

A MORE COMPLETE MODEL OF COMMUNICATING PROCESSES

E.C.R. HEHNER

Computer Systems Research Group, University of Toronto, Toronto M5S 1A1, Canada

C.A.R. HOARE

Oxford University Computing Laboratory, Programming Research Group, Oxford OX2 6PE, United Kingdom

Communicated by M. Nivat

Received January 1982

Revised October 1982

Abstract. A previous paper by Hoare gives axioms and proof rules for communicating processes that provide a calculus of total correctness. This paper gives explicit definitions of communicating processes as predicates. The former axioms and proof rules become theorems, proved using the explicit definitions. The defining predicates are more powerful than the proof rules for reasoning about processes, but less often useful for their construction. An implementation of the processes using partial recursive functions is given.

Introduction

Axioms and proof rules for communicating processes have been given by Hoare [1]. They serve three purposes:

- (a) they implicitly define communicating processes;
- (b) they can be used to reason about communicating processes;
- (c) they can be used to construct communicating processes for a given purpose.

In this paper, processes are given explicit definitions as predicates. The axioms and proof rules in [1] become theorems about processes, proved using the explicit definitions in this paper. This proves the soundness of the calculus presented in [1].

The proof rules are incomplete as a means of reasoning about communicating processes; an example of an expressible but not provable truth is given in [1]. The 'honest toil' of giving explicit predicate definitions has the advantage that it is complete, relative to data types. Its disadvantage is that the explicit definitions are less often useful for the construction of processes to fulfil a given purpose than the former proof rules (now theorems).

An axiomatic definition requires a model to demonstrate the consistency of the axioms. In [2], a set-theoretic model of communicating processes is given. From that, we know that at least one mathematical object satisfies the axioms. We should

like to know that at least one computable mathematical object satisfies the axioms. To that end, an implementation using partial recursive functions is presented in this paper.

1. Axiomatic basis

A message is a pair $\langle \text{channel}, \text{value} \rangle$ over an alphabet of channels and values. If direction of communication is important, a message can be taken instead to be a triple $\langle \text{channel}, \text{direction}, \text{value} \rangle$, or channels can be considered directed; however, we shall not need to refer to the direction of communication. (This is entirely realistic; if you tap a phone line, you can record 'values', but not direction.) To denote a message, we shall use infix '!' instead of angle brackets. For example $c!1$ is a message with channel c and value 1.

A particular process at a particular moment has a past, which is the sequence of messages it has communicated up to that moment. It also has a present, which is the set of messages that it can communicate at the next step. Which message in the set it actually does communicate next may be determined in part (or completely, or not at all) by the environment of the process (those other processes connected to it by channels). If the set is empty, the process cannot communicate further, and is said to be deadlocked.

Let *past* be a free variable representing a finite sequence of messages. Let *present* be a free variable representing a finite set of messages. A predicate R in *past* and *present* is said to describe a process P if, at all times (before and after each communication) during any evolution (execution) of P , R is true of P 's past and present. This is denoted

$$P \text{ sat } R$$

meaning " P is described by R ", or " P satisfies description R ".

Hoare proposes four 'healthiness conditions' that are reasonably required of the **sat** relation. For all processes P and all predicates $R, S, R(n)$,

(H1) $P \text{ sat true}$.

(H2) $\neg(P \text{ sat false})$.

(H3) If $R \Rightarrow S$ is a theorem, then $(P \text{ sat } R) \Rightarrow (P \text{ sat } S)$ is a theorem.

(H4) $(\forall n \in N. P \text{ sat } R(n)) \equiv (P \text{ sat } \forall n \in N. R(n))$.

On this basis, an axiom or proof rule is introduced in [1] for each construct (an axiom is just a proof rule with no premises).

For example, the process STOP can be defined by the axiom

$$(\text{STOP sat } R) \equiv R[\text{past}: \langle \rangle; \text{present}: \{\}].$$

The right side denotes the predicate obtained from R by simultaneously substituting, for all free occurrences of *past*, the empty sequence of messages, and for all free occurrences of *present*, the empty set of messages.

2. Explicit basis

" $P \text{ sat } R$ " means that predicate R describes process P , but the description may be a weak one. In the extreme, according to (H1), **true** is a (weak) description of every process. By (H4), the conjunction of all descriptions of a process P is also a description of P ; it is the strongest, and therefore an exact, description of P . If that description can be written explicitly, then it is an explicit definition of P .

Notationally, there is no reason to distinguish between a process and the predicate that defines it. Therefore we shall define a process as a predicate in the free variables *past* and *present* over a finite alphabet of messages. (The finiteness of the alphabet will be required later for computability.) For example, STOP can be defined as

$$\text{STOP: } \text{past} = \langle \rangle \wedge \text{present} = \{ \}.$$

The colon means "is defined as". (It is appropriate to use the same symbol for definition and substitution because definition is simply permission to substitute.)

For any process P , the possible initial communications are described exactly by the predicate $P[\text{past}: \langle \rangle]$, and the sequences of communications that lead to deadlock are described exactly by the predicate $P[\text{present}: \{ \}]$. If we wish to prove that a process cannot deadlock, we must prove

$$\forall \text{past}. \neg P[\text{present}: \{ \}].$$

With this way of defining processes, " $P \text{ sat } R$ " means

$$\forall \text{past}, \text{present}. P \Rightarrow R.$$

The healthiness conditions degenerate nicely.

(H1) $\forall \text{past}, \text{present}. P \Rightarrow \text{true}$ is a tautology.

(H2) $\neg(\forall \text{past}, \text{present}. P \Rightarrow \text{false})$ asserts that P is satisfiable.

(H3) If $R \Rightarrow S$ is a theorem, then $(\forall \text{past}, \text{present}. P \Rightarrow R) \Rightarrow (\forall \text{past}, \text{present}. P \Rightarrow S)$ is a theorem. This is provable in the predicate calculus.

(H4) $(\forall n \in N. \forall \text{past}, \text{present}. P \Rightarrow R(n)) \equiv (\forall \text{past}, \text{present}. P \Rightarrow \forall n \in N. R(n))$ is a tautology, assuming P does not mention n .

The word 'healthy' has been reduced to the word 'satisfiable', and indeed a predicate should be true of some *past* and *present* if it is to be a reasonable definition of a process. That excludes

MIRACLE: false

but is not sufficient as a characterization of processes. Consider

$$\text{NOWAY1: } \text{past} = \langle c!1 \rangle \wedge \text{present} = \{ \}.$$

Although NOWAY1 is clearly satisfiable, execution cannot begin, because NOWAY1[*past*: $\langle \rangle$] is not satisfiable; there is no initial message. It is reasonable to require of a process that when something cannot happen, it will not happen. A

different problem arises in

$$\text{NOWAY2: } \text{past} = \langle \rangle \wedge \text{present} = \{c!1\}.$$

Initially, the message $c!1$ can be communicated, but *past* remains forever empty. It is reasonable to require that when something can happen, something will happen.

Definition. A process P is a predicate in the free variables *past* and *present* such that

$$(P0) \quad \exists \text{past}, \text{present}. P,$$

$$(P1) \quad \forall \text{past}, \text{message}. (\exists \text{present}. P \wedge \text{message} \in \text{present}) \\ \equiv (\exists \text{present}. P[\text{past}: \text{past}^\wedge \langle \text{message} \rangle])$$

where $^\wedge$ is catenation of message sequences.

Perhaps (P1) is best explained as two implications, slightly rearranged.

$$\forall \text{past}, \text{present}, \text{message}.$$

$$P \wedge \text{message} \in \text{present} \Rightarrow$$

$$\exists \text{newpresent}. P[\text{past}: \text{past}^\wedge \langle \text{message} \rangle; \text{present}: \text{newpresent}]$$

says that *past* can be extended by any *message* in *present*, and P will still be satisfiable (it still describes the process).

$$\forall \text{past}, \text{present}, \text{message}.$$

$$P[\text{past}: \text{past}^\wedge \langle \text{message} \rangle] \Rightarrow$$

$$\exists \text{oldpresent}. P[\text{present}: \text{oldpresent}] \wedge \text{message} \in \text{oldpresent}$$

says that a nonempty *past* can be shortened by removing its final *message*, and P must have been satisfied by that shorter *past* and a *present* containing that *message*.

According to the definition of a process, neither MIRACLE nor NOWAY1 nor NOWAY2 is a process. But STOP is, and so is

$$\text{CHAOS: true}.$$

The next section will show why this process merits the name "CHAOS".

Having proposed a definition of a process which we named STOP, we must prove that the definition of STOP is consistent with the axiom proposed for STOP in [1]. The axiom must become a theorem in the predicate calculus. To prove it, we need the following.

Substitution fact. If P is a predicate, e is an expression, and x is a variable not appearing in e , then

$$P[x: e] \equiv \forall x. x = e \Rightarrow P$$

where the quantifier ranges over the type of x .

Proof. In [3], Shoenfield proves

$$P[x:e] \equiv \exists x. x = e \wedge P$$

from which we know both

$$\neg(P[x:e]) \equiv \neg\exists x. x = e \wedge P,$$

$$(\neg P)[x:e] \equiv \exists x. x = e \wedge \neg P.$$

But

$$(\neg P)[x:e] \equiv \neg(P[x:e]).$$

Hence

$$\begin{aligned} P[x:e] &\equiv \neg\neg P[x:e] \\ &\equiv \neg\exists x. x = e \wedge \neg P \\ &\equiv \forall x. x = e \Rightarrow P. \quad \square \end{aligned}$$

STOP Theorem. $(\text{STOP sat } R) \equiv R[\text{past}: \langle \rangle; \text{present}: \{\}].$

Proof. Using first the interpretation of “ $P \text{ sat } R$ ”, then the definition of “STOP”, then the substitution fact generalized to two variables, we have

$$\begin{aligned} \text{STOP sat } R &\equiv \forall \text{past}, \text{present}. \text{STOP} \Rightarrow R \\ &\equiv \forall \text{past}, \text{present}. \text{past} = \langle \rangle \wedge \text{present} = \{\} \Rightarrow R \\ &\equiv R[\text{past}: \langle \rangle; \text{present}: \{\}]. \end{aligned}$$

3. Implementation

An implementation of a process is a partial recursive function from sequences of messages to sets of messages. For each past within its domain, the function determines (computes) a present that contains the possible next messages. If all of the processes connected by a channel are simultaneously willing to communicate several different messages, i.e. the intersection of their presents contains more than one message, then an arbitrary one of those messages will be communicated. In that way, the system of processes (not the individual processes) is nondeterministic.

There is another kind of nondeterminism. A process predicate P can be satisfied by a particular past and present, say s and M , and also by s and N ; the same past can have more than one corresponding present. This process nondeterminism is a freedom for the implementor. An implementation, being deterministic, will deliver one of the corresponding presents for each past. In effect, for a function f to be an implementation of process P , it must be an implementation of some deterministic subprocess Q .

Definition. The pasts of a process P are defined as

$$\text{pasts}(P): \{past \mid \exists present. P\}.$$

Definition. Q is a subprocess of P if Q is a process ((P0) and (P1) are true of Q) and $Q \Rightarrow P$.

Definition. Process P is deterministic if $\forall past \in \text{pasts}(P). \exists ! present. P$.

Definition. An implementation of process P is a partial recursive function f from sequences to sets of messages, such that for some deterministic subprocess Q ,

$$\forall s \in \text{pasts}(Q). Q[past: s; present: f(s)].$$

We shall use $\#P\#$ as the name of an implementation of process P . For example,

$$\#STOP\#(s): \text{if } s = \langle \rangle \text{ then } \{ \}$$

is a partial recursive function defined only for the one sequence $\langle \rangle$. It is a correct implementation of STOP because

- (a) $\text{pasts}(\text{STOP}) = \{ \langle \rangle \},$
- (b) $\text{STOP}[past: \langle \rangle; present: \#STOP\#(\langle \rangle)]$
 $= \langle \rangle = \langle \rangle \wedge \{ \} = \{ \}$
 $= \text{true}.$

All other correct implementations of STOP must agree with this one for the sequence $\langle \rangle$.

CHAOS is the most nondeterministic process. Every deterministic process is a subprocess of CHAOS, so that an implementation of any process is also an implementation of CHAOS. The name "CHAOS" is deserved because it is not at all determined what CHAOS will do.

4. Process construction

Processes can be constructed from other processes in prescribed ways. Suppose that F is a one-place process constructor, i.e. from process P we construct process $F(P)$, from Q , $F(Q)$, etc. We require F to obey these laws.

- (PC0) If P is a process, then $F(P)$ is a process.
- (PC1) (*monotonicity*): $\forall P, Q$. If $P \Rightarrow Q$ is a theorem, then $F(P) \Rightarrow F(Q)$ is a theorem.
- (PC2) (*continuity*): If $P_0 P_1 P_2 \dots$ is a strengthening chain of processes, i.e. if $\forall i. P_{i+1} \Rightarrow P_i$ is a theorem, then $(\forall i. F(P_i)) \equiv F(\forall i. P_i)$ is a theorem.

By the usual trick known as 'Currying', a multi-place constructor can be regarded as the composition of several one-place constructors.

Each of Sections 5 to 14 presents one way of constructing processes from other processes. Each constructor is defined, a theorem (formerly axiom) is stated, and an implementation of it is displayed. For each constructor, it should be proved that (PC0)–(PC2) are satisfied, the theorem should be proved, and the correctness of the implementation should be proved. Most of these proofs are either tedious or obvious, and are omitted; only three of the more interesting proofs of theorems are included.

In addition to the standard notations of sets, predicates, and functions, the following notations will be used:

$\text{value}(m)$	is the value component of message m ,
$\text{channel}(m)$	is the channel component of message m ,
$\#s$	is the length of message sequence s ,
$r\hat{\ }s$	is the catenation of message sequences r and s ,
$\text{first}(s)$	is the first message in nonempty message sequence s ,
$\text{rest}(s)$	is the subsequence of nonempty message sequence s that excludes the first message,
$s.c$	is the subsequence of message sequence s that includes those messages having channel component c ,
$s \setminus c$	is the subsequence of message sequence s that excludes those messages having channel component c ,
$M.c$	is the subset of message set M that includes those messages having channel component c ,
$M \setminus c$	is the subset of message set M that excludes those messages having channel component c .

5. Output

A process that communicates the value of expression e on channel c (i.e. it communicates the message $c!e$) and then behaves like process P can be defined as

$$c!e \rightarrow P: \text{past} = \langle \rangle \wedge \text{present} = \{c!e\} \\ \vee \text{past} \neq \langle \rangle \wedge \text{first}(\text{past}) = c!e \wedge P[\text{past}: \text{rest}(\text{past})].$$

The output theorem, which will now be proved, is

$$(c!e \rightarrow P) \text{ sat } R \equiv R[\text{past}: \langle \rangle; \text{present}: \{c!e\}] \\ \wedge P \text{ sat } R[\text{past}: \langle c!e \rangle \wedge \text{past}].$$

The proof has three parts.

- (a) $(\forall \text{past}, \text{present}. \text{past} = \langle \rangle \wedge \text{present} = \{c!e\} \Rightarrow R) \\ \equiv R[\text{past}: \langle \rangle; \text{present}: \{c!e\}]$

by the substitution fact of Section 2.

(b) The predicate

$$\forall \text{past, present. } \text{past} \neq \langle \rangle \wedge \text{first}(\text{past}) = c!e \wedge P[\text{past} : \text{rest}(\text{past})] \Rightarrow R$$

can be simplified by considering the two cases $\text{past} = \langle c!e \rangle^r$ for some sequence r , and $\text{past} \neq \langle c!e \rangle^r$ for any sequence r . In the latter case, the antecedent is clearly false, so the implication is clearly true. The implication must hold for all values of past , so it reduces to the former case,

$$\forall r, \text{present. } P[\text{past} : r] \Rightarrow R[\text{past} : \langle c!e \rangle^r].$$

This, by a change of bound variable, is just

$$\forall \text{past, present. } P \Rightarrow R[\text{past} : \langle c!e \rangle^{\text{past}}]$$

which proves

$$\begin{aligned} & (\forall \text{past, present. } \text{past} \neq \langle \rangle \wedge \text{first}(\text{past}) = c!e \wedge P[\text{past} : \text{rest}(\text{past})] \Rightarrow R) \\ & \equiv (\forall \text{past, present. } P \Rightarrow R[\text{past} : \langle c!e \rangle^{\text{past}}]). \end{aligned}$$

(c) Part (a) has proved something of the form

$$(\forall p. A \Rightarrow R) \equiv B.$$

Part (b) has proved something of the form

$$(\forall p. C \Rightarrow R) \equiv D.$$

From these two, follows

$$(\forall p. A \vee C \Rightarrow R) \equiv B \wedge D$$

which is the desired result.

An implementation of output is

$$\begin{aligned} & c!e \rightarrow P^*(s) : \text{if } s = \langle \rangle \\ & \quad \text{then } \{c!e\} \\ & \quad \text{else if } \text{first}(s) = c!e \\ & \quad \quad \text{then } P^*(\text{rest}(s)). \end{aligned}$$

6. Input

Let c be a channel, M a finite set of messages for channel c , and $P(x)$ a process with parameter x . Then $c?x:M \rightarrow P(x)$ is a process that can initially communicate any message in M , say $c!x$, and then behave like $P(x)$. It is defined as

$$\begin{aligned} & c?x:M \rightarrow P(x) : \text{past} = \langle \rangle \wedge \text{present} = M \\ & \quad \vee \text{past} \neq \langle \rangle \wedge \text{first}(\text{past}) \in M \\ & \quad \wedge P(\text{value}(\text{first}(\text{past}))) [\text{past} : \text{rest}(\text{past})]. \end{aligned}$$

The input theorem

$$\begin{aligned}
 & (c?x:M \rightarrow P(x)) \text{ sat } R \\
 & \equiv R [past: \langle \rangle; present: M] \\
 & \wedge \forall m \in M. P(\text{value}(m)) \text{ sat } R[past: \langle m \rangle^{\wedge} past]
 \end{aligned}$$

can be proved in a manner similar to the proof of the output theorem. In fact, according to their semantics, output is just the special case of input in which M is a unitset.

An implementation of input is

$$\begin{aligned}
 & \#c?x:M \rightarrow P(x) \#(s): \text{ if } s = \langle \rangle \\
 & \quad \text{then } M \\
 & \quad \text{else if first}(s) \in M \\
 & \quad \quad \text{then } \#P(\text{value}(\text{first}(s))) \#(\text{rest}(s)).
 \end{aligned}$$

7. Recursion

If F is a one-place process constructor, then the recursive process definition

$$P: F(P)$$

means

$$P: \forall i. F^i(\text{CHAOS})$$

where

$$F^0(P) = P, \quad F^{i+1}(P) = F(F^i(P)).$$

Starting with CHAOS, we create an ordered sequence of processes $F^i(\text{CHAOS})$. The ordering is reverse implication ('is implied by'). The sequence is one of ever more deterministic subprocesses, and P is its limit. In other words, P is the least fixed point of F , where 'least' means 'weakest' or 'least deterministic'.

The theorem for recursively defined P is

$$(\forall i. (F^i(\text{CHAOS}) \text{ sat } R \uparrow i) \Rightarrow (F^{i+1}(\text{CHAOS}) \text{ sat } R \uparrow i + 1)) \Rightarrow P \text{ sat } R$$

where $R \uparrow i \equiv (\#past < i \Rightarrow R)$. To prove it, notice first that $R \uparrow 0 \equiv \text{true}$, hence $F^0(\text{CHAOS}) \text{ sat } R \uparrow 0$. With that basis, the major antecedent, by induction, implies

$$\begin{aligned}
 & \forall i. F^i(\text{CHAOS}) \text{ sat } R \uparrow i \\
 & \forall i. \forall past, present. F^i(\text{CHAOS}) \Rightarrow R \uparrow i \\
 & \forall past, present. \forall i. F^i(\text{CHAOS}) \Rightarrow R \uparrow i \\
 & \forall past, present. (\forall i. F^i(\text{CHAOS})) \Rightarrow (\forall i. R \uparrow i) \\
 & \forall past, present. P \Rightarrow R \\
 & P \text{ sat } R
 \end{aligned}$$

By choosing $\nexists \text{CHAOS} \nexists (s)$ to be the universe of messages for the channels of the construction, P can be implemented as

$$\nexists P \nexists (s): \bigcap_{i=0}^{\infty} \nexists F^i(\text{CHAOS}) \nexists (s)$$

assuming we know how to implement $F(P)$ in terms of an implementation of P . By choosing $\nexists \text{CHAOS} \nexists (s)$ to be the empty set, P can be implemented as

$$\nexists P \nexists (s): \bigcup_{i=0}^{\infty} \nexists F^i(\text{CHAOS}) \nexists (s).$$

8. Channel renaming

If process P does not mention channel d , then

$$P[c: d]$$

denotes a process like P except that all occurrences of channel c are replaced by channel d . Its theorem,

$$P[c: d] \text{ sat } R[c: d] \equiv P \text{ sat } R$$

where R does not mention c , is obvious. So is its implementation

$$\nexists P[c: d] \nexists (s): \nexists P \nexists [c: d](s).$$

9. Disjoint parallelism

Let P be a process that communicates on the set of channels C , and let Q be a process that communicates on the set of channels D , where $C \cap D = \{\}$. Then $P \parallel Q$ is a process that behaves like P and Q in parallel, not communicating with each other. It is defined as

$$\begin{aligned} P \parallel Q: & P[\text{past}: \text{past}.C; \text{present}: \text{present}.C] \\ & \wedge Q[\text{past}: \text{past}.D; \text{present}: \text{present}.D] \end{aligned}$$

where the dot notation has been extended in an obvious way.

From this definition, the theorem

$$(P \text{ sat } S) \wedge (Q \text{ sat } T) \Rightarrow P \parallel Q \text{ sat } (S \wedge T)$$

is easily proved.

Disjoint parallelism can be implemented as

$$\nexists P \parallel Q \nexists (s): \nexists P \nexists (s.C) \cup \nexists Q \nexists (s.D).$$

10. Channel connection

The process formed from process P by connecting its channels c and d and naming the connected channel b (a new channel name) is denoted

$$b = c \leftrightarrow d \text{ in } P.$$

The new process can communicate a message on the new channel b iff P can communicate a similar message on both channels c and d . For message $b!x$ to occur in the past of the new process, there must be a corresponding past of P in which $c!x$ and $d!x$ occur together. In fact, there must be two corresponding possible pasts of P , one in which $c!x$ and $d!x$ occur together in that order, and one in which they occur together in the other order. Without loss of generality, we can restrict our attention to one of the two orders. The definition can be stated as

$$\begin{aligned} b = c \leftrightarrow d \text{ in } P: & \exists s, M. P[\text{past}: s; \text{present}: M] \\ & \wedge \exists t. \forall i, x. (t_i = b!x) \equiv (s_i = c!x) \equiv (s_{i+1} = d!x) \\ & \wedge t \setminus b = s \setminus c \wedge \text{past} = t \setminus d \\ & \wedge \text{present} = \{b!x \mid \exists N. P[\text{past}: s \wedge \langle c!x \rangle \wedge \langle d!x \rangle; \text{present}: N]\} \\ & \cup M \setminus c \setminus d \end{aligned}$$

where t_i and s_i are the i th messages in message sequences t and s .

Its theorem is

$$\begin{aligned} (P \text{ sat } R) \Rightarrow & (b = c \leftrightarrow d \text{ in } P) \text{ sat } \exists x, y. R[c:x; d:y] \\ & \wedge \text{present}.b = \text{present}.x \cap \text{present}.y \\ & \wedge \text{past}.b = \text{past}.x = \text{past}.y. \end{aligned}$$

An implementation is

$$\#b = c \leftrightarrow d \text{ in } P \#(s): \#P \#(g(s, \langle \rangle))$$

where

$$\begin{aligned} g(s, t): & \text{ if } s = \langle \rangle \\ & \text{ then } t \\ & \text{ else if first}(s) = b!x \\ & \quad \text{ then } g(\text{rest}(s), t \wedge \langle c!x \rangle \wedge \langle d!x \rangle) \\ & \quad \text{ else } g(\text{rest}(s), t \wedge \langle \text{first}(s) \rangle) \end{aligned}$$

reconstructs the past with $c!x$ and $d!x$ replacing each occurrence of $b!x$.

11. Hiding

If P is a process that communicates on channel b (and possibly others) then

chan b in P

is a process like P , except that channel b is local, and not visible to (available for communication with) its environment. Channel b is used for internal communication in the new process. Hiding is quite a delicate subject, as the discussion will show, and we do not yet have an entirely satisfactory definition. A first attempt is

chan b in P :

$$\exists s, M. P[past: s; present: M] \wedge past = s \setminus b \wedge present = M \setminus b.$$

A variation is obtained by adding the conjunct

$$\begin{aligned} \forall r, t, N. s = r \hat{=} t \wedge P[past: r; present: N] \wedge N.b \neq \{ \} \\ \Rightarrow t \neq \langle \rangle \wedge \text{channel}(\text{first}(t)) = b \end{aligned}$$

within the scope of $\exists s$. Without this conjunct, a process is interruptible, i.e. there is the possibility of external communication as an alternative to internal communication. But with this conjunct, a process is uninterruptible while engaged in internal communication; when, in process P , communication is possible on either channel b or another channel, then, in process **chan b in P** , channel b will be chosen. The interruptible version includes the presents of unstable states, excluding the internal communications (hidden messages) that make the states unstable. It does so in order to include the possibility of external communications (interrupts) at those states. Unfortunately, this means that an implementation can deliver the present of an unstable state, ignoring forever the internal communication, waiting forever for an interrupt. The uninterruptible version includes only the presents of stable states. It is unsatisfactory because it excludes the possibility of external communication at an unstable state.

The hiding theorem is

$$(P \text{ sat } (R \wedge \exists f. \#past.b \leq f(past \setminus b)))$$

$$\Rightarrow ((\text{chan } b \text{ in } P) \text{ sat } (\exists s. R[past: s] \wedge past = s \setminus b \wedge present.b = \{ \}))$$

and can be proved from the uninterruptible version without using the theorem's antecedent. The purpose of the antecedent is to allow a definition of channel hiding that is weaker than the one we have given in the following way. In our uninterruptible version, a process which can engage in an infinite sequence of communications on the hidden channel does not communicate further with its environment; livelock (infinite internal chatter) is equivalent to deadlock. The hiding theorem, as an axiom in [1], allows livelock to be considered differently, in any way desired, even as equivalent to CHAOS; it does not allow anything to be proven about a livelocked process.

Hiding can be implemented, uninterruptibly, as follows.

$$\# \text{chan } b \text{ in } P\#(s): \#P\#(h(s, \langle \rangle))$$

where

$$\begin{aligned} h(s, t): & \text{ if } \#P\#(t).b = \{ \} \\ & \text{ then if } s = \langle \rangle \\ & \quad \text{ then } t \\ & \quad \text{ else } h(\text{rest}(s), t^{\wedge} \langle \text{first}(s) \rangle) \\ & \text{ else } h(s, t^{\wedge} \langle \text{choose}(\#P\#(t).b) \rangle). \end{aligned}$$

The function h constructs a new sequence t from the given sequence s such that $s = t \setminus b$ by inserting messages on channel b whenever possible. When there is a choice of messages that can be inserted, the choose function chooses an arbitrary one of them. Because *present* is always a finite set of messages, this choice will be from a finite set, and choose is therefore implementable. The function h is not necessarily terminating; this is in accord with the possibility of livelock.

The two kinds of nondeterminism mentioned in Section 3 can be called 'external' and 'internal'. External nondeterminism, having more than one message in a present, is resolved by the availability of external communication, and if that does not fully resolve it, then by other external forces. Internal nondeterminism, having more than one present, is resolved internally according to the implementation. But the two become mixed during channel hiding, when things external become internal. Separating them again properly is a delicate operation.

12. Nondeterministic union

The process that can behave either like process P or like process Q is defined as

$$P \text{ or } Q: P \vee Q.$$

Its theorem

$$(P \text{ or } Q) \text{ sat } R \equiv (P \text{ sat } R) \wedge (Q \text{ sat } R)$$

is easily proven.

One implementation is

$$\#P \text{ or } Q\#(s): \#P\#(s).$$

13. Conditional

The conditional process is defined as

$$\text{if } e \text{ then } P \text{ else } Q: e \Rightarrow P \wedge \neg e \Rightarrow Q$$

where e is a boolean expression, and P and Q are processes. Its theorem is

$$\begin{aligned} & (\text{if } e \text{ then } P \text{ else } Q) \text{ sat } R \\ & \equiv e \Rightarrow (P \text{ sat } R) \wedge \neg e \Rightarrow (Q \text{ sat } R). \end{aligned}$$

It can be implemented as

$$\begin{aligned} & * \text{if } e \text{ then } P \text{ else } Q * (s): \text{if } e \text{ then } *P * (s) \\ & \quad \text{else } *Q * (s). \end{aligned}$$

14. Alternation

Let I and J be input processes, as in Section 6. Then $I \parallel J$ is a process that behaves either like I or like J , depending on the availability of input for I and J . If, in its environment, an initial communication for I is possible but for J impossible, then it behaves like I . Conversely, if an initial communication for J is possible and for I impossible, it behaves like J . If both are possible, it behaves like either I or J . If neither is possible, it is deadlocked.

$$\begin{aligned} & c?x:M \rightarrow P(x) \parallel d?y:N \rightarrow Q(y): \\ & \quad \text{past} = \langle \rangle \wedge \text{present} = M \cup N \\ & \quad \vee \text{past} \neq \langle \rangle \wedge \text{first}(\text{past}) \in M \\ & \quad \quad \wedge P(\text{value}(\text{first}(\text{past}))) [\text{past}: \text{rest}(\text{past})] \\ & \quad \vee \text{past} \neq \langle \rangle \wedge \text{first}(\text{past}) \in N \\ & \quad \quad \wedge Q(\text{value}(\text{first}(\text{past}))) [\text{past}: \text{rest}(\text{past})]. \end{aligned}$$

The alternation theorem is

$$\begin{aligned} & (c?x:M \rightarrow P(x) \parallel d?y:N \rightarrow Q(y)) \text{ sat } R \\ & \equiv R[\text{past}: \langle \rangle; \text{present}: M \cup N] \\ & \quad \wedge (\forall m \in M. P(\text{value}(m)) \text{ sat } R[\text{past}: \langle m \rangle^{\wedge} \text{past}]) \\ & \quad \wedge (\forall m \in N. P(\text{value}(m)) \text{ sat } Q[\text{past}: \langle m \rangle^{\wedge} \text{past}]). \end{aligned}$$

An implementation of alternation is

$$\begin{aligned} & *c?x:M \rightarrow P(x) \parallel d?y:N \rightarrow Q(y) * (s): \\ & \quad \text{if } s = \langle \rangle \\ & \quad \text{then } M \cup N \\ & \quad \text{else if first}(s) \in M \end{aligned}$$

then $\nexists P(\text{value}(\text{first}(s)))\nexists(\text{rest}(s))$

else if $\text{first}(s) \in N$

then $\nexists Q(\text{value}(\text{first}(s)))\nexists(\text{rest}(s)).$

15. Conclusion

To program in the notation of communicating processes, one must begin with a specification, which is a predicate to be satisfied, and then choose an appropriate construction. An axiom or proof rule should say what the components of the construction must satisfy in order that the construction will satisfy the original predicate. Thus one creates specifications for subcomponents in the usual 'top-down' fashion.

The intended use of an object should guide the design of the object. Just as a specification, via the proof rules, guides the design of a process, so the desire for such proof rules guides the design of the range of process constructs. It is therefore appropriate for an axiomatic, implicit definition of communicating processes to precede an explicit definition. The explicit definition serves both as a guide for implementors, and as a more powerful tool for reasoning about processes. For example, defining

$$P: b!0 \rightarrow P$$

the predicate

$$(\text{chan } b \text{ in } P) \text{ sat true}$$

is cited in [1] as an expressible but unprovable truth. With the definitions in this paper, its proof is easy.

When processes are defined axiomatically, a model should be built to show that the axioms describe at least one object of interest; that was first done in [2]. The partial recursive functions in this paper serve as a second model, and a demonstration of computability. With more models, one gains confidence that at least some of the objects one wanted to define are in the class of objects defined.

How does one prove that the class defined includes all and only the desired objects? A Turing Machine simulation can show that all computable functions are expressible as processes, but the aptness of the process constructions is not amenable to proof in the usual sense. However, when a second, independent definition proves to be consistent with the first, confidence in the definitions is greatly increased. That is a contribution of this paper. In one sense, the predicate definitions are not independent: they were intended to be consistent with the axiomatic definitions. Nonetheless, they provide the reader with another look at communicating processes.

Acknowledgment

This paper was improved by the helpful comments of the referees, particularly Dr. Ernst-Rüdiger Olderog.

References

- [1] C.A.R. Hoare, A calculus of total correctness for communicating processes, *Sci. Comput. Programming* **1**(1) (1981) 49-72.
- [2] C.A.R. Hoare, S.D. Brookes and A.W. Roscoe, A theory of communicating sequential processes, Oxford University Programming Research Group Technical Monograph PRG-16 (1981).
- [3] J.R. Shoenfield, *Mathematical Logic* (Addison-Wesley, Reading, MA, 1967) 36.