe, f := 0, 0, 0, 1 \,\fatsemi\, (C \parallel\!\mid Y \parallel\!\mid X) \,]\!]_F$$
$$= \qquad\qquad\qquad\qquad\qquad\qquad\qquad\text{standard reasoning and Laws (14), (19), (20) and (24)}$$

$$(x := xin \parallel\!\mid (y := yin \,\fatsemi\, s := (y \neq 0))) \,\fatsemi$$
$$\mu Z \ ((((x := x + 1 \,{}_r\!\oplus\, \mathbf{skip}) \,\fatsemi\, X) \parallel\!\mid y := y - 1) \,\fatsemi\, s := (y \neq 0) \,\fatsemi\, Z$$
$$\qquad \vartriangleleft s \vartriangleright$$
$$\qquad out := x) \,.$$

To establish termination (of the tail recursion on Z that is a mildly disguised loop) it suffices to take variant function y.

On each iteration the choice between incrementing and leaving x unchanged is binary with fixed probability r. So after i iterations the value of x is binomially distributed in the interval $[xin, xin + i]$. But the number of iterations is $i = yin - y$ and so we take as loop invariant $0 \leq y \leq yin$ and

$$Inv := \forall j : [0, yin - y] \cdot x = xin + j \text{ with probability } \binom{yin - y}{j} r^j \,\bar{r}^{\,yin - y - j} \,.$$

Indeed Inv is established: initially $y = yin$ and so $x = xin$ with probability $\binom{0}{0} r^0 \,\bar{r}^{\,0-0} = 1$. It is also seen to be maintained, by the expected but slightly longer calculation.

Thus, localising also s

$$_{F \cup \{s\}}[\![\, ie, op, oe, f := 0, 0, 0, 1 \, \semi \, (C \parallel\!\!\!\mid Y \parallel\!\!\!\mid X)\,]\!]_{F \cup \{s\}}$$
$$=$$
$$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\text{standard reasoning and Laws (14) and (20)}$$
$$(x, y := xin, yin) \, \semi \, out := x \, ,$$

where the distribution of values for $out = x$ satisfies the invariant $Inv[out/x]$ and the negation of the guard ($y \neq 0$):

$$\forall j : [0, yin] \cdot out = xin + j \text{ with probability } \binom{yin}{j} r^j \, \overline{r}^{\,yin-j} \, .$$

In other words, as expected, out is binomially distributed over the interval $[xin, xin + yin]$ of integers, with 'weight' r.

Finally, localising also x and y gives (with E defined at the end of Section 5.1)

$$_{E}[\![\, ie, op, oe, f := 0, 0, 0, 1 \, \semi \, (C \parallel\!\!\!\mid Y \parallel\!\!\!\mid X)\,]\!]_{E}$$
$$=$$
$$out := inx + iny \, ,$$

with binomial distribution of values for out, as above.

6   Conclusions

The language PTSC of Zhu, Qin, He and Bowen [26], for probability, time and shared-variable concurrency, has been extended in this paper to incorporate interleaving and local block. The resulting language permits programs to be combined non-probabilistically in parallel, variables to be localised (essential if the complexities of shared-variable programs have any hope of being subdued) and an implementation to be compared, after localising appropriate low-level variables, with its more abstract specification. Laws have been given for the new operators and a register-transfer-level design has been shown to meet its probabilistic specification. In general, subtle methods are required to reason about shared-variable concurrency [24]. Use of law-based reasoning, as done in this case study, seems to offer an alternative approach that may even be able to be automated [9].

The next step confronting this work is, as for the original PTSC, construction of a denotational model. We should like to have a normal form, and sufficiently many laws to transform any program into it.

Experience with CSP has played an important part in the intuitions behind PTSC, particularly for the guarded constructs. The proximity of the PTSC and CSP realisations of our case study

confirms that, and suggests that it may be interesting to translate from (simple) designs in CSP that use channel communications to avoid races, to PTSC designs which require semaphores or equivalents for the same purpose. With more experience we will presumably gain some appreciation of the difference between algebraic reasoning and the various techniques for reasoning about shared-variable interactions. A Galois connection between the models would afford a systematic start. It would be interesting to compare PTSC with pGCL, the probabilistic version of Dijkstra's guarded-command language [13] under both the expectation-transformer semantics and distributional semantics.

The case study contains just concurrency, locality and probability, overlooking time. However it could easily be extended to contain time: the registers could take different times (modelled in PTSC using delay) to perform their updates and in particular the faulty increment register could take different times depending on whether or not a fault occurs. Perhaps register delays could even be configured to avoid races and hence obviate the need for semaphores (as in hardware optimisation). That constitutes further work. The treatment of time in PTSC is a little simplistic: it behaves just like a global variable. It would be interesting to explore alternatives.

References

[1] http://www.verilog.com

[2] J. C. M. Baeten and W. P. Weijland. *Process Algebra*. Cambridge Tracts in Theoretical Computer Science, **18**. Cambridge University Press, 1990.

[3] S. D. Brookes. A fully abstract semantics of a shared-variable parallel language. In *Proc. 8th Annual IEEE Symposium on Logic in Computer Science*, 98–109. IEEE Computer Society Press, 1993.

[4] E. W. Dijkstra. Solution of a problem in concurrent programming control. *Communications of the ACM*, **8**(9):569, 1965.

[5] E. W. Dijkstra. Guarded commands, nondeterminacy and the formal derivation of programs. *Communications of the ACM*, **18**:453–457, 1975.

[6] He Jifeng and H. Zhu. Formalising Verilog. In *Proc. ICECS 2000: 3rd IEEE International Conference on Electronics, Circuits and Systems*, 412–415. IEEE Computer Society Press, 2000.

[7] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall International. 1985.

[8] C. A. R. Hoare and He Jifeng. *Unifying Theories of Programming*. Prentice-Hall Series in Computer Science, 1998.

[9] P. Höfner and G. Struth. Can refinement be automated? *Electronic Notes in Theoretical Computer Science (ENTCS)*, **201**:197–222, 2008.

[10] C. Jones. *Probabilistic Nondeterminism. Ph.D.* thesis, LFCS, The University of Edinburgh, 1990. Monograph ECS-LFCS-90-105.

[11] C. B. Jones. *Development Methods for Computer Programs Including a Notion of Interference. PhD.* thesis, OUCL, The University of Oxford, 1981.

[12] D. Kozen. Semantics of probabilistic programs. *Journal of Computer and System Sciences*, **27**:333–354, 1983.

[13] A. K. McIver and C. C. Morgan. *Abstraction, Refinement and Proof for Probabilistic Systems.* Monographs in Computer Science, Springer Verlag, 2005.

[14] J. Manson, W. Pugh and S. Adve. The Java memory model. In *Principles of Programming Languages (POPL)*, 378–391, 2005.

[15] R. Milner. *Communication and Concurrency.* Prentice-Hall International, 1989.

[16] J. Misra and K. M. Chandy. Proof of networks of processes. *IEEE Transactions on Software Engineering*, **7**(7):417–426, 1981.

[17] C. C. Morgan. *Programming from Specifications.* Prentice-Hall International, second edition, 1994.

[18] C. C. Morgan, A. K. McIver, K. Seidel and J. W. Sanders. Refinement-oriented probability for CSP. *Formal Aspects of Computing*, **8**(6):617–647, 1996.

[19] C. C. Morgan. Proof rules for probabilistic loops. In He Jifeng, J. Cooke and P. Wallis, editors, *Proceedings of the BCS-FACS 7th Refinement Workshop.* Workshops in Computing, Springer Verlag, 1996.

[20] R. Motwani and R. Raghavan. *Randomized Algorithms.* Cambridge University Press, 1995.

[21] J. Ouaknine and S. A. Schneider. Timed CSP: a retrospective. Proceedings of the workshop *Essays on Algebraic Process Calculi, (APC 25). Electronic Notes in Theoretical Computer Science*, **162**:273–276, September 2006.

[22] S. Owicki and D. Gries. An axiomatic proof technique for parallel programs. *Acta Informatica*, **6**:319–340, 1976.

[23] G. D. Plotkin A structural approach to operational semantics. *The Journal of Logic and Algebraic Programming*, **60,61**:17–139, 2004. Originally: Technical Report 19, University of Aarhus Denmark, 1981.

[24] W.-P. de Roever, F. de Boer, U. Hannemann, J. Hooman, Y. Lakhnech, M. Poel and J. Zwiers. *Concurrency Verification: Introduction to Compositional and Noncompositional Methods.* Cambridge Tracts in Theoretical Computer Science **54**, Cambridge University Press, 2001.

[25] A. W. Roscoe. *The Theory and Practice of Concurrency.* Prentice-Hall International, 2005.

[26] H. Zhu, S. Qin, Jifeng He and J. P. Bowen. Integrating probability with time and shared-variable concurrency. *Proceedings of the 30th Annual IEEE/NASA Software Engineering Workshop*, IEEE Computer Society Press, 179–189, 2006.