

A CALCULUS
of
TOTAL CORRECTNESS
for
COMMUNICATING PROCESSES.

by
C.A.R. HOARE

April 1981

May 1981 (with minor revisions).

/// means "new paragraph please"

For Communicating Processes.

C. A. R. Hoare

April 1981


Dedication: to my son Matthew 1967-1981.

Summary.


A process communicates with its environment and with other processes by synchronised output and input on named channels. The current state of a process is defined by the sequences of messages which have passed along each of the channels up to some moment in time, and also by the sets of messages that may next be passed on each channel. A process satisfies an assertion if the assertion is at all times true of all possible states of the process.

We present a calculus for proving that a process satisfies the assertion describing its intended behaviour.


The following constructs are axiomatised: output; input; simple recursion; disjoint parallelism; channel renaming, connection and hiding; process chaining; nondeterminism; conditional; alternation; and mutual recursion. The calculus is illustrated by proof of a number of simple buffering protocols.




A ... a process with alphabet $\{\text{left}, \text{right}\}$



B a process $(P \parallel Q)$ with alphabet $\{\text{left}, c, d, \text{right}\}$



C the process $(b = (c \leftrightarrow d) \text{ in } B)$ with alphabet $\{\text{left}, b, \text{right}\}$



D is the process (chan b in C) with alphabet $\{\text{left}, \text{right}\}$

Figure 1

1. ASSERTIONS.

A process communicates with its environment by sending and receiving messages on named channels. (Fig. 1A).

The names of these channels constitute the alphabet of the process.

A process may be constructed from a group of subprocesses, intercommunicating on a network of named channels, (Fig. 1B, C). A message output

by one process along a channel is received instantaneously by all other processes connected to that channel, provided that all these processes are simultaneously prepared to input that message.

close up

// On each named channel, it is possible to keep a record of all messages passing along it. (For simplicity, we ignore direction of communication: if desired, this could be recorded as part of each message.) At any given moment, the record of all messages that have passed ^(so far) on a channel c is a finite sequence, which will be denoted by the variable " $c.past$ ".

no new para

At the very beginning, the value of $c.past$ (for each channel c) is the empty sequence $\langle \rangle$. During the evolution of a process, whenever a message m is communicated on channel c , the value of $c.past$ is extended on the right by m , and the new value is $(c.past \langle m \rangle)$.

/// At any given moment, the set of messages which a process is prepared to communicate on channel c is denoted by the variable " $c.ready$ ". When the process is not prepared to communicate at all on channel c , the value of $c.ready$ is the empty set \emptyset .

When a process is prepared to input on channel c , the value of $c.ready$ is the set M of all possible messages _{for that channel}. (When a process is prepared to output some message value m (selected from M), then the value of $c.ready$ is the ^{unit} set $\{m\}$, which has m as its only member.

/// Variables of the form $c.past$, $c.ready$ are known as channel variables. Since we do not wish to be concerned with the internal states and transitions of a process; we shall identify the current externally observable state of a process with the current values of its channel variables. $c.past$, for arranging the process.

4

/// An assertion with a given alphabet is a normal sentence of logic and mathematics, which may contain free channel variables of the form "c.past" and "c.ready", where c is a channel name in the alphabet of the assertion. The assertion describes certain possible states of some process at certain moments of time.

For example, the following are assertions, with informal explanations of their meaning.

(a) left.past = right.past

"The sequence of messages which has passed so far along the left channel is the same as the sequence that has passed along the right channel"

(b) left.ready = M

"The left channel is ready for input of any message in the set M"

(c) right.past < left.past

"The messages passed on the right channel form a proper initial subsequence of the messages that have passed on the left"

(d) right.ready = {first(left.past - right.past)}

"The right channel is ready for output of the earliest message on the left which has not yet been transmitted on the right"

/// Assertions may be readily combined by the familiar connectives of logic. For example, we define for future use the assertion:

$$\text{BUFF} \triangleq \text{left.past} = \text{right.past} \ \& \ \text{left.ready} = M \\ \vee \ \text{right.past} < \text{left.past} \ \& \\ \text{right.ready} = \{\text{first}(\text{left.past} - \text{right.past})\}.$$

This assertion describes all possible states of a buffering process (or transparent communications protocol), which outputs on its right channel the same sequence of messages which it inputs from the left, though possibly after some delay. When $\text{left.past} = \text{right.past}$, the process has an empty buffer, and it must then be prepared to input any message from the left. In the alternative case, the buffer is nonempty; it contains the sequence $(\text{left.past} - \text{right.past})$ of messages which are awaiting output on the right; and now the buffering process must be prepared to output the first element of this buffer. The assertion BUFF does not say whether or not input on the left is possible when the buffer is nonempty; and thus it does not specify any particular bound on the size of the buffer.

/// Let P be a process and let R be an assertion with the same alphabet as P .

Then P is said to satisfy R if at all times during the evolution of P , (before and after each communication) the assertion R correctly describes the observable state of P , i.e., the sequences of messages that have passed along its named channels, and the sets of messages that are ready to be communicated on the very next step.

This relation between processes and assertions is abbreviated:

$$P \text{ sat } R.$$

For example any process P which is to serve as a buffer or transparent communications protocol must satisfy the assertion BUFF . There are many processes that do so; - for example, a bounded buffer of any finite size or even an unbounded buffer; examples will be given later.

/// It follows from the intended interpretation of the relation "satisfies" that the following properties should be true for all processes P , and all predicates, R, S .

$$(H1) \quad P \text{ sat } \text{TRUE}$$

omit TRUE is a predicate which is always true of everything; it must therefore always be true of the behaviour of every process.

$$(H2) \quad \neg (P \text{ sat } \text{FALSE}) \equiv \dots$$

FALSE is the predicate that is always false of anything; it cannot therefore correctly describe the behaviour of any process.

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

If $(R \Rightarrow S)$ is a theorem, every state in which R is true is also a state in which S is true.

If all states of P are correctly described by R , they must also be correctly described by S , and hence

$$((P \text{ sat } R) \Rightarrow (P \text{ sat } S)) \text{ is also true!} \quad (H3) \text{ is}$$

a useful proof rule, known as the "rule of consequence"

Corollary:
$$\frac{R \equiv S}{(P \text{ sat } R) \equiv (P \text{ sat } S)}$$

(H4) If n is not a channel variable, and does not occur in P :

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

If, for each n in some set N , P satisfies $R(n)$, then each state of P is correctly described by $R(n)$, for all n in N . The converse implication follows from (H3).

Corollary: $(P \text{ sat } R) \& (P \text{ sat } S) \equiv (P \text{ sat } (R \& S))$

/// These four conditions are rather

similar to the healthiness conditions introduced by E.W. Dijkstra [1] to check the validity of each clause in the definition of his weakest precondition for sequential programming. Unfortunately, our calculus is not strong enough to prove healthiness in all cases; so we have to introduce the (conditions) as independent axioms, which must at least be consistent with the other proof rules of the calculus.

/// Let R be an assertion not containing the variable n ;
 then we define $R \uparrow n$ ^(R restricted to n) as the assertion satisfied by
 a process which behaves as described by R for
 at least $n-1$ steps, i.e., at least until the total
 number of communications on all channels reaches n .

Let $\{a, \dots, z\}$ be the

alphabet of R . Let $\#s$ stand for the length of
 the sequence s . Then we can define:

$$R \uparrow n \triangleq (\#a.\text{past} + \dots + \#z.\text{past} \geq n) \vee R$$

Example: $\text{BUFF} \uparrow n \triangleq (\#\text{left.past} + \#\text{right.past} \geq n) \vee \text{BUFF}$

Theorem 1. For any assertion R

(a) $R \uparrow 0$ is a theorem

$$(b) (\forall n: \text{NAT}. R \uparrow n) \equiv R$$

Proof: $c.\text{past}$ is a finite sequence for each channel c .

So $\#c.\text{past}$ is a natural number. R does not
 contain n , so

$$\begin{aligned} (\forall n: \text{NAT}. R \uparrow n) &\equiv (\forall n: \text{NAT}. \#a.\text{past} + \dots + \#z.\text{past} \geq n) \vee R \\ &\equiv R \end{aligned}$$

// Let R be an assertion possibly containing a variable x ,
 and let e be an expression of the same type as x .
 Then we define $R[e/x]$ as the assertion formed
 from R by substituting e for every free occurrence
 of x . (If any free variable of e would thereby
 become bound to a bound variable in R , the
 collision must be averted by systematic change of
 the offending bound variable.) For example, we define
 $BUFF' \triangleq (BUFF \uparrow (n+1)) [\langle x \rangle \text{left.past} / \text{left.past}]$
 $BUFF'' \triangleq BUFF' [\langle x \rangle \text{right.past} / \text{right.past}]$

After performing the substitutions, $BUFF''$ expands to:

$$\begin{aligned}
 & \# \langle x \rangle \text{left.past} + \# \langle x \rangle \text{right.past} \geq n+1 \\
 & \vee \langle x \rangle \text{left.past} = \langle x \rangle \text{right.past} \ \& \ \text{left.ready} = M \\
 & \vee \langle x \rangle \text{right.past} < \langle x \rangle \text{left.past} \\
 & \quad \& \ \text{right.ready} = \{\text{first}(\langle x \rangle \text{left.past} - \langle x \rangle \text{right.past})\}
 \end{aligned}$$

The following theorem is typical of the lengthy but shallow truths required in proofs of correctness of programs:

Theorem 2. $BUFF \uparrow n \Rightarrow (\forall x: M. BUFF'')$

Proof. each clause of the LHS implies the corresponding clause on the RHS.

11
//Let R be an assertion with alphabet $a..z$.

We introduce the convention that

$R[\langle \rangle / \text{past}]$

is the result of substituting the empty sequence $\langle \rangle$ for every occurrence of any of the channel variables $a.\text{past}, \dots, z.\text{past}$. For example

$\text{BUFF}[\langle \rangle / \text{past}] \equiv \langle \rangle = \langle \rangle \ \& \ \text{left.ready} = M$
 $\vee \langle \rangle \langle \rangle \ \& \ \dots$

which is equivalent to "left.ready = M". If $P \text{ sat } R$, then $R[\langle \rangle / \text{past}]$ describes all the possible states of P at its very beginning, before it has engaged in communication on any of its channels. These states are defined in terms of $a.\text{ready}, \dots, z.\text{ready}$, which specify the sets of communications for which P should be ready on its very first step. Thus if any process is to satisfy the assertion BUFF , it must at the beginning be ready to input on its left channel any value in the set M .

/// By a similar convention

$$R[\emptyset/\text{ready}]$$

is the result of substituting the empty set \emptyset for every occurrence of any of the channel variables $a.\text{ready}, \dots, z.\text{ready}$. For example:

$$\text{BUFF}[\emptyset/\text{ready}] = \text{left.past} = \text{right.past} \ \& \ \phi = M \\ \vee \text{right.past} < \text{left.past} \ \& \ \phi = \{\dots\}$$

which is always false. If $P \text{ sat } R$, then $R[\emptyset/\text{past}]$ describes all possible states of P , in which it is not ready for communication along any of its channels. These states are known as deadlock states; and it is usually desired to prove that they cannot occur. The states are defined in terms of the variables $a.\text{past}, \dots, z.\text{past}$; and therefore we only need to prove that $R[\emptyset/\text{ready}]$ is false for all values of these variables. For example, any process that satisfies BUFF can never deadlock, (unless the set M of all possible messages is empty, — a possibility which we can realistically ignore).

/// As a final convention, we allow successive substitutions to be separated by commas; for example:

$$R[\langle \rangle/\text{past}, \emptyset/\text{ready}] = (R[\langle \rangle/\text{past}])[\emptyset/\text{ready}].$$

// One of the simplest processes ^(with alphabet A) is the process $STOP_A$ which is already deadlocked at its start. Clearly, it is never ready to do anything, so $c.ready = \emptyset$ for all c in A . Furthermore, the sequence of messages transmitted along each channel remains forever empty, i.e. $c.past = \langle \rangle$. In summary, the process $STOP_A$ has only this single state; consequently, it satisfies an assertion R if and only if R correctly describes its only state, i.e. if R is true when all the variables of the form $c.ready$ take the value \emptyset , and all the variables of the form $c.past$ take the value $\langle \rangle$. This informal reasoning justifies the axiom:

$$(STOP_A \text{ sat } R) \equiv R[\emptyset / ready, \langle \rangle / past]$$

Examples. The following are theorems:

$$STOP_A \text{ sat } (c.ready \neq \{x\} \ \& \ \#c.past \neq 3)$$

$$\rightarrow (STOP_{LR} \text{ sat } BUFF)$$

where $LR = \{left, right\}$

// $STOP_A$ is rather a useless process; it has been introduced here only to provide a simple example of an axiom, i.e. and how it can be informally justified.

In the remainder of this monograph, we introduce a number of programming constructs suitable for the programming of communicating processes.

Each construct is given a syntax, and an informal explanation of its semantics. The semantics is formalised by an axiom or proof rule which is illustrated by application to some simple example. Treatment of each example is spread over several consecutive subsections.

2.1 Output

Let P be a process;

let c be a channel name in the alphabet of P ; and

let e be an expression (not containing channel variables).

Then we use the notation

$$(c!e \rightarrow P)$$

to denote the process which first outputs the value of e on channel c and then behaves like P .

In its initial state, when the

past of all its channels is empty, this process is prepared to communicate the value of e on channel c , so that $c.\text{ready} = \{e\}$. It is not prepared to communicate ^{on} any other channel, so initially $d.\text{ready} = \phi$ for all channels d other than c .

An assertion R is true of this initial state if and only if it is true when the channel variables of R take their initial values, as described above. This may be expressed by substituting these values in R , giving

$$R[\langle \rangle / \text{past}, \{e\} / c.\text{ready}, \phi / \langle \text{ready} \rangle]$$

(The use of the expression e to stand for its value is justified only in a programming notation which excludes assignment of new values to variables.

occurring in e)

The subsequent states of $(c!e \rightarrow P)$ are very similar to the states of P ; the only difference is in the value of $c.past$. If in a state of P $c.past$ has value s , then in the corresponding state of $(c!e \rightarrow P)$, $c.past$ has the value $\langle e \rangle s$. In order to prove

$$(c!e \rightarrow P) \underline{\text{sat}} R$$

it is the process P that must ensure, not that its own states satisfy R , but rather that the corresponding states of $(c!e \rightarrow P)$ are correctly described by R . In other words, R must be true when the value of $c.past$ is replaced by $\langle e \rangle c.past$; or more formally:

$$P \underline{\text{sat}} (R [\langle e \rangle c.past / c.past])$$

To prove that all states of a process are correctly described by R , it is sufficient to prove that the initial state satisfies R , and that the subsequent states do so too. The preceding paragraphs deal with these two cases; putting them together we get the rules:

$$((c!e \rightarrow P) \underline{\text{sat}} R) \equiv (R[\langle e \rangle / past, \{e\} / c.ready, \emptyset / ready] \& P \underline{\text{sat}} (R[\langle e \rangle c.past / c.past]))$$

Example.

$$((\text{right! } x \rightarrow p) \text{ sat } \text{BUFF}') \equiv$$

$$S \ \& \ (p \text{ sat } \text{BUFF}' [\langle x \rangle / \text{right.past} / \text{right.past}])$$

where $S \triangleq \text{BUFF}' [\langle \rangle / \text{past}, \{x\} / \text{right.ready}, \emptyset / \text{ready}]$

On performing the substitutions, S expands to

$$\# \langle x \rangle + \# \langle \rangle \geq n+1$$

$$\forall \langle x \rangle = \langle \rangle \ \& \ \emptyset = M$$

$$\forall \langle \rangle < \langle x \rangle \ \& \ \{x\} = \{\text{first}(\langle x \rangle - \langle \rangle)\}.$$

The last clause makes S a trivial theorem.

Theorem 3.

$$((\text{right! } x \rightarrow p) \text{ sat } \text{BUFF}') \equiv (p \text{ sat } \text{BUFF}'')$$

Proof. The theorem S can be omitted from a conjunction, and the definition of BUFF'' is used.

The axiom for output has the same apparent "backwards" quality as the axiom of assignment in sequential programming. Readers who have become familiar with the latter may note that the command $(c!e \rightarrow P)$ has the same apparent effect on $c.\text{past}$ as the command

$$(P; c.\text{past} := \langle e \rangle c.\text{past})$$

provided that P contains no assignments. Thus the second term of the axiom of output is derivable from (the axiom of assignment).

2.2 Input

Let $P(x)$ be a process whose behaviour (but not alphabet) possibly depends on the value of the free variable x . Let c be a channel in the alphabet of $P(x)$, and let M be a finite ~~nonempty~~ set of message values which can be communicated on channel c . Then

$$(c?x: M \rightarrow P(x))$$

is the process which is initially prepared to input on channel c any value in the set M . The newly input value is given the local name x , and the process subsequently behaves like $P(x)$. The variable x is regarded as a bound variable, so

$$(c?x: M \rightarrow P(x))$$

is the same process as

$$(c?y: M \rightarrow P(y)).$$

Example.

$$\text{COPYSTEP} \triangleq (\text{left}?x: M \rightarrow (\text{right}!x \rightarrow p))$$

COPYSTEP first inputs a value from the left, then outputs this same value to the right, and then behaves like p .

/// The input command is similar to the output command except in two respects. Firstly, the initial value of $c.ready$ is not just a single value, but the whole of the set M . Secondly, the subsequent behaviour $P(x)$ may depend on the input value x , which is not known in advance; and therefore $P(x)$ must be proved to meet its specification for all values of x ranging over the set M . This reasoning informally justifies the axiom:

Let R be an assertion not containing x .

$$((c?x: M \rightarrow P(x)) \underline{sat} R) \equiv (R[\langle \rangle / \text{past}, M / c.ready, \emptyset / \text{ready}] \ \& \ \forall x: M. (P(x) \underline{sat} R[\langle x \rangle c.past / c.past]))$$

Example.

$$(COPYSTEP \underline{sat} (BUFF \uparrow_{n+1})) \equiv S \ \& \ (\forall x: M. (\text{right}!x \rightarrow p) \underline{sat} BUFF')$$

where $S \triangleq (BUFF \uparrow_{n+1})[\langle \rangle / \text{past}, M / \text{left.ready}, \emptyset / \text{ready}]$
 $\equiv (\# \langle \rangle + \# \langle \rangle \geq n+1) \vee (\langle \rangle = \langle \rangle \ \& \ M = M) \vee (\langle \rangle \langle \rangle \ \& \dots)$

The second clause makes S a theorem.

Theorem 4.

$$(COPYSTEP \underline{sat} (BUFF \uparrow_{n+1})) \equiv (p \underline{sat} (\forall x: M. BUFF''))$$

Proof. Theorem 3, definition of $BUFF'$ and (H4)

2.3. Recursion

Let p be a variable standing for a process with a given alphabet. Let $F(p)$ be the description of a process (with the same alphabet) containing none or more occurrences of the variable p . Then

$$\mu p. F(p)$$

is the recursively defined process, which starts off behaving like $F(p)$, and on encountering an occurrence of p , behaves like $(\mu p. F(p))$ again.

Example.

$$\text{COPY} \triangleq \mu p. (\text{left}?x: M \rightarrow (\text{right}!x \rightarrow p))$$

The process COPY is an infinitely repeating cycle, each iteration of which inputs a message from the left and outputs the same message to the right.

A recursively defined process is intended to be a "fixed point" of its defining function, F , i.e.,

$$\mu p. F(p) \text{ is the same process as } F(\mu p. F(p)) \dots \quad (1)$$

Let R be an assertion, and suppose for an arbitrary process p we can prove

$$(p \text{ sat } (R \wedge n)) \Rightarrow (F(p) \text{ sat } R \wedge (n+1)). \quad (2)$$

From theorem 1(a) and (H1) it follows that

$$(\mu p. F(p)) \text{ sat } (R \wedge 0)$$

By substituting $\mu p F(p)$ for p in (2), and using (1) we get

$$(\mu p. F(p) \text{ sat } R \uparrow n) \Rightarrow (\mu p. F(p) \text{ sat } R \uparrow (n+1))$$

(close up)

By the obvious induction on n we get

$$\forall n. (\mu p. F(p) \text{ sat } (R \uparrow n))$$

By (H4) and theorem 1(b), we conclude

$$(\mu p. F(p)) \text{ sat } R.$$

This reasoning serves as

(an informal justification of the

following proof rule:

$$\frac{(\mu p \text{ sat } (R \uparrow n)) \Rightarrow (F(p) \text{ sat } (R \uparrow (n+1)))}{\mu p. F(p) \text{ sat } R}$$

Theorem 5. COPY sat BUFF.

By the rule given above, it is sufficient to prove

$$(p \text{ sat } (\text{BUFF} \uparrow n)) \Rightarrow (\text{COPYSTEP } \text{sat} (\text{BUFF} \uparrow_{n+1}))$$

By theorem 4, this is equivalent to

$$(p \text{ sat } \text{BUFF} \uparrow n) \Rightarrow (p \text{ sat } (\forall c: M. \text{BUFF}''))$$

which follows from theorem 2 by (H3).

/// Now at last we see the motivation for the choice of assertions used in the previous examples. Of course, a proof would normally be presented in the reverse order, with proof requirements for the component processes being derived by formal manipulation from the proof requirement of the whole process. The reader is invited to use this top-down method to prove the obvious fact:

$$(\mu p. (b!0 \rightarrow p)) \text{ sat } (b.\text{ready} \neq \emptyset).$$

2.4. Channel renaming.

Let P be a process, with channel c in its alphabet, and let d be a ^{channel} name not in its alphabet.

Then $P[d/c]$ is taken to denote a process that behaves just like P , except that

c is removed from its alphabet

d is included in its alphabet

whenever P would have used channel c for input or output, $P[d/c]$ uses d instead.

$P[d/c]$ can clearly be derived from the definition of the process P by replacing each occurrence of the name c by an occurrence of d .

Example.

$$\text{COPY}[d/\text{right}] \triangleq \mu p. (\text{left}?x: M \rightarrow (d!x \rightarrow p)).$$

A similar transformation may be made to any assertion satisfied by P , in accordance with the following convention:

$$R[d/c] \triangleq R[d.\text{past}/c.\text{past}, d.\text{ready}/c.\text{ready}].$$

The appropriate axiom is quite obvious:

$$(P[d/c] \text{ sat } R[d/c]) \equiv (P \text{ sat } R)$$

2.5. Disjoint parallelism

Let P and Q be processes with disjoint alphabets. Since they have no channel name in common, they are unconnected, and therefore cannot communicate or interact with each other in any way.

The notation $(P \parallel Q)$ denotes a process which behaves like P and Q evolving in parallel; its alphabet is clearly the union of the alphabets of P and Q . Channel renaming can be used when needed to achieve disjointness of alphabets.

Example.

$$\text{PROT} \triangleq (\text{COPY}[d/\text{right}]) \parallel (\text{COPY}[c/\text{left}])$$

This combination is illustrated in Fig 1 B.

\parallel The states of $(P \parallel Q)$ correspond to elements of the cartesian product space of the set of states of P and the set of states of Q . If P satisfies S , then S has the same alphabet as P ; it therefore correctly describes the current values of those channels in the state of $(P \parallel Q)$ which are in the alphabet of P ; and hence

$$(P \parallel Q) \text{ sat } S.$$

Similarly, if Q sat T it follows that

$$(P \parallel Q) \text{ sat } T.$$

Hence by (H4, corollary), we justify the proof rule.

$$\frac{(P \text{ sat } S) \& (Q \text{ sat } T)}{(P \parallel Q) \text{ sat } (S \& T)}$$

Example.

Let $\text{BUFF}(c, d) \triangleq \text{BUFF}[d/\text{right}] \& \text{BUFF}[c/\text{left}]$

Theorem. 6.

$\text{PROT} \text{ sat } \text{BUFF}(c, d)$

Proof. immediate from Theorem 5 and the proof rules for renaming and disjoint parallelism.

2.6. Channel Connection.

Let P be a process with channels c and d in its alphabet. We may wish to connect together these two channels, so that messages passed on either of them are simultaneously passed on the other. For technical reasons, we give a new name b to the newly connected channel, and eliminate the names c and d from the alphabet of P . The process resulting from this connection and renaming will be denoted

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

Example.

$$\text{PROT} \triangleq (b = c \leftrightarrow d \text{ in } \text{PROT})$$

This is illustrated in Figure 1C. The intended effect of this connection is that everything output on the right channel by the left operand of \parallel gets simultaneously input on the left channel of the right operand of \parallel , (and vice-versa).

~~The sequence of messages so passed is recorded as $b.\text{past}$, and the readiness of the joined channel is indicated by $b.\text{ready}$. The values of these new channel variables are defined in the following~~

/// When two channels $\left. \begin{matrix} c \text{ and } d \end{matrix} \right\}$ are connected, a message can be passed on the connecting channel b if and only if both of the connected channels are ready for that communication; *is* at all times:

$$b.\text{ready} = (c.\text{ready} \wedge d.\text{ready})$$

As a consequence, whenever c is ready for output and d for input, $d.\text{ready}$ is the universal set Ω , and the connected channel b ready for output of the same value as c . Similar remarks apply when d is ready for output and c for input. When both c and d are ready for input, so is b . When either of c or d is unready then so is b . There remains the case that both c and d are ready for output, and the readiness of b depends $\left. \begin{matrix} \text{on} \end{matrix} \right\}$ whether the values output are the same. This case is not very useful, and should probably be excluded in a practical programming notation.

/// Each message transmitted on either of the connected channels $\left. \begin{matrix} c \text{ and } d \end{matrix} \right\}$ is instantaneously passed by the connecting channel b to the other one. The sequences of messages so transmitted must therefore always be the same:

$$b.\text{past} = c.\text{past} = d.\text{past}.$$

20

/// It is the duty of an implementation of the connection operator to ensure that $b.ready$ and $b.past$ have the right values, as described in the above paragraphs. The programmer can just assume that this has been done. Thus we derive the proof rule

$$P \text{ sat } R$$

$$(b = c \leftrightarrow d \text{ in } P) \text{ sat } (b.ready = c.ready \wedge d.ready \\ \wedge b.past = c.past = d.past \\ \wedge R)$$

Unfortunately, the assertion in the consequent of this rule contains the channel names c and d , which are not supposed to be in the alphabet of the process concerned. This problem is easily solved by the valid technique of weakening the consequent. (H3); it is easy to check that the following proof rule is a logical consequence of the one justified above.

28a

P sat R

$$(b=c \leftrightarrow d \text{ in } P) \text{ sat } (\exists x, y. b.\text{ready} = x \wedge y \\ \& R[b.\text{past}/c.\text{past}, b.\text{past}/d.\text{past}, \\ x/c.\text{ready}, y/d.\text{ready}])$$

Theorem 7.

$$\text{PROTOL sat } \exists x, y. (b.\text{ready} = x \wedge y \& \text{BB})$$

where $\text{BB} \triangleq \text{BUFF}(c, d)[b.\text{past}/c.\text{past}, b.\text{past}/d.\text{past}, \\ x/c.\text{ready}, y/d.\text{ready}]$.

Proof. immediate from theorem 6.

Here is BB written out in full:

$$\begin{aligned} & (\text{left.past} = b.\text{past} \& \text{left.ready} = M \\ & \vee b.\text{past} < \text{left.past} \& y = \{\text{first}(\text{left.past} - b.\text{past})\}) \\ & \& (b.\text{past} = \text{right.past} \& x = M \\ & \vee \text{right.past} < b.\text{past} \& \text{right.ready} = \{\text{first}(b.\text{past} - \text{right.past})\}) \end{aligned}$$

2.7. Hiding

Let P be a process with channel b in its alphabet. Suppose that b is a channel which connects two or more component subprocesses of P , as described in the previous section. Since b is still in the alphabet of P , it can still be used for communication with the environment of P . Indeed, no communication can take place on channel b without the knowledge and consent of the environment. However, in the design of any mechanism, we usually wish to conceal its internal workings from its environment, and this is especially important for electronic devices, which can work millions of times faster than the environment could keep pace with. We therefore wish to hide from the environment ^(of P) all communications passing between subprocesses of P along channel b . Each such communication is intended to occur automatically and instantaneously as soon as all the processes connected by the channel are ready for it. And, of course, channel b must be removed from the alphabet of P . The required effect is denoted:

(chan b in P)

which declares the name b as a local channel in P .

As with other local variables, we postulate,
(chan b in P) is the same as (chan c in P [c/b])
where c is not in the alphabet of P.

Example:

PROTOCOL \triangleq (chan b in (PROTOC))

In this example, the channel b connects ^{the} two parallel subprocesses of the process PROTOC. One of the processes acts like a trivial transmitter of a protocol, and the other as a trivial receiver.

The channel b serves as the transmission line between them.

The user of the mechanism is not concerned with the nature, number, or content of the messages passing along the transmission line, which are therefore concealed from him, as shown in Fig 1D.

//// A state of the process (chan b in P) is said to be stable if there is no further possibility of communication on channel b, i.e.,

$$b.ready = \emptyset$$

In an unstable state, when communication is possible on channel b, we want that communication to take place invisibly at high speed; and this will bring the process to a new and usually different state.

Of course, if one of the "other" channels is ready and the environment is prepared to communicate on at the same time as b,

that channel, the external communication can occur instead
- but ^(this) cannot be relied upon. If the environment is not
prepared to communicate on any of the other ready channels,
a (we insist that a ready

internal communication must sooner or later occur -

and preferably sooner; Thus the unstable states

are evanescent, and cannot be relied upon; in

specifying the externally ^(visible) behaviour of processes, it

seems sensible to ignore them. In

other words we choose to interpret

$P \text{ sat } R$

as a claim that R is true of all stable states of P

For each stable state of (chan b in P), there exists a state
of P in which $b.\text{ready} = \emptyset$ and in which $b.\text{past}$ has some
value of no further interest. This informal reasoning suggests
a proof rule:

$(P \text{ sat } R)$ is equivalent to

$(\text{chan } b \text{ in } P) \text{ sat } (\exists b.\text{past}. R[\emptyset / b.\text{ready}])$

(Here we have quantified over a channel variable as

if it were an ordinary variable. The meaning

is the same as if an ordinary variable s had been

substituted, i.e.,

$\exists s. R[s / b.\text{past}, \emptyset / b.\text{ready}]$)

/// Unfortunately this proof rule leads to a contradiction.
Consider the process

$$P \triangleq \mu p. b!0 \rightarrow p$$

P outputs an unbounded sequence of zeroes on channel b , and is always prepared to output another; we can prove:

$$P \text{ sat } (b.\text{ready} \neq \emptyset)$$

From this, using the incorrect rule given above, we deduce

$$(\text{chan } b \text{ in } P) \text{ sat } \exists b.\text{past} ((b.\text{ready} \neq \emptyset) [\emptyset / b.\text{ready}])$$

The assertion here reduces to $\emptyset \neq \emptyset$, which violates the condition (H2) (counterexample due to W.A. Roscoe).

The trouble here is that we have tried to hide an infinite sequence of internal communications, with disastrous consequences for our theory. The consequences in practice could be equally unfortunate, because the resulting process might expend all its energies on internal communication, and never pay any further attention to its environment. This phenomenon is known as "livelock" or "infinite chatter", and there are sound theoretical and practical reasons for requiring a programmer to prove it cannot occur. A simple way of doing this is to prove that the number of messages which can be passed along the hidden channel b is bounded by some

total function f of the state of the other non-hidden channels:

$$\#b.past \leq f(c.past, \dots, z.past)$$

where c, \dots, z are all the other channels in the alphabet of the process.

Summarising the discussion above, we formulate the proof rules:

$$\frac{P \text{ sat}(R \& (\#b.past \leq f(c.past, \dots, z.past))) \& b.ready \text{ is finite}}{(\text{chan } b \text{ in } P) \text{ sat}(\exists b.past. R[\phi/b.ready])}$$

Theorem 8.

$$\text{PROTOCOL } \text{sat}(\exists b.past, x, y. (\phi = xny \& BB))$$

$$\text{Proof. } BB \Rightarrow (BB \& \#b.past \leq \#left.past)$$

The conclusion follows from theorem 7 and (H3)

54

We are at last ready to prove

Theorem 9. PROTOCOL sat BUFF

Proof. We prove the assertion of Theorem 8. implies BUFF.

Expanding the assertion, BB we get four cases:

$$\text{left.past} = \text{b.past} = \text{right.past} \ \& \ \text{left.ready} = \infty = M$$

$$\vee \text{right.past} < \text{b.past} = \text{left.past} \ \& \ \text{right.ready} = \{\text{first}(\text{left.past} - \text{b.past})\} \ \& \ \dots$$

$$\vee \text{right.past} = \text{b.past} < \text{left.past} \ \& \ \infty = M \ \& \ y = \{\dots\}$$

$$\vee \text{right.past} < \text{b.past} < \text{left.past} \ \& \ \text{right.ready} = \text{first}(\text{b.past} - \text{right.past}) \ \& \ y = \dots$$

where irrelevant phrases are replaced by ellipses.

The first two clauses obviously imply the corresponding clauses of BUFF. The third clause describes an unstable state, and contradicts the term $(\phi = \infty \cap y)$;

this case is therefore eliminated. The fourth clause also implies the corresponding clause of BUFF, using transitivity of $<$ and the fact that

$$r < b < l \Rightarrow \text{first}(b-r) = \text{first}(l-r).$$

2.8. Process chaining.

The connection of processes in a series by their right and left channels is such a useful operation that it deserves a special notation:

$$(P \Leftrightarrow Q) \triangleq \text{chan } b \text{ in } (b = c \Leftrightarrow d \text{ in } ((P [d/\text{right}]) ||| (Q [c/\text{left}])))$$

where b, c, d are fresh channel names

Example. $\text{PROTOCOL} = (\text{COPY} \Leftrightarrow \text{COPY})$

Unfortunately, the proof rule for this defined construct is hardly simpler than its definition.

Let s, x , and y be fresh variables.

$$\text{Let } S' = S [s/\text{right.past}] [x/\text{right.ready}]$$

$$\text{Let } T' = T [s/\text{left.past}] [y/\text{left.ready}]$$

Let f be a total function of pairs of sequences.

$$P \text{ sat } S, \quad Q \text{ sat } T,$$

$$S' \& T' \Rightarrow s \leq f(\text{left.past}, \text{right.past})$$

$$(P \Leftrightarrow Q) \text{ sat } (\exists s, x, y. (x \cap y = \emptyset \ \& \ S' \& T'))$$

Theorem 10 If $P \text{ sat } \text{BUFF}$ and $Q \text{ sat } \text{BUFF}$.

then $(P \Leftrightarrow Q) \text{ sat } \text{BUFF}$

Proof. essentially the same as given for theorem 9.

Corollaries: $(\text{PROTOCOL} \Leftrightarrow \text{COPY}) \text{ sat } \text{BUFF}$

$(\text{PROTOCOL} \Leftrightarrow \text{PROTOCOL}) \text{ sat } \text{BUFF}$

etc.

2.9. Nondeterministic union.

Let P and Q be process descriptions with the same alphabet. Then the notation

$$(P \text{ or } Q)$$

stands for a process that behaves either like P or like Q . The choice between the alternatives is left completely unspecified, and may be made arbitrarily as the process $(P \text{ or } Q)$ evolves, or may be fixed by its implementor before the starts. The choice cannot be influenced by the environment of the process, and is undetectable at the time it is made — though it may be deducible from the subsequent behaviour of the process.

Example. (PROTOCOL or COPY)

This behaves either like a one-place buffer or a two-place buffer, the choice being unspecified and unknown.

If, during the life of this process, the length of left.past ever exceeds the length of right.past by two, then we can deduce that the choice has ^(fallen on) the first alternative.

If we want to be sure that $(P \text{ or } Q)$ satisfies R , since we do not know which of P or Q will be selected, we had better prove that they both satisfy R

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

Theorem 11. (PROTOCOL or COPY) sat BUFF

Proof. from Theorems 9 and 5.

2.10 Conditional

Let e be a Boolean-valued expression not containing any channel variables. Let P and Q be processes with the same alphabet. Then the process

is one that behaves like P if e evaluates to true, or behaves like Q if e evaluates to false.

The proof rule is correspondingly simple:

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

An example will be given in 2.12.

2.11. Alternation.

Let $P(x)$ and $Q(y)$ be processes whose behaviour possibly depends on the values of the free variables x and y respectively; but all of them have the same alphabet.

Let c and d be distinct channel names in this alphabet. Let M be the set of messages that can be communicated on c , and let N be the set for d . Then the notation

$$(c?x:M \rightarrow P(x) \square d?y:N \rightarrow Q(y))$$

denotes a process which behaves as follows. Initially, it is prepared to input either on channel c or on channel d ; in the first case its subsequent behaviour is defined by $P(x)$, where x stands for the value input on c ; and in the second case, its subsequent behaviour is ~~similarly~~ defined by $Q(y)$, where y is the value input on d .

(Only one of the two inputs can take place; but in contrast to nondeterministic union, the choice can be influenced by the other processes connected to the channels c and d . If the process (or processes) connected to one of them remains forever unprepared for communication, then communication can still occur, but only on the other channel. But if ^(all) the processes connected to each of the channels ~~all become ready~~ for communication, then it is nondeterministic ~~on which~~

(will take place.)

on which channel the communication) (An efficient implementation should select the first to become ready; but such considerations of efficiency rightly cannot be formalised in a calculus of correctness; and a programmer clearly must not rely on them, since he has delegated to the implementor all control over the relative speeds of the processes. Example.

$$\text{MERGESTEP} \triangleq (\text{left1?}x: M \rightarrow \text{right}!(1,x) \rightarrow p \quad \square \text{left2?}x: M \rightarrow \text{right}!(2,x) \rightarrow p)$$

This process has alphabet {left1, left2, right}. It inputs a message x on either left1 or left2, tags it with a 1 or 2 to indicate its source, and outputs the tagged message on the right, after which it behaves like p.

In the initial state of a process described using \square , both the channels involved are ready for input, and all the other channels are unready. Each subsequent state corresponds either to a state of $P(x)$ or to a state of $Q(y)$; and both cases must be proved correct. The proof rule ^(therefore) is modelled on that for simple input.

If c and d are distinct channel names:

$$(c?x:M \rightarrow P(x) \parallel d?y:N \rightarrow Q(y)) \text{ sat } R$$

$$\equiv R[\langle \rangle / \text{past}, M / c.\text{ready}, N / d.\text{ready}, \emptyset / \text{ready}]$$

$$\& \forall x:M. P(x) \text{ sat } R[\langle x \rangle c.\text{past} / c.\text{past}]$$

$$\& \forall y:N. Q(y) \text{ sat } R[\langle y \rangle d.\text{past} / d.\text{past}].$$

Example.

Let $\text{sel}(n, s)$ be a sequence formed from s by selecting only those items tagged with n , and then removing the tags; or, more formally:

$$\text{sel}(n, s) = \begin{cases} \text{if } s = \langle \rangle \text{ then } \langle \rangle \\ \text{else if first}(s) = (n, x) \text{ then } \langle x \rangle \text{ sel}(n, \text{rest}(s)) \\ \text{else } \text{sel}(n, \text{rest}(s)). \end{cases}$$

$$\text{Let MERGED} \triangleq \text{sel}(1, \text{right.past}) \leq \text{left1.past}$$

$$\& \text{sel}(2, \text{right.past}) \leq \text{left2.past}$$

$$\& (\text{left1.ready} = \text{left2.ready} = M \& EQ$$

$$EQ = \text{sel}(1, \text{right.past}) = \text{left1.past} \& \text{sel}(2, \text{right.past}) = \text{left2.past} \vee \text{right.ready} \neq \emptyset) \& \neg EQ$$

$$\text{Theorem 12. MERGESTEP sat (MERGED} \Gamma_{n+1}) \equiv$$

$$\forall x:M. (\text{right}!(1, x) \rightarrow p) \text{ sat (MERGED} \Gamma_{n+1})[\langle x \rangle \text{left1.past} / \text{left1.past}]$$

$$\& \forall x:M (\text{right}!(2, x) \rightarrow p) \text{ sat (MERGED} \Gamma_{n+1})[\langle x \rangle \text{left2.past} / \text{left2.past}]$$

Proof. The omitted terms are trivial theorems.

The reader may care to complete the proof that

$$(\text{np. MERGESTEP}) \text{ sat MERGED.}$$

/// The notation and proof rule for alternation can clearly be adapted for more than two alternatives; and since $\langle c!e \rightarrow P \rangle$ is the same as $\langle c?x:\{e\} \rightarrow P \rangle$, output can be readily

substituted for input.

2.12 General Recursion.

The method of defining processes by recursion can be generalised to allow mutual recursion. A set of processes defined by mutual recursion constitute a solution to a set of simultaneous fixed point equations, just as $\mu p. F(p)$ is a solution for p in the single equations:

$$p \triangleq F(p).$$

A pair of mutually recursive equations take the form

$$p \triangleq F(p, q)$$

$$q \triangleq G(p, q)$$

where $F(p, q)$ and $G(p, q)$ are process descriptions, which in general contain the process variables p and q .

/// The method of mutual recursion generalises even further to infinite sets of simultaneous equations, one for each member s in some counting set S :

$$p(s) = F(p, s) \quad \text{for all } s \text{ in } S.$$

The solutions to all these simultaneous equations constitute an array p , with an element $p(s)$ for each s in S .

This array of processes is denoted by the formula:

$$\mu p(s:S). F(p, s).$$

However, it is often clearer to write the definitions in the equational form shown above

Example.

Let M^* be the set of all finite sequences of elements of M .

Let $IN \triangleq (\text{left? } x: M \rightarrow p(\langle x \rangle))$

Let $INOROUT \triangleq (\text{left? } x: M \rightarrow p(s \langle x \rangle)$
 $\quad \quad \quad \square \text{right! first}(s) \rightarrow p(\text{rest}(s))$
 $\quad \quad \quad)$

Let $STEP \triangleq \text{if } s = \langle \rangle \text{ then } IN \text{ else } INOROUT$

Let $B \triangleq \mu p(s: M^*). STEP$

For each s in M^* , $B(s)$ behaves like an unbounded buffer with current content s . If s is empty, $B(s)$ is prepared only to input on the left any value x in M , and then behave like $B(\langle x \rangle)$, that is, like a buffer containing only the value x . But if s is nonempty, $B(s)$ is prepared: ^{newline} either (1) to input a new element x , which is appended to the stored buffer, so that its subsequent behaviour is $B(s \langle x \rangle)$; or (2) to output the first element of its buffer, which is then removed, so that its subsequent behaviour is $B(\text{rest}(s))$.

The same definition can be written out more clearly in the form of an equation in B :

$B(s) \triangleq \text{if } s = \langle \rangle \text{ then } \text{left? } x: M \rightarrow B(\langle x \rangle)$
 $\quad \quad \quad \text{else } (\text{left? } x: M \rightarrow B(s \langle x \rangle)$

$\quad \quad \quad \square \text{right! first}(s) \rightarrow B(\text{rest}(s))$

$\quad \quad \quad)$ For all s in M^*

/// The proof rule for generalised recursion is similar to that for simple recursion, except that the formulae are quantified over all s in the counting set S .

$$\frac{(\forall s: S. p(s) \text{ sat } (R(s) \Gamma_n) \Rightarrow \forall s: S. F(p, s) \text{ sat } (R(s) \Gamma_{n+1}))}{\forall s: S. ((\mu p(s: S) F(p, s))(s)) \text{ sat } R(s)}$$

Example. Let us define

$$\text{BUFF}(s) \triangleq \text{BUFF}[s \text{ left.past} / \text{left.past}]$$

$\text{BUFF}(s)$ describes the behaviour of a buffer that has input the sequence s , but not yet output it.

$\text{BUFF}(s)$ therefore should describe the future behaviour of the process $B(s)$, as stated in the following theorem.

Theorem 13. $\forall s: S. B(s) \underline{\text{sat}} \text{BUFF}(s)$

Proof.

By the rule of recursion, we can assume

$$\forall s: M^*. p(s) \underline{\text{sat}} (\text{BUFF}(s) \uparrow n) \quad \text{--- (0)}$$

and must prove

$$\text{STEP } \underline{\text{sat}} (\text{BUFF}(s) \uparrow_{n+1}) \quad \text{for } s \in M^*$$

which by the conditional rule, splits in two:

$$s = \langle \rangle \Rightarrow \text{IN } \underline{\text{sat}} (\text{BUFF}(s) \uparrow_{n+1}) \quad \text{--- (1)}$$

$$\text{and } s \neq \langle \rangle \Rightarrow \text{INOROUT } \underline{\text{sat}} (\text{BUFF}(s) \uparrow_{n+1}) \quad \text{--- (2)}$$

For (1), we assume $s = \langle \rangle$ and need to prove

$$(\text{BUFF}(s) \uparrow_{n+1}) [\langle \rangle / \text{past}, M / \text{left.ready}, \phi / \text{ready}] \quad \text{--- (1a)}$$

$$p(\langle \rangle) \underline{\text{sat}} (\text{BUFF}(s) \uparrow_{n+1}) [\langle \rangle \text{left.past} / \text{left.past}] \quad \text{--- (1b)}$$

(1a) is a trivial theorem, and the assertion of (1b) is equivalent to

$$p(\langle \rangle) \underline{\text{BUFF}} [s \langle \rangle \text{left.past} / \text{left.past}] \uparrow n$$

which by definition is $\text{BUFF}(s \langle \rangle) \uparrow n$

So (1b) follows directly from the assumption (0)

and the condition $s = \langle \rangle$.

For (2) we assume $s \neq \langle \rangle$ and need to prove

$$(\text{BUFF}(s) \uparrow_{n+1}) [\langle \rangle / \text{past}, M / \text{left.ready}, \{\text{first}(s)\} / \text{right.ready}, \phi / \text{ready}] \quad (2a)$$

$$\& \forall x: M. p(s \langle x \rangle) \underline{\text{sat}} (\text{BUFF}(s) \uparrow_{n+1}) [\langle x \rangle \text{left.past} / \text{left.past}] \quad (2b)$$

$$\& p(\text{rest}(x)) \underline{\text{sat}} (\text{BUFF}(s) \uparrow_{n+1}) [\langle \text{first}(s) \rangle \text{right.past} / \text{right.past}] \quad (2c)$$

(2a) is a trivial theorem. The assertion of (2b) is equivalent to $\text{BUFF}(s \langle x \rangle) \uparrow n$, and the assertion of (2c) is equivalent to $\text{BUFF}(\text{rest}(s)) \uparrow n$; so both (2b) and (2c) follow from the assumption (0)

/// To check the above claims of trivial theoremhood or equivalence, it is necessary only to expand the abbreviations. For example

$$(BUFF(s) \uparrow_{n+1}) [\langle \text{first}(s) \rangle \text{right.past} / \text{right.past}] \equiv$$

$$\equiv \# \text{left.past} + \# (\langle \text{first}(s) \rangle \text{right.past}) \geq n+1$$

$$\vee (s \text{ left.past}) = \langle \text{first}(s) \rangle \text{right.past} \ \& \ \text{left.ready} = M$$

$$\vee \langle \text{first}(s) \rangle \text{right.past} < (s \text{ left.past})$$

$$\ \& \ \text{right.ready} = \{ \text{first}(s \text{ left.past} - \langle \text{first}(s) \rangle \text{right.past}) \}$$

$$BUFF(\text{rest}(s)) \uparrow_n \equiv$$

$$\# \text{left.past} + \# \text{right.past} \geq n$$

$$\vee (\text{rest}(s) \text{ left.past}) = \text{right.past} \ \& \ \text{left.ready} = M$$

$$\vee \text{right.past} < (\text{rest}(s) \text{ left.past})$$

$$\ \& \ \text{right.ready} = \{ \text{first}((\text{rest}(s) \text{ left.past}) - \text{right.past}) \}$$

When $s \neq \langle \rangle$,
 (these are clearly equivalent, clause by clause.)

Theorem 14. $B(\langle \rangle) \underline{\text{sat}} BUFF$

Proof. Put $s = \langle \rangle$ in Theorem 13.