

Kent Academic Repository

Full text document (pdf)

Citation for published version

Schneider, Steve A. and Bryans, Jeremy W. and Davies, Jim (1994) A Failures Semantics for ET-LOTOS. Technical report. , Presented at the COST 247 - WG 6.1 working group meeting in Brighton

DOI

Link to record in KAR

<https://kar.kent.ac.uk/21190/>

Document Version

UNSPECIFIED

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

A failures semantics for ET-LOTOS¹

Steve Schneider
Royal Holloway, University of London

Jeremy Bryans and Jim Davies
University of Reading

Abstract A denotational semantics is presented for ET-LOTOS [LeL94] in the style of semantics for timed CSP, in terms of timed failures with additional information about internal events. This semantics is consistent with the original operational semantics, and is shown to be the weakest congruence stronger than trace equivalence—the weakest useful congruence.

1 Introduction

The LOTOS specification language was developed to support specification and design of distributed and communicating systems. Recently, a number of proposals have arisen for incorporating time into the language [LeL94, QFA93, BLT94], enabling the modelling of quantitative timing behaviour.

A LOTOS specification consists primarily of a description of desirable behaviour. But we also require a notion of how closely any proposed implementation should resemble that ideal. Many approaches require that an implementation should be at least weakly bisimilar to the specification process. Bisimulation equivalence has the pleasing property that it can often be efficiently checked mechanically, and further it usually has a complete set of characterising algebraic laws. Proposals for timed extensions to LOTOS are generally accompanied by appropriate notions of bisimulation.

On the other hand, the kind of distinctions made by bisimulation equivalence are finer than can generally be made by interacting agents, which do not have access to the internal state of their communication partners. In fact, processes which are not bisimilar may still be indistinguishable in any context. This paper takes the view that if no context can tell a specification process from a proposed implementation, then the implementation should be considered to be satisfactory. Since bisimulation congruence is strictly stronger than this notion of equivalence, it is useful as a sufficient condition for equivalence, but should not be considered as a necessary condition.

¹Presented at the COST 247 workshop on extending process algebras, Brighton July 1994

We follow the CSP approach of [Sch94], and see how it adapts to the LOTOS operators. For definiteness, we will consider the timed LOTOS proposal of [LeL94], as this appears to be the most amenable to this treatment. We begin by requiring that processes which have different traces should be considered distinct, and then search for the weakest congruence over the language that preserves this distinction. We find that the timed failures used in CSP are not quite enough, and that we also require information about the times at which certain internal events are possible, because some LOTOS operators are sensitive to internal events. The result is a fully abstract denotational semantics for timed LOTOS .

Equivalence and congruence

A congruence on a language is an equivalence relation that is preserved by all of the operators of the language. Hence it is preserved by all contexts built from those operators. Given a notion \equiv of equivalence, the corresponding congruence \cong is the weakest congruence (i.e. the one which makes as few distinctions as possible) which implies \equiv -equivalence. In other words

$$\forall P, Q \bullet P \cong Q \Rightarrow P \equiv Q$$

and if we have another congruence \cong' for which

$$\forall P, Q \bullet (P \cong' Q \Rightarrow P \equiv Q)$$

then

$$\forall P, Q \bullet P \cong' Q \Rightarrow P \cong Q$$

For example, untimed trace equivalence on processes identifies two processes if any finite sequence of events (any untimed trace) which can be observed during an execution of one may also be observed of the other. Untimed trace congruence (hereafter called trace congruence) for timed CSP turns out to be timed failures equivalence, where two processes are identified if they have the same set of timed failures (see [Sch94] for details).

In this paper we are interested in the weakest congruence stronger than (untimed) trace equivalence. We regard this as the weakest congruence of any practical use—anything weaker will equate processes which have different traces! For timed CSP, this turns out to be the timed failures semantics, and this is in turn equivalent to may testing equivalence [Hen88] on the operational semantics (where two processes are equivalent if the set of tests each may pass is the same). In general however, may testing equivalence might not be a congruence, and in fact it is not a congruence for timed versions of LOTOS.

Basic Timed LOTOS

We focus on Leduc and Léonard's proposal for Basic ET-LOTOS [LeL94]. They present an extension to the language of Basic LOTOS to include timing operators, and give an operational semantics for the language, in terms of event transitions and delay transitions. Processes are defined in terms of those transitions that may be performed.

An event transition is of the form $P \xrightarrow{a} Q$. This indicates that process P may immediately perform event a and then behave as process Q . A delay transition is of the form $P \xrightarrow{d} Q$. This indicates that process P may evolve in d units of time to process Q .

LOTOS assumes a universal set of gates G . It also makes use of special events $i, \delta \notin G$, and we will introduce other such special events $\omega \notin G, v \in G$ in this paper. The event a ranges over $G \cup \{i, \delta, \omega\}$. A set of gates Γ is a subset of G . The variables t, u range over a set of time variables; d ranges over the set of time values, which is taken to be $[\theta, \infty) \cup \{\infty\}$, though it is noted in [LeL94] that any commutative monoid with an absorbent (zero) element ∞ for $+$ will serve as a time domain.

The formal semantics of Basic ET-LOTOS is given in [LeL94]; we reproduce it here to keep the paper self-contained.

The syntax is as follows: A process P is defined by an equational definition with a where clause binding a vector of process variables to a vector of process terms:

$$P := Q \quad \text{where} \quad \tilde{Y} := \tilde{Q}$$

Every Y occurring free in some Q must appear in the vector of process variables \tilde{Y} . As discussed below, the vector of process terms must be time-guarded in the vector of process variables. This constraint is not imposed in [LeL94], but is necessary for the semantic approach taken in this paper.

The terms Q are defined as follows:

$$\begin{aligned} Q := & \text{stop} \mid \text{exit}\{d\} \mid a@t\{d\};Q \mid \Delta^d Q \mid Q \square Q \\ & \mid Q \mid[\Gamma] \mid Q \mid \text{hide } \Gamma \text{ in } Q \mid Q \gg Q \mid Q [> Q \mid Y \end{aligned}$$

These denote respectively: deadlock; termination at time d ; life reducer (or event timeout); delay; choice; concurrency; abstraction; enabling (or sequential composition); disabling; process variable.

If t does not appear free in Q then the term $a\{d\};Q$ may be used as an abbreviation for $a@t\{d\};Q$. The term $a;Q$ abbreviates $a\{\infty\};Q$ for $a \neq i$, and $i;Q$ is used for $i\{0\};Q$.

The rules for the operators are given as follows:

$$\frac{}{\text{stop} \xrightarrow{d} \text{stop}}$$

$$\frac{}{\text{a@t}\{d_1\}; P \xrightarrow{a} [0/t]P}$$

$$\frac{}{\text{a@t}\{d_1 + d\}; P \xrightarrow{d} \text{a@t}\{d_1\}; [t + d/t]P}$$

$$\frac{}{\text{a@t}\{d_1 + d\}; P \xrightarrow{d} \text{stop}} \quad [a \neq i, d > d_1]$$

$$\frac{P \xrightarrow{a} P'}{\Delta^0 P \xrightarrow{a} P'}$$

$$\frac{}{\Delta^{d_1+d} P \xrightarrow{d} \Delta^{d_1} P}$$

$$\frac{P \xrightarrow{d} P'}{\Delta^0 P \xrightarrow{d} P'}$$

$$\frac{}{\text{exit}\{d_1\} \xrightarrow{\delta} \text{stop}}$$

$$\frac{}{\text{exit}\{d_1 + d\} \xrightarrow{d} \text{exit}\{d_1\}}$$

$$\frac{}{\text{exit}\{d_1\} \xrightarrow{d} \text{stop}} \quad [d > d_1]$$

$$\frac{P \xrightarrow{a} P'}{P \square Q \xrightarrow{a} P' \quad Q \square P \xrightarrow{a} P'}$$

$$\frac{P \xrightarrow{d} P' \quad Q \xrightarrow{d} Q'}{P \square Q \xrightarrow{d} P' \square Q'}$$

$$\frac{P \xrightarrow{a} P'}{P \mid[\Gamma] \mid Q \xrightarrow{a} P' \mid[\Gamma] \mid Q \quad Q \mid[\Gamma] \mid P \xrightarrow{a} Q \mid[\Gamma] \mid P'} \quad [a \notin \Gamma \cup \{\delta\}]$$

$$\frac{P \xrightarrow{a} P' \quad Q \xrightarrow{a} Q'}{P \mid[\Gamma] \mid Q \xrightarrow{a} P' \mid[\Gamma] \mid Q'} \quad [a \in \Gamma \cup \{\delta\}]$$

$$\frac{P \xrightarrow{d} P' \quad Q \xrightarrow{d} Q'}{P \mid[\Gamma] \mid Q \xrightarrow{d} P' \mid[\Gamma] \mid Q'}$$

$$\frac{P \xrightarrow{a} P'}{\text{hide } \Gamma \text{ in } P \xrightarrow{a} \text{hide } \Gamma \text{ in } P'} \quad [a \notin \Gamma]$$

$$\frac{P \xrightarrow{a} P'}{\text{hide } \Gamma \text{ in } P \xrightarrow{i} \text{hide } \Gamma \text{ in } P'} \quad [a \in \Gamma]$$

$$\frac{P \xrightarrow{d} P' \quad \forall a \in \Gamma \bullet \neg(P \xrightarrow{a})}{\text{hide } \Gamma \text{ in } P \xrightarrow{d} \text{hide } \Gamma \text{ in } P'}$$

$$\frac{P \xrightarrow{a} P'}{P \gg Q \xrightarrow{a} P' \gg Q} \quad [a \neq \delta]$$

$$\frac{P \xrightarrow{\delta} P'}{P \gg Q \xrightarrow{i} Q}$$

$$\frac{P \xrightarrow{d} P' \quad \neg(P \xrightarrow{\delta})}{P \gg Q \xrightarrow{d} P' \gg Q}$$

$$\frac{P \xrightarrow{a} P'}{P [> Q \xrightarrow{a} P' [> Q} \quad [a \neq \delta]$$

$$\frac{Q \xrightarrow{a} Q'}{P [> Q \xrightarrow{a} Q'}$$

$$\begin{array}{c}
P \xrightarrow{\delta} P' \\
\hline
P \text{ [> } Q \xrightarrow{\delta} P'
\end{array}$$

$$\begin{array}{c}
P \xrightarrow{d} P' \qquad Q \xrightarrow{d} Q' \\
\hline
P \text{ [> } Q \xrightarrow{d} P' \text{ [> } Q'
\end{array}$$

$$\begin{array}{c}
[g_1/h_1, \dots, g_n/h_n]P \xrightarrow{a} P' \\
Q[h_1, \dots, h_n] := P \\
\hline
Q[g_1, \dots, g_n] \xrightarrow{a} P'
\end{array}$$

$$\begin{array}{c}
[g_1/h_1, \dots, g_n/h_n]P \xrightarrow{d} P' \\
Q[h_1, \dots, h_n] := P \\
\hline
Q[g_1, \dots, g_n] \xrightarrow{d} P'
\end{array}$$

Operationally, there is no problem with instant or unguarded recursions; their transitions may be calculated using the rules in exactly the same way as guarded recursions. A definition such as $S := S$ is perfectly legitimate, and defines a time-stop process, one that has no event or delay transitions (in contrast to `stop` which can always allow time to pass). A definition such as $S := a@t\{1\}; S$ is guarded by the event a , and can do as many copies of the event a at time θ as the environment is prepared to allow.

However, the standard denotational approach for timed CSP (which is being followed here) would need a significant alteration in order to deal with such process definitions. Recursive definitions are required to take some time to reach their next invocation. This makes the semantics more straightforward, at the expense of expressibility. Thus any recursive definition must be *time-guarded* in the sense that there is some non-zero delay, imposed by a Δ^d construct, guarding any recursive call. (In a mutual recursion, there must be some universal non-zero time guard which holds for each recursive definition.) As a result, no time-stopping process can be written; the examples in the preceding paragraph are all considered to be unguarded in the timed sense.

Traces

The traces of a process may be extracted from its operational semantics by use of the rules below.

$$\frac{}{\langle \rangle \in \text{traces}(P)}$$
$$\frac{P \xrightarrow{a} P' \quad tr \in \text{traces}(P')}{\langle a \rangle \frown tr \in \text{traces}(P)} \quad [a \neq i]$$
$$\frac{P \xrightarrow{i} P' \quad tr \in \text{traces}(P')}{tr \in \text{traces}(P)}$$
$$\frac{P \xrightarrow{d} P' \quad tr \in \text{traces}(P')}{tr \in \text{traces}(P)}$$

The set $\text{traces}(P)$ is defined to be the smallest set closed under these inference rules.

Timed failures

The set of events Σ includes every gate in G , and the events δ and ω . The internal event i is not in Σ .

A timed event (t, a) is an element of $[\theta, \infty) \times \Sigma$.

A timed failure is a particular view of an execution of a process, recording what might have been observed by its environment, or by another process. It will contain the events that were performed, the times at which they were performed, and the times at which events were refused by the process. The approach of CSP is to associate each process with the set of those timed failures that are consistent with

some run of the process. Throughout this paper, the term *failure* means timed failure.

A timed trace is a finite sequence of timed events in which the times are non-decreasing. The set of timed traces is denoted TT .

A timed refusal is a set of timed events.

A timed failure is a (timed trace, timed refusal) pair. This is considered a record of an execution; the timed trace records those timed events that were performed, and the timed refusal records those timed events that were refused by the process, during the execution.

Since process executions are finitely variable, in that only finitely many changes of state may occur in a finite interval, it is sufficient to consider timed refusal sets which are finitely variable; where the set of events being refused changes only finitely often. Timed refusals thus have a particular structure: a finite union of refusal tokens.

$$\begin{aligned} FRTOK &= \{[b, e) \times A \mid 0 \leq b < e < \infty, A \subseteq^{fn} \Sigma\} \\ FRSET &= \{\bigcup B \mid B \subseteq^{fn} FRTOK\} \end{aligned}$$

Notation

We use s to range over TT , the set of timed traces, and X to range over $FRSET$, the set of timed refusals; t, u range over the finite times $[0, \infty)$; I ranges over subsets of $[0, \infty)$ (usually intervals).

We use the following operations on sequences of timed events: $s_1 \frown s_2$ denotes the concatenation of s_1 and s_2 ; $end(s \frown \langle (t, a) \rangle) = t$.

The following projections on sequences are defined by list comprehension:

$$\begin{aligned} s \uparrow I &= \langle (t, a) \mid (t, a) \leftarrow s, t \in I \rangle \\ s \setminus A &= \langle (t, a) \mid (t, a) \leftarrow s, a \notin A \rangle \\ s \upharpoonright A &= \langle (t, a) \mid (t, a) \leftarrow s, a \in A \rangle \\ s + u &= \langle (t + u, a) \mid (t, a) \leftarrow s \rangle \\ events(s) &= \langle a \mid (t, a) \leftarrow s \rangle \\ \sigma(s) &= \{a \mid s \upharpoonright \{a\} \neq \langle \rangle\} \end{aligned}$$

We also define a number of projections on refusal sets:

$$\begin{aligned} X \triangleleft u &= \{(t, a) \mid (t, a) \in X, t < u\} \\ X \uparrow I &= \{(t, a) \mid (t, a) \in X, t \in I\} \end{aligned}$$

$$\begin{aligned}
X - u &= \{(t - u, a) \mid (t, a) \in X, t \geq u\} \\
\sigma(X) &= \{a \mid (u, a) \in X\} \\
\text{end}(X) &= \text{sup}\{u \mid (u, a) \in X\}
\end{aligned}$$

We define $\text{end}(s, X) = \max\{\text{end}(s), \text{end}(X)\}$.

Example

The timed failure

$$(\langle (\mathcal{B}, a), (\mathcal{B}, b) \rangle, [\mathcal{0}, \mathcal{B}] \times \{b\} \cup [\mathcal{B}, \mathcal{5}] \times \{a\})$$

corresponds to an execution in which the event **b** was offered to the process but not accepted (i.e. refused) up until time \mathcal{B} ; the event **a** was then observed at time \mathcal{B} ; a further copy of the event **a** was refused over the interval $[\mathcal{B}, \mathcal{5}]$; and finally the event **b** was performed at time $\mathcal{8}$.

This would be a possible observation of the process $\mathbf{a} ; \Delta^2 \mathbf{b} ; \text{stop}$. The event **a** can occur at any time, or not at all. If it occurs at time \mathcal{B} , then no further copies of it are possible, and so will be refused over the interval $[\mathcal{B}, \mathcal{5}]$. The event **b** is not enabled in this execution until time $\mathcal{5}$, so it can be refused over the interval $[\mathcal{0}, \mathcal{B}]$. It may be performed at any time after time $\mathcal{5}$; in this instance, it is performed at time $\mathcal{8}$.

Extracting failures

We can define inductively what it means for a timed failure to be consistent with a given execution. The relation **exhibits** is a predicate on processes and timed failures drawn from $TT \times FRSET$. The events that may appear in timed traces are $G \cup \{\delta, \omega\}$ —any action of interest apart from i .

$$\begin{array}{c}
\hline
P \text{ exhibits } (\langle \rangle, \{\}) \\
\hline
\\
\\
\frac{P \xrightarrow{a} Q \wedge Q \text{ exhibits } (s, X)}{P \text{ exhibits } (\langle (\mathcal{0}, a) \rangle \frown s, X)} \quad [a \neq i] \\
\\
\\
\frac{P \xrightarrow{i} Q \wedge Q \text{ exhibits } (s, X)}{P \text{ exhibits } (s, X)} \\
\hline
\end{array}$$

$$\frac{
\begin{array}{l}
P \xrightarrow{t} Q \\
Q \text{ exhibits } (s, X - t) \\
\forall a \in \sigma(X \triangleleft t) \bullet \neg(P \xrightarrow{a})
\end{array}
}{
P \text{ exhibits } (s + t, X)
}$$

The relation **exhibits** is defined to be the smallest relation closed under these inference rules—it holds only in cases that can be established using these rules.

The last consistency rule exploits a variant of the reverse persistency result 4.2.4 of [LeL94]: that if $P \xrightarrow{t} Q$ and $\neg(P \xrightarrow{a})$ then

$$\forall t' < t \bullet P \xrightarrow{t'} P' \Rightarrow \neg(P' \xrightarrow{a})$$

In other words, if P is able to refuse a initially, then at every point before the end of an evolution a may still be refused. Hence the refusal set $X \triangleleft t$ recorded during the evolution is consistent with the execution if P is unable to perform any of the events in that set (at the beginning of that evolution).

The failures of P can now be extracted:

$$\text{failures}(P) = \{(s, X) \mid P \text{ exhibits } (s, X)\}$$

Example

The following sequence of executions is possible for $a; \Delta^2 b; \text{stop}$, and this is consistent with the failure $(\langle (\mathcal{B}, a), (\mathcal{B}, b) \rangle, [\emptyset, \mathcal{B}] \times \{b\} \cup [\mathcal{B}, 5] \times \{a\})$.

$$\begin{array}{rcl}
0 & a; \Delta^2 b; \text{stop} & \\
& \downarrow \mathcal{B} & \\
\mathcal{B} & a; \Delta^2 b; \text{stop} & \xrightarrow{a} \Delta^2 b; \text{stop} \\
& & \downarrow \mathcal{B} \\
5 & & \Delta^0 ; b; \text{stop} \\
& & \downarrow \mathcal{B} \\
8 & & b; \text{stop} \xrightarrow{b} \text{stop}
\end{array}$$

The inference rules may be used with this sequence of transitions (considered in reverse order) to derive the claim above:

$$\begin{array}{l}
\text{stop exhibits } (\langle \rangle, \{\}) \\
\vdash \quad \mathbf{b}; \text{stop exhibits } (\langle (\theta, b) \rangle, \{\}) \\
\vdash \quad \Delta^0 \mathbf{b}; \text{stop exhibits } (\langle (\mathcal{B}, b) \rangle, \{\}) \\
\vdash \quad \Delta^2 \mathbf{b}; \text{stop exhibits } (\langle (\mathcal{S}, b) \rangle, [\theta, 2] \times \{a\}) \\
\vdash \quad \mathbf{a}; \Delta^2 \mathbf{b}; \text{stop exhibits } (\langle (\theta, a), (\mathcal{S}, b) \rangle, [\theta, 2] \times \{a\}) \\
\vdash \quad \mathbf{a}; \Delta^2 \mathbf{b}; \text{stop exhibits } (\langle (\mathcal{B}, a), (\mathcal{S}, b) \rangle, [\theta, \mathcal{B}] \times \{b\} \cup [\mathcal{B}, \mathcal{S}] \times \{a\})
\end{array}$$

The traces of a process may be deduced from its failures:

Theorem 1.1 $\text{traces}(P) = \{\text{events}(s) \mid (s, X) \in \text{failures}(P)\}$ □

Proof Each execution of P gives rise to a trace tr , and is consistent with a set of failures of P each of whose timed trace s has $\text{events}(s) = tr$. □

Corollary 1.2 Timed failures equivalence is stronger than trace equivalence □

Distinguishing processes with different failures

The next theorem gives the converse relationship between trace congruence and timed failures equivalence. If failures equivalence were a congruence, then the previous corollary together with this theorem would yield that failures equivalence is trace congruence. Alas, as we shall see later, failures equivalence is not a congruence for the language Basic ET-LOTOS.

Theorem 1.3 If two processes are trace congruent, then they have the same timed failures □

Proof Given a particular timed failure (s_θ, X_θ) , with a particular representation for the refusal set X_θ

$$X_\theta = [u_1, u'_1) \times A_1 \cup [u_2, u'_2) \times A_2 \cup \dots \cup [u_m, u'_m) \times A_m$$

we can construct a context that detects (s_θ, X_θ) .

$$\begin{aligned}
P'_{\langle \rangle} &= \omega; \text{Stop} \\
P'_{\langle (t, a) \rangle} \frown_s &= \Delta^t \mathbf{a} @ t \{0\}; P'_{s-t} \\
Q'_i &= \coprod_{a \in A_i} \Delta^u \mathbf{a} @ t \{u' - u\}; \text{Stop} \quad \square \quad \Delta^{u'} \mathbf{i}; \omega; \text{Stop} \\
Q' &= \coprod_{\omega} Q'_i
\end{aligned}$$

In this construction, ω is a special ‘success’ event that is not contained in the set of gates G available to processes. Note the use of generalised choice in the definition of \mathbf{Q}_i ; since the set A_i is finite, this is simply multiple applications of the \square operator.

The timed LOTOS context for (s_θ, X_θ) is

$$\mathbf{C}_{(s_\theta, X_\theta)}(\mathbf{R}) = (\mathbf{R} \llbracket \Sigma \rrbracket (\mathbf{P}'_{s_\theta} \llbracket \omega \rrbracket \mathbf{Q}'_{X_\theta})) \setminus \Sigma$$

The process $\mathbf{C}_{(s_\theta, X_\theta)}(\mathbf{P})$ is able to perform ω if and only if $(s_\theta, X_\theta) \in \text{failures}(\mathbf{P})$: the \mathbf{P}'_{s_θ} component can perform ω only after performing the events in the trace s_θ , and the \mathbf{Q}'_{X_θ} component can perform ω only if everything in X_θ is refused.

If (s_θ, X_θ) is a failure of \mathbf{P} , then $\text{traces}(\mathbf{C}_{(s_\theta, X_\theta)}(\mathbf{P})) = \{\langle \rangle, \langle \omega \rangle\}$. If (s_θ, X_θ) is not a failure of \mathbf{P} , then $\text{traces}(\mathbf{C}_{(s_\theta, X_\theta)}(\mathbf{P})) = \{\langle \rangle\}$. Thus if two processes have the same traces in all contexts, then they must have the same timed failures. \square

This result also holds for timed CSP. There it is also true that failures equivalence is a congruence, and so it follows that it is trace congruence for the language of timed CSP.

Failures equivalence is not a congruence

In the case of ET-LOTOS, however, failures equivalence is not enough to yield a congruence. Failures contain no information about internal activity. This does not matter in timed CSP, since none of the operators of the language respond to internal events. In ET-LOTOS, there are two operators that are sensitive to hiding, and which therefore may not preserve failures equivalence: choice and disabling.

Consider the two processes

$$\begin{aligned} \mathbf{P}_1 &= \mathbf{i}; \mathbf{a}\{5\}; \mathbf{Stop} \\ \mathbf{P}_2 &= \mathbf{a}\{5\}; \mathbf{Stop} \end{aligned}$$

These two processes have the same timed failures:

$$\begin{aligned} &\{(\langle \rangle, X) \mid a \notin \sigma(X \uparrow [0, 5])\} \\ &\cup \{(\langle (u, a) \rangle, X) \mid u \leq 5 \wedge a \notin \sigma(X \uparrow [0, u])\} \end{aligned}$$

But \mathbf{P}_1 will always perform an internal event at time θ , whereas \mathbf{P}_2 performs no event before its initial \mathbf{a} .

A context $\mathbf{C}(\mathbf{R}) = \mathbf{R} \square \Delta^1 \mathbf{b}\{5\}; \mathbf{Stop}$, involving choice may distinguish \mathbf{P}_1 from \mathbf{P}_2 : the process $\mathbf{C}(\mathbf{P}_1)$ is unable to perform a \mathbf{b} event at time 1—the internal event of \mathbf{P}_1 will resolve the choice at time θ . On the other hand, the process $\mathbf{C}(\mathbf{P}_2)$ is able to perform event \mathbf{b} at time 1. Thus $\langle b \rangle$ is a trace of $\mathbf{C}(\mathbf{P}_2)$ but not of $\mathbf{C}(\mathbf{P}_1)$.

A context $D(R) = \Delta^1 \mathbf{b}\{5\}; \text{Stop} [\triangleright R$ involving disabling may distinguish P_1 from P_2 : $D(P_1)$ is unable to perform \mathbf{b} since the interrupt will occur at time θ . Yet $D(P_2)$ is able to perform \mathbf{b} , since the interrupt need not occur at all, and certainly need not occur before time l when \mathbf{b} becomes available.

These examples illustrate that the timed failures information is not sufficient to predict all the traces of a process in any context.

The extra information required is the time of the first event (internal or external) during a particular execution corresponding to a timed failure. This will be the time at which a choice is resolved, or at which a disabling occurs. It is only the first event that can be detected by these two operators, so information concerning subsequent internal activity is not required.

We will associate a time as well as a timed failure with a given execution. A number of time values could be associated with any particular timed failure. A timed failure triple is a triple:

$$(t, s, X)$$

which is associated with a process iff the process has an execution where the first (internal or external) event occurs at time t , s records the trace of visible events, and X records the visible events that were refused. If $s = \langle \rangle$ then t can take a finite or infinite value; otherwise it must be finite (in fact no greater than the beginning of s). The infinite value in $(\infty, \langle \rangle, X)$ indicates that there is an execution consistent with $(\langle \rangle, X)$ in which no event occurs—in other words, no event need occur before $\text{end}(X)$.

These triples can be extracted directly from the operational semantics in much the same way as the timed failures were earlier:

$$\frac{}{P \text{ exhibits}' (\infty, \langle \rangle, \{\})}$$

$$\frac{P \xrightarrow{a} Q \quad Q \text{ exhibits}' (t, s, X)}{P \text{ exhibits}' (\theta, \langle (\theta, a) \rangle \frown s, X)} \quad [a \neq i]$$

$$\frac{P \xrightarrow{i} Q \quad Q \text{ exhibits}' (t, s, X)}{P \text{ exhibits}' (\theta, s, X)}$$

$$\frac{\begin{array}{l} P \xrightarrow{d} Q \\ Q \text{ exhibits}' (u, s, X - d) \\ \forall a \in \sigma(X \triangleleft d) \bullet \neg(P \xrightarrow{a}) \end{array}}{P \text{ exhibits}' (u + t, s + t, X)}$$

We associate a set of failure triples with a process P

$$\text{failuretriples}(P) = \{(t, s, X) \mid P \text{ exhibits}' (t, s, X)\}$$

The following theorem is immediate:

Theorem 1.4 $\text{traces}(P) = \{\text{events}(s) \mid (t, s, X) \in \text{failuretriples}(P)\}$ □

Proof Each execution of P gives rise to a trace tr , and is consistent with a set of failure triples of P each of whose timed trace s has $\text{events}(s) = tr$. □

Corollary 1.5 Failure triple equivalence is stronger than trace equivalence □

The next theorem points in the opposite direction.

Theorem 1.6 Trace congruence is stronger than failure triple equivalence □

Proof We have two cases to consider: t finite and t infinite.

If t is finite then let $\mathbb{C}_{(t,s,X)}(\mathbb{R})$ be the context

$$\mathbb{C}_{(\langle t, v \rangle) \frown_{s, X \cup [t, t+1) \times \{v\}} (v ; v ; \text{Stop} \ [\triangleright \mathbb{R})}$$

Observe that $t \leq \text{begin}(s)$, so $\text{trace}(t, v) \frown s$ is well-formed. Then for v to be performed and then refused at time t , we must have that the disabling occurs at time t : the first event from the execution of \mathbb{R} which corresponds to (s, X) . Hence $\mathbb{C}_{(t,s,X)}(P)$ will have $\langle \omega \rangle$ as a trace if and only if (t, s, X) is a failure triple of P .

Now consider the case that t is infinite. We first require a lemma:

Lemma 1.7 If $P \text{ exhibits}' (t, \langle \rangle, X)$ and $t \geq \text{end}(X)$ then $P \text{ exhibits}' (\infty, \langle \rangle, X)$ □

This is easily established by examining the proof of \mathbf{P} **exhibits'** $(t, \langle \rangle, X)$, and considering the prefix of it which has inferences based only on delay transitions.

Now if $(\infty, \langle \rangle, X)$ is a behaviour of \mathbf{P} but not of \mathbf{Q} , then there are two possibilities: either $(\langle \rangle, X)$ is not a failure of \mathbf{Q} , in which case there is a context $\mathbf{C}_{(\langle \rangle, X)}$ to distinguish \mathbf{P} from \mathbf{Q} ; or else it is a failure of \mathbf{Q} , in which case any associated time value must be less than $\mathit{end}(X)$, by the lemma above. In this latter case, consider the context

$$\mathbf{D}(\mathbf{R}) = \mathbf{C}_{(\langle \langle \mathit{end}(X), v \rangle \rangle, X)}(\mathbf{R} \square \Delta^{\mathit{end}(X)} v \{0\}; \mathbf{Stop})$$

Then $\mathbf{D}(\mathbf{P})$ can perform an ω , since it is possible that the choice will not be resolved by \mathbf{P} by time $\mathit{end}(X)$. On the other hand, by the lemma above, any execution of \mathbf{Q} exhibiting $(\langle \rangle, X)$ must perform some event strictly earlier than $\mathit{end}(X)$, and so the choice can never be resolved against \mathbf{Q} during such an execution. Thus $\mathbf{D}(\mathbf{Q})$ cannot perform ω . So there is a context that distinguishes \mathbf{P} from \mathbf{Q} .

We conclude that trace congruence implies failure triple equivalence; processes with different timed failure triples can be placed in contexts that yield different traces. \square

Now we can show that failure triple equivalence is a congruence. It will then follow from Theorem 1.6 and Corollary 1.5 that failure triple equivalence is trace congruence.

Theorem 1.8 Failure triple equivalence is a congruence \square

Proof We show this by giving the semantic equations for timed LOTOS in terms of timed failure triples. We must prove that these equations correspond to the result of extracting the timed failure triples directly from the operational semantics. Since the equations are given in a compositional way, they automatically define a congruence.

We establish the following by structural induction on \mathbf{P} , where $\llbracket \cdot \rrbracket$ is the semantic function from ET-LOTOS to the semantic domain TM_L defined below:

$$\forall \mathbf{P} \bullet \mathit{failuretriples}(\mathbf{P}) = \llbracket \mathbf{P} \rrbracket$$

This completes the proof. \square

Denotational semantics for ET-LOTOS

The semantic domain TM_L is defined to be those sets

$$S \subseteq ([0, \infty) \times TT \times FRSET) \cup (\{\infty\} \times \{\langle \rangle\} \times FRSET)$$

of failure triples satisfying the following condition:

$$(C) (t, \langle \rangle, X) \in S \wedge t \geq \text{end}(X) \Rightarrow (\infty, \langle \rangle, X) \in S$$

The semantic function $\llbracket \cdot \rrbracket$ is a mapping $\text{ET-LOTOS} \rightarrow \text{TM}_L$, defined by the following equations:

$$\begin{aligned} \llbracket \text{stop} \rrbracket &= \{(\infty, \langle \rangle, X) \mid X \in \text{RSET}\} \\ \llbracket \text{exit}\{d\} \rrbracket &= \{(\infty, \langle \rangle, X) \mid ([\theta, d] \times \{\delta\}) \cap X = \{\}\} \\ &\cup \\ &\{(t, \langle (t, \delta) \rangle, X) \mid ([\theta, t] \times \{\delta\}) \cap X = \{\} \wedge t \leq d\} \\ \llbracket \text{a@t}\{d\}; \mathbb{Q} \rrbracket &= \{(\infty, \langle \rangle, X) \mid ([\theta, d] \times \{a\}) \cap X = \{\}\} \\ &\cup \\ &\{(u, \langle (u, a) \rangle \frown (s+u), X) \mid ([\theta, u] \times \{a\}) \cap X = \{\} \wedge u \leq d \\ &\quad \wedge \exists v \bullet (v, s, X-u) \in \llbracket \mathbb{Q}[u/t] \rrbracket\} \\ \llbracket \text{i@t}\{d\}; \mathbb{Q} \rrbracket &= \{(\infty, \langle \rangle, X) \mid \text{end}(X) \leq d\} \\ &\cup \\ &\{(u, s+u, X) \mid u \leq d \wedge \exists v \bullet (v, s, X-u) \in \llbracket \mathbb{Q}[u/t] \rrbracket\} \\ \llbracket \Delta^d \mathbb{Q} \rrbracket &= \{(u+d, s+d, X) \mid (u, s, X-d) \in \llbracket \mathbb{Q} \rrbracket\} \\ \llbracket \mathbb{P} \sqcap \mathbb{Q} \rrbracket &= \{(u, s, X) \mid (u, s, X) \in \llbracket \mathbb{P} \rrbracket \wedge \\ &\quad \exists v \geq u \bullet (v, \langle \rangle, X \triangleleft u) \in \llbracket \mathbb{Q} \rrbracket\} \\ &\cup \\ &\{(u, s, X) \mid (u, s, X) \in \llbracket \mathbb{Q} \rrbracket \wedge \\ &\quad \exists v \geq u \bullet (v, \langle \rangle, X \triangleleft u) \in \llbracket \mathbb{P} \rrbracket\} \\ \llbracket \mathbb{P} \llbracket \Gamma \rrbracket \mathbb{Q} \rrbracket &= \{(u, s, X) \mid \exists (u_P, s_P, X_P) \in \llbracket \mathbb{P} \rrbracket, (u_Q, s_Q, X_Q) \in \llbracket \mathbb{Q} \rrbracket \bullet \\ &\quad u = \min\{u_P, u_Q\} \wedge s \in s_P \llbracket \Gamma \rrbracket s_Q \\ &\quad \wedge X \uparrow \Gamma = X_P \uparrow \Gamma \cup X_Q \uparrow \Gamma \\ &\quad \wedge X \setminus \Gamma = X_P \setminus \Gamma \cap X_Q \setminus \Gamma\} \\ \llbracket \text{hide } \Gamma \text{ in } \mathbb{Q} \rrbracket &= \{(u, s \setminus \Gamma, X) \mid (u, s, X \cup [\theta, \text{end}(s, X)]) \times \Gamma \in \llbracket \mathbb{Q} \rrbracket\} \\ \llbracket \mathbb{P} \gg \mathbb{Q} \rrbracket &= \{(u, s, X) \mid (u, s, X \cup [\theta, \text{end}(s, X)]) \times \{\delta\} \in \llbracket \mathbb{P} \rrbracket \wedge \delta \notin \sigma(s)\} \\ &\cup \\ &\{(u, s \frown (s'+t), X) \mid (u, s \frown \langle (t, \delta) \rangle, X \triangleleft t \cup [\theta, t] \times \{\delta\} \in \llbracket \mathbb{P} \rrbracket \\ &\quad \wedge \delta \notin \sigma(s) \\ &\quad \wedge \exists v \bullet (v, s', X-t) \in \llbracket \mathbb{Q} \rrbracket\} \end{aligned}$$

$$\begin{aligned}
[[\mathbf{P} \triangleright \mathbf{Q}]] &= \{(u, s, X) \mid (u, s, X) \in [[\mathbf{P}]] \wedge \\
&\quad \exists v \geq u, v \geq \text{end}(s, X) \bullet (v, \langle \rangle, X) \in [[\mathbf{Q}]]\} \\
&\cup \\
&\{(u, s \widehat{\ } s', X) \mid \exists (u_P, s, X_P) \in [[\mathbf{P}]], (u_Q, s', X) \in [[\mathbf{Q}]] \bullet \\
&\quad X_P = X \triangleleft u_Q \wedge \text{end}(s) \leq u_Q \wedge u = \min(u_P, u_Q)\}
\end{aligned}$$

To give a semantics to recursively defined processes, we impose a metric space structure on TM_L . Define

$$\begin{aligned}
S \triangleleft t &= \{(u, s, X) \mid (u, s, X) \in S, u < t, \text{end}(s, X) < t\} \\
&\cup \\
&\{(\infty, \langle \rangle, X) \mid (t', \langle \rangle, X) \in S, \text{end}(X) < t \leq t'\} \\
d(S_1, S_2) &= \inf\{2^{-t} \mid S_1 \triangleleft t = S_2 \triangleleft t\}
\end{aligned}$$

Theorem 1.9 The distance function d is a metric on TM_L , and (d, TM_L) is a complete metric space. \square

The body of a recursive definition $\mathbf{Y} := \mathbf{Q}$ is a function from \mathbf{Y} to \mathbf{Q} which we require to be time guarded. The mapping on TM_L corresponding to this function is a contraction mapping. By Banach's fixed point theorem [Sut75] this mapping has a unique fixed point, since TM_L is a complete metric space. Full details of this argument are presented in [DaS93].

The semantics of $\mathbf{P} := \mathbf{Y}$ where $\mathbf{Y} = \mathbf{Q}$ is defined to be the unique fixed point of the contraction mapping corresponding to \mathbf{Q} .

A note on the structure of timed refusals

Timed refusals are carefully defined to be left-closed and right-open, so refusal information has the form $[b, e) \times A$. This decision reflects the way in which processes make events available. For a context to discover that events are being refused by a process, the context must itself offer them, and the transition system defined for Basic ET-LOTOS has the property that events are always made available at some first point in time, so we may as well record refusal information as being closed at the bottom. In fact, if we allow refusal information based on open intervals as well, then we may end up distinguishing processes that cannot be told apart in any context. Consider the following:

$$\begin{aligned}
P_1 &:= (i; a\{0\}; \text{stop} \sqcap i; b; \text{stop} \sqcap i; b\{0\}; \text{stop}) \\
P_2 &:= i; ((a\{0\}; \text{stop} \parallel [a] \parallel i\{\infty\}; a; \text{stop}) \sqcap b; \text{stop})
\end{aligned}$$

Each of these processes performs its first event (internal) at time θ . Each can perform **a** at time θ only, and can perform **b** at any time. Each can refuse **b** at time θ provided **a** is not refused at that time; and each can refuse **b** at any time after time θ . In fact, they have the same failure triples, and no context can distinguish them. However, if we were to allow refusal information also in terms of open intervals, then we would expect $(\theta, \langle, [\theta, 1) \times \{a\} \cup (\theta, 1) \times \{b\})$ to be exhibited by the P1, but not by P2, which can exhibit only $(\theta, \langle, [\theta, 1) \times \{a\} \cup [\epsilon, 1) \times \{b\})$ for arbitrarily small $\epsilon > \theta$. Thus we would distinguish processes that no context can distinguish.

The reason for using right-open intervals is more pragmatic: no new information would be gained by including right-closed intervals, again because of the property that events are always enabled at first instants. If a refusal $[b, e] \times A$ is possible for a process, then there will always be some $\epsilon > 0$ for which $[b, e + \epsilon) \times A$ is also a refusal; the events in A will not be enabled until some positive time after e . Hence including refusal information based on closed intervals yields no new information about a process.

If events need not become enabled at particular instants, but may become enabled strictly after some instant, then we will need to record other kinds of refusal information. In Full ET-LOTOS, a selection predicate is permitted to restrict the times at which events may occur, so for example the construction $\mathbf{b} @ \mathbf{t}[t > 0]; \mathbf{stop}$ allows **b** to occur at any time after time θ ; there is no moment at which it becomes available. In fact, depending on the freedom one is allowed in the predicate, intervals may not be suitable at all (for example, if **a** is permitted to occur only at a rational time) and it seems that the appropriate refusal sets will have significantly less structure, perhaps ranging over arbitrary sets of timed events. In any case, contexts can now be constructed that make finer distinctions. For example, the context for Y

$$\text{hide } \{a, b\} \text{ in } ((a; \mathbf{stop} \sqcap \mathbf{b} @ \mathbf{t}[t > 0]; \mathbf{stop} \sqcap \Delta^2 \omega; \mathbf{stop}) \parallel [a, b] \parallel Y) \\ \parallel \{\} \parallel \Delta^1 i; \mathbf{stop}$$

can distinguish P1 (which will allow an ω to be performed) from P2 (which will not).

Comparison with weak bisimulation

Weak timed bisimulation is an equivalence relation defined on processes to abstract away from internal events. The definition in [LeL94] is as follows:

$$P \xrightarrow{*}^d Q \text{ iff } P \xrightarrow{d_1} \xrightarrow{d_2} \dots \xrightarrow{d_n} Q \quad \text{where } d = \sum_{i=1}^n d_i \\ P \xRightarrow{d} Q \text{ iff } P(\xrightarrow{i})^* \xrightarrow{*}^{d_1} (\xrightarrow{i})^* \dots \xrightarrow{*}^{d_n} (\xrightarrow{i})^* Q \quad \text{where } d = \sum_{i=1}^n d_i$$

$$P \xRightarrow{a} Q \text{ iff } P(\overset{i}{\rightarrow})^* \xrightarrow{a} (\overset{i}{\rightarrow})^* Q$$

$$P \xRightarrow{\varepsilon} Q \text{ iff } P(\overset{i}{\rightarrow})^* Q$$

Then a relation R is a weak timed bisimulation if and only if whenever $B_1 R B_2$ then for any α :

1. if $B_1 \xRightarrow{\alpha} B'_1$ then there is some B'_2 such that $B_2 \xRightarrow{\alpha} B'_2$ and $B'_1 R B'_2$
2. if $B_2 \xRightarrow{\alpha} B'_2$ then there is some B'_1 such that $B_1 \xRightarrow{\alpha} B'_1$ and $B'_1 R B'_2$

Weak timed bisimulation equivalence is incomparable with failure triple congruence. For example, the processes $\mathbf{a}; \mathbf{stop}$ and $\mathbf{i}; \mathbf{a}; \mathbf{stop}$ are weak timed bisimulation equivalent, but have different failure triples: for example the failure triple $(1, \langle(1, a)\rangle, \{\})$ is a possible behaviour of the first process, but not of the second (since the time value must be θ with such a trace).

However, weak timed bisimulation *congruence* will be stronger than failure triple equivalence, as will be illustrated by the following examples, where pairs of processes are exhibited which have the same failure triple sets, but are not bisimilar.

Consider the following pair of processes:

$$\begin{aligned} \mathbf{P1} & := (\mathbf{i}; \Delta^3 \mathbf{a}\{5\}; \mathbf{stop}) \sqcap (\mathbf{i}; \Delta^3 \mathbf{b}\{5\}; \mathbf{stop}) \\ \mathbf{P2} & := \mathbf{i}; \Delta^3 ((\mathbf{i}; \mathbf{a}\{5\}; \mathbf{stop}) \sqcap (\mathbf{i}; \mathbf{b}\{5\}; \mathbf{stop})) \end{aligned}$$

Each of $\mathbf{P1}$ and $\mathbf{P2}$ will nondeterministically offer either an \mathbf{a} or a \mathbf{b} over the interval $[3, 8]$, so they have the same failure triple semantics. But $\mathbf{P1}$ resolves the choice at time θ , whereas $\mathbf{P2}$ does not resolve the choice until time 3 , so it can allow three units of time to pass (performing only its internal transition at time θ) and reach a state from which both \mathbf{a} and \mathbf{b} are possible; this delay cannot be matched by $\mathbf{P1}$.

This next pair of processes also have the same failure triple sets—no context can distinguish them—but they are not bisimilar.

$$\begin{aligned} \mathbf{P3} & = \mathbf{i}; \mathbf{stop} \sqcap \mathbf{i}; \mathbf{a}; \mathbf{stop} \\ \mathbf{P4} & = \mathbf{P3} \sqcap \mathbf{i}; \mathbf{a}\{2\}; \mathbf{stop} \end{aligned}$$

The process $\mathbf{P3}$ nondeterministically chooses either to deadlock immediately, or to offer the event \mathbf{a} until it is accepted. Process $\mathbf{P4}$ has these same choices, but may also choose to offer \mathbf{a} for two units of time only. This is a transition that cannot be matched by $\mathbf{P3}$.

A standard example from untimed process algebra also illustrates the fact that failures equivalence is weaker than weak timed bisimulation congruence. Consider the following pair of processes:

$$\begin{aligned} P5 &= i; a\{4\}; \text{stop} \sqcap i; b\{4\}; \text{stop} \\ P6 &= P5 \sqcap i; (a\{4\}; \text{stop} \sqcap b\{4\}; \text{stop}) \end{aligned}$$

The first process offers either an **a** or a **b** event for four units of time. The second may offer either of these, or else may offer its environment the choice of both. They are not bisimilar: **P6** can reach a state in which both **a** and **b** are offered, and **P5** can reach no similar state. Again, no context can distinguish these processes.

Urgency

The model TM_L does not handle the phenomenon of ‘urgency’ or forcing of events, where events must occur by some particular time. Urgent events can be modelled in ET-LOTOS by judicious use of timesteps: the process $a;P \sqcap \text{timestep}$ where $\text{timestep} := \text{timestep}$ forces the event **a** to occur at time θ . However, Leduc and Léonard do not encourage the use of timesteps, and take the view that timesteps are unnecessary in specification.

For the denotational semantics, we were careful to forbid non-time-guarded recursions, and no timestopping processes can be constructed from the operators of Basic ET-LOTOS without such recursive definitions. The handling of forced events and time-stopping processes within a denotational model is a non-trivial extension of the model presented here. In the document [BDS94] an adaptation of an earlier *signals* model is proposed, which includes the modelling of urgent events.

As an example of the sorts of difficulties that arise when internal events can be made urgent, we will examine the urgency operator $\text{urge } a \text{ in } P$ of [BLT94], which makes all occurrences of **a** urgent in process **P**. This may be used on internal or external actions.

Consider the two processes below. They have the same failure triples. Each may offer nondeterministically either **a** or **b** sometime between time 1 and time 3 .

$$\begin{aligned} P7 &:= i; (\Delta^1 ; i\{2\}; a; \text{stop} \sqcap i\{2\}; \Delta^1 ; b; \text{stop}) \\ P8 &:= i; (\Delta^1 ; i\{2\}; b; \text{stop} \sqcap i\{2\}; \Delta^1 ; a; \text{stop}) \end{aligned}$$

If the internal event is made urgent, then the resulting processes are different:

$$\text{urge } i \text{ in } P7 \neq \text{urge } i \text{ in } P8$$

The left-hand process can offer only **b**, and the right-hand can offer only **a**. So failure triples is no longer a congruence in the presence of this operator; more information is required in order to obtain a congruence. However, we conjecture that the **urge** operator restricted to external actions will preserve traces-and-timestops congruence over the language of Basic ET-LOTOS.

Acknowledgements

The authors would like to thank the members of the CONCUR2 project for comments on an earlier version of this material presented in Brighton, March 1994, and also the participants of the COST 247 WG1 workshop on extended process algebras, also in Brighton, July 1994.

Steve Schneider is grateful to the CONCUR2 project and the COST 247 project for providing funding towards this research. Jim Davies is grateful to the DRA for funding.

References

- [BDS94] J. Bryans, J.W. Davies, and S.A. Schneider, *Denotational semantic models for real-time LOTOS*, Brighton COST 247 WG1 workshop on extended process algebras, July 1994.
- [BLT94] T. Bolognesi, F. Lucidi and S. Trigila, *Converging towards a Timed LOTOS Standard*, Computer Standards and Interfaces, Vol. 16 1994.
- [DaS93] J.W. Davies and S.A. Schneider, *Recursion induction for real-time processes*, Formal Aspects of Computing, 5(6), 1993.
- [Hen88] M. Hennessy, *Algebraic theory of processes*, MIT Press, 1988.
- [LeL94] G. Leduc and L. Leonard, *A Formal Definition of Time in LOTOS*, University of Liege, 1994.
- [QFA93] J. Quemada, D. Frutos and A. Azcorra, *TIC: A Timed Calculus*, Formal Aspects of Computing, 1993.
- [Sch94] S.A. Schneider, *An operational semantics for timed CSP* Information and Computation, to appear, 1994
- [Sut75] W. A. Sutherland, *Introduction to Metric and Topological Spaces*, Oxford University Press 1975.