An alternative order for the failures model

by A.W. Roscoe[1]

Oxford University Computing Laboratory,
8-11 Keble Road, Oxford OX1 3QD, U.K.

0. Introduction. This paper, which is a revised version of one which appeared in [R4], introduces an alternative, coarser, partial order on the (improved) failures model of [BR] (sometimes called the failures-divergences model). The new order gives *exactly* the same semantics to CSP as the old one. As well as being of intrinsic interest for this reason the new order allows one to establish some interesting new results about the semantics of CSP and also (if desired) to extend the model to encompass certain extra forms of unbounded nondeterminism.

The present failures model $\mathcal{N}$, with its explicit treatment of divergence, was independently introduced in [R1,B] to overcome various technical difficulties in the "pure" failures model of [HBR, BHR]. In the earlier model we had based our partial order on nondeterminism

$$P \sqsupseteq Q \;\Leftrightarrow\; P \subseteq Q$$

and so it was natural that this idea should be extended to the improved model:

$$P \sqsupseteq Q \;\Leftrightarrow\; \mathcal{D}[\![P]\!] \subseteq \mathcal{D}[\![Q]\!] \wedge \mathcal{F}[\![P]\!] \subseteq \mathcal{F}[\![Q]\!]\ .$$

This order works very well, of course. $P \sqsupseteq Q$ just when P "improves" Q, or when P is more deterministic than Q. $\sqsupseteq$ is a complete partial order, though when the underlying alphabet of communications is infinite the compactness condition

$$(\forall X \in \wp(Y).(s,X) \in \mathcal{F}[\![P]\!]) \;\Rightarrow\; (s,Y) \in \mathcal{F}[\![P]\!]$$

[1] The author gratefully acknowledges that the work reported in this paper was supported by ONR grant N00014-87-G-0242.

is necessary for this. ($\wp(Y)$ denotes the *finite* subsets of Y.) (This assumption turns out to be somewhat weaker than an assumption of finite nondeterminism.) In this paper we will sometimes refer to this order as the *nondeterminism* order.

There are good intuitive reasons for expecting all CSP operations to be monotonic and continuous with respect to $\sqsupseteq$, and so they are. That the order produces the right semantics for recursion is demonstrated by the congruence of the failures denotational semantics and its operational semantics [BRW] as well as the intuitively attractive idea that the recursive term $\mu P.F(P)$ must, in this order, always denote the most nondeterministic solution to

$$P \;=\; F(P)\;.$$

Summary of definitions and notation. The reader should consult earlier works, particularly [H], for the syntax of CSP and the meaning of its constructs. This section contains a summary of the technical details of the model $\mathcal{N}$ and the semantics of CSP in that model. It assumes a knowledge of standard notations concerning traces.

$\mathcal{N}$ is defined relative to some non-empty alphabet Σ of communications which is fixed for all processes. This parameter will generally be understood rather than mentioned explicitly. Each syntactic CSP process P will be identified with an element of $\mathcal{N}_\Sigma$. Thus we will not use process alphabets (αP) as part of the semantics, unlike [H]. The use of process alphabets produces a trivially isomorphic theory where each process P is a member of $\mathcal{N}_{\alpha P}$, where αP is the process' own alphabet. The advantage of this, essentially typed, theory is that the parallel operator does not require the explicit mention of alphabets. Its disadvantages for theoretical work like that in this paper are that more mathematical housekeeping is required and that the theory of mutual recursion becomes rather messy.

$\mathcal{N}$ consists of all pairs $P = \langle F, D\rangle$ such that $F \subseteq \Sigma^* \times \mathcal{P}(\Sigma)$ and $D \subseteq \Sigma^*$, and such that for s, $t \in \Sigma^*$ and X, $Y \subseteq \Sigma$

$$\begin{array}{lrcl}
(1) & (F \neq \emptyset) & \wedge & ((st,\emptyset) \in F \Rightarrow (s,\emptyset) \in F) \\
(2) & (t,X) \in F \wedge Y \subseteq X & \Rightarrow & (t,Y) \in F \\
(3) & (t,X) \in F \wedge \forall a \in Y.(t\langle a\rangle,\emptyset) \notin F & \Rightarrow & (t, X \cup Y) \in F \\
(4) & (\forall X' \in \wp(X).(t,X') \in F) & \Rightarrow & (t,X) \in F \\
(5) & s \in D & \Rightarrow & st \in D \\
(6) & s \in D & \Rightarrow & (st,X) \in F.
\end{array}$$

Let $P = \langle F, D\rangle \in \mathcal{N}$ be a process. Define

$$\begin{array}{rcl}
\mathcal{F}[\![P]\!] & = & F \\
\mathcal{D}[\![P]\!] & = & D \\
\mathit{traces}(P) & = & \{t \mid (t,\emptyset) \in F\} \\
\mathit{refusals}(P) & = & \{X \mid (\langle\rangle, X) \in F\}
\end{array}$$

and for $s \in \mathit{traces}(P)$ define

$$P\,\mathit{after}\,s = \langle\{(t,B) \mid (st,B) \in F\}, \{t \mid st \in D\}\rangle$$

and for $s \in \Sigma^*$ define

$$\mathcal{R}[\![P]\!]s = \{X \mid (s, X) \in \mathcal{F}[\![P]\!]\}$$

so that $\mathcal{R}[\![P]\!]s = \mathit{refusals}(P \mathbin{\mathit{after}} s)$ when $s \in \mathit{traces}(P)$. For $s \in \mathit{traces}(P)$, $P \mathbin{\mathit{after}} s$ is a process.

The determinism order $\sqsubseteq$ is defined on $\mathcal{N}$ by

$$P \sqsubseteq Q \;\Leftrightarrow\; \mathcal{D}[\![P]\!] \supseteq \mathcal{D}[\![Q]\!] \wedge \mathcal{F}[\![P]\!] \supseteq \mathcal{F}[\![Q]\!].$$

$\langle \mathcal{N}, \sqsubseteq \rangle$ is a complete partial order: $\perp = \langle \Sigma^* \times \mathcal{P}(\Sigma), \Sigma^* \rangle$ is the least element, and if Δ is a directed subset of $\mathcal{N}$ the least upper bound $\bigsqcup \Delta$ is

$$\langle \bigcap \{F \mid \langle F, D \rangle \in \Delta\}, \bigcap \{D \mid \langle F, D \rangle \in \Delta\} \rangle \; .$$

CSP is given a denotational semantics over $\mathcal{N}$ as follows.

The atomic processes STOP and SKIP are defined

$$\mathit{STOP} = \langle \{(\langle\,\rangle, X) \mid X \subseteq \Sigma\}, \emptyset \rangle$$

$$\mathit{SKIP} = \langle \{(\langle\,\rangle, X) \mid \surd \notin X\} \cup \{(\langle \surd \rangle, X) \mid X \subseteq \Sigma\}, \emptyset \rangle \; .$$

Let $P = \langle F, D\rangle$, $P' = \langle F', D'\rangle$ and for $b \in B \subseteq \Sigma$, $P_b = \langle F_b, D_b\rangle$ all be processes. Then

$$
\begin{array}{rcl}
\mathcal{D}[\![a \to P]\!] & = & \{\langle a\rangle s \mid s \in D\} \\
\mathcal{F}[\![a \to P]\!] & = & \{(\langle\,\rangle, X) \mid a \notin X\} \cup \{(\langle a\rangle s, X) \mid (s, X) \in F\} \\
\mathcal{D}[\![x : B \to P_x]\!] & = & \{\langle b\rangle s \mid b \in B \wedge s \in D_b\} \\
\mathcal{F}[\![x : B \to P_x]\!] & = & \{(\langle\,\rangle, X) \mid B \cap X = \emptyset\} \cup \{(\langle b\rangle s, X) \mid b \in B \wedge (s, X) \in F_b\} \\
\mathcal{D}[\![P \sqcap P']\!] & = & D \cup D' \\
\mathcal{F}[\![P \sqcap P']\!] & = & F \cup F' \\
\mathcal{D}[\![P \Box P']\!] & = & D \cup D' \\
\mathcal{F}[\![P \Box P']\!] & = & \{(\langle\,\rangle, X) \mid (\langle\,\rangle, X) \in F \cap F'\} \cup \{(s, X) \mid s \neq \langle\,\rangle \wedge (s, X) \in F \cup F'\} \\
 & & \cup\{(s, X) \mid s \in \mathcal{D}[\![P \Box P']\!]\} \\
\mathcal{D}[\![P\ {}_B\|_C\ P']\!] & = & \{st \mid s \in (B \cup C)^* \wedge s\upharpoonright B \in D \wedge s\upharpoonright C \in \mathit{traces}(P')\} \\
 & & \cup\{st \mid s \in (B \cup C)^* \wedge s\upharpoonright B \in \mathit{traces}(P) \wedge s\upharpoonright C \in D'\} \\
\mathcal{F}[\![P\ {}_B\|_C\ P']\!] & = & \{(s, (X \cap B) \cup (Y \cap C) \cup Z) \mid Z \cap (B \cup C) = \emptyset \wedge s \in (B \cup C)^* \\
 & & \wedge (s\upharpoonright B, X) \in F \wedge (s\upharpoonright C, Y) \in F'\} \\
 & & \cup\{(s, X) \mid s \in \mathcal{D}[\![P\ {}_B\|_C\ P']\!]\} \\
\mathcal{D}[\![P \mid\mid\mid P']\!] & = & \bigcup\{\mathit{merge}\langle s, t\rangle \mid s \in D \wedge t \in \mathit{traces}(P')\} \\
 & & \cup\bigcup\{\mathit{merge}\langle s, t\rangle \mid s \in D' \wedge t \in \mathit{traces}(P)\} \\
\mathcal{F}[\![P \mid\mid\mid P']\!] & = & \{(s, X) \mid \exists t, u.\, s \in \mathit{merge}\langle t, u\rangle \wedge (t, X) \in F \wedge (u, X) \in F'\} \\
 & & \cup\{(s, X) \mid s \in \mathcal{D}[\![P \mid\mid\mid P']\!]\} \\
\mathcal{D}[\![P; P']\!] & = & \{st \mid s \in D \wedge s \text{ tick-free}\} \\
 & & \cup\{st \mid s\langle\surd\rangle \in \mathit{traces}(P) \wedge t \in D' \wedge s \text{ tick-free}\} \\
\mathcal{F}[\![P; P']\!] & = & \{(s, X) \mid (s, X \cup \{\surd\}) \in F \wedge s \text{ tick-free}\} \\
 & & \cup\{(st, X) \mid s\langle\surd\rangle \in \mathit{traces}(P) \wedge s \text{ tick-free} \wedge (t, X) \in F'\} \\
 & & \cup\{(s, X) \mid s \in \mathcal{D}[\![P; P']\!]\} \\
\mathcal{D}[\![P \backslash a]\!] & = & \{(s\backslash a)t \mid s \in D\} \cup \{(s\backslash a)t \mid \forall n.s\langle a\rangle^n \in \mathit{traces}(P)\} \\
\mathcal{F}[\![P \backslash a]\!] & = & \{(s\backslash a, X) \mid (s, X \cup \{a\}) \in F\} \cup \{(s, X) \mid s \in \mathcal{D}[\![P\backslash a]\!]\} \\
\mathcal{D}[\![f[P]]\!] & = & \{(f(s))t \mid s \in D\} \\
\mathcal{F}[\![f[P]]\!] & = & \{(f(s), X) \mid (s, f^{-1}(X)) \in F\} \cup \{(s, X) \mid s \in \mathcal{D}[\![f[P]]\!]\} \\
\mathcal{D}[\![f^{-1}[P]]\!] & = & \{s \mid f(s) \in D\} \\
\mathcal{F}[\![f^{-1}[P]]\!] & = & \{(s, X) \mid (f(s), f(X)) \in F\} \cup \{(s, X) \mid s \in \mathcal{D}[\![f^{-1}[P]]\!]\}
\end{array}
$$

Each of these operations is continuous with respect to $\sqsubseteq$, and so we may define the meaning of a recursion (single or mutual) by the least fixed point in the usual way.

1. The new order. The failures model includes the assumption that once a process *can* diverge (on trace s, say) then we are not interested in what it can do (or fail to do) on any extension of s. This assumption, which we made for various technical and philosophical reasons, essentially corresponds to an assumption that a divergent process is *undefined*. The new order (which we will call the *definedness* order) is based solely on this principle. P will be weaker than Q just when Q's divergences are a subset of P's and all of P's convergent behaviour is copied exactly in Q. Included in P's convergent behaviour are its *minimal* divergence traces, because the process has completed these traces before it can diverge. We demand that these be included among Q's traces but not necessarily among its divergences. If X is any set of traces we define $\mu(X)$ to be the minimal elements of X: $\{s \in X \mid \not\exists t \in X . t < s\}$.

$$\begin{aligned} P \leq Q \;\Leftrightarrow\; & \mathcal{D}[\![Q]\!] \subseteq \mathcal{D}[\![P]\!] \wedge \\ & s \notin \mathcal{D}[\![P]\!] \Rightarrow \mathcal{R}[\![P]\!]s = \mathcal{R}[\![Q]\!]s \wedge \\ & \mu(\mathcal{D}[\![P]\!]) \subseteq \mathit{traces}(Q) \end{aligned}$$

Several points are immediately apparent about this order. First, it is *coarser* than the old one, in that

$$P \leq Q \;\Rightarrow\; P \sqsubseteq Q\ ,$$

and *strictly* so, in that there are many pairs of processes P, Q such that

$$P \sqsubseteq Q \quad \text{but not} \quad P \leq Q\ .$$

Secondly, the divergence-free processes (i.e., the ones with $\mathcal{D}[\![P]\!] = \emptyset$) are all incomparable and maximal in the order. (No process is strictly above a non-divergent one.) Thirdly, the new order has the same least element as the old one ($\perp$, the immediately divergent process).

The new order is complete and, where appropriate, has identical least upper bounds for directed sets as the old one. These and some other useful facts about the new order are established in the next Lemma.

Lemma 1.1

a) $P \leq Q \Rightarrow P \sqsubseteq Q$

b) $\perp = \langle \Sigma^* \times \mathcal{P}(\Sigma), \Sigma^* \rangle$ is the least element of $\mathcal{N}$ for both orders. The $\leq$-maximal elements are the divergence-free processes.

c) If $P \leq R$ and $P \sqsubseteq Q \sqsubseteq R$, then $P \leq Q$.

d) Any nonempty subset S of $\mathcal{N}$ has greatest lower bounds with respect to both $\leq$ and $\sqsubseteq$. In general, $\sqcap_{\leq} S \sqsubseteq \sqcap_{\sqsubseteq} S$.

e) Any subset of $\mathcal{N}$ with any $\leq$-upper bound has a least upper bound.

f) Each $\leq$-directed set has a least upper bound.

g) If $\bigsqcup_{\leq} S$ is defined then so is $\bigsqcup_{\sqsubseteq} S$ and the two are equal. Furthermore $\bigsqcup_{\leq} S = P^* = \langle F^*, D^* \rangle$, where $F^* = \bigcap\{F \mid \langle F, D \rangle \in S\}$ and $D^* = \bigcap\{D \mid \langle F, D \rangle \in S\}$.

Proof. (a) and (b) and are trivial. For (c), we observe that $P \leq Q$ if and only if $P \sqsubseteq Q$ and

(i) $(s, X) \in \mathcal{F}[\![P]\!] \wedge s \notin \mathcal{D}[\![P]\!] \Rightarrow (s, X) \in \mathcal{F}[\![Q]\!]$, and

(ii) $\mu(\mathcal{D}[\![P]\!]) \subseteq \mathit{traces}(Q)$, so to prove the result it will be sufficient to prove (i) and (ii). If $(s,X) \in \mathcal{F}[\![P]\!] \wedge s \notin \mathcal{D}[\![P]\!]$ then, since $P \leq R$, we know $(s,X) \in \mathcal{F}[\![R]\!]$. Hence $(s,X) \in \mathcal{F}[\![Q]\!]$ as $Q \sqsubseteq R$. Exactly the same argument applies for (ii).

The $\sqsubseteq$-greatest lower bound of nonempty $S \subseteq \mathcal{N}$ is always given by $\langle \overline{F^*}, D^* \rangle$, where $F^* = \bigcup\{F \mid \langle F, D\rangle \in S\}$, $D^* = \bigcup\{D \mid \langle F, D\rangle \in S\}$ and where $\overline{F}$ is the closure of F with respect to the compactness axiom (4) (i.e., a pair (s,X) is included if all (s,Y) are for finite subsets Y of X). The application of this closure operator is not necessary when S or Σ is finite. This bound does not work in general for the definedness order $\leq$, however, since one does not in general have $P \in S \Rightarrow \bigsqcap S \leq P$. The greatest lower bound of $S = \{\langle F_i, D_i\rangle \mid i \in \Lambda\}$ is constructed so that it diverges as soon as the behaviour of any two elements of S starts to differ, either by having a different refusal set or a different next communication. We define $\bigsqcap_{\leq} S$ to be $\langle F, D\rangle$, where

- $D = \bigcup\{D_i \mid i \in \Lambda\} \cup \{st \mid \exists i, j.(\exists Y.(s,Y) \in F_i \setminus F_j) \vee (\exists a.(s\langle a\rangle, \emptyset) \in F_i \setminus F_j)\}$
- $F = \bigcup\{F_i \mid i \in \Lambda\} \cup \{(s,X) \mid s \in D\}$

It is easy to show that this process is in $\mathcal{N}$ and is indeed the $\leq$ greatest lower bound of S. That this greatest lower bound is $\sqsubseteq$-less than the other one follows trivially from the fact that $\leq$ is coarser than $\sqsubseteq$.

(e) follows because, as is fairly well-known, *any* partial order which has greatest lower bounds for nonempty sets has this property. The usual argument is repeated here. If S is a set with an upper bound, then U_S, the set of upper bounds of S is nonempty and so $x = \bigsqcap U_S$ exists. Since $y \leq z$ whenever $y \in S$ and $z \in U_S$ it follows that each $y \in S$ is a lower bound for U_S. As x is the *greatest* lower bound for U_S it follows that $x \geq y$ for all $y \in S$ and therefore that $x \in U_S$. Plainly x is the least element of U_S and is therefore the least upper bound of S.

We now turn to the proof of (f). Notice that if Δ is a $\leq$-directed set it is also $\sqsubseteq$-directed. Its $\sqsubseteq$-least upper bound

$$\bigsqcup \Delta \;=\; \langle \bigcap\{\mathcal{F}[\![P]\!] \mid P \in \Delta\}, \bigcap\{\mathcal{D}[\![P]\!] \mid P \in \Delta\}\rangle$$

is now shown to be an $\leq$-upper bound. If $P \in \Delta$ then $\mathcal{D}[\![P]\!] \supseteq \mathcal{D}[\![\bigsqcup \Delta]\!]$ by the properties of $\sqsubseteq$. If $s \notin \mathcal{D}[\![P]\!]$ and $Q \in \Delta$, it follows from the existence of R such that $R \geq P$ and $R \geq Q$ that $s \in \mathcal{D}[\![Q]\!]$ or $\mathcal{R}[\![P]\!]s = \mathcal{R}[\![Q]\!]s$. Hence $\mathcal{R}[\![\bigsqcup \Delta]\!]s = \mathcal{R}[\![P]\!]s$. Finally, if s is minimal in $\mathcal{D}[\![P]\!]$ and $Q \in \Delta$ it again follows from the existence of $R \geq P, Q$ that $s \in \mathit{traces}(Q)$. This shows that $s \in \mathit{traces}(\bigsqcup \Delta)$, which completes the proof that $P \leq \bigsqcup \Delta$. Obviously $\bigsqcup \Delta$ is the $\sqsubseteq$-least $\leq$-upper bound.

It follows by part (e) from the fact that Δ has one $\leq$-upper bound that it has a least one $\bigsqcup_{\leq} \Delta$. Since $\leq$ is coarser than $\sqsubseteq$ it follows that $\bigsqcup_{\leq} \Delta \sqsubseteq \bigsqcup \Delta$. Combining this with the last observation in the previous paragraph proves that the two are the same, which is what we wanted.

The first sentence of (g) follows easily from the second, which is what we prove. We show first that if $P' = \langle F', D'\rangle$ is the actual least upper bound on S then $D^* = D'$, where D^* is as defined in the statement of the Lemma. Trivially $D' \subseteq D^*$ and, since D' satisfies axiom (5), if $D' \neq D^*$

there is $s \in (\mu D^*) \setminus D'$. As $s \in \mu D^*$ there must be some $P_1 = \langle F_1, D_1 \rangle \in S$ such that $s \in \mu D_1$. Since $P_1 \leq P'$ we therefore know that $s \in \mathit{traces}(P')$. It follows that $P'' = \langle F'', D'' \rangle$ defined by

$$\begin{array}{rcl} F'' & = & F' \cup \{(st, X) \mid t \in \Sigma^* \wedge X \subseteq \Sigma\} \\ D'' & = & D' \cup \{st \mid t \in \Sigma^*\} \end{array}$$

is a process. But it easy to show that $P \leq P''$ for all $P \in S$ (for example, by part (c) above) and that $P'' < P'$. It follows that P' cannot be the least upper bound on S, a contradiction.

It is easy to show that P^* defined in the statement of the Lemma satisfies axioms (1), (2), (5) and (6). We next note that $F^* \supseteq F'$ since F' is a $\sqsubseteq$-upper bound for S. Now by the above paragraph those parts of F^* and F' implied by divergence and axiom (5) are equal. Suppose that $s \notin D' = D^*$. Then there is $P = \langle F, D \rangle \in S$ such that $s \notin D$. Necessarily $\mathcal{R}[\![P]\!]s = \mathcal{R}[\![P']\!]s$ as $P \leq P'$. It follows that $\mathcal{R}[\![P']\!]s \supseteq \mathcal{R}[\![P^*]\!]s$. Putting these facts together yields $F' \supseteq F^*$, proving that in fact $F' = F^*$. Since $D' = D^*$ and $F' = F^*$ we have thus proved that P^* is the actual least upper bound, as claimed. This completes the proof of the Lemma. $\Box$

It is easy to construct functions which are monotonic in either of our two orders without being monotonic in the other. However all the usual CSP constructs, with their standard definitions over $\mathcal{N}$, turn out to be monotonic in the new order as well as the old. Of course one can prove each case separately (and easily), but this is unenlightening. The underlying reason why they all work is described in the next paragraph.

With the exception of recursion, a special case which will be dealt with later, every CSP operator over $\mathcal{N}$ is defined by mapping the behaviours (failures and divergences) of the argument process(es) to the behaviours to which they correspond in the image. An examination of these definitions will reveal that, for each operator F, all convergent behaviours of $F(P)$ (its behaviours on non-divergent traces plus the facts that its minimal divergences are traces) are consequences of the convergent behaviours of P. Hence, if $P \leq Q$ and $s \notin \mathcal{D}[\![F(P)]\!]$, $\mathcal{R}[\![F(P)]\!]s$ is derived only from the non-divergent behaviours of P, each of which is also possible for Q by definition of $\leq$. Thus

$$\mathcal{R}[\![F(P)]\!]s \ \subseteq \mathcal{R}[\![F(Q)]\!]s \ .$$

The reverse inclusion follows from $\sqsubseteq$-monotonicity, as does

$$\mathcal{D}[\![F(P)]\!] \ \supseteq \ \mathcal{D}[\![F(Q)]\!] \ .$$

If s is minimal in $\mathcal{D}[\![F(P)]\!]$ then the fact that it is a trace of $F(P)$ is derived from some convergent behaviour of P, necessarily present also in Q; hence s is a trace of $F(Q)$, completing the proof that $F(P) \leq F(Q)$.

The most interesting operator to look at in conjunction with the preceding paragraph is sequential composition (;). This is the only operator where a minimal divergence trace of an argument can contribute to a non-divergent trace of the result. This happens in $P; Q$ if P can diverge immediately after terminating ($\surd$). The reader who inspects the semantics given earlier will be able to show that, for example $(\surd \rightarrow \ \bot); P = P$ for all P. Operationally, the fact that the left-hand process would diverge after $\surd$ does not matter since the occurrence of $\surd$ means that it can be 'turned off'.

This non-dependence of the convergent behaviour of $F(P)$ upon the divergent behaviour of P is closely related to a property of the operational semantics. The first-step behaviour of $F(P)$

may depend on the behaviour of P or not (the latter is the case with the prefixing operators, for example). However, if it does and P can perform an internal action to become, say, P', $F(P)$ can perform an internal action and become $F(P')$. Thus, when an operated-on process is active, its internal actions occur independently of F: they are outside its control. An immediate consequence is that, if F has brought P to a point where it can diverge but is still interested in what it can do, then $F(P)$ can diverge also.

When a function is $\leq$-monotonic and $\sqsubseteq$-continuous, it is $\leq$-continuous. For if $\mathcal{D}$ is $\leq$-directed, then $\{F(P) \mid P \in \mathcal{D}\}$ is also, with least upper bound $\bigsqcup\{F(P) \mid P \in \mathcal{D}\}$ as described above. However, since F is $\sqsubseteq$-continuous and $\mathcal{D}$ is necessarily $\sqsubseteq$-directed, we have

$$F(\bigsqcup \mathcal{D}) \;=\; \bigsqcup\{F(P) \mid P \in \mathcal{D}\}\,.$$

This completes the proof as the least upper bound of $\mathcal{D}$ is the same in both orders.

This means that all CSP constructs other than recursion are continuous in the new order, and so we may deal with recursion using least fixed points in the usual way. As is well known, this makes recursion itself a continuous operation in that if $F(P, Q)$ is continuous, then $\mu P.F(P, Q)$ represents a continuous function of Q. Therefore all CSP terms represent continuous functions of their free process variables.

The fact that the two orders yield the same semantics for CSP (i.e., ascribe the same value in $\mathcal{N}$ to each term) follows very easily by structural induction once we observe that the only place where the orders are used is recursion, and the two orders yield identical least upper bounds to all sequences that can appear there. (Of course, the fact that the two orders have the same bottom is also used here.)

2. Consequences for proof rules. The existence of the definedness order has a number of striking consequences for proofs by recursion induction and unique fixed point rules. I had previously proved some of the following results from the theorem (of [BRW]) expressing the congruence of the operational and denotational semantics for CSP. These earlier proofs, though interesting, were far more difficult and less natural than the following.

Theorem 2.1. Suppose the recursive CSP term $\mu P.F(P)$ is divergence-free (i.e., the least solution to $P = F(P)$ has $\mathcal{D}[\![P]\!] = \emptyset$). Then it is the *only* solution to $P = F(P)$.

Proof. We know that any solution Q to this equation must satisfy

$$Q \;\geq\; \mu P.F(P)$$

but, being divergence-free, $\mu P.F(P)$ is maximal in $\mathcal{N}$ under $\leq$. Hence $Q = \mu P.F(P)$. $\Box$

Of course, this result is equally true of mutual recursions. Indeed, one does not need to know that all the mutually defined processes are divergence-free before one can draw useful conclusions. A mutual recursion can be identified with a function from the product space $\mathcal{N}^\Lambda$ to itself, where Λ is an indexing set with one element for each process being simultaneously defined. Λ might be infinite, for example $\mathbf{N}$ or T^*, or finite. Elements of $\mathcal{N}^\Lambda$ are written in vector notation: $\underline{P}$ is a typical element of $\mathcal{N}^\Lambda$ and the λ-component of $\underline{P}$ is P_λ.

Theorem 2.2. Suppose $\underline{P} = F(\underline{P})$ is a CSP mutual recursion indexed by some set Λ, and that in the least solution the λ-component is divergence-free. Then all solutions have identical λ-components.

Proof. If $\underline{P}_\lambda$ is divergence-free then

$$\underline{P} \le \underline{P'} \Rightarrow P_\lambda = P'_\lambda$$

where $\le$ has been extended to the product space in the usual way (i.e., co-ordinatewise). □

The above results are useful, for they allow one to extend the use of the unique fixed point rule to any recursive definition which is known to have a divergence-free least solution. One application is to the analysis of networks where internal communication is hidden. These can often be proved divergence-free by specific techniques. In certain circumstances one can show that the semantic value of the network is the same as the appropriate component of the solution to the mutual recursion over its state space obtained by applying suitable "expansion theorems". The above theorems show that this "state-recursion" has only one solution, which is extremely useful when proving equality between the (state-space of) the network and some other system. More details of these applications will be found in [R3].

The rest of the results of this section all have the flavour of the above Theorem: extending to general divergence-free recursions results which were already known for constructive recursions. (The definition of a constructive recursion will be given shortly.) They are all couched in terms of functions that are $\le$-monotonic or $\le$-continuous rather than directly as theorems about CSP, for reasons that will become apparent in the next section.

Lemma 2.3. Suppose F is a continuous function from $\mathcal{N}^\Lambda$ to itself, and that in the least solution to $\underline{P} = F(\underline{P})$ the λ-component P_λ is divergence-free. Then for each trace s there exists a natural number n such for *any* $\underline{Q} \in P^\Lambda$ and $m \ge n$,

$$s \notin \mathcal{D}[\![F^m(\underline{Q})_\lambda]\!] \quad \text{and} \quad \mathcal{R}[\![F^m(\underline{Q})_\lambda]\!]s = \mathcal{R}[\![P_\lambda]\!]s\,.$$

Proof. Suppose $\underline{Q}$ and s are as above. We know that the least solution is $\bigsqcup\{F^n(\perp^\Lambda) \mid n \in \mathbf{N}\}$. Thus there is some n such that $s \notin \mathcal{D}[\![F^n(\perp^\Lambda)_\lambda]\!]$.

An easy induction shows that $F^k(\perp^\Lambda) \le F^k(\underline{Q})$ for all k, and that

$$F^n(\perp^\Lambda) \le F^m(\underline{Q}) \quad \text{when } n \le m.$$

The definition of $\le$ and the fact that $F^n(\perp^\Lambda)_\lambda \le P_\lambda$ then easily gives

$$s \notin \mathcal{D}[\![F^m(\underline{Q})_\lambda]\!] \quad \text{and} \quad \mathcal{R}[\![F^m(\underline{Q})_\lambda]\!]s = \mathcal{R}[\![F^n(\perp^\Lambda)_\lambda]\!]s = \mathcal{R}[\![P_\lambda]\!]s$$

for $m \ge n$, which establishes the Lemma. □

For a number of years I have been attracted to fixed point induction rules of the following form.

Proforma rule. If R is a predicate on (vectors of) processes which is satisfiable (i.e., $R(P)$ holds for some P) and satisfies some "continuity" condition, and $\mu P.F(P)$ is a recursive definition, satisfying some "well-definedness" condition, such that

$$\forall P.(R(P) \Rightarrow R(F(P)))$$

then infer $R(\mu P.F(P))$. $\Box$

These have been discussed before in, for example, [R1,BHR,Re,RR1,RR2,R3]. Two slots need to be filled in: the conditions on predicates and recursions. In valid versions of the rule one naturally finds that the stronger one condition is, the weaker the other needs to be.

Up to now, by far the most useful versions of the rule have been based on cases where F is a contraction map and $\{P \mid R(P)\}$ a closed subset in some complete metric space. This applies, for example, over untimed models like $\mathcal{N}$ when F is constructive, that is

$$\forall P.\forall n.F(P \downarrow n) \downarrow n+1 \;=\; F(P) \downarrow n+1$$

and R satisfies

$$(\alpha) \qquad\qquad \neg R(P) \;\Rightarrow\; (\exists n.\forall Q.\, Q \downarrow n = P \downarrow n \Rightarrow \neg R(Q))$$

where $P \downarrow n$ represents the n-step behaviour of P. That is, over $\mathcal{N}$,

$$\mathcal{D}[\![P \downarrow n]\!] \;=\; \mathcal{D}[\![P]\!] \cup \{st \mid |s| = n \wedge s \in \mathit{traces}(P)\}$$

$$\mathcal{F}[\![P \downarrow n]\!] \;=\; \mathcal{F}[\![P]\!] \cup \{(s, X) \mid s \in \mathcal{D}[\![P \downarrow n]\!]$$

Over product spaces $\mathcal{N}^{\Lambda}$, $P \downarrow n$ is taken co-ordinatewise. See [BHR] for more details. Over the timed models of [RR1,RR2,Re] F is an arbitrary recursion (for there all recursions represent contraction maps) and R represents a closed set in the metric spaces used as models.

The proof of validity in these cases is very simple. We know, by the contraction mapping theorem, that F has a unique fixed point in the space of all processes; but, as F preserves R it is also a contraction mapping on the complete metric space $\{P \mid R(P)\}$. As F has *a* fixed point in the subspace, the *unique* fixed point in the whole space must lie in the subspace.

On the assumption that equality with a given process is to be an acceptable predicate R, it is clear that any acceptable "well-definedness" condition on recursions must at least imply uniqueness of fixed points. In the context of the *traces* model for CSP I was able to establish (in [R1]) a version of the rule which relied only upon unique fixed points, and not on anything more concrete, even continuity of F. Over any partial order we can define the *interval topology* to be the smallest topology in which all the "closed intervals" $[x, y] = \{z \mid x \le z \le y\}$ are closed sets. Now the interval topology over the traces model is compact (a very common property for models of computation – see the appendix) so that the following theorem applies. (For one topology, $\mathcal{T}$, to *extend* another, $\mathcal{U}$, simply means every set open (closed) in $\mathcal{U}$ is open (closed) in $\mathcal{T}$. If $\mathcal{T}$ extends $\mathcal{U}$ one sometimes says that $\mathcal{T}$ is *finer* than $\mathcal{U}$. If, in this case, $\mathcal{T}$ is compact then so must be $\mathcal{U}$: hence the conditions of the following theorem require the compactness of the interval topology. Fortunately this topology is always compact under the conditions of the theorem, as is proved in the appendix.)

Theorem 2.4 [R1]. Suppose $\langle X, \le \rangle$ is a complete lattice and that $\mathcal{T}$ is a compact topology on X which extends the interval topology. Suppose also that the monotonic function $f : X \to X$ has a unique fixed point x. Then if $Y \subseteq X$ is nonempty, closed in $\mathcal{T}$ and such that $f(Y) \subseteq Y$, we can deduce $x \in Y$.

Proof. A straightforward transfinite induction establishes that each of the sets $[f^{\alpha}(\bot), f^{\alpha}(\top)] \cap Y$ (all of which are closed) are all non-empty, where $\bot$ and $\top$ are respectively the least and greatest elements of X. (The limit ordinal case uses the compactness of $\mathcal{T}$.) But the fact that f has a unique fixed point implies that $[f^\alpha(\bot), f^\alpha(\top)]$ is eventually (i.e., for large enough α) the singleton set containing the fixed point. The result follows immediately. $\Box$

This proof depends crucially on the existence of a top element. Indeed the theorem is not true of more general partial orders. One need only consider the flat truth value domain $\{true, false, \bot\}$ and the continuous function $\neg$. This has unique fixed point $\bot$ and maps the closed set $\{true, false\}$ into itself.

Having observed this it once seemed to me to be unlikely that one could get an analogue of Theorem 2.4 for the failures model. However, the new order will allow us to prove two useful partial analogues, at least. They both have as preconditions the fact that the unique fixed point is divergence-free, which is unlikely to be much of a practical handicap since almost all processes one will tolerate in practice will satisfy this constraint.

First we define three more conditions on predicates.

(β) If $\neg R(P)$ then there is a finite set T of traces such that, for any Q, if

$$s \in \mathcal{D}[\![P]\!] \Leftrightarrow s \in \mathcal{D}[\![Q]\!] \quad \text{and} \quad \mathcal{R}[\![P]\!]s = \mathcal{R}[\![Q]\!]s$$

for all $s \in T$, then $\neg R(Q)$.

(γ) If $\neg R(P)$ then there are finite sets T of traces and F of failures such that, for any Q, if

$$s \in \mathcal{D}[\![P]\!] \Leftrightarrow s \in \mathcal{D}[\![Q]\!]$$

for all $s \in T$ and

$$(s, X) \in \mathcal{F}[\![P]\!] \Leftrightarrow (s, X) \in \mathcal{F}[\![Q]\!]$$

for all $(s, X) \in F$ then $\neg R(Q)$.

(δ) If $\neg R(P)$ then there are finite sets T of traces and F of *finite* failures (failures (s, X) with X finite) such that, for any Q, if

$$s \in \mathcal{D}[\![P]\!] \Leftrightarrow s \in \mathcal{D}[\![Q]\!]$$

for all $s \in T$ and

$$(s, X) \in \mathcal{F}[\![P]\!] \Leftrightarrow (s, X) \in \mathcal{F}[\![Q]\!]$$

for all $(s, X) \in F$ then $\neg R(Q)$.

These conditions simply say that any failure to satisfy the predicate is detectable from a process' behaviour on a finite set of traces, or on a finite set of possible divergences and failures, or on a finite set of possible divergences and finite failures. (α) says that failure is detectable by some finite length of trace. (In each case the sets or length can depend on the failing process.)

The definitions of the conditions β, γ, δ can all be extended to predicates of product spaces $\mathcal{N}^\Lambda$. One simply insists that failure is detectable from a finite set (or sets) as before from a finite selection of components of the vector of processes. This is different from the extension of α to product spaces, where failure need only be detectable from *all* components up to some fixed length.

(β) If $\neg R(\underline{P})$ then there are finite sets T of traces and Φ of indices such that, for any $\underline{Q}$, if

$$s \in \mathcal{D}[\![P_\lambda]\!] \Leftrightarrow s \in \mathcal{D}[\![Q_\lambda]\!] \text{ and } \mathcal{R}[\![P_\lambda]\!]s = \mathcal{R}[\![Q_\lambda]\!]s$$

for all $s \in T$ and $\lambda \in \Phi$, then $\neg R(\underline{Q})$.

(γ) If $\neg R(\underline{P})$ then there are finite sets Φ of indices, T of traces and F of failures such that, for any $\underline{Q}$, if whenever $\lambda \in \Phi$

$$s \in \mathcal{D}[\![P_\lambda]\!] \Leftrightarrow s \in \mathcal{D}[\![Q_\lambda]\!]$$

for all $s \in T$ and

$$(s, X) \in \mathcal{F}[\![P_\lambda]\!] \Leftrightarrow (s, X) \in \mathcal{F}[\![Q_\lambda]\!]$$

for all $(s, X) \in F$ then $\neg R(\underline{Q})$.

(δ) If $\neg R(\underline{P})$ then there are finite sets Φ of indices, T of traces and F of finite failures such that, for any $\underline{Q}$, if whenever $\lambda \in \Phi$

$$s \in \mathcal{D}[\![P_\lambda]\!] \Leftrightarrow s \in \mathcal{D}[\![Q_\lambda]\!]$$

for all $s \in T$ and

$$(s, X) \in \mathcal{F}[\![P_\lambda]\!] \Leftrightarrow (s, X) \in \mathcal{F}[\![Q_\lambda]\!]$$

for all $(s, X) \in F$ then $\neg R(\underline{Q})$.

The conditions α, β, γ, δ are clearly successively stronger: for example every β-predicate is an α-predicate. Each of them generates a topology on $\mathcal{N}$ or $\mathcal{N}^\Lambda$ where a set is closed precisely when it has the form $\{P \mid R(P)\}$ for an allowable predicate R. Several later results depend on facts about these topologies, but rather than analyse them in detail here we will state results as required and relegate proofs to the appendix. The following proposition lists some of these.

Proposition 2.5.

a) The conditions β, γ and δ are equivalent if the underlying alphabet Σ is finite. If it is infinite δ is strictly stronger than γ and γ is strictly stronger than β. α is equivalent to β if the alphabet is finite and the index-set Λ is finite. Otherwise β is strictly stronger than α.

b) For all $\xi \in \{\beta, \gamma, \delta\}$ the ξ-topology on $\mathcal{N}^\Lambda$ defined above is well-defined. The topology generated by ξ on $\mathcal{N}^\Lambda$ is the same as the product topology where each copy of $\mathcal{N}$ is given the topology generated by ξ.

c) The δ-topology on $\mathcal{N}^\Lambda$ is compact, Hausdorff and zero-dimensional.

d) If $\underline{P} \in \mathcal{N}^\Lambda$, each of the sets $\{\underline{Q} \mid \underline{Q} \geq \underline{P}\}$, $\{\underline{Q} \mid \underline{Q} \leq \underline{P}\}$, $\{\underline{Q} \mid \underline{Q} \sqsupseteq \underline{P}\}$ and $\{\underline{Q} \mid \underline{Q} \sqsubseteq \underline{P}\}$ are closed in the δ-topology (and hence in the others as well). $\Box$

Having made all these definitions and stated the above proposition we are now in a position to fulfil the promise made earlier and prove some analogues of Theorem 2.4. The *lower topology* on a partial order is defined to be the smallest topology on which all sets of the form $x\!\uparrow = \{y \mid x \leq y\}$ are closed (see the appendix for more details).

Theorem 2.6. Let X be a complete partial order and let $\mathcal{T}$ be a compact extension of its lower topology. Then, if $f : X \to X$ is a monotonic function whose least fixed point x is maximal (and hence unique) and Y, closed in $\mathcal{T}$, is such that $f(Y) \subseteq Y$, we have $x \in Y$.

Proof. This is very similar to that of Theorem 2.4. Transfinite induction establishes that the (closed) sets $(f^\alpha(\bot))\!\uparrow \cap Y$ are all non-empty, where $\bot$ is the least element of X. But the fact that the least fixed point of f is maximal implies that $(f^\alpha(\bot))\!\uparrow$ eventually (i.e., for large enough α) contains only the fixed point. The result follows immediately. $\Box$

Corollary. If $F : \mathcal{N}^\Lambda \to \mathcal{N}^\Lambda$ is monotonic with a divergence-free least fixed point $\underline{P}$ and R is a satisfiable δ-predicate such that $\forall \underline{Q}.R(\underline{Q}) \Rightarrow R(F(\underline{Q}))$ then $R(\underline{P})$.

Proof. This follows trivially from Theorem 2.6 and Proposition 2.5 (c) and (d). $\Box$

This corollary shows that the proforma rule is valid in any case where the least fixed point can be shown to be divergence free and the predicate satisfies δ. This and the other version given below show that we can do without constructiveness for many applications.

When the function F is continuous one can, because of Lemma 2.3, weaken δ to β.

Theorem 2.7. If $F : \mathcal{N}^\Lambda \to \mathcal{N}^\Lambda$ is continuous with a divergence-free least fixed point $\underline{P}$ and R is a satisfiable β-predicate such that $\forall \underline{Q}.(R(\underline{Q}) \Rightarrow R(F(\underline{Q})))$ then $R(\underline{P})$.

Proof. If $R(\underline{P})$ did not hold then we could find a finite sets Φ of indices and T of traces such that for any $\underline{Q}$, if for all $\langle \lambda, s\rangle \in \Phi \times T$ we had $\mathcal{R}[\![P_\lambda]\!]s = \mathcal{R}[\![Q_\lambda]\!]s$ and $s \notin \mathcal{D}[\![Q_\lambda]\!]$, then $R(\underline{Q})$ would not hold. And we know that there exists $\underline{Q} \in \mathcal{N}^\Lambda$ such that $R(\underline{Q})$, and hence $R(F^n(\underline{Q}))$ for all n. However, choosing n to be greater than all the ns chosen by Lemma 2.3 for the $\langle \lambda, s\rangle \in \Phi \times T$, this gives a contradiction. $\Box$

There is a striking similarity between the proof of Theorem 2.7 and the corresponding result for metric spaces and contraction mappings discussed above. Both rely on the fact that iterations begun from an arbitrary point in the space (the process known to satisfy the predicate) converge to the fixed point. For a consequence of Lemma 2.3 is that, if a continuous function F has a divergence-free least fixed point, then every sequence of the form $\langle F^n(Q) \mid n \in \mathbf{N}\rangle$ converges to the fixed point in the β-topology.

It is mathematically pleasing to have characterisations of when the recursion induction rule is valid that are independent of a specific metric. Even though the great majority of recursions included in CSP programs are constructive, it proves useful to have these abstract forms of the rule when the function is derived from the process and not the other way round: this idea is developed in [R3].

3. Unbounded nondeterminism; hiding infinite sets. As stated in the introduction, the completeness of the determinism order $\sqsubseteq$ depends on the compactness axiom (4). As an illustration of this, consider the "processes"

$$P_n = \prod_{n \leq m} m \rightarrow STOP$$

where $\prod$ represents the nondeterministic choice over a possibly infinite set of indexed processes. The natural interpretation of P_n as a set of failures (assuming its alphabet is $\mathbf{N}$) is

$$\{(\langle\rangle, X) \mid \{n, n+1, n+2, \ldots\} \not\subseteq X\} \cup \{(\langle m \rangle, X) \mid m \geq n\}$$

which does not satisfy the compactness axiom. If these processes were allowed we would, as one would expect from their definitions, have $P_n \sqsubseteq P_{n+1}$. However the chain $\langle P_n \mid n \in \mathbf{N}\rangle$ would have no upper bound, for if $P \sqsupseteq P_n$ for all n then P can neither communicate on its first step nor refuse the whole alphabet – a clear contradiction to axiom (3).

The compactness axiom is not as restrictive as one of finite nondeterminism. The latter would state that, for each s, $\mathcal{R}[\![P]\!]s$ contains a finite number of maximal elements, with every $Y \in \mathcal{R}[\![P]\!]s$ contained in one of them. Axiom (4) allows infinite nondeterminism provided it is finitely presented in some sense, for example it allows sets of refusals such as

$$\{X \subseteq \mathbf{N} \mid \forall i.\{2i, 2i+1\} \not\subseteq X\}$$

where there are uncountably many maximal elements. However, as we see from the processes P_i above it prevents us from expressing many natural processes properly.

It is clear that if one wants to model unbounded nondeterminism properly, then we are far better off without this compactness axiom. Fortunately the definedness order $\leq$ does not need it for completeness, and so it may be discarded. (Note that the P_i as defined above are not comparable under $\leq$.) I should remark at this point that an alternative, slightly simpler, form of axiom (3) which is sometimes seen

$$(3') \qquad (s, X) \in \mathcal{F}[\![P]\!] \wedge (s\langle a \rangle, \emptyset) \notin \mathcal{F}[\![P]\!] \Rightarrow (s, X \cup \{a\}) \in \mathcal{F}[\![P]\!]$$

requires axiom (4) in the sense that $(3') + (4) \equiv (3) + (4)$ but $(3')$ alone is strictly weaker than (3). $(3')$ allows such "processes" as $P = \langle \{(<>, X) \mid X \subseteq \mathbf{N} \wedge X \text{ is finite}\}, \emptyset\rangle$ (which cannot refuse the whole set of integers or communicate any). This P is not acceptable as $P\backslash\mathbf{N}$ would not have any traces at all. (See the definition of the infinite hiding operator below.) Thus $(3')$ is not adequate if axiom (4) is omitted.

We will denote by $\mathcal{N}'_\Sigma$ the failures model defined without axiom (4) over some alphabet Σ. (As before, we will usually suppress the subscript.) Of course $\mathcal{N}'$ only differs from $\mathcal{N}$ when Σ is infinite. The orders $\sqsubseteq$ and $\leq$ are defined exactly as before. $\leq$ is complete, but if Σ is infinite $\sqsubseteq$ is not (as is illustrated by the example above).

All parts of Lemma 1.1 continue to hold in $\mathcal{N}'$ except that it is now unnecessary (and wrong) to apply the closure operator $\overline{F}$ when finding the $\sqsubseteq$-greatest lower bound of a nonempty set. It is interesting to note that, over $\mathcal{N}'$, $\sqsubseteq$-greatest lower bounds are found by union and $\leq$-least upper bounds (where they exist) are found by intersection.

All the operators that were used over $\mathcal{N}$ (and defined earlier) may be defined over $\mathcal{N}'$ using exactly the same clauses. In proving the operators well-defined over $\mathcal{N}$, axiom (4) is never used except when proving axiom (4). It follows that they are all well-defined over $\mathcal{N}'$. We now prove that they are all monotonic and continuous.

If $P \in \mathcal{N}'$, define $P^{\#}$ as follows:

$$\mathcal{D}[\![P^{\#}]\!] \;=\; \mathcal{D}[\![P]\!]$$

$$\mathcal{F}[\![P^{\#}]\!] \;=\; \mathit{traces}(P) \times \mathcal{P}(\Sigma).$$

This is a process with the same traces as P but which can refuse anything at any time. Clearly $P^{\#} \in \mathcal{N}$ and the function thus represented from $\langle \mathcal{N}', \leq \rangle$ to $\langle \mathcal{N}, \sqsubseteq \rangle$ is continuous.

The fact that all operators other than recursion are $\sqsubseteq$-monotonic follows because in all cases the behaviour of the result of an operator depends in a direct and positive way on the behaviour of its arguments. The proof of $\leq$-monotonicity is completed by exactly the same argument we used earlier over $\mathcal{N}$.

The operators' continuity can be proved by using the # operator. If F is one of the operations and $\Delta \subseteq \mathcal{N}'$ is directed, we know by a standard argument using monotonicity that

$$\bigsqcup \{F(P) \mid P \in \Delta\} \leq F(\bigsqcup \Delta)\,.$$

We also have, since # is continuous from $\mathcal{N}'$ to $\mathcal{N}$ and F is continuous over $\mathcal{N}$:

$$\bigsqcup \{F(P^{\#}) \mid P \in \Delta\} = F((\bigsqcup \Delta)^{\#})\,.$$

It follows (using the same observation as in the last paragraph) that

$$\mathcal{D}[\![\bigsqcup \{F(P) \mid P \in \Delta\}]\!] = \mathcal{D}[\![F(\bigsqcup \Delta)]\!]\,.$$

But if $P \leq Q$ and $\mathcal{D}[\![P]\!] = \mathcal{D}[\![Q]\!]$ it is easy to see that $P = Q$; continuity follows immediately.

All the usual algebraic laws of CSP (from [BR,H], for example) remain true over the extended model $\mathcal{N}'$. However we will see below that several natural extensions of these do not extend to the infinite hiding operator defined below.

As we know, all of the operators dealt with above map $\mathcal{N}$ to itself, so there is little reason to move to the extended model if they are all one wants to use. However, as hinted earlier, there are interesting operators which we can define over $\mathcal{N}'$ that could not be defined before. Firstly we can now define the nondeterministic composition of an arbitrary nonempty set of processes by union. If $\mathcal{S}$ is any nonempty set of processes

$$\mathcal{D}[\![\sqcap \mathcal{S}]\!] = \bigcup \{\mathcal{D}[\![P]\!] \mid P \in \mathcal{S}\}$$

$$\mathcal{F}[\![\sqcap \mathcal{S}]\!] = \bigcup \{\mathcal{F}[\![P]\!] \mid P \in \mathcal{S}\}\,.$$

This is continuous in each process argument individually, but not over infinite vectors of processes. (In considering the monotonicity and continuity of $\sqcap$ it is easier to think of it as a function of vectors rather than sets of processes, for the latter have no obvious partial order.) For example, if we define vectors $\underline{P}^n$ by

$$P^n_m = \begin{cases} \perp & \text{if } n \leq m \\ \mathit{STOP} & \text{otherwise} \end{cases}$$

we have $\underline{P}^n \leq \underline{P}^{n+1}$ but $\sqcap \underline{P}^n = \perp$ for all n while $\sqcap(\bigsqcup\{\underline{P}^n \mid n \in \mathbf{N}\}) = \sqcap(\mathit{STOP}^{\mathbf{N}}) = \mathit{STOP}$.

The existence of the above operator implies that any satisfiable specification defined purely as a property of the traces/failures/divergences of a process (every individual behaviour satisfies some property) has a most nondeterministic solution: simply take the nondeterministic composition of all solutions. (Some properties of this form, like "deadlock-free", do not have most nondeterministic solutions over $\mathcal{N}$.)

Perhaps the most annoying restriction imposed by the original model on the language was the impossibility of hiding an infinite set of events. If we allow infinite alphabets at all, it is strange only to be able to hide finite sets. There is no reason not to define it over $\mathcal{N}'$, however.

$$\begin{array}{rcl} \mathcal{D}[\![P \backslash X]\!] & = & \{(s \backslash X)t \mid s \in \mathcal{D}[\![P]\!]\} \\ & & \cup\{(s \backslash X)t \mid \exists s_1 < s_2 < \cdots . s = s_1 \wedge \forall i.(s_i \backslash X = s \backslash X \wedge s_i \in \mathit{traces}(P))\} \\ \mathcal{F}[\![P \backslash X]\!] & = & \{(s \backslash X, Y) \mid (s, X \cup Y) \in \mathcal{F}[\![P]\!]\} \cup \{(s, Y) \mid s \in \mathcal{D}[\![P \backslash X]\!]\} \end{array}$$

Notice that in this definition we insist that the sequence of traces yielding divergence are linearly ordered. This is often not done when X is finite because then König's Lemma proves that, if there are infinitely many s_i with $s_i \backslash X \leq s$ for all i then there is an infinite ordered sequence of the type above.

It is straightforward to show that this operator is well-defined and monotone with respect to $\leq$ and $\sqsubseteq$. However it is not continuous, again because of the introduction of unbounded nondeterminism. Consider the sequence $\langle P_n \mid n \in \mathbf{N}\rangle$, where

$$P_n \;=\; (x : \{k \mid k \leq n\} \rightarrow \mathit{STOP}) \Box (x : \{k \mid k > n\} \rightarrow \perp).$$

These form a $\leq$-chain with limit

$$P \;=\; x : \mathbf{N} \rightarrow \mathit{STOP}.$$

$P \backslash \mathbf{N} = \mathit{STOP}$ but, for each n, $P_n \backslash \mathbf{N} = \perp$.

Thus we have succeeded in modelling unboundedly nondeterministic operators, but with the seemingly inevitable loss of continuity. This means that in general the least fixed point might not be reached after only ω iterations of the underlying function. Indeed, one can give a simple if contrived example to show that any ordinal can be required. Given an ordinal α we take it (i.e., $\{\beta \mid \beta < \alpha\}$) as our alphabet and define

$$P \;=\; \beta : \alpha \rightarrow (P \parallel \gamma : \beta \rightarrow \mathit{STOP}) \backslash \alpha\,.$$

A little thought will reveal that the least fixed point here is

$$\beta : \alpha \rightarrow \mathit{STOP}$$

since every execution generates a strictly decreasing sequence of ordinals, which cannot be infinite. It takes precisely α iterations to reach this fixed point, unless α is finite in which case it takes $\alpha + 1$.

Unbounded nondeterminism seems to appear in two distinct forms: it manifests itself in a finite time, as in

$$Q_1 \ = \sqcap\{n \to STOP \mid n \in \mathbf{N}\}$$

where the initial refusals of $\mathcal{N}'$ capture it, or takes an infinite amount of time to appear as in

$$Q_2 \ = \sqcap\{P_n \mid n \in \mathbf{N}\} \quad \textit{where} \quad P_0 = STOP,\ P_{n+1} = a \to P_n\,.$$

Q_2 appears to be able to perform any finite number of as, but not to be able to perform an infinite sequence of them. Thus, intuitively, we would not expect $Q_2\backslash a$ to diverge. (Note that this form of unbounded nondeterminism can appear where there is a finite alphabet.) Unfortunately, in the $\mathcal{N}'$ semantics, $Q_2\backslash a$ *can* diverge because it interpolates a process' infinite behaviours from increasing sequences of its finite behaviours. $\mathcal{N}'$ has no means of describing this second, and perhaps more subtle, form of unbounded nondeterminism. (This example illustrates the fact that, in $\mathcal{N}'$, hiding (even finite) is not infinitely distributive over nondeterministic choice.)

The law $P\backslash X\backslash Y = P\backslash Y\backslash X$ fails to hold in general for essentially the same reason. To see this consider the process

$$Q = x : \mathbf{N} \to P_n$$

where P_n are as defined above. If a is hidden before $\mathbf{N}$ there is no divergence as there is no infinite execution sequence of as following any natural number: $Q\backslash a\backslash\mathbf{N} = STOP$. However if $\mathbf{N}$ is hidden first this information is lost and so divergence is predicted: $Q\backslash\mathbf{N}\backslash a = \ \bot$.

A way around this problem is to introduce a new component into the model to represent the infinite traces of a process. Thus a process will be described as a triple $\langle F, D, I\rangle$, where F and D have the same structure as before and I is the set of infinite traces it can perform. See the sequel [R2] for the construction of this model. In the sequel we will see that the present model and semantics in general give a pessimistic view of CSP, in that, if P is any CSP term, P_1 is the value ascribed to it by the semantics above and P_2 is the operationally natural value then $P_1 \leq P_2$. Equality is achieved once one introduces the infinite traces component. We also recover the desirable algebraic properties that were shown to have been lost above.

Thus arguably the model $\mathcal{N}'$ is not really a "proper" model for unboundedly nondeterministic CSP, rather an approximation to the correct one. Nevertheless the fact that it exists and is a complete partial order is certainly interesting. (In the sequel we will see that when it is extended by infinite traces it is not only an incomplete partial order but has no complete partial order consistent with CSP. The existence of the fixed points of recursions becomes much harder to prove.)

Proof rules over $\mathcal{N}'$. Perhaps paradoxically we find that the model $\mathcal{N}'$ which was made possible by $\leq$ has slightly fewer of the attractive properties proved in Section 2 (using $\leq$) than does $\mathcal{N}$.

Theorems 2.1 and 2.2 and Lemma 2.3 hold over $\mathcal{N}'$ for exactly the same reasons as before. One can still define the properties α, β, γ and δ exactly as was done there. Parts (a) and (b) of Proposition 2.5 still hold. The δ-topology on $\mathcal{N}'$ is still compact, but since every open (closed) set clearly contains $\langle \overline{F}, D\rangle$ if and only if it contains $\langle F, D\rangle$ for each $\langle F, D\rangle \in \mathcal{N}'$, it fails even to be T_0. (Recall that $\overline{F}$ denotes the closure of F under the compactness axiom.) For this reason it fails to be an extension of either lower topology. The γ-topology is Hausdorff and satisfies all parts of Proposition 2.5 (d). Unfortunately it is not compact, unless of course Σ is finite.

All this means that none of $\alpha - \delta$ satisfies the conditions of Theorem 2.6. The lower topology is itself compact, as it always is over a consistently complete cpo (see the appendix), so at least Theorem 2.6 is not vacuous. Whether there are any compact extensions of it that are practically useful will be a subject of further research.

Since Lemma 2.3 still holds in the extended model it follows that Theorem 2.7 does as well. However, unlike over $\mathcal{N}$, care is necessary in applying this theorem over $\mathcal{N}'$ since now by no means all functions are continuous. In particular it will be unsafe to use it whenever the function in question involves infinite nondeterminism or infinite hiding.

4. Conclusions. We have seen that it is possible for two different orders to give exactly the same semantics to CSP, at first sight a rather surprising result. Since I discovered the new order several of my colleagues have pointed out to me that a very similar alternative order based on definedness has recently been discovered by Greg Nelson for Dijkstra's language of guarded commands [N]. The semantic model used represents each process as a relation between initial and final states, with the additional final state $\perp$ representing the possibility of nontermination. As in this paper he replaced the usual refinement order by a new one which treats all nondivergent behaviour as incomparable: two processes P, Q are ordered just when P can diverge whenever Q can and all of P's nondivergent behaviour is reflected exactly in Q:

$$(\sigma, \perp) \notin P \Rightarrow P(\sigma) = Q(\sigma)\,.$$

(σ is any initial state.) Just as was the case in CSP, the new order there provided exactly the same semantics as before but increases the range of constructs that can be considered.

In both of these languages the nondeterminism order must remain the order of refinement and will therefore continue to play a most important role. However just because it has that function does not means that it is the best order on which to base the semantic fixed point theory: we have just seen how by moving to $\leq$ we were able to prove a wide range of new results about fixed points and to extend the model and language.

As I remarked above, the extension $\mathcal{N}'$ of the failures model is only able to cope properly with certain sorts of unbounded nondeterminism. It is perhaps best viewed as a stepping stone to the model incorporating infinite traces that will be introduced in [R2], though where it is adequate it is certainly simpler and easier to analyse than the new one.

Now that we have two different orders on the failures model we are faced with the question of which should be presented as the "standard" order for recursion to those meeting the model for the first time. If we agree that the restriction to finitely nondeterministic operators is undesirable there seem to be two options. One could present this new order (over $\mathcal{N}'$), mentioning the important role of $\sqsubseteq$ in refinement. On the other hand one could continue to use $\sqsubseteq$ over $\mathcal{N}'$ and appeal to the existence of $\leq$ and the results of this paper to assert the existence of fixed points, even though the order would not be complete. This is an important question but at this stage it is too early to decide.

Appendix: the topology of $\mathcal{N}$ and $\mathcal{N}'$. In this section we justify some of the claims that were made about the topological properties of the various conditions on predicates and about the lower and interval topologies. Of necessity we assume that the reader has a basic knowledge of general topology: essentially up to product spaces and Tychonoff's theorem.

First we establish the various parts of 2.5, with a few observations about $\mathcal{N}'$. The fact that β, γ and δ are equivalent when Σ is finite is obvious because there are only finitely many refusal sets and all refusal sets are finite. The case when Σ is infinite is illustrated by setting $\Sigma = \mathbf{N}$. Then

$$R_\beta(P) \equiv \exists n.(<>, \{n\}) \in \mathcal{F}[\![P]\!]$$

is a β-predicate but not a γ-predicate, while

$$R_\gamma(P) \equiv \forall s.(s, \mathbf{N}) \notin \mathcal{F}[\![P]\!]$$

is a γ-predicate but does not satisfy δ. (R_γ is deadlock-freedom, an important predicate: δ is the only one of our conditions that does not allow *all* predicates that are formed by specifying that all behaviours must satisfy some condition on divergences and failures.)

The fact that α is equivalent to β (and hence to γ and δ) if the alphabet and Λ are both finite follows easily from the fact that there are then only finitely many traces of any length and Φ (in the definition of β) can be equal to Λ. If Λ is infinite the predicate

$$R^1_\alpha(\underline{P}) \equiv \exists \lambda. P_\lambda = STOP$$

satisfies α but not β. The following predicate of a single process also has these properties

$$R^2_\alpha(P) \equiv P \neq STOP$$

when the alphabet Σ is infinite.

The fact that β, γ and δ, as defined in section 2 (over either model and over a simple or product space), generate topologies is easily demonstrated. In each case it is trivial that the whole space and the empty set are closed. That an arbitrary intersection of closed sets $\bigcap\{C_\lambda \mid \lambda \in I\}$ is closed follows because if some P fails to be in here then $P \notin C_\lambda$, say. This last fact is witnessed by some finite set(s) (of traces, or traces and failures, etc.). Exactly the same sets witness that P cannot be in the intersection. If C and D are closed then so is $C \cup D$, for if $P \notin C \cup D$ then there are finite sets for each of C and D witnessing this. The union(s) of these witness that P cannot be in C or in D, and therefore not in $C \cup D$.

To prove that the topologies produced on $\mathcal{N}^\Lambda$ (and, indeed on $\mathcal{N}'^\Lambda$) are just the products of the topologies on the individual spaces it is helpful to consider bases for the *open* sets of the topologies. We deal here with δ, the others (β and γ) being very similar indeed. An examination of δ (defined respectively over a single and a product space) reveals that a set U is open (the complement or negation of an open set) if and only for all $P \in U$ there the $\langle T, F\rangle$-ball (respectively $\langle T, F, \Phi\rangle$-ball) about P is contained in U for some finite sets $T, F(, \Phi)$. The $\langle T, F\rangle$-ball is defined to be those processes which cannot be distinguished from P by their divergences in T and their finite failures in F. The $\langle T, F, \Phi\rangle$-ball about $\underline{P} \in \mathcal{N}^\Lambda$ is similarly defined to be all those vectors which cannot be distinguished from $\underline{P}$ by inspecting the divergences in T and finite failures in F of their λ-components for $\lambda \in \Phi$. Denote these balls by

$$N_{\langle T,F\rangle}(P) \quad \text{and} \quad N_{\langle T,F,\Phi\rangle}(\underline{P})$$

respectively. It is easy to see that if $Q \in N_{\langle T,F \rangle}(P)$ then

$$N_{\langle T,F \rangle}(Q) = N_{\langle T,F \rangle}(P)$$

and that the same holds over $\mathcal{N}^{\Lambda}$. It follows that the balls are in each case a basis for the respective topology.

A basis for the product topology on $\mathcal{N}^{\Lambda}$ generated by δ on each component is formed by taking a finite set S of pairs (λ, U) where U is a basis element on $\mathcal{N}$ (and, without loss of generality, $(\lambda, U), (\lambda, U') \in S \Rightarrow U = U'$) and letting $U_S = \{\underline{P} \mid (\lambda, U) \in S \Rightarrow P_\lambda \in U\}$. Given $P \in U_S$, we know that for each $(\lambda, U) \in S$ there are T_λ and F_λ such that $N_{\langle T_\lambda, F_\lambda \rangle}(P_\lambda) \subseteq U)$. Define

$$\begin{array}{rcl} \Phi & = & \{\lambda \mid (\lambda, U) \in S\} \\ T & = & \bigcup\{T_\lambda \mid (\lambda, U) \in S\} \\ F & = & \bigcup\{F_\lambda \mid (\lambda, U) \in S\}\,. \end{array}$$

It is easy to see that $N_{\langle T,F,\Phi \rangle}(\underline{P}) \subseteq U_S$. It follows that the δ-topology on $\mathcal{N}^{\Lambda}$ is finer (an extension of) the product topology. The other way round is easier: given $N_{\langle T,F,\Phi \rangle}(\underline{P})$, set $S = \{(\lambda, N_{\langle T,F \rangle}(P_\lambda)) \mid \lambda \in \Phi\}$. It is easy to see that $U_S = N_{\langle T,F,\Phi \rangle}$.

The compactness axiom (4) in the definition of $\mathcal{N}$ means that each element is completely determined by its divergences and finite failures: there is a unique element of $\mathcal{N}$ corresponding to each pair $\langle \mathit{ff}, D \rangle$ of finite failures and divergences satisfying axioms (1-3),(5),(6). This means that the natural projection from $\mathcal{N}$ to the set $\mathcal{N}_{\mathit{fin}}$ of all such pairs is a bijection. If the δ-topology is defined on $\mathcal{N}_{\mathit{fin}}$ in the same way as over $\mathcal{N}$ it is clear that this map in a homeomorphism. Thus if we could prove that $\mathcal{N}_{\mathit{fin}}$ is compact under δ we would also have proved that $\mathcal{N}$ is.

It is clear that the δ-toplogy on $\mathcal{N}_{\mathit{fin}}$ is just the subspace topology inherited from the topology equivalently defined on the set A of arbitrary pairs $\langle \mathit{ff}, D \rangle$ of finite failure and divergence sets (i.e., not necessarily satisfying the axioms). It is easy to prove that the latter is just the product of the two topologies on sets of finite failures and traces respectively where bases are given by

$$\begin{array}{rcll} N_X(\mathit{ff}) & = & \{\mathit{ff}' \mid \mathit{ff}' \cap X = \mathit{ff} \cap X\} & \quad X \text{ a finite set of finite failures} \\ N_X(D) & = & \{D' \mid D' \cap X = D \cap X\} & \quad X \text{ a finite set of traces} \end{array}$$

which are in turn homeomorphic to the product of $\Sigma^* \times \wp(\Sigma)$ and Σ^* copies respectively of the discrete topology on $\{0, 1\}$ – compact by Tychonoff's theorem. It follows that $\mathcal{N}$ is compact if and only if it is a closed subset of A. Since the intersection of closed subsets is closed it will be sufficient to check that the sets of pairs satisfying each of (1-3),(5),(6) individually are closed. This is implied by the fact that, in each case, if a pair $\langle \mathit{ff}, D \rangle$ does not satisfy the axiom then this fact is discernable from checking if $\langle \mathit{ff}, D \rangle$ contains each of some finite set of behaviours. For example, if it fails (1) then *either* $(<>, \emptyset) \notin \mathit{ff}$ (one behaviour) or there are $s < t$ such that $(s, \emptyset) \notin \mathit{ff}$ and $(t, \emptyset) \notin \mathit{ff}$ (two behaviours). The most interesting case is axiom (3). Here, because we are restricting ourselves to finite failures, the set Y must (implicitly) be finite itself. Failure to satisfy this axiom would manifest itself by $(s, X) \in \mathit{ff}$, $(s, X \cup Y) \notin \mathit{ff}$ and $(s\langle a \rangle, \emptyset) \notin \mathit{ff}$ for each $a \in Y$. (It is the need to make (3) closed that has forced us down into finite failures.)

Since the topology on A has been established to be homeomorphic to a Hausdorff one it follows trivially that the toplogies on $\mathcal{N}_{\mathit{fin}}$, and hence on $\mathcal{N}$, are Hausdorff also.

Since $\mathcal{N}$ is compact Hausdorff under δ it follows from the equivalence of the δ-topology on $\mathcal{N}^\Lambda$ and the product topology that the δ-topology on $\mathcal{N}^\Lambda$ is compact Hausdorff (Tychonoff's theorem again). The topology is zero-dimensional because it has a basis of closed and open (clopen) sets: the ball $N_{\langle T,F,\Phi\rangle}(\underline{P})$ is closed because (as is easily checked), the complement of this ball is equal to

$$\bigcup\{N_{\langle T,F,\Phi\rangle}(\underline{Q}) \mid \underline{Q} \notin N_{\langle T,F,\Phi\rangle}(\underline{Q})\}$$

which is an open set. This completes the proof of 2.5 (c).

Each of the predicates described in (d) have the property that their failure can be demonstrated by at most two behaviours from a single component of Q. For example, the failure of $\underline{Q} \geq \underline{P}$ is either demonstrated by some divergence of Q_λ not being present in P_λ, $(s, \emptyset)$ not being in $\mathcal{F}[\![Q]\!]$ for some minimal divergence of P or $(s, X) \in \mathcal{F}[\![Q]\!] \not\Leftrightarrow (s, X) \in \mathcal{F}[\![P]\!]$ for some $s \notin \mathcal{D}[\![P]\!]$ and X finite. (Here we use the fact that $\mathcal{F}[\![P]\!]$ and $\mathcal{F}[\![Q]\!]$ are determined by their sets of finite failures.)

The δ-topology on $\mathcal{N}'$ is not even T_0 because when two processes have the same set of divergences and finite failures there is no open set that contains one without the other. In fact it inherits compactness from $\mathcal{N}_{\mathit{fin}}$ in the same way as above, since its open sets are precisely the inverse images of $\mathcal{N}_{\mathit{fin}}$'s open sets under the natural projection.

When Σ is infinite, none of the stronger topologies on $\mathcal{N}$ can be compact because of the theorem that no proper extension of a compact Hausdorff topology can be compact. The γ-topology on $\mathcal{N}'$ is easily shown to be Hausdorff (any pair of distinct processes differ on some behaviour) but turns out not to be compact, essentially because the infinitary version of axiom (3) is then not closed in the sense described above.

This concludes our discussion of the α, β, γ and δ topologies. We now turn our attention more generally to the subject of the interval and lower topologies. (For a much deeper discussion of these, in particular over lattices, see [CCL].)

The interval topology is always T_1, for each singleton set is closed ($[x, x] = \{x\}$) and the lower topology is always T_0. The interval topology often satisfies stronger separation properties than this (for example compact Hausdorff) but need not, as is demonstrated by the lattice with only an infinite set of incomparable points between top and bottom. The lower topology never satisfies T_1 in a nontrivial partial order, for if $x < y$ then every open set containing y contains x.

Over a consistently complete cpo X (i.e., one where each finite set of elements with an upper bound have a least upper bound) the lower topology is always compact. To prove this we appeal to a theorem of Alexander (see, for example [K]) that a topology is compact if each cover by elements from some sub-basis has a finite subcover. Thus in our case it sufficient to prove that if S is a subset of X such that for all finite $F \subseteq S$ the set $\bigcap\{x\uparrow \mid x \in F\}$ is nonempty, then $\bigcap\{x\uparrow \mid x \in S\}$ is nonempty. The assumptions imply that each finite $F \subseteq S$ has a least upper bound $\bigsqcup F$. The set of these as F varies forms a directed set, the limit of which is the least upper bound of S, proving that $\bigcap\{x\uparrow \mid x \in S\}$ is nonempty as desired, for it must contain this limit.

A simple extension of this argument shows that the interval topology over a consistently complete cpo is compact also. (In fact, one proves that the perhaps finer topology with sub-basis of closed sets $\{x\uparrow, x\downarrow \mid x \in X\}$ is compact, where $x\downarrow = \{y \mid y \leq x\}$.) For in that case any subset G of this sub-basis with the finite intersection property splits into two parts $\{x\uparrow \mid x \in S\} \cup \{y\downarrow \mid y \in T\}$.

For as before, the least upper bound of S exists, and since $x\uparrow \cap y\downarrow \neq \emptyset$ when $x \in S$ and $y \in T$ we know that each $y \in T$ is an upper bound for S, so it follows that $\bigsqcup S \leq y$ for all $y \in T$ and hence that $\bigsqcup S \in \bigcap G$.

Acknowledgements. I would like to thank Peter Collins for help in discovering the Alexander sub-basis theorem used in the Appendix. A number of people, in particular Tom Verhoeff and a referee, have helped me by pointing out minor errors and stylistic improvements. Several people, including Carroll Morgan, have pointed out the similarity of this work with that of Nelson which was referred to in the conclusion.

.

References.

[B] Brookes, S.D., *A Model for Communicating Sequential Processes*, Oxford University D.Phil. thesis, 1983.

[BHR] Brookes, S.D., Hoare, C.A.R., and Roscoe, A.W., *A Theory of Communicating Sequential Processes*, Journal of the Association for Computing Machinery, vol. 31, no. 3, 560-599.

[BR] Brookes, S.D., and Roscoe A.W., *An improved failures model for communicating processes* , Springer Lecture Notes in Computer Science, vol. 197, 281-305.

[BRW] Brookes, S.D., Roscoe A.W., and Walker, D.J., *An operational semantics for CSP*, Submitted for publication.

[CCL] Gierz, G., Hofmann, K.H., Keimel, K., Lawson, J.D., Mislove, M., and Scott, D.S., *A compendium of continuous lattices*, Springer-Verlag (1980).

[H] Hoare, C.A.R., *Communicating sequential processes*, Prentice-Hall, 1985

[HBR] Hoare, C.A.R., Brookes, S.D., and Roscoe, A.W., *A theory of communicating sequential processes*, Oxford University Computing Laboratory, Programming Research Group, Technical Report PRG-16.

[K] Kelley, J.L., *General Topology*, Springer GTIM27 (1975).

[N] Nelson, G., *A generalisation of Dijkstra's calculus*, Research report 16, Digital Systems Research Center, 1987.

[Re] Reed, G.M., *A uniform mathematical theory for real-time distributed computing*, Oxford University D.Phil. thesis, 1988.

[RR1] Reed, G.M., and Roscoe, A.W., *A timed model for communicating sequential processes*, Proceedings of ICALP'86, Springer LNCS 226 (1986), 314-323.

[RR2] Reed, G.M., and Roscoe, A.W., *Metric spaces as models for real-time concurrency*, in Proceedings of MFPS87 (Springer LNCS 298).

[R1] Roscoe, A.W., *A mathematical theory of communicating processes*, Oxford University D.Phil. thesis, 1982.

[R2] Roscoe, A.W., *Unbounded nondeterminism in CSP*, **To be filled in when we know where this is to appear**

[R3] Roscoe, A.W., *Topology, computer science, and the mathematics of convergence*, in Topology and category theory in computer science (Reed, Roscoe and Wachter, eds) OUP 1991.

[R4] Roscoe, A.W., *Two papers on CSP*, Techical monograph PRG-67, Oxford University Computing Laboratory, 1988.