

Timewise Refinement for Communicating Processes

Steve Schneider

Royal Holloway, University Of London
Egham, Surrey, TW20 0EX, UK
steve@dcs.rhnc.ac.uk

Abstract

A theory of timewise refinement is presented. This allows the translation of specifications and proofs of correctness between semantic models, permitting each stage in the verification of a system to take place at the appropriate level of abstraction. The theory is presented within the context of CSP. A denotational characterisation of refinement is given in terms of relations between behaviours at different levels of abstraction. Various properties for the preservation of refinement through parallel composition are discussed. An operational characterisation is also given in terms of timed and untimed tests, and observed to coincide with the denotational characterisation.

Keywords: Concurrency, Real-time systems, Refinement, Timed CSP, Timewise refinement, Verification.

1 Introduction and general theory

Verification of time-critical systems requires the application of necessarily complicated and detailed techniques, reflecting the complex nature of such systems and the detailed and precise requirements upon them. Yet it is often the case that a significant proportion of specifications on timed systems will be concerned with logical behaviour rather than timing behaviour, and proposed implementations will often be correct with respect to these parts of the specification by virtue of their functional properties, independently of their timing properties.

Specifications often split naturally into parts that are dependent upon time, and time-independent parts. For example a timed buffer must meet certain functional constraints: it can output only what has been previously input, and in the same order; it should be deadlock-free, and will eventually become ready to output when non-empty. There might also be requirements on the capacity of the buffer. In addition to these requirements, there might be some desired timing properties, concerning throughput, maximum rate of input, minimum rate of output, and so on. The techniques required to establish these two kinds of properties will often be quite different, and the timing behaviour of a proposed implementation may be irrelevant in verifying the functional correctness. In such cases, we would hope to perform each part of the verification at the appropriate level of abstraction.

Development of a timed system could also follow this pattern. A system may be developed initially with respect to the time-independent aspects of its specification. A subsequent stage of development may introduce precise timed behaviour into the system to ensure that the timing aspects of the specification are also met. It is important to ensure that the correctness of the first stage is preserved when time is introduced.

This paper proposes a framework in the context of CSP for system development and verification at various levels of abstraction, in such a way that consistency is ensured between results obtained at different levels. We investigate refinement relations between processes in different models of the CSP hierarchy [Ree88]. It is important to identify which properties (such as deadlock-freedom or determinism) can be translated between models, since only for such properties can

verifications be mapped up the hierarchy. The more mature and powerful techniques available in the more abstract models (such as model-checking, algebraic techniques, specialised theories for deadlock-freedom, and simply more abstract reasoning) may then be used in conjunction with the more cumbersome and difficult methods required for the more detailed timed aspects of the verification.

This paper investigates two refinement relations in detail, both between an untimed and a timed model of CSP. The first untimed model is concerned only with safety specification. The second is also able to address fairness and (untimed) liveness requirements. The relationships between these two models and the timed infinite model for timed CSP [Sch92, MRS95] will be presented.

General framework

The various models for CSP all embody a particular approach to denotational semantics. A process is modelled in terms of the observations that may be made of it, which may also be considered as the behaviours it may exhibit. If we have a set \mathcal{O} of all possible observations, then a process is identified with a subset of \mathcal{O} . The corresponding semantic model \mathcal{M} consists of those subsets of \mathcal{O} that may be considered to represent some process. A set of healthiness conditions, or axioms, for \mathcal{M} are used to characterise these subsets of \mathcal{O} .

A programming language \mathcal{L} is used for describing processes. Each term in \mathcal{L} together with a binding $\rho \in \mathcal{E} = \mathcal{V} \rightarrow \mathcal{M}$ of variables to values is associated with an element of \mathcal{M} , called its semantics or meaning, by means of a semantic function $\mathcal{F} : \mathcal{L} \rightarrow \mathcal{E} \rightarrow \mathcal{M}$. This meaning consists of the set of possible behaviours for the term. This function is compositional, in the sense that the process associated with any particular term depends only on the processes associated with its components, and on how these components are composed. A program is a term with no free variables. The process associated with such a program is independent of the environment ρ used to evaluate it.

Following the notation introduced in [Hoa85], we give specifications in terms of predicates upon observations. A term P meets a specification $S(o)$ in environment ρ if all of the observations in the semantics of P in ρ meet the corresponding predicate. In this case, we write $P \text{ sat}_\rho S(o)$.

Definition 1.1

$$P \text{ sat}_\rho S(o) \Leftrightarrow \forall o : \mathcal{O} \bullet (o \in \mathcal{F}[[P]]\rho \Rightarrow S(o))$$

□

A term P_1 is refined by another term P_2 when every possible behaviour of P_2 is also a possible behaviour of P_1 in any environment. In this case we write $P_1 \sqsubseteq P_2$, and consider P_2 to be more deterministic than P_1 , since P_1 can do everything P_2 can, and possibly more. If $P_1 \sqsubseteq P_2$, and $P_1 \text{ sat}_\rho S(o)$, then it follows that $P_2 \text{ sat}_\rho S(o)$; refining a process maintains correctness with respect to specifications. This approach also allows terms P to act as specifications: P_2 meets specification P if it is a refinement of P .

The nature of the semantic model is dependent upon the nature of the observation set \mathcal{O} . Observations describe executions of systems at a particular level of abstraction. For example, the use of traces as observations provides only the sequences of events that a system may perform; refusals provide information about contexts in which a system may deadlock; and timed traces also provide information about the times at which events may occur. The use of a particular kind of observation depends on the kind of specification we wish to consider, and the level of abstraction at which we need to consider the system in order to establish correctness.

If we have two different semantic models \mathcal{M}_A and \mathcal{M}_C , based upon different sets of observations \mathcal{O}_A and \mathcal{O}_C respectively, then we are able to analyse systems at two different levels of abstraction; and we may ask when a description at the level of \mathcal{M}_C refines a description at the level of \mathcal{M}_A .

We firstly employ a relation ${}_{A}\mathcal{R}_C \subseteq \mathcal{O}_A \times \mathcal{O}_C$ to relate observations at the different levels of abstraction. The intention is that if $b_A {}_{A}\mathcal{R}_C b_C$ then b_A and b_C are both descriptions, at different levels of abstraction, of the same execution; or alternatively, that b_A is an abstract description of b_C . There is of course no guarantee that the relation ${}_{A}\mathcal{R}_C$ captures a useful relationship between behaviours; this depends upon the intended application of the theory.

An environment $\rho_A : \mathcal{V} \rightarrow \mathcal{M}_A$ is refined by an environment $\rho_C : \mathcal{V} \rightarrow \mathcal{M}_C$ if $\rho_A(Y)$ contains every abstract behaviour related to each concrete behaviour in $\rho_C(Y)$. This reflects the idea that the concrete refines the abstract in that it contains fewer behaviours:

Definition 1.2

$$\rho_A \sqsubseteq_{{}_{A}\mathcal{R}_C} \rho_C \Leftrightarrow \forall Y : \mathcal{V} \bullet {}_{A}\mathcal{R}_C^{-1}(\rho_C(Y)) \subseteq \rho_A(Y)$$

□

The refinement relation may now be given between terms from \mathcal{M}_A and terms from \mathcal{M}_C with respect to the relation ${}_{A}\mathcal{R}_C$ and a pair of environments:

Definition 1.3

$$P_A \sqsubseteq_{{}_{A}\mathcal{R}_C} P_C \Leftrightarrow \forall \rho_A, \rho_C \bullet \rho_A \sqsubseteq_{{}_{A}\mathcal{R}_C} \rho_C \Rightarrow {}_{A}\mathcal{R}_C^{-1}(\mathcal{F}_C[P_C]\rho_C) \subseteq \mathcal{F}_A[P_A]\rho_A$$

□

The following lemma states that refinement between models is monotonic with respect to refinement within models:

Lemma 1.4

$$P_{A1} \sqsubseteq_A P_{A2} \sqsubseteq_{{}_{A}\mathcal{R}_C} Q_{C1} \sqsubseteq_C Q_{C2} \Rightarrow P_{A1} \sqsubseteq_{{}_{A}\mathcal{R}_C} Q_{C2}$$

□

A verification of P_A may be translated into a verification of P_C by use of the following inference rule, whose soundness follows from the definitions above:

$$\frac{\begin{array}{l} \rho_A \sqsubseteq_{{}_{A}\mathcal{R}_C} \rho_C \\ P_A \mathbf{sat}_{\rho_A} S_A(o_A) \\ P_A \sqsubseteq_{{}_{A}\mathcal{R}_C} P_C \end{array}}{P_C \mathbf{sat}_{\rho_C} \forall o_A \bullet (o_A {}_{A}\mathcal{R}_C o_C \Rightarrow S_A(o_A))}$$

We may thus consider the specification

$$S_C(o_C) = \forall o_A \bullet (o_A {}_{A}\mathcal{R}_C o_C \Rightarrow S_A(o_A))$$

to be the translation of $S_A(o_A)$.

Observe that if $\mathcal{M}_A = \mathcal{M}_C$, and the relation ${}_{A}\mathcal{R}_C$ is the identity relation, then the refinement relation $\sqsubseteq_{{}_{A}\mathcal{R}_C}$ is simply refinement under the non-deterministic ordering; and the rule states that if a program meets a specification, then so too does any refinement of it. The refinement relation within a single model \mathcal{M}_A is defined as

$$P_1 \sqsubseteq_A P_2 \Leftrightarrow \forall \rho_A \bullet \mathcal{F}_A[P_2]\rho_A \subseteq \mathcal{F}_A[P_1]\rho_A$$

The above rule is sound for any refinement relation ${}_{A}\mathcal{R}_C$, by virtue of the definitions of the operators involved. We shall later take advantage of the dual property of *completeness*, which states that whenever the conclusion to the above rule is true, then a term P_A and environment ρ_A can be found such that the antecedents to the rule are true. In general, this property will not hold, but it may hold for some particular refinement relations ${}_{A}\mathcal{R}_C$, and in fact we will establish that it holds for the relations presented in this paper.

Definition 1.5 A refinement relation ${}_{A}\mathcal{R}_C$ is said to be *complete* if whenever the conclusion of the above rule holds for any S_A , then there is some P_A and ρ_A for which the three antecedents hold. \square

The following lemma provides a necessary and sufficient condition for a refinement relation to be complete, in the case where \mathcal{F}_A is surjective for any ρ_A (as is the case for the untimed semantic models used in this paper).

Lemma 1.6 The relation ${}_{A}\mathcal{R}_C$ is complete if and only if ${}_{A}\mathcal{R}_C^{-1}(T)$ is an element of \mathcal{M}_A whenever T is an element of \mathcal{M}_C \square

Proof Assume that ${}_{A}\mathcal{R}_C^{-1}(T)$ is an element of \mathcal{M}_A whenever T is an element of \mathcal{M}_C . Consider the conclusion $P_C \mathbf{sat}_{\rho_C} \forall o_A \bullet (o_A \mathcal{A}\mathcal{R}_C o_C \Rightarrow S_A(o_A))$. Then it follows that ${}_{A}\mathcal{R}_C^{-1}(\mathcal{F}_C[P_C]\rho_C) \mathbf{sat} S_A(o_A)$. Define ρ_A by defining it on each Y as $\rho_A(Y) = {}_{A}\mathcal{R}_C^{-1}(\rho_C(Y))$. Then ρ_A is well-defined since ${}_{A}\mathcal{R}_C^{-1}(\rho_C(Y))$ is always in \mathcal{M}_A . Since ${}_{A}\mathcal{R}_C^{-1}(\mathcal{F}_C[P_C]\rho_C) \in \mathcal{M}_A$ and \mathcal{F}_A is surjective there is some P_A such that $\mathcal{F}_A[P_A]\rho_A = {}_{A}\mathcal{R}_C^{-1}(\mathcal{F}_C[P_C]\rho_C)$. Hence all three antecedents of the inference rule hold.

If on the other hand there is some P_C for which ${}_{A}\mathcal{R}_C^{-1}(\mathcal{F}_C[Q_C]\rho_C) \notin \mathcal{M}_A$, then consider P_A and ρ_A for which $P_A \sqsubseteq_{\mathcal{A}\mathcal{R}_C} P_C$. The term P_A will not meet the specification $S_A(o_A)$ defined by $S_A(o_A) = o_A \in {}_{A}\mathcal{R}_C^{-1}(\mathcal{F}_C[P_C]\rho_C)$. This is because ${}_{A}\mathcal{R}_C^{-1}(\mathcal{F}_C[P_C]\rho_C)$ is a proper subset of $\mathcal{F}_A[P_A]\rho_A$, and so will contain some behaviour o_A which breaks the specification $S_A(o_A)$; yet the term Q_C in ρ_C meets the translation of that specification:

$$\forall o_A \bullet o_A \mathcal{A}\mathcal{R}_C o_C \Rightarrow S_A(o_A)$$

Hence there is some specification S_A for which the conclusion holds, but for which the second and third antecedents of the rule cannot both hold. \square

The definitions given above applied to terms, of which programs are a special case. In the case of programs (but not of general terms) the environment information does not influence the semantics, and environmental considerations may be elided from the definitions and rules below. The satisfaction relation for programs may be given independently of environmental information:

Definition 1.7

$$P \mathbf{sat} S(o) \Leftrightarrow \exists \rho \bullet P \mathbf{sat}_{\rho} S(o)$$

\square

The inference rule given above simplifies to the following in the case of programs:

$$\frac{\begin{array}{l} P_A \mathbf{sat} S_A(o_A) \\ P_A \sqsubseteq_{\mathcal{A}\mathcal{R}_C} P_C \end{array}}{P_C \mathbf{sat} \forall o_A \bullet (o_A \mathcal{A}\mathcal{R}_C o_C \Rightarrow S_A(o_A))}$$

and is complete in the same circumstances.

Completeness of the refinement relation allows the following rule for programs:

$$\frac{P_C \mathbf{sat} \forall o_A \bullet (o_A \mathcal{A}\mathcal{R}_C o_C \Rightarrow S_A(o_A))}{\exists P_A \bullet P_A \sqsubseteq_{\mathcal{A}\mathcal{R}_C} P_C \wedge P_A \mathbf{sat} S_A(o_A)}$$

The discussion so far has all been independent of the details of the programming language. In order to prove that one program refines another using the above theory, it is necessary to

calculate the semantics of each program in the appropriate model, and then check that the refinement relation holds between them. Compositionality often plays a critical role in breaking down verification obligations on large systems to manageable components. We aim to exploit the compositional nature of program semantics, and so we investigate when refinements established between components of abstract and concrete systems mean that the entire abstract system is refined by the entire concrete system.

Our aim is to find relationships concerning the operators of the language \mathcal{L} so that refinement between terms and programs may be established without resorting to explicit calculation of their semantics, by reasoning at the syntactic level.

A syntactic operator \oplus' of \mathcal{L} is a refinement of operator \oplus if \oplus' -combinations of refinements of processes refine \oplus -combinations of the original processes:

Definition 1.8 An operator \oplus' of the language \mathcal{L} with arity α refines operator \oplus with the same arity, if

$$(\forall i < \alpha \bullet P_i \sqsubseteq_{\mathcal{A}\mathcal{R}\mathcal{C}} P'_i) \Rightarrow \oplus(P_i \mid i < \alpha) \sqsubseteq_{\mathcal{A}\mathcal{R}\mathcal{C}} \oplus'(P'_i \mid i < \alpha)$$

□

In general this definition applies to two languages \mathcal{L} and \mathcal{L}' , but for the purposes of this paper we will work within a single language.

The framework presented above is very well known. But to go further, we must focus on particular models, languages, and refinement relations. We are interested in conditions for refinement relations to exist between programs (which will vary from relation to relation), and how specifications translate between models.

In this paper we are concerned with mapping results up the hierarchy of untimed and timed models for CSP. We will concentrate on two relations in detail, both from an untimed to a timed model: one from the untimed traces model, which is used for analysis of safety properties; and one from the untimed infinite traces model [Ros88] (which also contains failures and divergences), a more sophisticated untimed model supporting consideration of liveness issues.

2 Communicating Sequential Processes

Syntax

The language of Communicating Sequential Processes (CSP) is given as follows:

$$P ::= Loop \mid Stop \mid Skip \mid P ; P \mid P \triangleright^t P \mid P \square P \mid a : A \longrightarrow P_a \mid \prod_{i \in I} P_i \\ \mid P \mathbin{\|}_A P \mid P \parallel P \mid P \setminus A \mid f(P) \mid f^{-1}(P) \mid Y \mid \mu Y \circ P$$

Here the set A is a subset of the universal set of events Σ ; I is a subset of the set of indexes \mathcal{I} ; f is a function $\Sigma \rightarrow \Sigma$; Y is drawn from the set of process variables \mathcal{V} ; and t is drawn from the set of times, the non-negative real numbers. The fixed point operation $\mu Y \circ P$ binds free occurrences of Y in P . The *programs* of the language are those terms with no free process variables. Observe that the requirement that arguments to an arbitrary internal choice be indexed from \mathcal{I} ensures that the size of the choice is bounded by some cardinal, which is required to ensure that the language is well-defined by this syntax. For more details of the technicalities, see [MRS95].

The constructors given above represent respectively: the most non-deterministic process; deadlock; successful termination; sequential composition; timeout; external choice; prefix choice; non-deterministic choice; synchronised parallel; interleaving parallel; interface abstraction or hiding;

two forms of alphabet renaming; process variable; and recursion. For a more detailed discussion of the language, the reader is referred to [DaS95].

The following abbreviations often prove useful:

$$\begin{aligned}
Wait\ t &= Stop \triangleright^t Skip \\
b \longrightarrow P &= a : \{b\} \longrightarrow P(a) \text{ (where } P(b) \text{ is defined to be } P) \\
b \xrightarrow{t} P &= b \longrightarrow Wait\ t ; P \\
P \parallel Q &= P \parallel_{\Sigma} Q \\
P \sqcap Q &= \prod_{i \in \{1,2\}} P_i \text{ (where } P_1 \text{ is defined to be } P \text{ and } P_2 \text{ is defined to be } Q) \\
P \parallel_A Q &= h_A(f_A(P) \parallel_{AU0.\Sigma} g_A(Q))
\end{aligned}$$

where

$$\begin{aligned}
f_A(x) &= \begin{array}{ll} x & \text{if } x \in A \\ 0.x & \text{otherwise} \end{array} \\
g_A(x) &= \begin{array}{ll} x & \text{if } x \in A \\ 1.x & \text{otherwise} \end{array} \\
h_A(x) &= \begin{array}{ll} y & \text{if } x \in 0.y \\ y & \text{if } x \in 1.y \\ x & \text{otherwise} \end{array}
\end{aligned}$$

When modelling timed processes, we must take care to ensure that recursive calls are time guarded, so that a minimum delay must elapse between successive recursive calls. This is achieved by ensuring that every instance of the process variable of a recursive term should appear in the right-hand argument of a non-zero timeout. A set of rules for determining when a term is time guarded is detailed in [DaS93].

Notation

The set Σ is the set of visible events. Variables a, b, c are taken to range over Σ . If $M \subseteq \Sigma$ then $c.M$ is shorthand for the set of events $\{c.m \mid m \in M\} \subseteq \Sigma$. The variable t ranges over \mathbf{R}_+ , the set of non-negative real numbers. Variable tr ranges over Σ^* , finite sequences of events from Σ ; u ranges over Σ^ω , infinite sequences of events from Σ ; $X \subseteq \Sigma$ denotes a set of events; s ranges over $(\mathbf{R}_+ \times \Sigma)^* \cup (\mathbf{R}_+ \times \Sigma)^\omega$, the (finite and infinite) sequences of timed visible events; we use $\aleph \subseteq \mathbf{R}_+ \times \Sigma$ to represent a timed refusal, a set of timed visible events. In fact the timed refusals used within the timed model have a particular structure, as defined later in *IRSET*.

We use the following operations on (untimed and timed) sequences of events: $\#w$ is the length of the sequence w ; $w_1 \widehat{\ } w_2$ denotes the concatenation of w_1 and w_2 ; $last(w \widehat{\ } x) = x$ yields the last item in the list. The notation $w_1 \preceq w_2$ means that w_1 is a (not necessarily contiguous) subsequence of w_2 ; the notation $w_1 \leq w_2$ means that w_1 is a prefix of w_2 ; the notation $w_1 \leq_i w_2$ means that w_1 is a prefix of w_2 such that $\#w_2 \leq \#w_1 + i$.

The following projections are defined on untimed sequences by list and set comprehension:

$$\begin{aligned}
tr \upharpoonright A &= \langle a \mid a \leftarrow tr, a \in A \rangle \\
tr \setminus A &= \langle a \mid a \leftarrow tr, a \notin A \rangle \\
tr \downarrow c &= \langle x \mid a \leftarrow tr, a = c.x \rangle \\
\sigma(tr) &= \{a \mid tr \upharpoonright \{a\} \neq \langle \rangle\}
\end{aligned}$$

For timed sequences, we define the beginning and end of a sequence in the following way: $begin(\langle(t, a)\rangle \widehat{s}) = t$, $end(s \widehat{\langle(t, a)\rangle}) = t$, $end(s) = \infty$ when s is infinite, and for convenience $begin(\langle\rangle) = \infty$ and $end(\langle\rangle) = 0$. The following projections on timed sequences are defined by list and set comprehension:

$$\begin{aligned}
s \triangleleft t &= \langle(t', a) \mid (t', a) \leftarrow s, t' \leq t\rangle \\
s \triangleleft t &= \langle(t', a) \mid (t', a) \leftarrow s, t' < t\rangle \\
s \uparrow t &= \langle(t', a) \mid (t', a) \leftarrow s, t' = t\rangle \\
s \upharpoonright A &= \langle(t', a) \mid (t', a) \leftarrow s, a \in A\rangle \\
s \setminus A &= \langle(t', a) \mid (t', a) \leftarrow s, a \notin A\rangle \\
s - t &= \langle(t' - t, a) \mid (t', a) \leftarrow s, t' \geq t\rangle \\
strip(s) &= \langle a \mid (t', a) \leftarrow s \rangle \\
\sigma(s) &= \{a \mid s \upharpoonright \{a\} \neq \langle\rangle\}
\end{aligned}$$

We also define a number of projections on timed refusal sets:

$$\begin{aligned}
\aleph \triangleleft t &= \{(t', a) \mid (t', a) \in \aleph, t' < t\} \\
\aleph \triangleright t &= \{(t', a) \mid (t', a) \in \aleph, t' \geq t\} \\
\aleph \upharpoonright A &= \{(t', a) \mid (t', a) \in \aleph, a \in A\} \\
\aleph - t &= \{(t' - t, a) \mid (t', a) \in \aleph, t' \geq t\} \\
\sigma(\aleph) &= \{a \mid (t', a) \in \aleph\} \\
end(\aleph) &= sup\{t' \mid (t', a) \in \aleph\}
\end{aligned}$$

We will use $(s, \aleph) - t$ as an abbreviation for $(s - t, \aleph - t)$, and $end(s, \aleph)$ for $max\{end(s), end(\aleph)\}$.

Semantic models

The hierarchy of models presented in [Ree88] (see Figure 1) supports reasoning at a number of levels of abstraction, allowing aspects of behaviour dependent upon refusal information, stability information, or timing information to be included as required. In addition to Reed's hierarchy of models, we have the infinite timed model \mathcal{M}_{TI} , presented in [Sch92] and [MRS95]; and the untimed infinite traces model \mathcal{M}_{UI} of [Ros88], which is an extension of the failures-divergences model of [BrR85]. In this paper we will focus on the three models which yield the most general results concerning refinement: the untimed traces model \mathcal{M}_{UT} , and the two infinite models. These three models are presented in full, together with their corresponding semantic functions, in Appendix A. The semantic functions \mathcal{F}_{UT} and \mathcal{F}_{UI} are surjective in the sense required for completeness in Lemma 1.6; this follows directly from similar results presented in [Sch89].

The untimed traces model

Observations in the model \mathcal{M}_{UT} are simply finite sequences of events, or *traces*. A trace tr of a system is a record of the events performed during some (partial) execution of the system. Thus the observation set \mathcal{O}_{UT} is defined to be Σ^* , where Σ is the universal set of events.

The model \mathcal{M}_{UT} is the set of nonempty prefix closed subsets S of \mathcal{O}_{UT} .

The untimed infinite traces model

This model is first described in [Ros88]. In other presentations, processes consist of three components, modelling the three kinds of observation that may be made: a failure set $F \subseteq \Sigma^* \times \mathbf{P}(\Sigma)$;




Figure 1: Reed's hierarchy and additional models

a divergence set $D \subseteq \Sigma^*$; and an infinite traces set $I \subseteq \Sigma^\omega$. A divergence $tr \in D$ is a sequence of events such that after some prefix of tr the system may perform an infinite sequence of internal actions. A failure $(tr, X) \in F$ is an observation of a system if either the sequence of external events tr may be observed during an execution, after which no further internal progress may be made and the process refuses to engage in any event from the set X ; or else tr is a divergent trace. An infinite trace u is an infinite sequence of actions such that either the system may perform the whole trace during a single execution, or else some prefix of it is a divergent trace.

For the sake of uniformity within this presentation, we consider a process to consist of a single set S of pairs, where the first component is a label from the set $\{f, d, i\}$, and the second component is a behaviour from the corresponding behaviour set. Thus S is a subset of

$$\{f\} \times (\Sigma^* \times \mathbf{P}(\Sigma)) \cup \{d\} \times \Sigma^* \cup \{i\} \times \Sigma^\omega$$

The timed infinite traces model

In this model, the times at which events are performed and refused are recorded. This model assumes that systems are finitely variable: an infinite sequence of internal and external actions may not be performed in a finite time. Thus the only infinite traces that may be observed must take infinitely long to occur. Furthermore, since a change in the set of events made available to the environment is considered to correspond to an internal action, this model needs to consider only those refusal sets which contain finitely many changes in any finite interval.

The set of traces $T\Sigma_{\leq}^\omega$ and refusal sets $IRSET$ are adequate for capturing all possible observations

of finitely variable systems:

$$\begin{aligned}
T\Sigma_{\leq}^{\omega} &= \{s \in (\mathbf{R}^+ \times \Sigma)^* \cup \mathbf{R}^+ \times \Sigma\}^{\omega} \mid \langle (t_1, a_1), (t_2, a_2) \rangle \preceq s \Rightarrow t_1 \leq t_2 \\
&\quad \wedge \#s = \infty \Rightarrow \text{end}(s) = \infty\} \\
RTOK &= \{[b, e) \times A \mid 0 \leq b < e < \infty \wedge A \subseteq \Sigma\} \\
RSET &= \{\bigcup R \mid R \subseteq RTOK \wedge R \text{ is finite}\} \\
IRSET &= \{\bigcup R \mid R \subseteq RTOK \wedge \forall t \bullet (\bigcup R) \triangleleft t \in RSET\}
\end{aligned}$$

Behaviours consist of (trace, refusal) pairs. In contrast to the untimed case, the refusal is observed during the occurrence of the trace, rather than simply afterwards.

For example, the behaviour $(\langle (3, a), (3, d), (8, b) \rangle, [0, 20) \times \{c\})$ is a record indicating that the process was observed to refuse event c beginning at time 0, that while it was continuing to do so, it performed event a and then d at time 3, and then event b at time 8. Finally, the observer stopped watching c being refused at time 20.

As usual, a process consists of the set of possible behaviours that may be observed of it. A detailed discussion of the model appears in [Sch92, MRS95].

Example

Define the program AB as follows:

$$AB = \mu Y \circ (a \xrightarrow{3} Y \triangleright^5 b \longrightarrow \text{Stop})$$

Then $\mathcal{F}_{UT}\llbracket AB \rrbracket$ contains both the traces $\langle a \rangle$ and $\langle a, a, b \rangle$, but not trace $\langle b, a \rangle$. The untimed infinite semantics $\mathcal{F}_{UI}\llbracket AB \rrbracket$ contains failures $(f, (\langle a, a \rangle, \{a\}))$ and $(f, (\langle a, b \rangle, \{b, c\}))$, but not $(f, \langle a \rangle, \{b\})$ or $(f, \langle b, a \rangle, \{c\})$; it contains the infinite trace $(i, \langle a, a, a, \dots \rangle)$; and it contains no divergences.

The timed behaviours $\mathcal{F}_{TI}\llbracket AB \rrbracket$ include $(\langle (2, a), (9, a) \rangle, [0, 6) \times \{b\})$: the process may perform event a at time 2, and again at time 9, while refusing to perform b between times 0 and 6. The behaviour $(\langle \rangle, [0, 5) \times \{b\} \cup [5, \infty) \times \{a\})$ is also possible: if no external events are performed, then b will be refused for the first 5 units of time, after which the timeout will occur, and a will be refused thereafter. Neither $(\langle (2, a) \rangle, [0, 1) \times \{a\})$ nor $(\langle \rangle, [0, 10) \times \{b\})$ are possible timed behaviours of $\mathcal{F}_{UI}\llbracket AB \rrbracket$.

3 Timewise refinement

3.1 Trace refinement

We consider an untimed trace to be an abstract description of a timed failure if the trace corresponds to the sequence of events in the timed trace. We thus define the refinement relation between untimed traces \mathcal{O}_{UT} and timed failures \mathcal{O}_{TI} as follows:

$$tr \mathrel{_{UT}\mathcal{R}_{TI}} (s, \aleph) \iff tr = \text{strip}(s)$$

For a timed trace s , the sequence $\text{strip}(s)$ is the trace s with the times removed from the events.

Theorem 3.1 This refinement relation is complete □

Proof By Lemma 1.6 it is enough to show that $T = \mathrel{_{UT}\mathcal{R}_{TI}}^{-1} (U)$ is a well-defined process for any timed process U . But $\mathrel{_{UT}\mathcal{R}_{TI}}^{-1} (U) = \{\text{strip}(s) \mid (s, \aleph) \in U\}$, and this set is clearly non-empty (since U is) and prefix closed (since U is), and hence it is a well-defined process, meeting the definition in Appendix A. □

It turns out that all of the CSP operators preserve this refinement relation:

Theorem 3.2 Given any CSP operator \oplus , and two vectors of terms of length $\text{arity}(\oplus)$ \underline{P} and \underline{Q} such that $\underline{P} \sqsubseteq_{\text{UT}\mathcal{R}_{\text{TI}}} \underline{Q}$, then

$$\oplus(\underline{P}) \sqsubseteq_{\text{UT}\mathcal{R}_{\text{TI}}} \oplus(\underline{Q})$$


□

Proof By a structural induction over the CSP syntax. The timed and untimed semantics (given in Appendix A) of each CSP operator need to be considered in turn. The case for all operators except recursion follows the same pattern. As an example, we provide the case of the parallel operator. The proof for recursion follows the structure of that given for failure refinement in Lemma 3.5 below, except that in the case here continuity of the CSP operators in \mathcal{M}_{UT} means that a standard induction rather than a transfinite induction is sufficient.

The proof for the parallel operator runs as follows. Assume that $P_i \sqsubseteq_{\text{UT}\mathcal{R}_{\text{TI}}} Q_i$ for $i = 1, 2$. Then we aim to prove that $P_1 \parallel P_2 \sqsubseteq_{\text{UT}\mathcal{R}_{\text{TI}}} Q_1 \parallel Q_2$.

Consider some $(s, \aleph) \in \mathcal{F}_{\text{TI}}[Q_1 \parallel Q_2]\rho_C$. We must prove that $\text{strip}(s) \in \mathcal{F}_{\text{UT}}[P_1 \parallel P_2]\rho_A$. By the timed semantics of the parallel operator there are \aleph_i , $i = 1, 2$ such that $(s, \aleph_i) \in \mathcal{F}_{\text{TI}}[Q_i]\rho_C$. Since $P_i \sqsubseteq_{\text{UT}\mathcal{R}_{\text{TI}}} Q_i$ (by the inductive hypothesis) we have $\text{strip}(s) \in \mathcal{F}_{\text{UT}}[P_i]\rho_A$ for each i , and hence $\text{strip}(s) \in \mathcal{F}_{\text{UT}}[P_1 \parallel P_2]\rho_A$ by the untimed trace semantics of the parallel operator. □

Corollary 3.3 For any program P , $P \sqsubseteq_{\text{UT}\mathcal{R}_{\text{TI}}} P$ □



The payoff from this result is that whenever a trace specification may be verified of a CSP program in the untimed traces model, then it follows immediately that its translation into the timed model will hold for the same program on its more complicated semantics. Also $\text{Skip} \sqsubseteq_{\text{UT}\mathcal{R}_{\text{TI}}} \prod_{t \in I} \text{Wait } t$ for any set of times I , so arbitrary delays can be introduced into programs while still preserving refinement, since if $P \sqsubseteq_{\text{UT}\mathcal{R}_{\text{TI}}} Q$, then it follows that $\text{Skip}; P = P \sqsubseteq_{\text{UT}\mathcal{R}_{\text{TI}}} \prod_{t \in I} \text{Wait } t; Q$. Thus an untimed verification can be carried out and delays inserted subsequently.

The translation of a specification $S(\text{tr})$ on traces will be $S'(s, \aleph)$, given by

$$S'(s, \aleph) \Leftrightarrow \forall \text{tr} \in \Sigma^* \bullet (\text{tr} \text{ UT}\mathcal{R}_{\text{TI}}(s, \aleph) \Rightarrow S(\text{tr}))$$

If S is admissible (i.e. $(\forall \text{tr} < u \bullet S(\text{tr})) \Rightarrow S(u)$ for every infinite trace u) then the translation is equivalent on processes to the specification $S''(s, \aleph) = S(\text{strip}(s))$. Thus admissible specifications may be translated to timed specifications by a simple substitution of the free variable. Since most safety specifications are admissible, this does not amount to a practical limitation.

As an example, consider the safety requirement that no event should be performed after a b . This is given by

$$S(\text{tr}) = \forall \text{tr}_0, \text{tr}_1 \bullet (\text{tr} = \text{tr}_0 \frown \langle b \rangle \frown \text{tr}_1) \Rightarrow \text{tr}_1 = \langle \rangle$$

A verification in the traces model that program AB satisfies this specification would be quite straightforward. We may translate this verification to the timed model, and conclude that $AB \text{ sat } S(\text{strip}(s))$ in that model. This may then be used in a timed verification. For example, consider the timed specification that event a should never be performed within 8 time units of any b :

$$\langle (t, b) \rangle \preceq s \Rightarrow a \notin \sigma(s \uparrow (t-8, t+8))$$

This specification reads as follows: if (t, b) is recorded in the trace s , then a does not appear in the set of events recorded in s during the interval $(t-8, t+8)$. Then the untimed specification tells us that a cannot occur after b , i.e. in the interval $(t, t+8)$, so the only cases to consider in the timed model are a occurring before b , or at the same time, i.e. the interval $(t-8, t]$. For this case, a timed analysis on AB is required.

In general (even when $S(\text{tr})$ is inadmissible), $S(\text{tr})$ is translated to $(\#s < \infty \Rightarrow S(\text{strip}(s)))$.

3.2 Failures Refinement

We think of a process refusing a particular set, in the untimed sense, if it eventually reaches a state after which no event from that set is possible. In the timed world, this corresponds to the information that there is some time after which the set may be continuously refused. Thus for a timed behaviour (s, \aleph) with finite timed trace s , an abstract view of this behaviour would be an untimed version $\text{strip}(s)$ of the trace, and for any set X , if there is some t for which $[t, \infty) \times X$ is contained in \aleph , then \aleph is evidence that X may eventually be refused forever.

$$(f, (tr, X)) \text{ UI } \mathcal{R}_{TI} (s, \aleph) \Leftrightarrow tr = \text{strip}(s) \wedge \exists t \bullet [t, \infty) \times X \subseteq \aleph$$

Relating timed infinite traces to untimed ones, we obtain the following refinement relation between \mathcal{O}_{UI} and \mathcal{O}_{TI} :

$$(i, u) \text{ UI } \mathcal{R}_{TI} (s, \aleph) \Leftrightarrow u = \text{strip}(s)$$

Observe that there is no timed version of divergence in this model.

Theorem 3.4 This refinement relation is complete. □

Proof We need to show that $T = \text{UI } \mathcal{R}_{TI}^{-1}(U)$ is a well-defined process for any timed process U (i.e. T meets axioms 1–8 given in Appendix A).

Axiom 2 for $U \in \mathcal{M}_{TI}$ yields axioms 1, 2 and 6 for T ; axioms 4, 5 and 7 are trivial for T since T has no behaviours of the form (d, tr) .

Axiom 3 for U yields axiom 3 for T : if $(f, (tr, X)) \in T \wedge \forall a \in Y \bullet (f, (tr \frown \langle a \rangle, \{\})) \notin T$ then $(s, [t_0, \infty) \times X) \in U$ for some $t_0 > \text{end}(s)$, and also for any $a \in Y$ and any t we have $(s \frown \langle (t, a) \rangle, \{\}) \notin U$ (since their untimed counterparts are not in T). Now by axiom 3 for U there is some $\aleph' \supseteq \aleph$ which contains all unperformable events. Assume for a contradiction that $(t, a) \notin \aleph'$ for some $t \geq t_0$ and $a \in Y$. Then $((s \triangleleft t) \frown \langle (t, a) \rangle, \aleph' \triangleleft t) \in U$, and so $(s \frown \langle (t, a) \rangle, \{\}) \in U$ by axiom 2 (since $s = s \triangleleft t$), yielding the contradiction. Hence $(t, a) \in \aleph'$ whenever $t \geq t_0$, $a \in Y$, yielding $[t_0, \infty) \times Y \subseteq \aleph'$, and also $[t_0, \infty) \times (X \cup Y) \subseteq \aleph'$, and so $(s, [t_0, \infty) \times (X \cup Y)) \in U$, and hence $(f, (tr, X \cup Y)) \in S$, establishing that axiom 3 for S holds.

It remains only to establish axiom 8 for S . Consider a behaviour $(f, (s, \{\})) \in T$. Then there must be some timed trace s_0 such that $(s_0, \{\}) \in U$ and $\text{strip}(s_0) = s$. Now by axiom 4 for \mathcal{M}_{TI} there must be some process $U' \in \mathcal{CL}$ such that $(s_0, \{\}) \in U'$ and $U' \subseteq U$.

Let $c : \mathbf{P}(\mathbf{R}) \rightarrow \mathbf{R}$ be a choice function. Then define:

$$\begin{aligned} V_0 &= \{s_0\} \\ V_{n+1} &= \{s \widehat{\ } \langle (t_a, a) \rangle \mid s \in V_n \wedge a \in \Sigma \\ &\quad \wedge \{t \mid (s \widehat{\ } \langle (t, a) \rangle, \{\}) \in U'\} \neq \{\} \\ &\quad \wedge t_a = c(\{t \mid (s \widehat{\ } \langle (t, a) \rangle, \{\}) \in U'\}) \end{aligned}$$

$$V = \overline{\bigcup_{i \in \mathbf{N}} T_i}$$

Since U' is finitely variable and closed, any infinite trace in V comes from a legitimate infinite trace in U' . And since U' meets axiom 3 for \mathcal{M}_{TI} , the set of all events that do not extend a given finite trace in V must be refusable for all time after the corresponding timed trace in U' , since there is no time after that timed trace at which any of those events is possible. It follows that V is a set of traces that establishes that axiom 8 holds for T . \square

A study of the CSP operators reveals the following:

Theorem 3.5 Given any CSP operator \oplus except parallel composition, and two vectors of terms of length $\text{arity}(\oplus)$ \underline{P} and \underline{Q} such that $\underline{P} \sqsubseteq_{\text{vT}\mathcal{R}_{TI}} \underline{Q}$, then

$$\oplus(\underline{P}) \sqsubseteq_{\text{vT}\mathcal{R}_{TI}} \oplus(\underline{Q})$$

\square

Proof The structure of the cases of this proof is entirely similar to that of Lemma 3.2. All operators follow the same pattern, with the exception of the proof of recursion, which we now give.

Assume that $\underline{P} \sqsubseteq_{\text{vT}\mathcal{R}_{TI}} \underline{Q}$. Now consider a particular pair of environments $\rho_A \sqsubseteq_{\text{vT}\mathcal{R}_{TI}} \rho_C$. Let $U = (\mathcal{F}_{TI}[\mu Y \circ \underline{Q}]\rho_C)$ and $T = {}_{\text{vT}\mathcal{R}_{TI}}^{-1}(U)$. We must prove $T \subseteq \mathcal{F}_{UI}[\mu Y \circ \underline{P}]\rho_A$. Observe that $T \in \mathcal{M}_{UI}$ by the proof of Theorem 3.4

We use the fact that $\mathcal{F}_{UI}[\mu Y \circ \underline{P}]\rho_A = F^\alpha(\perp)$ for some ordinal α , where $\perp = \mathcal{F}_{UI}[\text{Loop}]\rho_A$, and

$$F(S) = \mathcal{F}_{UI}[\underline{P}](\rho[S/Y])$$

Hence it is enough to prove by transfinite induction that $T \subseteq F^\alpha(\perp)$ for all ordinals α .

Base case: trivial, since any element of \mathcal{M}_{UI} is contained in \perp .

Successor case: if $T \subseteq F^\alpha(\perp)$, then

$$\begin{aligned} T &= {}_{\text{vT}\mathcal{R}_{TI}}^{-1}(U) \\ &\quad U \text{ is a fixed point of } \lambda S \bullet \mathcal{F}_{TI}[\underline{Q}]\rho_C[S/Y] \\ &= {}_{\text{vT}\mathcal{R}_{TI}}^{-1}(\mathcal{F}_{TI}[\underline{Q}](\rho_C[U/Y])) \\ &\quad \text{structural inductive hypothesis: } \underline{P} \sqsubseteq_{\text{vT}\mathcal{R}_{TI}} \underline{Q}, \rho_A[T/Y] \sqsubseteq_{\text{vT}\mathcal{R}_{TI}} \rho_C[U/Y] \\ &\subseteq \mathcal{F}_{UI}[\underline{P}](\rho_A[T/Y]) \\ &\quad \text{transfinite inductive hypothesis, and monotonicity of all CSP operators} \\ &\subseteq \mathcal{F}_{UI}[\underline{P}](\rho_A[F^\alpha(\perp)/Y]) \\ &\quad \text{definition} \\ &= F^{\alpha+1}(\perp) \end{aligned}$$

as required.

Limit case: if γ is a limit ordinal, then the inductive hypothesis states that $T \subseteq F^\alpha(\perp)$ for all $\alpha < \gamma$. Then $F^\gamma(\perp) = \bigsqcup_{\alpha < \gamma} F^\alpha(\perp)$. Since T is an upper bound for each element of the chain, and $T \in \mathcal{M}_{UI}$, then the least upper bound of the chain $F^\gamma(\perp)$ is defined, and $T \subseteq F^\gamma(\perp)$, as required.

This establishes the case for recursion. □

We again obtain that $Skip \sqsubseteq_{UI} \mathcal{R}_{TI} \prod_{t \in I} Wait t$ for any set of times I , so arbitrary delays can be introduced into programs while still preserving timewise failures refinement.

Unfortunately, parallel composition does not preserve refinement in general. One example where it fails is in the case of two programs Q_1 and Q_2 , illustrated in Figure 3.2. They are always willing to perform an event at some time in the future, by offering it periodically (so neither will eventually always refuse it), but they are unable to find any time on which they can synchronise, so their combination is able to refuse the offer forever.

$$\begin{aligned} Q_1 &= \mu Y \circ (a \longrightarrow Stop) \triangleright^1 Wait 3 ; Y \\ Q_2 &= Wait 2 ; Q_1 \end{aligned}$$




Figure 2: alternating offers

Each of Q_1 and Q_2 are refinements of $P = a \longrightarrow Stop$, but $Q_1 \parallel Q_2$ is not a refinement of $P \parallel P$, since it may refuse a forever, as Q_1 and Q_2 can never synchronise on a ; yet $P \parallel P$ is unable to refuse a .

However, we will later obtain the following theorem (Theorem 4.9)

Theorem 3.6 Every CSP program P has that $P \sqsubseteq_{UI} \mathcal{R}_{TI} P$ □

Observe that this theorem holds for $Q_1 \parallel Q_2$, since the untimed semantics reveals a divergence in each of Q_1 and Q_2 : a possible infinite sequence of internal events (repeatedly timing out and passing round the loop without the performance of event a). In the failures model, the untimed semantics is that of *Loop* which is refined by any process. In fact, any program which could repeatedly make and retract offers in the timed model would be divergent in the failures model (though as in the case of Q_1 and Q_2 , it could also refine a non-divergent program).

Parallel composition and refinement

The importance of the parallel operator moves us to investigate conditions under which it does preserve refinement. The results presented below concern programs, but they generalise to terms in particular environments, in the obvious way.

Non-Retraction

The example above illustrates one of the ways in which refinement may be lost by parallel composition: the periodic withdrawal of offers. One way to ensure synchronisation is to maintain offers until they are accepted.

A process which does not withdraw offers (though it may make new ones) until it next performs a visible event is termed *non-retracting*. This is similar to (though slightly weaker than) the notion of *nonpre-emptive* given in [CIZ92], although that definition is given in operational terms.

Definition 3.7 A program Q is *non-retracting* if

$$(s, \aleph) \in \mathcal{F}_{TI}[\![Q]\!] \Rightarrow (s, \aleph \cup \{(t, a) \mid \exists t' \bullet (t', a) \in \aleph \wedge \text{end}(s) \leq t < t'\}) \in \mathcal{F}_{TI}[\![Q]\!])$$

□

If an event may be refused at a time t' , then it must be possible that it was continuously refused since the occurrence of the last visible event, at time $\text{end}(s)$. Thus once an event is guaranteed to have been offered, it must be continually offered thereafter.

As expected, we obtain that parallel composition preserves refinement for non-retracting programs:

Lemma 3.8 If $P_1 \sqsubseteq_{UI\mathcal{R}_{TI}} Q_1$ and $P_2 \sqsubseteq_{UI\mathcal{R}_{TI}} Q_2$, and Q_1 and Q_2 are both non-retracting, then $(P_1 \parallel P_2) \sqsubseteq_{UI\mathcal{R}_{TI}} (Q_1 \parallel Q_2)$ □

Proof This is a special case of Lemma 3.10 below, with $A_1 = A_2 = \Sigma$, and the fact that non-retraction is stronger than eventual non-retraction on Σ . □

This idea of non-retraction may be generalised, so that it is concerned only with particular events rather than all events, and with the fact that a time after which offers should be maintained is reached eventually rather than immediately.

Definition 3.9 A program Q is *eventually non-retracting on A* if for any trace s there is some time $t(s)$ such that

$$(s, \aleph) \in \mathcal{F}_{TI}[\![Q]\!] \Rightarrow (s, \aleph \cup \{(t, a) \mid a \in A \wedge \exists t' \bullet (t', a) \in \aleph \wedge t(s) \leq t < t'\}) \in \mathcal{F}_{TI}[\![Q]\!])$$

□

Observe that if S is eventually non-retracting on A , and $B \subseteq A$, then it is also eventually non-retracting on B .

This form of eventual non-retraction allows a period of unstable behaviour before settling down. This permits some timeout behaviour disallowed by a non-retraction requirement. For example, $a \longrightarrow P \overset{3}{\triangleright} b \longrightarrow Q$ is eventually non-retracting (if P and Q are) although the offer of a will be retracted at time 3.

Lemma 3.8 may be generalised to interface parallel, by considering programs that are non-retracting on their common interface.

Lemma 3.10 If $P_1 \sqsubseteq_{UI\mathcal{R}_{TI}} Q_1$ and $P_2 \sqsubseteq_{UI\mathcal{R}_{TI}} Q_2$, and Q_1 and Q_2 are both eventually non-retracting on $A_1 \cap A_2$, then

$$(P_1 \parallel_{A_1 \parallel_{A_2}} P_2) \sqsubseteq_{UI\mathcal{R}_{TI}} (Q_1 \parallel_{A_1 \parallel_{A_2}} Q_2)$$

□

Proof If $(i, u) \text{ UI}\mathcal{R}_{TI} (s, \aleph)$ and $(s, \aleph) \in \mathcal{F}_{TI}[\![Q_1 \parallel_{A_1 \parallel_{A_2}} Q_2]\!]$, then it follows immediately that $(i, u \upharpoonright A_1) \in \mathcal{F}_{UT}[\![P_1]\!]$ and $(i, u \upharpoonright A_2) \in \mathcal{F}_{UT}[\![P_2]\!]$, and hence that $(i, u) \in \mathcal{F}_{UI}[\![P_1 \parallel_{A_1 \parallel_{A_2}} P_2]\!]$.

Consider $(f, (tr, X))$ related via $\text{UI}\mathcal{R}_{TI}$ to $(s, \aleph) \in \mathcal{F}_{TI}[\![Q_1 \parallel_{A_1 \parallel_{A_2}} Q_2]\!]$, with $(f, (tr, X)) \text{ UI}\mathcal{R}_{TI} (s, \aleph)$. Then $tr = \text{strip}(s)$, and $\exists t \bullet [t, \infty) \times X \subseteq \aleph$. Now by the semantics of the parallel

operator, there are \aleph_1 and \aleph_2 such that $(s \upharpoonright A_1, \aleph_1) \in \mathcal{F}_{TI}[\![Q_1]\!]$, $(s \upharpoonright A_2, \aleph_2) \in \mathcal{F}_{TI}[\![Q_2]\!]$, and $\aleph \upharpoonright (A_1 \cup A_2) = (\aleph_1 \upharpoonright A_1) \cup (\aleph_2 \upharpoonright A_2)$. Define

$$\begin{aligned} X_1 &= \{a \in A_1 \mid \forall t \bullet a \in \sigma(\aleph_1 \triangleright t)\} \\ X_2 &= \{a \in A_2 \mid \forall t \bullet a \in \sigma(\aleph_2 \triangleright t)\} \end{aligned}$$

Then $X_1 \cup X_2 = X \upharpoonright (A_1 \cup A_2)$. Furthermore, by the eventual non-retraction of Q_1 and Q_2 , it follows that there are t_1 and t_2 such that

$$\begin{aligned} (s \upharpoonright A_1, \aleph_1 \cup [t_1, \infty) \times X_1) &\in \mathcal{F}_{TI}[\![Q_1]\!] \\ (s \upharpoonright A_2, \aleph_2 \cup [t_2, \infty) \times X_2) &\in \mathcal{F}_{TI}[\![Q_2]\!] \end{aligned}$$

Since each P_i is refined by the corresponding Q_i , it follows that

$$\begin{aligned} (f, (\text{strip}(s) \upharpoonright A_1, X_1)) &\in \mathcal{F}_{UI}[\![P_1]\!] \\ (f, (\text{strip}(s) \upharpoonright A_2, X_2)) &\in \mathcal{F}_{UI}[\![P_2]\!] \end{aligned}$$

and so $(f, (\text{strip}(s), X_1 \cup X_2)) \in \mathcal{F}_{UI}[\![P_1 \parallel_{A_1} P_2]\!]$ by the semantics of the parallel operator. Hence the parallel operator preserves refinement for eventually non-retracting programs. \square

However, if only one of the programs is non-retracting, then the refinement need not be preserved through a parallel combination. For example, consider the following programs, illustrated in Figure 3.2 (where succ is an alphabet renaming which maps n to $n + 1$).

$$\begin{aligned} Q_1 &= \text{Wait } 2; \mu Y \circ (0 \longrightarrow \text{Stop} \square \text{Wait } 1; \text{succ}(Y)) \\ Q_2 &= \mu Y \circ (n : \mathbf{N} \longrightarrow \text{Stop}) \stackrel{1}{\triangleright} \text{succ}(Y) \end{aligned}$$




Figure 3: non-synchronising offers

The process described by Q_1 makes natural numbers available, one at a time. Q_1 is non-retracting, and it is also a refinement of $P_1 = n : \mathbf{N} \longrightarrow \text{Stop}$; on the empty trace, nothing may be refused forever. The process described by Q_2 begins with all natural numbers available, and retracts them one at a time. Q_2 is a refinement of

$$P_2 = \prod_{F \subseteq \mathbb{N}} n : (\mathbf{N} \setminus F) \longrightarrow \text{Stop}$$

for which all events are possible, and any finite set of events may eventually be refused forever. The parallel combination $Q_1 \parallel Q_2$ is semantically equivalent to Stop , since there is no event that Q_1 and Q_2 may cooperate on: Q_1 is prepared to perform event m from time $m + 2$ onwards, but Q_2 is not prepared to perform it beyond time $m + 1$. On the other hand, $P_1 \parallel P_2$ is equivalent to P_2 , which cannot deadlock before any events have been performed. Hence $Q_1 \parallel Q_2$ is not a refinement of $P_1 \parallel P_2$, even though Q_1 is non-retracting.

Promptness

The above example highlights other ways in which parallel combination can fail to preserve refinement. If Q_1 had made all of its offers by some time t , then the counterexample would not work, since the non-retraction of Q_1 would ensure that all of the offers, made by time t , must remain on offer until acceptance occurs.

We define a program to be t -prompt if it must make its offers within t time units of the end of the trace: if a set may be refused for time t , then it may be refused forever.

Definition 3.11 A program Q is t -prompt if

$$(s, \aleph) \in \mathcal{F}_{TI}[\![Q]\!] \wedge [t', t' + t) \times A \subseteq \aleph \wedge t' \geq \text{end}(s) \Rightarrow (s, \aleph \cup [t', \infty) \times A) \in \mathcal{F}_{TI}[\![Q]\!]$$

A program is *prompt* if it is t -prompt for some t . □

We then obtain the following alternative result, which places no constraints upon Q_2 :

Lemma 3.12 If $P_1 \sqsubseteq_{UI\mathcal{R}_{TI}} Q_1$ and $P_2 \sqsubseteq_{UI\mathcal{R}_{TI}} Q_2$, Q_1 is non-retracting and t -prompt for some t , then $(P_1 \parallel P_2) \sqsubseteq_{UI\mathcal{R}_{TI}} (Q_1 \parallel Q_2)$ □

Proof This is a special case of Lemma 3.14 below. □

Promptness may be generalised to apply only to a particular set of events A .

Definition 3.13 A program Q is t -prompt on A if

$$(s, \aleph) \in \mathcal{F}_{TI}[\![Q]\!] \wedge B \subseteq A \wedge [t', t' + t) \times B \subseteq \aleph \wedge t' \geq \text{end}(s) \\ \Rightarrow (s, \aleph \cup [t', \infty) \times B) \in \mathcal{F}_{TI}[\![Q]\!]$$

□

Then the condition for parallelism to preserve refinement may be correspondingly generalised: if the interface between two programs may be split into two parts, and each program is prompt and non-retracting on a different part, then refinement will be preserved by parallel composition.

Lemma 3.14 If $P_1 \sqsubseteq_{UI\mathcal{R}_{TI}} Q_1$, $P_2 \sqsubseteq_{UI\mathcal{R}_{TI}} Q_2$, $B_1 \cup B_2 = A_1 \cap A_2$, Q_1 is eventually non-retracting on B_1 and prompt on B_1 , and Q_2 is eventually non-retracting on B_2 and prompt on B_2 , then

$$(P_1 \parallel_{A_1} \parallel_{A_2} P_2) \sqsubseteq_{UI\mathcal{R}_{TI}} (Q_1 \parallel_{A_1} \parallel_{A_2} Q_2)$$

□

Proof It is clear that infinite traces of the timed parallel combination will appear in untimed form in the untimed combination. So we need only show that

$$(f, (tr, X)) \text{ UI}\mathcal{R}_{TI} (s, \aleph) \wedge (s, \aleph) \in \mathcal{F}_{TI}[\![Q_1 \parallel_{A_1} \parallel_{A_2} Q_2]\!] \\ \Rightarrow (f, (tr, X)) \in \mathcal{F}_{UI}[\![P_1 \parallel_{A_1} \parallel_{A_2} P_2]\!]$$

Assume the antecedent, and consider $(s, \aleph) \in \mathcal{F}_{TI}[\![Q_1 \parallel_{A_1} \parallel_{A_2} Q_2]\!]$, built from behaviours $(s \upharpoonright A_1, \aleph_1) \in \mathcal{F}_{TI}[\![Q_1]\!]$ and $(s \upharpoonright A_2, \aleph_2) \in \mathcal{F}_{TI}[\![Q_2]\!]$, where we have $(\aleph_1 \upharpoonright A_1) \cup (\aleph_2 \upharpoonright A_2) = \aleph \upharpoonright (A_1 \cup A_2)$. We have by the relation $\text{UI}\mathcal{R}_{TI}$ that $tr = \text{strip}(s)$, and also there is some t such that

$t \geq \text{end}(s) + \max\{t(s \upharpoonright A_1, Q_1), t(s \upharpoonright A_2, Q_2)\}$ for some t , such that $[t, \infty) \times X \subseteq \aleph$. Now Q_1 is t_1 -prompt for some t_1 , and Q_2 is t_2 -prompt for some t_2 . Define

$$\begin{aligned} X_1 &= \sigma((\aleph_1 \upharpoonright (X \cap A_1 \cap B_1)) \triangleright (t_1 + t)) \\ Y_1 &= \{a \in X \cap A_1 \cap B_2 \mid [t_2 + t, \infty) \times \{a\} \subseteq \aleph_1\} \\ X_2 &= \sigma((\aleph_2 \upharpoonright (X \cap A_2 \cap B_2)) \triangleright (t_2 + t)) \\ Y_2 &= \{a \in X \cap A_2 \cap B_1 \mid [t_1 + t, \infty) \times \{a\} \subseteq \aleph_2\} \end{aligned}$$

Then $X_1 \cup Y_1 \cup X_2 \cup Y_2 = X \cap (A_1 \cup A_2)$. Now since Q_1 is eventually non-retracting on B_1 , it follows (also using subset closure of refusals for Q_1) that

$$(s \upharpoonright A_1, \aleph_1 \cup [t, t + t_1) \times X_1) \in \mathcal{F}_{TI}[\![Q_1]\!]$$

And hence it follows by promptness that

$$(s \upharpoonright A_1, \aleph_1 \cup [t, \infty) \times X_1) \in \mathcal{F}_{TI}[\![Q_1]\!]$$

Also, observe that $[t_2 + t, \infty) \times Y_1 \subseteq \aleph_1$. Since Q_1 is a refinement of P_1 it follows that $(f, (tr \upharpoonright A_1, X_1 \cup Y_1)) \in \mathcal{F}_{UI}[\![P_1]\!]$. Similarly we obtain that $(f, (tr \upharpoonright A_2, X_2 \cup Y_2)) \in \mathcal{F}_{UI}[\![P_2]\!]$. Thus from the semantics of parallel, we obtain that $(f, (tr, X)) \in \mathcal{F}_{UI}[\![P_1 \parallel_{A_1} P_2]\!]$, yielding the result. \square

This result is particularly useful, as it applies immediately to systems such as those described in terms of input/output automata [LyV92] where input is always possible, and so components are always non-retracting and prompt on input. In these systems, parallel composition connects outputs from one program to corresponding inputs of the other. Thus the interface in a parallel composition may be partitioned into those events input by one component, and those input by the other. The two programs will be prompt and non-retracting on these two sets respectively.

This result may also be applied to CSP descriptions of occam programs, since in such programs output guards are not permitted. Once a program is prepared to perform an output, it remains ready to perform it until it occurs. Consequently, program are always non-retracting on output, so parallel composition will preserve refinement for prompt components.

Compactness

The final condition we will present here concerns the nature of the untimed programs P_1 and P_2 . A process is compact if its refusals are determined (in a particular way) by the finite refusal sets. If the untimed processes are compact, then it turns out that only one of the timed processes need be non-retracting for refinement to be preserved by the parallel operator. In the example above, P_2 fails this condition; any finite subset of \mathbf{N} may be refused, but infinite subsets may not.

Definition 3.15 A program P is *compact* if for any $tr \in \Sigma^*$, $Y \subseteq \Sigma$ we have

$$(\forall X \subseteq^{fn} Y \bullet (f, (tr, X)) \in \mathcal{F}_{UI}[\![P]\!]) \Rightarrow (f, (tr, Y)) \in \mathcal{F}_{UI}[\![P]\!]$$

\square

Lemma 3.16 If $P_1 \sqsubseteq_{UI \mathcal{R}_{TI}} Q_1$, $P_2 \sqsubseteq_{UI \mathcal{R}_{TI}} Q_2$, Q_1 is eventually non-retracting on $A_1 \cap A_2$, and P_1, P_2 are compact, then

$$(P_1 \parallel_{A_1} P_2) \sqsubseteq_{UI \mathcal{R}_{TI}} (Q_1 \parallel_{A_1} Q_2)$$

\square

Proof It is clear that infinite traces of the timed process will appear in the untimed process. So we need only show that

$$\begin{aligned} (f, (tr, X)) \text{ }_{UI\mathcal{R}_{TI}} (s, \aleph) \wedge (s, \aleph) \in \mathcal{F}_{TI}[[Q_1 \upharpoonright_{A_1} \parallel_{A_2} Q_2]] \\ \Rightarrow (f, (tr, X)) \in \mathcal{F}_{UI}[[P_1 \upharpoonright_{A_1} \parallel_{A_2} P_2]] \end{aligned}$$

Consider $(s, \aleph) \in \mathcal{F}_{TI}[[Q_1 \upharpoonright_{A_1} \parallel_{A_2} Q_2]]$, built from behaviours $(s \upharpoonright_{A_1}, \aleph_1) \in \mathcal{F}_{TI}[[Q_1]]$ and $(s \upharpoonright_{A_2}, \aleph_2) \in \mathcal{F}_{TI}[[Q_2]]$, where $\aleph_1 \upharpoonright_{A_1} \cup \aleph_2 \upharpoonright_{A_2} = \aleph \upharpoonright_{(A_1 \cup A_2)}$. We have by the relation $_{UI\mathcal{R}_{TI}}$ that $tr = strip(s)$, and also that there is some $t \geq end(s) + t(s \upharpoonright_{A_1}, Q_1)$ such that $[t, \infty) \times X \subseteq \aleph$. Define

$$\begin{aligned} X_1 &= \{a \in X \cap A_1 \mid \forall t' \bullet a \in \sigma(\aleph_1 \triangleright t')\} \\ X_2 &= \{a \in X \cap A_2 \mid \exists t' \bullet [t', \infty) \times \{a\} \subseteq \aleph_2\} \end{aligned}$$

Since $\aleph_1 \upharpoonright_{A_1} \cup \aleph_2 \upharpoonright_{A_2} = \aleph \upharpoonright_{(A_1 \cup A_2)}$, we have that $X_1 \cup X_2 = X$. Since Q_1 is eventually non-retracting, we obtain that $(s \upharpoonright_{A_1}, \aleph_1 \cup [t, \infty) \times X_1) \in \mathcal{F}_{TI}[[Q_1]]$. Since P_1 is refined by Q_1 , we have that $(f, (tr \upharpoonright_{A_1}, X_1)) \in \mathcal{F}_{UI}[[P_1]]$.

Now consider a finite set $\{a_1, a_2 \dots a_n\} = Y \subseteq X_2$. For each a_i there is a corresponding time t_i such that $[t_i, \infty) \times \{a_i\} \subseteq X_2$. Thus there is a time $t_0 = \max\{t_i\}$ such that $[t_0, \infty) \times Y \subseteq X_2$. Since P_2 is refined by Q_2 , it follows that $(f, (tr \upharpoonright_{A_2}, Y)) \in \mathcal{F}_{UI}[[P_2]]$. This is true for all finite subsets Y of X_2 , so by compactness of P_2 we have that $(f, (tr \upharpoonright_{A_2}, X_2)) \in \mathcal{F}_{UI}[[P_2]]$. Hence $(f, (tr, X)) \in \mathcal{F}_{UI}[[P_1 \upharpoonright_{A_1} \parallel_{A_2} P_2]]$ as required. \square

Compactness is often easy to check, since it will be present in any program not containing infinite non-determinism. Thus any program not containing any infinite choice or infinite-to-one renaming will automatically be compact.

Specification

In the untimed infinite traces model specifications may be considered as consisting of three components, dealing with the failures, divergences and infinite traces. In other words, for any given $S(l, b)$ there are S_f, S_d and S_i such that

$$\begin{aligned} S(l, b) &\Leftrightarrow (l = f \wedge b = (tr, X)) \Rightarrow S_f(tr, X) \\ &\quad \wedge (l = d \wedge b = tr) \Rightarrow S_d(tr) \\ &\quad \wedge (l = i \wedge b = u) \Rightarrow S_i(u) \end{aligned}$$

Then a specification $(S_f(tr, X), S_d(tr), S_i(u))$ translates to the timed specification

$$\begin{aligned} S'(s, \aleph) &\Leftrightarrow \#s < \infty \wedge [t, \infty) \times X \subseteq \aleph \Rightarrow S_f(strip(s), X) \\ &\quad \wedge \#s = \infty \Rightarrow S_i(strip(s)) \end{aligned}$$

For example, the specification ‘deadlock-free’ constrains only the possible failure set, with $S_f(tr, X) \Leftrightarrow X \neq \Sigma$. The translation is equivalent to

$$\#s < \infty \Rightarrow \neg \exists t \bullet [t, \infty) \times \Sigma \subseteq X$$

which is the timed version of deadlock-freedom. Thus an untimed verification of deadlock-freedom for a system remains valid under timewise refinement.

The untimed specification of a buffer which passes messages of type M may be given simply as a predicate S_f :

$$\begin{aligned} S_f(tr, X) &\Leftrightarrow \quad tr \downarrow out \leq tr \downarrow in \\ &\quad \wedge \quad tr \downarrow out = tr \downarrow in \Rightarrow X \cap in.M = \{\} \\ &\quad \wedge \quad tr \downarrow out < tr \downarrow in \Rightarrow out.M \not\subseteq X \end{aligned}$$

The translation is equivalent to the following timed specification:

$$\begin{aligned}
& strip(s) \downarrow out \leq strip(s) \downarrow in \\
\wedge & strip(s) \downarrow out = strip(s) \downarrow in \Rightarrow \neg \exists t, m \bullet [t, \infty) \times \{in.m\} \subseteq \mathbb{N} \\
\wedge & strip(s) \downarrow out < strip(s) \downarrow in \Rightarrow \neg \exists t \bullet [t, \infty) \times out.M \subseteq \mathbb{N}
\end{aligned}$$

which is the specification of a timed buffer.

As an example of an application of the theory, consider Roscoe's first (untimed) buffer law presented in [Hoa85], which tells us that the chaining together of two buffers is again a buffer. The chaining operator is defined in terms of parallel, hiding, and renaming (where $swap_{a,b}$ renames channel a to b and vice versa):

$$P_1 \gg P_2 \hat{=} (swap_{out,c}(P_1) \parallel_{\{in,c\}} \parallel_{\{out,c\}} swap_{in,c}(P_2)) \setminus c$$

However, this law does not hold in general in the timed model. As we have seen, B_1 and B_2 might fail to agree on a time to synchronise on their common internal channel, resulting in their combination refusing ever to output.

In order to establish conditions under which the law does hold, we will make use of the fact that untimed buffers are compact (since if some input can be refused, then so too can all possible inputs), and also of the fact that the refinement relations in this paper are complete, which yields that every timed buffer is a refinement of some untimed buffer. We may then obtain conditions under which a chain of buffers again yields a buffer.

For example, if every timed buffer B_i is eventually non-retracting on input, then the chain $B_1 \gg B_2 \gg \dots \gg B_n$ is again a buffer, eventually non-retracting on input. This follows from the fact that each B_i is a refinement of some untimed buffer A_i (by completeness of the refinement relation); that we have a condition which may be applied at every step of building up the chain to ensure that the timed chain refined the untimed chain (in the general parallel case, we require only non-retraction on the interface); and that the chain $A_1 \gg A_2 \gg \dots \gg A_n$ is an untimed buffer (from Roscoe's law), from which it follows that any refinement of it is a timed buffer. A similar result holds if each B_i is non-retracting on output; or if odd (or even) numbered buffers are non-retracting on both input and output. It follows that the combination $B_1 \gg COPY \gg B_2 \dots COPY \gg B_n$ is a buffer, for any timed buffers B_i .

4 An operational characterisation

A more immediately intuitive semantic approach often employed in the theory of process algebra is that of operational semantics: processes are defined in terms of transitions that they may perform and subsequent states that may be reached. Within this framework, equivalence between processes may be characterised in terms of bisimulation relations [Mil89], or by means of equivalence under some notion of testing [Hen88].

In the testing approach, a test is defined to be a process T which also has the capacity to perform a special success event ω , which is considered to be distinct from the set of synchronisation events Σ . An execution of a process P is a maximal (finite or infinite) sequence of transitions starting from P . Then we say that P may T if there is some execution of $(P \parallel T) \setminus \Sigma$ which passes through a state from which ω is a possible transition; and P must T if every execution of $(P \parallel T) \setminus \Sigma$ passes through such a state. Then P is equivalent to Q under may testing if for any test T , P may $T \Leftrightarrow Q$ may T ; and P and Q are equivalent under must testing if P must $T \Leftrightarrow Q$ must T for any test T .

An operational semantics has been given for CSP in [BRW92] and [Ros88]. Transitions are given as $P \xrightarrow{\mu} P'$, indicating that a process P may perform a μ event (i.e. an internal or visible event)

and then behave as P' . In this section we will subscript the transition with a u to indicate that this is an untimed transition. Equivalence in the untimed traces model \mathcal{M}_{UT} is exactly the same as equivalence under may testing using the transitions given in [Ros88]; and equivalence in the untimed infinite traces model \mathcal{M}_{UI} is exactly the same as equivalence under must testing using those transitions. More details may be found in [Hen88, BRW92, Ros88]. The important properties from our point of view is that each trace of P predicted by the traces model corresponds to an execution of P in which that sequence of visible events is performed (as well as possibly some internal events); that any divergence corresponds to an execution in which some prefix of the divergent trace is performed, followed by an infinite sequence of internal τ steps; any failure (tr, X) corresponds either to a divergence (i.e. an infinite sequence of τ steps after some prefix of the trace) or to an execution in which the entire sequence tr of events is performed, and a state is reached from which no internal progress can be made, and from which no event in the refusal set X is possible; and for every infinite trace u there is an execution that either diverges after some prefix of u or performs the entire sequence of events u . And conversely, any execution given by the operational semantics is recorded appropriately in the denotational semantics.

An operational semantics has also been provided for timed CSP in [Sch91, Sch95], where processes may undergo timed transitions: $P \xrightarrow{(t,\mu)} P'$ indicates that the process P may perform event μ at time t , and subsequently behave as P' . We will subscript timed transitions with t to distinguish them from untimed transitions. Evolutions, or time passing transitions, were also provided in the operational semantics. Equivalence in the infinite timed failures model \mathcal{M}_{TI} is the same as equivalence under must testing using the transitions given in [Sch95]. Again, timed failures (s, \aleph) are present in the denotational semantics of a process P precisely when there is some execution of P in which events are performed at the times recorded in s , passing through states in which the events recorded in the refusal set \aleph were not possible.

Both may and must forms of testing can be defined for both untimed and timed CSP. We will say $P \underline{\text{may}}_u T$ if some execution of $(P \parallel T) \setminus \Sigma$ in terms of untimed \xrightarrow{u} transitions passes through a state in which ω is possible; and we will use $P \underline{\text{must}}_u T$, $P \underline{\text{may}}_t T$ and $P \underline{\text{must}}_t T$ in a similar fashion.

We may interpret a timed program Q also at the untimed level. Given a timed program Q , we may define an untimed program $\Theta(Q)$ whose untimed transitions are derived from the timed ones:

$$\frac{Q \xrightarrow{(t,\mu)}_t Q'}{\Theta(Q) \xrightarrow{\mu}_u \Theta(Q')}$$

Then $\mathcal{F}_{UT}[\Theta(Q)]$, $\mathcal{F}_{UD}[\Theta(Q)]$, $\mathcal{F}_{UF}[\Theta(Q)]$, and $\mathcal{F}_I[\Theta(Q)]$ can all be defined using the resulting transition system.

In the traces model, $P \sqsubseteq_{UT} Q$ is true exactly when $\forall T \bullet (Q \underline{\text{may}} T \Rightarrow P \underline{\text{may}} T)$. By analogy, we may characterise an operational version of timed refinement, where if Q may pass a timed test T , then P may pass the same T considered as an untimed test.

Definition 4.1 $P \sqsubseteq_t Q$ is defined by

$$P \sqsubseteq_t Q \Leftrightarrow \forall T \bullet Q \underline{\text{may}}_t T \Rightarrow P \underline{\text{may}}_u \Theta(T)$$

□

It turns out that this notion of refinement is the same as the denotational version of traces refinement.

Theorem 4.2 $P \sqsubseteq_t Q \Leftrightarrow P \sqsubseteq_{UT} \mathcal{R}_{TI} Q$

□

Proof “ \Rightarrow ” Assume $P \not\sqsubseteq_{UI} \mathcal{R}_{TI} Q$. Then there is some trace s of Q such that $strip(s)$ is not a trace of P . Let $s = \langle (t_1, a_1), \dots, (t_n, a_n) \rangle$. Then define the test T by

$$T = a_1 \longrightarrow \dots \longrightarrow a_n \longrightarrow \omega \longrightarrow Stop$$

Since the timed operational semantics are equivalent to the denotational semantics, there is some execution of Q giving rise to trace s , so there is some execution of $(Q \parallel T) \setminus \Sigma$ which reaches a state in which T can perform ω . However, there is no such execution of $(P \parallel \Theta(T)) \setminus \Sigma$, since if there were then this would correspond to P performing the events in $strip(s)$, which would mean that $strip(s)$ is a trace of P , yielding a contradiction. Thus $Q \underline{\text{may}}_t T$ but $\neg(P \underline{\text{may}}_u T)$, and so $\neg(P \sqsubseteq_t Q)$

“ \Leftarrow ” Assume $P \sqsubseteq_{UI} \mathcal{R}_{TI} Q$, and consider a test T for which $Q \underline{\text{may}}_t T$. Then there is some execution of $(Q \parallel T) \setminus \Sigma$ which leads to a state in which $\xrightarrow{\omega}_t$ is possible. The contribution of Q to this execution corresponds to some timed trace s . Then $strip(s)$ is a trace of P , so P has some execution giving rise to $strip(s)$. Thus $(P \parallel \Theta(T)) \setminus \Sigma$ has an execution which takes $\Theta(T)$ through states corresponding to the timed states that T passes through in the successful execution of $(Q \parallel T) \setminus \Sigma$, and so it reaches a state in which an $\xrightarrow{\omega}_u$ transition is possible. Thus $P \underline{\text{may}}_u T$. \square

In the failures/divergences model, $P \sqsubseteq_{UI} Q$ is equivalent to $P \underline{\text{must}} T \Rightarrow Q \underline{\text{must}} T$ for any T . Again by analogy, we characterise an operational version of timed refinement:

Definition 4.3 $P \sqsubseteq_f Q$ is defined by

$$P \sqsubseteq_f Q \Leftrightarrow \forall T \bullet P \underline{\text{must}}_u \Theta(T) \Rightarrow Q \underline{\text{must}}_t T$$

\square

This formulation of refinement is equivalent to the denotational version of failures refinement.

Theorem 4.4 $P \sqsubseteq_f Q \Leftrightarrow P \sqsubseteq_{UI} \mathcal{R}_{TI} Q$ \square

Proof “ \Rightarrow ” If $P \not\sqsubseteq_{UI} \mathcal{R}_{TI} Q$ then either (1) there is some $(s, [t', \infty) \times X) \in \mathcal{F}_{TI}[[Q]]$ with $(f, (strip(s), X)) \notin \mathcal{F}_{UI}[[P]]$, or (2) there is some infinite trace s such that $(s, \{\}) \in \mathcal{F}_{TI}[[Q]]$ and $strip(s) \notin \mathcal{F}_{UI}[[P]]$.

- 1 Let $s = \langle (t_1, a_1), \dots, (t_n, a_n) \rangle$, and let $t_0 = 0$. Assume without loss of generality that $t' \geq t_n$. Define

$$\begin{aligned} T_i &= Wait(t_i - t_{i-1}); ((a_i \longrightarrow T_{i+1}) \overset{0}{\triangleright} \omega \longrightarrow Stop) \quad 0 < i \leq n \\ T_{n+1} &= Wait(t' - t_n); x : X \longrightarrow \omega \longrightarrow Stop \end{aligned}$$

If $(P \parallel \Theta(T_1)) \setminus \Sigma$ has an execution that is not successful, then the contribution from P must correspond to the failure $(strip(s), X)$, yielding a contradiction. Thus $P \underline{\text{must}}_u T_1$. On the other hand, Q has an execution corresponding to $(s, [t', \infty) \times X)$, and so $(Q \parallel T_1) \setminus \Sigma$ does have an unsuccessful execution, thus $\neg(Q \underline{\text{must}}_t T_1)$.

- 2 Let $s = \langle (t_1, a_1), \dots, (t_i, a_i), \dots \rangle$. Then let the trace during an interval $[n, n+1)$ be given by $\langle (t_{n,1}, a_{n,1}), \dots, (t_{n,m}, a_{n,m}) \rangle$. This must be finite for any interval, since the trace s is finitely variable, i.e. its restriction to any finite interval is finite. Define

$$\begin{aligned} T_{n,i} &= Wait(t_{n,i} - t_{n,i-1}); \\ &\quad ((a_{n,i} \longrightarrow T_{n,i+1}) \overset{0}{\triangleright} \omega \longrightarrow Stop) \quad 0 < i \leq m \\ T_{n,m+1} &= Stop \\ T_n &= T_{n,1} \parallel Wait 1; T_{n+1} \end{aligned}$$

Each $T_{n,i}$ is well-defined, so each of the equations for the T_i is 1-guarded.) Then if $(P \parallel \Theta(T_0)) \setminus \Sigma$ has an unsuccessful execution, the contribution of P must correspond to $strip(s)$, yielding a contradiction; thus $P \underline{\text{must}}_u \Theta(T_0)$. However, $(Q \parallel T_0) \setminus \Sigma$ does have an unsuccessful execution, driven by an execution of Q corresponding to s . Thus $\neg(Q \underline{\text{must}}_i T_0)$.

“ \Leftarrow ” Assume that $P \sqsubseteq_{vI\mathcal{R}_{TI}} Q$, and that $\neg(Q \underline{\text{must}}_i T)$. It will be enough to prove that $\neg(P \underline{\text{must}}_u \Theta(T))$. Consider an unsuccessful execution of $(Q \parallel T) \setminus \Sigma$. There are a number of possibilities; we consider the events that were internalised by the $\setminus \Sigma$ abstraction:

- $Q \parallel T$ performs infinitely many events from Σ . Then there is a corresponding infinite trace s of both Q and T . Since $P \sqsubseteq_{vI\mathcal{R}_{TI}} Q$, the trace $strip(s)$ is an infinite trace of P . If P diverges at some point along $strip(s)$, then this will give rise to an unsuccessful execution of $(P \parallel T) \setminus \Sigma$. Otherwise there is an infinite untimed execution of $\Theta(T)$ performing the same events, and passing through corresponding states to those reached in the timed execution. Hence there is an infinite execution of $(P \parallel \Theta(T)) \setminus \Sigma$ where ω is not possible in any state (since the possibility of ω depends purely on the state reached by $\Theta(T)$), and so $\neg(P \underline{\text{must}}_u \Theta(T))$.
- $Q \parallel T$ performs finitely many events from Σ :
 - If T performs infinitely many timed τ transitions, then $\Theta(T)$ may perform infinitely many untimed ones, passing through corresponding states, so if P does not diverge (which itself leads to an unsuccessful execution) then this will yield an unsuccessful execution of $(P \parallel \Theta(T)) \setminus \Sigma$.
 - If T performs finitely many τ actions, then it will come to arrive in a final state T' . Any events that T' is able to perform are blocked by Q for all time, and so P (if it does not diverge) may reach a stable state P' in which none of those events are possible. Also, $\Theta(T)$ may by untimed transitions reach a state $\Theta(T')$, which is also unable to perform those events that T' was unable to perform, and so $P' \parallel \Theta(T')$ will be unable to progress. Thus the execution from $P \parallel \Theta(T)$ to $P' \parallel \Theta(T')$ is maximal, and unsuccessful.

□

These results may also be used to establish that any untimed description of a system in CSP is refined by the same description considered as a timed description.

We first provide an untimed operational semantics for the timeout operator. The timeout may always be resolved by its left-hand argument performing a visible action, but any internal progress made by that argument does not resolve the timeout.

$$\frac{P \xrightarrow{a}_{\lambda_u} P'}{P \triangleright^t P'' \xrightarrow{a}_{\lambda_u} P'} \qquad \frac{P \xrightarrow{\tau}_{\lambda_u} P'}{P \triangleright^t P'' \xrightarrow{\tau}_{\lambda_u} P' \triangleright^t P''}$$

Furthermore, the timeout may occur:

$$\frac{}{P \triangleright^t P'' \xrightarrow{\tau}_{\lambda_u} P''}$$

Lemma 4.5 Any timed process P has $\Theta(P) \sqsubseteq_{vT\mathcal{R}_{TI}} P$ and $\Theta(P) \sqsubseteq_{vI\mathcal{R}_{TI}} P$. □

Proof (sketch) This follows from the above-mentioned equivalence of the denotational and operational semantics in both the untimed cases and the timed case. In the first case, if there is

a timed trace s of P , then there is some execution of P which gives rise to this trace. But then every step of this execution can be matched by an untimed step, so there is an equivalent untimed execution of $\Theta(P)$, which corresponds to the trace $strip(s)$. Since the untimed operational and denotational semantics are equivalent, the trace $strip(s)$ appears in the trace set of $\Theta(P)$.

In the case of failures refinement, similar reasoning shows that infinite timed traces will be matched by infinite untimed ones; and a timed failure $(s, [t, \infty) \times X)$ with finite trace s will correspond to some execution of P . After the trace s has been performed there are two possibilities. The execution may contain an infinite sequence of internal events; these can be matched by $\Theta(P)$, leading to a divergence and the inclusion of $(f, (strip(s), X))$ as a failure of $\Theta(P)$. The other possibility is that a final state is reached from which no event in X , or any further internal progress, is possible (since X is refused from that point onwards); in this case a corresponding untimed state in which X may be refused is reachable from $\Theta(P)$ by means of a corresponding execution, and the failure $(f, (strip(s), X))$ again appears as a failure of $\Theta(P)$. \square

Definition 4.6 Define relation R to be a *failures simulation* if whenever $R(P_1, P_2)$, then

1. If $P_2 \xrightarrow{\mu} P'_2$ then there is some P'_1 such that $P_1 \xrightarrow{\mu} P'_1$ and $R(P'_1, P'_2)$
2. If $P_1 \xrightarrow{\mu} P'_1$ then there are processes P'_1, P'_2 , such that $P_1 \xrightarrow{\mu} P'_1$ and $P_2 \xrightarrow{\mu} P'_2$ and $R(P'_1, P'_2)$.

We say that P_1 *failures simulates* P_2 if there is some failures simulation between them. \square

Lemma 4.7 If P_1 failures simulates P_2 then $P_1 \sqsubseteq_{UI} P_2$ \square

Proof It follows from the definition of failures simulation that any trace, infinite trace, or divergent trace of P_2 must also be a trace, infinite trace, or divergent trace respectively, of P_1 . Furthermore, any refusal of P_2 (not arising from divergence) is also a refusal of P_1 , since if $R(P_1, P_2)$ and $\neg P_2 \xrightarrow{\mu}$ then $\neg P_1 \xrightarrow{\mu}$; it follows that any failure of P_2 is a failure of P_1 . \square

Lemma 4.8 Any CSP process P has P failures simulates $\Theta(P)$. \square

Proof Let the relation R hold between two programs P and $\Theta(Q)$ if P and Q are syntactically identical up to the values of timeouts. A straightforward structural induction on P shows that R is a failures simulation.

As an illustration we will establish the case of the timeout operator.

Assume that P and Q are syntactically identical up to the values of timeouts. If $P = P_1 \triangleright^t P_2$, then $Q = Q_1 \triangleright^{t'} Q_2$, where P_i is syntactically identical to Q_i up to the values of timeouts for $i = 1, 2$.

The operational rules required to establish this case are the following

$$\frac{Q \xrightarrow{(t', a)}_t Q'}{Q \triangleright^t Q'' \xrightarrow{(t', a)}_t Q'} \quad [t' \leq t] \qquad \frac{Q \xrightarrow{(t', \tau)}_t Q'}{Q \triangleright^t Q'' \xrightarrow{(t', \tau)}_t Q' \triangleright^{t-t'} Q''} \quad [t' \leq t]$$

Furthermore, the timeout may occur if Q can evolve for t units of time:

$$\frac{Q \xrightarrow{t}}{Q \triangleright^t Q'' \xrightarrow{(t, \tau)}_t Q''}$$

There are two clauses that define failures simulation.

Clause 1: If $\Theta(Q) \xrightarrow{\mu}_{\lambda_u} R$ then $Q \xrightarrow{(t', \mu)}_t Q'$ for some t', Q' by the definition of Θ . There are three possibilities to be considered, one for each of the above operational rules. We will consider only the second; the other cases are entirely similar. In this case we have $Q_1 \xrightarrow{(t', \tau)}_t Q'_1$. The inductive hypothesis states that $P_1 \xrightarrow{\tau}_{\lambda_u} P'_1$ with P'_1 identical to Q'_1 up to the values of timeout, so it follows that $P'_1 \triangleright^t P_2$ is identical to $Q'_1 \triangleright^{t-t'} Q_2$ up to the value of timeouts, as required.

Clause 2: If $P \xrightarrow{\mu}_{\lambda_u}$ then there are two possibilities: either $\mu = \tau$ or $\mu \neq \tau$. In the latter case the proof is straightforward. In the case where $\mu = \tau$, it is a result from [Sch95] that any timed process can either evolve or perform a τ transition: either (1) $Q_1 \rightsquigarrow^t$ or (2) $Q_1 \xrightarrow{(t', \tau)}_t Q'_1$ for some $t' \leq t$.

In case (1), the third operational rule is applicable, and so $Q \xrightarrow{(t, \tau)}_t Q_2$. Hence $\Theta(Q) \xrightarrow{\tau}_{\lambda_u} \Theta(Q_2)$. The third untimed rule for timeout has that $P \xrightarrow{\tau}_{\lambda_u} P_2$, and by the inductive hypothesis on the relationship between P_2 and Q_2 it follows that $R(P_2, \Theta(Q_2))$ as required.

In case (2), the inductive hypothesis for P_1 and Q_1 states that there are processes P'_1 and Q'_1 such that $P_1 \xrightarrow{\tau}_{\lambda_u} P'_1$ and $\Theta(Q_1) \xrightarrow{\tau} Q'_1$ with $R(P'_1, Q'_1)$. Hence there are processes $P'_1 \triangleright^t P_2$ and $Q'_1 \triangleright^{t-t'} \Theta(Q_2)$ which are related by R , and for which $P \xrightarrow{\tau}_{\lambda_u} P'_1 \triangleright^t P_2$ and $\Theta(Q) \xrightarrow{\tau}_{\lambda_u} Q'_1 \triangleright^{t-t'} \Theta(Q_2)$, as required.

Thus the case is established.

Hence P failures simulates $\Theta(Q)$ whenever P and Q are identical up to the values of timeouts. The lemma follows from the special case $P = Q$. \square

Theorem 4.9 Any process P has $P \sqsubseteq_{UI \mathcal{R}_{TI}} P$ and $P \sqsubseteq_{UI \mathcal{R}_{TI}} P$. \square

Proof Lemmas 4.5, 4.7, 4.8 yield that

$$P \sqsubseteq_{UI} \Theta(P) \sqsubseteq_{UI \mathcal{R}_{TI}} P$$

and Lemma 1.4 yields that

$$P \sqsubseteq_{UI \mathcal{R}_{TI}} P$$

as required. \square

5 Examples

A stop and wait protocol

The well-known alternating bit protocol is a useful common example, since it has been treated by so many different formalisms that it provides a means of comparing and contrasting them. We will use it here simply to illustrate some of the techniques presented earlier.

The untimed alternating bit protocol consists of a sender and receiver communicating over two lossy channels. The nature of a generic lossy channel may be specified at the untimed level using the infinite traces model. The specification SM on a medium $M_{in, out}$ with input in and output out consists of three parts, two on the trace refusal observations, and one on the infinite traces:


$$\begin{aligned} SM1(tr, X) & \quad tr \downarrow out \preceq tr \downarrow in \\ SM2(tr, X) & \quad out.M \not\subseteq X \vee in.M \cap X = \{\} \\ SM3(u) & \quad \#(u \upharpoonright in) = \infty \Rightarrow \#(u \upharpoonright out) = \infty \end{aligned}$$

$SM1(tr, X)$ simply states that the sequence of messages passed on channel out should be a (not necessarily contiguous) subsequence of those passed on channel in , so messages may be lost but not corrupted; $SM2(tr, X)$ states that at least one of input and output should not be refused (where X is the refusal set); and $SM3(u)$ is a fairness condition that requires that output should not be lost infinitely often.

The requirement we have of the entire system is that it should behave as a (one-place) buffer. Our specification is

$$\begin{aligned} SPEC(tr, X) = & \quad tr \downarrow out \leq_1 tr \downarrow in \\ & \wedge \quad tr \downarrow out = tr \downarrow in \Rightarrow in.M \cap X = \{\} \\ & \wedge \quad tr \downarrow out < tr \downarrow in \Rightarrow out.M \not\subseteq X \end{aligned}$$

The network used is pictured as follows:



The basic idea of the protocol is to add an extra bit to each of the messages sent along the lossy channels which alternates between 0 and 1. The sending process sends multiple copies of each message until it receives an acknowledgement. As soon as the receiving process gets a new message it sends acknowledgements of it until the next message arrives. The two ends can always spot a new message or acknowledgement because of the alternating bit.

The two media are described as $M1 = M_{a,b}$ and $M2 = M_{c,d}$, passing messages from a to b , and from c to d

This strategy may be captured by the following CSP descriptions of the sender S and the receiver R . We set $R = R(0)$ and $S = S(0)$, where for $s \in \{0, 1\}$ and x in the set of messages M we define

$$\begin{aligned} S(s) &= in?x \longrightarrow S'(s, x) \\ S'(s, x) &= a!(s, x) \longrightarrow S'(s, x) \\ &\quad \square d?s \longrightarrow S(\bar{s}) \\ &\quad \square d?\bar{s} \longrightarrow a!(s, x) \longrightarrow S'(s, x) \\ R(s) &= b?(s, x) \longrightarrow out!x \longrightarrow c!s \longrightarrow R(\bar{s}) \\ &\quad \square b?(\bar{s}, x) \longrightarrow c!\bar{s} \longrightarrow R(s) \end{aligned}$$

The entire network consists of the parallel combination of the sender and receiver together with the two media; and the channels a, b, c , and d are all made internal.

$$NETWORK = ((S \upharpoonright_{\{in, a, d\}} \parallel_{\{out, b, c\}} R) \Sigma \parallel_{\{a, b, c, d\}} (M1 \upharpoonright_{\{a, b\}} \parallel_{\{c, d\}} M2)) \setminus \{a, b, c, d\}$$

The network is considered as the parallel combination of the protocol and the media.

An analysis at the untimed level establishes that the system is livelock-free, essentially because of the fairness of the media which cannot lose an infinite sequence of messages. It is also deadlock-free: if S cannot make progress, then it must be waiting for both media, which must therefore both be ready to interact with R , and so R is able to make progress. Finally, it is straightforward to show that it is functionally equivalent to a one-place buffer.

Timed descriptions of the alternating bit protocol commonly employ a timeout in the description of the sender process, since the intention is that the sender should wait for an acknowledgement and then retransmit a message if this does not arrive within a certain interval. But in fact there is no need to withdraw the capability of receiving a message on the acknowledgement channel simply because a retransmission has been enabled, and so at the untimed level this behaviour may be modelled as a choice.

To provide a timed refinement of the protocol, we wish to preserve correctness of the system. The most general form of correctness that could be preserved by a timewise refinement would be for timed versions of the media to meet simply the translations of the untimed specifications with no further constraints. Thus we prefer not to impose the restriction on the media that they are non-retracting.

A timed version TS of the sender process may be obtained simply by including a delay t before retransmission of a message. The length of this delay will be influenced by such factors as the length of time before an acknowledgement would be expected to arrive, and the reluctance to send unnecessary messages. The timed receiver process TR still behaves sequentially, and has no time-critical behaviour.

Some small delays ϵ are introduced to ensure that the recursive loops are time-guarded. (These play the role of the original δ delay enforced by event prefix in earlier versions of timed CSP [ReR86]).

$$\begin{aligned}
TS(s) &= in?x \longrightarrow TS'(s, x) \\
TS'(s, x) &= Wait\ t; a!(s, x) \longrightarrow TS'(s, x) \\
&\quad \square d?s \xrightarrow{\epsilon} TS(\bar{s}) \\
&\quad \square d?\bar{s} \xrightarrow{\epsilon} a!(s, x) \longrightarrow TS'(s, x) \\
TR(s) &= b?(s, x) \xrightarrow{\epsilon} out!x \longrightarrow c!s \longrightarrow TR(\bar{s}) \\
&\quad \square b?(\bar{s}, x) \xrightarrow{\epsilon} c!\bar{s} \longrightarrow TR(s)
\end{aligned}$$

Given two timed media $TM1$ and $TM2$ that meet the timed translation of SM , by completeness there are two untimed media $M1$ and $M2$ which meet SM and which are refined by $TM1$ and $TM2$. Then $M1 \upharpoonright_{\{a,b\}} \upharpoonright_{\{c,d\}} M2 \sqsubseteq_{\nu I \mathcal{R} \tau I} TM1 \upharpoonright_{\{a,b\}} \upharpoonright_{\{c,d\}} TM2$ by Lemma 3.10, since the intersection of the interface sets is empty, so both $TM1$ and $TM2$ are trivially non-retracting on it. Also, by Theorem 3.5, and since delays may be introduced into an untimed description to produce a timed refinement, and the intersection of the two interface sets is empty, we have that $S \upharpoonright_{\{in,a,d\}} \upharpoonright_{\{out,b,c\}} R \sqsubseteq_{\nu I \mathcal{R} \tau I} TS \upharpoonright_{\{in,a,d\}} \upharpoonright_{\{out,b,c\}} TR$. Furthermore, both the sender and the receiver are non-retracting and prompt, and so $TS \upharpoonright_{\{in,a,d\}} \upharpoonright_{\{out,b,c\}} TR$ is also non-retracting and prompt. Thus the timed network

$$TNETWORK = ((TS \upharpoonright_{\{in,a,d\}} \upharpoonright_{\{out,b,c\}} TR) \Sigma \upharpoonright_{\{a,b,c,d\}} (TM1 \upharpoonright_{\{a,b\}} \upharpoonright_{\{c,d\}} TM2)) \setminus \{a, b, c, d\}$$

is a timewise refinement of the untimed network, and so it must be a one-place buffer. Thus the functional correctness of the timed network may be deduced from an untimed analysis.

Of course, to do an analysis of the timing behaviour of the network it would be necessary to use the full power of the timed model. To consider the maximum time between input and output it is necessary to know for how long it is necessary to input messages into the media before output can be guaranteed; and to optimise the value of the timeout t it is necessary to know the expected delay in the media of a successfully transmitted message. The technique of timewise refinement cannot contribute to these concerns; its role is rather to complement them by allowing the appropriate use of more abstract methods for some analysis of aspects of a system's behaviour, even when other aspects require the use of the more complicated timed models.

A railroad crossing

This example was originally presented in [Lev87]. It gives an extremely simple model of a railroad crossing which is nevertheless complex enough to illustrate a number of aspects of the modelling and verification of timed systems.

The system is described as consisting of three components: a train, a gate, and a gate controller. The gate should be up to allow traffic to pass when no train is approaching, but should be lowered to obstruct traffic when a train is close to reaching the crossing. It is the task of the controller to monitor the approach of a train, and to instruct the gate to be lowered within the appropriate time. The train is modelled at a high level of abstraction: the only relevant aspects of the train's behaviour are when it is nearing the crossing, when it is entering it, when it is leaving it; and the delays between these events.

A number of safety conditions are formulated. These require the gate to be down when the train enters the crossing; the gate not to change state for ten time units before the train enters the crossing; and the train to have left the crossing by the time the gate goes up. We also require the liveness property that the crossing is deadlock-free. A system that deadlocked with the gate down would meet the safety conditions, but would not be satisfactory.

We begin with an untimed description and analysis of the system, to investigate which of these properties may be verified at the untimed level. We keep the process descriptions as simple as possible, including only those events that are relevant to consideration of these properties.

The gate controller *Controller* receives two types of signal from the crossing sensors: *near.ind*, which informs the controller that the train is approaching, and *out.ind*, which indicates that the train has left the crossing. It sends two types of signal to the crossing gate mechanism: *down.command*, and *up.command*, which instruct the gate to go down and up respectively. These four events form the alphabet *C* of the controller.

The gate, modelled by *Gate*, responds to the commands sent by the controller. We include additional events *up* and *down* to model the position of the gate. These two events, together with *up.command* and *down.command* form the alphabet *G* of the gate.

This leaves us with the following description of the crossing mechanism:

$$Crossing = Controller _C \parallel_G Gate$$

The controller responds to sensory inputs by issuing gate command signals:

$$\begin{aligned} Controller &= near.ind \longrightarrow down.command \longrightarrow Controller \\ &\square \\ &out.ind \longrightarrow up.command \longrightarrow Controller \end{aligned}$$

The gate process responds to the controller's signals by raising and lowering the gate.

$$\begin{aligned} Gate &= down.command \longrightarrow down \longrightarrow Gate \\ &\square \\ &up.command \longrightarrow up \longrightarrow Gate \end{aligned}$$

To reason about the behaviour of the system as a train approaches and reaches the crossing, we model the effect of such a happening via the crossing sensors. The train triggers the sensors by means of the *near.ind* and *out.ind* events. The events *train.near*, *enter.crossing* and *leave.crossing* model respectively the situations where the train is close to the crossing, the train enters the crossing, and the train leaves the crossing. These five events are all that are required for the sake of this analysis: they form the alphabet *T* of the train.

We will use the process *Train* to model the approach of the train, and its effect upon the crossing system.

$$\begin{aligned} \textit{Train} = & \textit{train.near} \longrightarrow \textit{near.ind} \longrightarrow \textit{enter.crossing} \longrightarrow \\ & \textit{leave.crossing} \longrightarrow \textit{out.ind} \longrightarrow \textit{Train} \end{aligned}$$

The crossing system, in conjunction with the train, is described as follows:

$$\textit{System} = \textit{Train} \mathit{T} \parallel_{\textit{CUG}} \textit{Crossing}$$

We are now able to express the properties we gave earlier in terms of the events we have chosen to model the system.

If the train enters the crossing, then the gate should have gone down more recently than it went up:

$$\textit{Safety1}(tr) = \textit{last}(tr) = \textit{enter.crossing} \Rightarrow \textit{last}(tr \upharpoonright \{\textit{down}, \textit{up}\}) = \textit{down}.$$

If the train enters the crossing at time t , then no *down* or *up* events should have occurred in the preceding ten time units; the projection of the trace to those events over that interval is empty:

$$\textit{Safety2}(s, \aleph) = \langle (t, \textit{enter.crossing}) \rangle \in s \Rightarrow (s \upharpoonright \{\textit{down}, \textit{up}\} \upharpoonright [t - 10, t]) = \langle \rangle$$

If the gate goes up, then the train must have left the crossing more recently than it entered it:

$$\begin{aligned} \textit{Safety3}(tr) = & \textit{last}(tr) = \textit{up} \Rightarrow \\ & \textit{last}(tr \upharpoonright \{\textit{enter.crossing}, \textit{leave.crossing}\}) = \textit{leave.crossing} \end{aligned}$$

Finally, the system must be deadlock-free.

$$\textit{Liveness1}(tr, \aleph) = \textit{deadlock-freedom}$$

Each of these properties has been expressed at the highest possible level of abstraction. In each case, the simplest model has been used to capture the required property. Safety properties 1 and 3 are expressible in the untimed traces model. Safety property 2 concerns explicit timing issues, so the timed model is required in order to express it. Deadlock-freedom is expressible using the untimed failures model.

Safety properties 1 and 3, and the liveness property, are candidates for being established by the untimed system description. They may be established by use of algebraic laws, or by use of proof rules. An alternative approach would be to use model-checking for these properties directly[Ros94]. The states of the system are shown in Figure 4 (where internal transitions corresponding to unwinding of recursions have been elided). Examination of the diagram reveals that at any point where *up* is possible, there must have been a *leave.crossing* event more recently than an *enter.crossing* event. Thus *Safety3* is satisfied. On the other hand, there are *enter.crossing* transitions where *up* is more recent than *down*, showing that in fact *Safety1* is not satisfied. Even though it is expressible as an untimed requirement, it turns out that its validity rests upon timing properties of the system (in particular, that the gate goes down in less time than it takes for the train to reach the crossing). Finally, every state has some transition out of it, so the system is deadlock-free, meeting the liveness requirement. Both *Safety3* and *Liveness1* are easily checked by the Failures Divergences Refinement checker FDR [FSE93].

Timewise refinement allows timing information to be added to the description of the system while preserving the properties already established. We firstly include the timing information

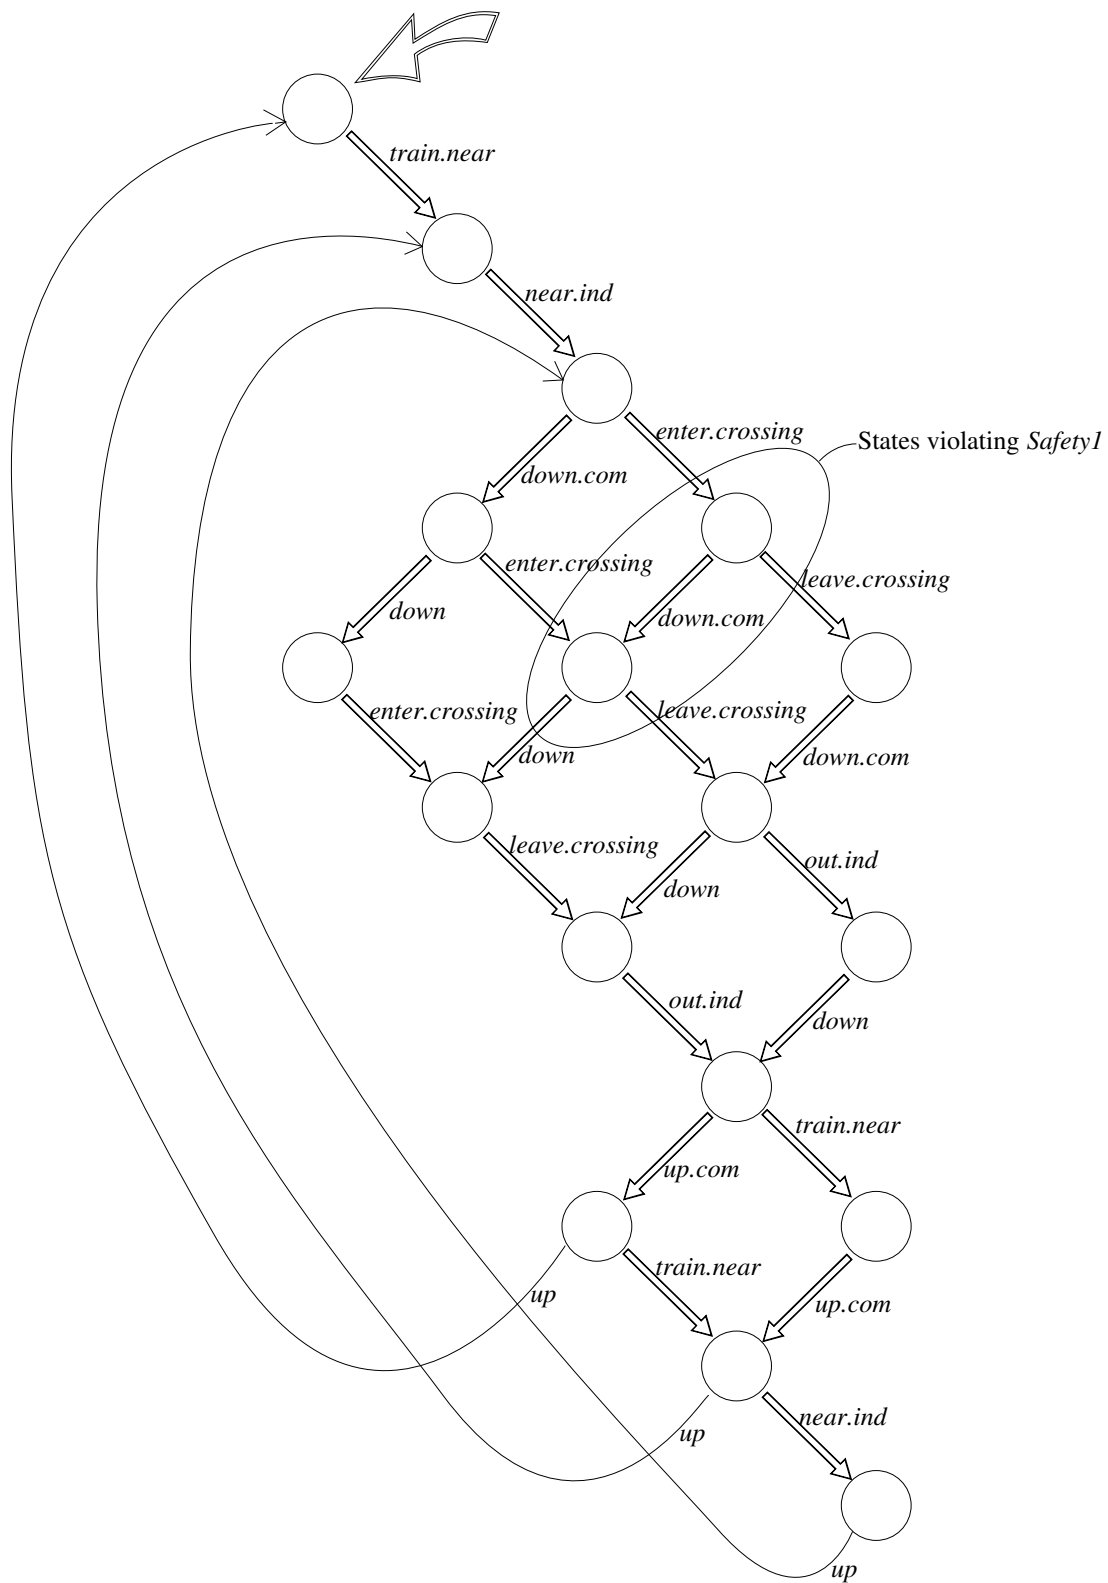


Figure 4: States and transitions for an untimed railway crossing

we have about the train: that it takes at least 5 minutes from triggering the *near.ind* sensor to reach the crossing; and that it takes at least 20 seconds to get across the crossing.

$$\begin{aligned} TTrain &= \text{train.near} \longrightarrow \text{near.ind} \xrightarrow{300} \text{enter.crossing} \xrightarrow{20} \\ &\quad \text{leave.crossing} \longrightarrow \text{out.ind} \longrightarrow TTrain \end{aligned}$$

The controller takes a negligible amount of time ϵ from receiving a signal from a sensor to relaying the corresponding instruction to the gate.

$$\begin{aligned} TController &= \text{near.ind} \xrightarrow{\epsilon} \text{down.command} \longrightarrow TController \\ &\quad \square \\ &\quad \text{out.ind} \xrightarrow{\epsilon} \text{up.command} \longrightarrow TController \end{aligned}$$

The timed gate process *TGate* process takes a non-negligible amount of time to get the gate into position following an instruction:

$$\begin{aligned} TGate &= \text{down.command} \xrightarrow{100} \text{down} \longrightarrow TGate \\ &\quad \square \\ &\quad \text{up.command} \xrightarrow{100} \text{up} \longrightarrow TGate \end{aligned}$$

However, this is still considerably less time than it takes for the train to reach the crossing, so the timed description is sufficiently detailed to establish *Safety1*, under the additional assumption that the events *up* and *down* are entirely under the control of the *Gate*. For further discussion of environmental assumptions, see [Sch96]. This environmental assumption is captured as $[0, \infty) \times \{\text{up}, \text{down}\} \subseteq \aleph$, and so the specification met by the system is

$$[0, \infty) \times \{\text{up}, \text{down}\} \subseteq \aleph \Rightarrow \text{Safety1}$$

The timed description is also sufficient to establish *Safety2*.

The timed process descriptions are simply the untimed process descriptions with times added. The results of Section 3.1 guarantee that the timed system trace refines the untimed one, ensuring that the timed system meets *Safety3*.

Finally, the three component timed processes are all non-retracting. Lemma 3.8 ensures that the system consisting of their parallel combination is a failures refinement of the untimed system, so it retains the untimed property *Liveness1*.

6 Discussion

We have seen how verifications of specifications can be mapped up the CSP hierarchy of models, and also an example of how general laws might be translated. Other properties (such as deterministic or compact) do not translate in general. For example, the deterministic untimed process $a \longrightarrow \text{Stop}$ is refined by the non-deterministic timed process $a \longrightarrow \text{Stop} \sqcap \text{Wait } 5 ; a \longrightarrow \text{Stop}$, which can perform or refuse to perform *a* at time 2.

Comparisons

There has also been some work in this area in the contexts of timed CCS and of timed ACP. Larsen and Yi [LaY93] have proposed a notion of *time-abstracting* bisimulation, which specifies when timed processes are equivalent modulo timing behaviour. Thus one process may be used to specify simply the functional behaviour of a system by requiring that any proposed implementation should be time-abstracting bisimilar to it. They prove that time-abstracting equivalence is

decidable for a timed CCS calculus [Wan90], in contrast to the refinement relation presented in this paper, which is not decidable. Interestingly, they also establish that time-abstracting congruence (i.e. equivalence in all contexts) is standard timed bisimulation. The corresponding result for this paper is that untimed traces congruence for timed processes is the same as (finite) timed failures equivalence.

Baeten and Bergstra [BaB92] have considered the embedding of untimed ACP into real time ACP. They propose a translation of untimed ACP into the timed setting, for example translating a to $\int_{t>0} a(t)$: an untimed a process specifies nothing about the time the a should occur, so it translates to the timed process that can perform an a at any time. This is also the philosophy of this paper. They also consider the translation of certain identities of ACP into the timed framework; this permits reasoning at a higher (untimed) level of abstraction to be incorporated when detailed reasoning about timing issues is also required.

Urgent events

Events in timed CSP are treated as non-urgent: processes are described in terms of when events become enabled and disabled. A process does not have complete *control* over the performance of (visible) events, since the cooperation of its environment is always required for an event to occur. The approach taken in timed CSP, and some other process algebras, is that the process therefore is not given *responsibility* for the performance of events; the most that can be expected of a process is that it is willing to go along with the occurrence of an event.

The other principal approach to events in timed process algebras, taken by Timed LOTOS [QuF87, BoL92] among others, is to assume that the environment of a process will permit certain events to occur when the process wishes to perform them. Thus actions of a process may be treated as *urgent*, in that they occur as soon as they are enabled by the process. However, since the environmental assumption is not always justified—the process may be placed in parallel with a process which is unwilling to allow an urgent action to occur—it is possible that a catastrophic timelock occurs and no progress can be made. It is the responsibility of the system designer to ensure that no timelocks occur in a particular design. Alternatively, the parallel operator may treat urgent events as observable signals. In this treatment non-urgent events may synchronise with urgent ones but may not block them. This approach was taken in [DJS90, Dav93].

In fact, urgent and non-urgent actions are incorporated within a single framework in [BoL92], where actions are initially non-urgent, but may be made urgent by an *urgency* operator.

The question arises: how do the techniques developed here apply in the case of urgency?

The natural timed denotational semantics for systems with urgent actions is in terms of timed traces. Thus we might hope that the traces refinement discussed in this paper would remain applicable in the presence of urgent actions. Indeed, the may testing characterisation of the refinement relation yields the same results, and all of the timed LOTOS operators preserve the refinement relation. Also, the urgency operator $\rho(U)$ preserves refinement, in the sense that if P is refined by Q , then it is also refined by $\rho(U)(Q)$; making events urgent does not remove any safety properties. We might expect to use timewise refinement techniques when the requirements on a system split naturally into untimed safety properties, and timing requirements.

The notion of a timed failure is inapplicable to an urgent action, since it is performed as soon as it becomes possible. All the information about when a process was ready to perform it is contained in the timed trace; refusal information is redundant (though it remains appropriate for non-urgent actions). Thus the failures refinement relation is not appropriate when dealing with urgent timed systems. This is reflected by the relation obtained by considering the relation \sqsubseteq_f defined in terms of must testing. If process Q has an urgent action a to be performed at time t , then the test $T = \text{Wait}(t+1); \omega \rightarrow \text{Stop}$ will yield a timelock at time t , since it will not cooperate with Q on the action a . However, any untimed process P (except an immediately

divergent one) has $P \text{ \underline{must} } \Theta(T)$. Thus no useful process P has $P \sqsubseteq_f Q$. The refinement relation collapses in the face of urgent actions, holding only when the untimed process is divergent.

In systems where both urgent and non-urgent events are present, it appears that the results for traces refinement apply as easily as in the pure non-urgent case of timed CSP. Failures refinement as discussed here will be useful only as far as the first urgent action. It is not even clear whether there is any more appropriate notion of failures style refinement, since timed liveness properties in urgent systems do not appear to have any untimed counterpart. The archetypical timed liveness property is that of timelock-freedom, and it is unlikely that this could be established by untimed analysis.

Current and future work

Earlier work [Sch89] investigated the relationship between the untimed models and the standard timed failures model of [Ree88]. The difficulties encountered in using that model to treat infinite behaviour led to the development of the infinite failures model, which supports a more natural treatment of timewise refinement from the untimed models.

Other refinement relations are also under investigation. In particular, a relation between the failures/divergences model and the timed failures stabilities model that treats instability as divergence has that all CSP operators preserve refinement; and this refinement relation is complete for stable processes. When stability considerations are important then this relation would be the natural one to use. Of particular interest is the relationship between the timed models and the timed probabilistic models for CSP developed by Lowe [Low91]. Work has already been done in this direction (see e.g. [Low92]), which it seems should fit into the framework presented in this paper.

The underlying theory presented here is of course more general than simply CSP, and should be applicable wherever processes are modelled in terms of the behaviours they may exhibit. It may for example be applicable to Gerth and Kuiper's interface refinement [GKS92]. I feel that the theory will be useful only if refinement relations can be established at the syntactic level, since if refinement can be shown only by examining the semantics directly, then verifying abstract specifications of processes via refinement is unlikely to be much easier than performing the verification directly.

Acknowledgements

Thanks are due to Mike Mislove, Bill Roscoe, Mike Reed, Jim Davies, Dave Jackson, Michael Goldsmith, and members of the ESPRIT SPEC, REACT and CONCUR projects for useful discussions on this material. In particular Bill Roscoe was responsible for the formulation of the alternating bit protocol presented here. Thanks are also due to the anonymous referees whose thorough reading and careful comments have been extremely useful.

This work was funded under SERC research fellowship B91/RFH/312. Additional funding was received from EC Basic Research Action 7166 CONCUR 2, and the U.S. Office of Naval Research.

References

- [BaB92] J.C.M. Baeten and J.A. Bergstra, *Discrete time process algebra*, University of Amsterdam, Report P9208b, 1992.
- [BoL92] T. Bolognesi and F. Lucidi, *Timed process algebras with urgent interactions and a unique powerful binary operator*, LNCS 600, 1992.

- [BrR85] S.D. Brookes and A.W. Roscoe, *An improved failures model for communicating sequential processes*, Proceedings, Seminar on Concurrency, LNCS 197, 1985.
- [BRW92] S.D. Brookes, A.W. Roscoe, and D.J. Walker, *An operational semantics for CSP*, submitted for publication, 1992.
- [CIZ92] R. Cleaveland and A.E. Zwarico, *A theory of testing for real-time*, North Carolina SU and Johns Hopkins, 1992.
- [DaS93] J.W. Davies and S.A. Schneider, *Recursion induction for real-time processes*, Formal Aspects of Computing, 5 (6), 1993.
- [DaS95] J.W. Davies and S.A. Schneider, *A brief history of timed CSP*, Theoretical Computer Science, 138, 1995.
- [Dav93] J.W. Davies, *Specification and proof in real-time CSP*, Cambridge University Press, 1993.
- [DJS90] J.W. Davies, D.M. Jackson, and S.A. Schneider, *Making things happen in Timed CSP*, PRG Technical Report 2-90, Oxford, 1990.
- [FSE93] Formal Systems (Europe) Ltd, *The Failures Divergences Refinement User Manual*, 1993.
- [GKS92] R. Gerth, R. Kuiper, and J Segers, *Interface refinement in reactive systems*, Proceedings of CONCUR '92, LNCS 630, 1992.
- [Hen88] M. Hennessy, *Algebraic Theory of Processes*, MIT Press, 1988.
- [Hoa85] C.A.R. Hoare, *Communicating Sequential Processes*, Prentice-Hall, 1985.
- [LaY93] K.G. Larsen and Wang Yi, *Time abstracted bisimulation: implicit specifications and decidability*, Proceedings of MFPS '93, LNCS, 1993.
- [Lev87] N. G. Leveson and J. L. Stolzy, *Safety Analysis Using Petri Nets' IEEE Transactions on Software Engineering*, 1987.
- [Low91] G. Lowe, *Prioritized and probabilistic models of timed CSP*, Oxford University Computing Laboratory Technical Report PRG-TR-24-91, 1991.
- [Low92] G. Lowe, *Relating the prioritized model of timed CSP to the timed failures model*, Oxford University Computing Laboratory, 1992.
- [LyV92] N. Lynch and F. Vaandrager, *Forward and backward simulations for timing-based systems*, Proceedings, Real-time: theory in practise, LNCS 600, 1992.
- [Mil89] R. Milner, *Communication and Concurrency*, Prentice-Hall, 1989.
- [MRS95] M.W. Mislove, A.W. Roscoe, and S.A. Schneider, *Fixed points without completeness*, Theoretical Computer Science, 138, 1995.
- [QuF87] J. Quemada and A. Fernandez *Introduction of quantitative relative time into LOTOS*, in "Protocol Specification, Testing and Verification VII" North Holland, 1987.
- [Ree88] G.M. Reed, *A Uniform Mathematical Theory for Real-Time Distributed Computing*, Oxford University DPhil thesis, 1988.
- [ReR86] G.M. Reed and A.W. Roscoe, *A timed model for communicating sequential processes*, Proceedings, 13th ICALP, LNCS 226, 1986.
- [Ros88] A.W. Roscoe, *Unbounded Nondeterminism in CSP*, Oxford University Computing Laboratory technical monograph PRG-67, 1988.
- [Ros94] A.W. Roscoe, *Model-checking CSP*, in *A Classical Mind* edited by A.W. Roscoe, Prentice-Hall, 1994.

- [Sch89] S.A. Schneider, *Correctness and Communication in Real-Time Systems*, Oxford University DPhil. thesis 1989.
- [Sch91] S.A. Schneider, *An operational semantics for timed CSP*, Proceedings of the Workshop on Concurrency, Report 63, Programming Methodology Group, Chalmers University, 1992.
- [Sch92] S.A. Schneider, *Unbounded nondeterminism for real-time processes*, Oxford University Technical Report 13-92, 1992.
- [Sch95] S.A. Schneider, *An operational semantics for timed CSP*, Information and Computation, 116 (2), 1995.
- [Sch96] S.A. Schneider, *Specification and verification in Timed CSP*, in *Real-time Systems: Specification, Verification and Analysis*, Prentice-Hall 1996
- [Wan90] Wang Yi, *Real-time behaviour of asynchronous agents*, Proceedings of CONCUR '90, LNCS 458, 1990.

A Semantic models and functions

Traces

The traces model \mathcal{M}_{UT} is defined to be those sets of traces that are non-empty, and closed under prefixing. They are ordered under set containment.

The semantic function \mathcal{F}_{UT}

The space of environments is defined as

$$\mathcal{E}_{UT} = \mathcal{V} \rightarrow \mathcal{M}_{UT}$$

The semantic function

$$\mathcal{F}_{UT} : CSP \longrightarrow \mathcal{E}_{UT} \longrightarrow \mathcal{M}_{UT}$$

is defined by the following set of equations:

$$\begin{aligned}
\mathcal{F}_{UT}[\mathit{Loop}]\rho &\hat{=} \{tr \in \Sigma^*\} \\
\mathcal{F}_{UT}[\mathit{Stop}]\rho &\hat{=} \{\langle \rangle\} \\
\mathcal{F}_{UT}[\mathit{Skip}]\rho &\hat{=} \{\langle \rangle, \langle \surd \rangle\} \\
\mathcal{F}_{UT}[P ; Q]\rho &\hat{=} \{tr \mid tr \in \mathcal{F}_{UT}[P]\rho \wedge tr \upharpoonright \{\surd\} = \langle \rangle\} \\
&\cup \\
&\{tr_P \frown tr_Q \mid tr_P \frown \langle \surd \rangle \in \mathcal{F}_{UT}[P]\rho \wedge \\
&\quad tr_P \upharpoonright \{\surd\} = \langle \rangle \wedge tr_Q \in \mathcal{F}_{UT}[Q]\rho\} \\
\mathcal{F}_{UT}[P \overset{t_0}{\triangleright} Q]\rho &\hat{=} \mathcal{F}_{UT}[P]\rho \cup \mathcal{F}_{UT}[Q]\rho \\
\mathcal{F}_{UT}[P \square Q]\rho &\hat{=} \mathcal{F}_{UT}[P]\rho \cup \mathcal{F}_{UT}[Q]\rho \\
\mathcal{F}_{UT}[a : A \longrightarrow P_a]\rho &\hat{=} \bigcup_{a \in A} \{\langle a \rangle \frown tr \mid tr \in \mathcal{F}_{UT}[P_a]\rho\} \\
\mathcal{F}_{UT}[\prod_{i \in I} P_i]\rho &\hat{=} \bigcup_{i \in I} \mathcal{F}_{UT}[P_i]\rho
\end{aligned}$$

$$\begin{aligned}
\mathcal{F}_{UT}[[P \parallel_B Q]]\rho &\hat{=} \{tr \in (A \cup B)^* \mid tr \upharpoonright A \in \mathcal{F}_{UT}[[P]]\rho \wedge \\
&\quad tr \upharpoonright B \in \mathcal{F}_{UT}[[Q]]\rho\} \\
\mathcal{F}_{UT}[[P \parallel Q]]\rho &\hat{=} \{tr \mid \exists tr_P \in \mathcal{F}_{UT}[[P]]\rho, tr_Q \in \mathcal{F}_{UT}[[Q]]\rho \bullet \\
&\quad tr \text{ interleaves } (tr_P, tr_Q)\} \\
\mathcal{F}_{UT}[[P \setminus A]]\rho &\hat{=} \{tr \setminus A \mid tr \in \mathcal{F}_{UT}[[P]]\rho\} \\
\mathcal{F}_{UT}[[f(P)]]\rho &\hat{=} \{f(tr) \mid tr \in \mathcal{F}_{UT}[[P]]\rho\} \\
\mathcal{F}_{UT}[[f^{-1}(P)]]\rho &\hat{=} \{tr \mid f(tr) \in \mathcal{F}_{UT}[[P]]\rho\} \\
\mathcal{F}_{UT}[[Y]]\rho &\hat{=} \rho(Y)
\end{aligned}$$

where *interleaves* is defined on sequences $m, m_1, m_2 : \Sigma^* \cup \Sigma^\omega$ as follows:

$$\begin{aligned}
m \text{ interleaves } (m_1, m_2) &\Leftrightarrow \#m = \#m_1 + \#m_2 \\
&\wedge \exists f_1 : (1..\#m_1) \mapsto_{<} (1..\#m), f_2 : (1..\#m_2) \mapsto_{<} (1..\#m) \bullet \\
&\quad \text{ran}(f_1) \cap \text{ran}(f_2) = \{\} \\
&\quad \wedge \text{ran}(f_1) \cup \text{ran}(f_2) = 1..\#m \\
&\quad \wedge \forall n \in 1..\#m_1 \bullet m_1 @ n = m @ f_1(n) \\
&\quad \wedge \forall n \in 1..\#m_2 \bullet m_2 @ n = m @ f_2(n)
\end{aligned}$$

where $1..n = \{j \mid 1 \leq j \leq n\}$ $1..\omega = \mathbf{N}$, $A \mapsto_{<} B$ is the set of monotonically increasing injective total functions from A to B , $\text{ran}(f)$ is the range of function f , and $m @ n$ is the n th element of sequence m .

The semantics $\mathcal{F}_{UT}[[\mu Y \circ P]]\rho$ of a recursive term is defined to be the least fixed point of the function

$$\lambda S \bullet \mathcal{F}_{UT}[[P]](\rho[S/Y]) : \mathcal{M}_{UT} \rightarrow \mathcal{M}_{UT}$$

where substitution in an environment is defined by

$$\begin{aligned}
\rho[S/Y](Z) &= \rho(Z) \quad Z \neq Y \\
\rho[S/Y](Y) &= S
\end{aligned}$$

Untimed infinite traces, failures, and divergences

The process axioms given in [Ros88] correspond to the following properties required of a set S for it to correspond to the set of observations of some process. Thus the semantic model \mathcal{M}_{UI} is the collection of sets

$$S \subseteq \{f\} \times (\Sigma^* \times \mathbf{P}(\Sigma)) \cup \{d\} \times \Sigma^* \cup \{i\} \times \Sigma^\omega$$

(ordered under reverse containment) that meet these eight axioms:

- (1) $(f, (tr_1 \hat{\ } tr_2, \{\})) \in S \Rightarrow (f, (tr_1, \{\})) \in S$
- (2) $(f, (tr, X)) \in S \wedge Y \subseteq X \Rightarrow (f, (tr, Y)) \in S$
- (3) $(f, (tr, X)) \in S \wedge \forall a \in Y \bullet (f, (tr \hat{\ } \langle a \rangle, \{\})) \notin S \Rightarrow (f, (tr, X \cup Y)) \in S$
- (4) $(d, tr_1) \in S \Rightarrow (d, tr_1 \hat{\ } tr_2) \in S$
- (5) $(d, tr_1) \in S \Rightarrow (f, (tr_1 \hat{\ } tr_2, X)) \in S$
- (6) $(i, tr \hat{\ } u) \in S \Rightarrow (f, (tr, \{\})) \in S$
- (7) $(d, tr) \in S \Rightarrow (i, tr \hat{\ } u) \in S$
- (8) $(f, (tr_1, \{\})) \in S \Rightarrow \exists T \bullet (\forall tr_2 \in T \bullet (f, (tr_1 \hat{\ } tr_2, \{a \mid tr_2 \hat{\ } \langle a \rangle \notin T\})) \in S \wedge \{(i, tr_1 \hat{\ } u) \mid u \in \overline{T}\} \subseteq S)$

Here $\overline{T} = \{u \in \Sigma^\omega \mid \forall tr < u \bullet tr \in T\}$, where T ranges over finite prefix closed sets of finite traces.

The semantic function \mathcal{F}_{UI}

The function \mathcal{F}_{UI} is defined in terms of three functions \mathcal{F}_{UD} , \mathcal{F}_{UF} , and \mathcal{F}_I , yielding divergences, failures, and infinite traces respectively. It is then given by

$$\begin{aligned} \mathcal{F}_{UI}[[P]]\rho &= \{(d, tr) \mid tr \in \mathcal{F}_{UD}[[P]]\rho\} \\ &\cup \{(f, (tr, X)) \mid (tr, X) \in \mathcal{F}_{UF}[[P]]\rho\} \\ &\cup \{(i, u) \mid u \in \mathcal{F}_I[[P]]\rho\} \end{aligned}$$

The semantic function \mathcal{F}_{UD}

Define $\mathcal{E}_{UD} = \mathcal{V} \rightarrow \mathcal{M}_{UD}$. The semantic function

$$\mathcal{F}_{UD} : CSP \rightarrow \mathcal{E}_{UD} \rightarrow \mathcal{M}_{UD}$$

is defined by the following set of equations:

$$\begin{aligned} \mathcal{F}_{UD}[[Loop]]\rho &\hat{=} \{tr \mid tr \in \Sigma^*\} \\ \mathcal{F}_{UD}[[Stop]]\rho &\hat{=} \{\} \\ \mathcal{F}_{UD}[[Skip]]\rho &\hat{=} \{\} \\ \mathcal{F}_{UD}[[P; Q]]\rho &\hat{=} \{tr \frown tr' \mid tr \in \mathcal{F}_{UD}[[P]]\rho \wedge \surd \notin \sigma(tr) \wedge tr' \in \Sigma^*\} \\ &\cup \\ &\{tr \frown tr' \mid (tr \frown \langle \surd \rangle, \{\}) \in \mathcal{F}_{UF}[[P]]\rho \wedge \surd \notin \sigma(tr) \\ &\quad \wedge tr' \in \mathcal{F}_{UD}[[Q]]\rho\} \\ \mathcal{F}_{UD}[[P \stackrel{t_0}{\triangleright} Q]]\rho &\hat{=} \mathcal{F}_{UD}[[P]]\rho \cup \mathcal{F}_{UD}[[Q]]\rho \\ \mathcal{F}_{UD}[[P \square Q]]\rho &\hat{=} \mathcal{F}_{UD}[[P]]\rho \cup \mathcal{F}_{UD}[[Q]]\rho \\ \mathcal{F}_{UD}[[a : A \longrightarrow P_a]]\rho &\hat{=} \bigcup_{a \in A} \{\langle a \rangle \frown tr \mid tr \in \mathcal{F}_{UD}[[P_a]]\rho\} \\ \mathcal{F}_{UD}[[\prod_{i \in I} P_i]]\rho &\hat{=} \bigcup_{i \in I} \mathcal{F}_{UD}[[P_i]]\rho \\ \mathcal{F}_{UD}[[P _A \parallel_B Q]]\rho &\hat{=} \{tr \frown tr' \mid tr \in (A \cup B)^* \\ &\quad \wedge \\ &\quad (tr \upharpoonright A, \{\}) \in \mathcal{F}_{UF}[[P]]\rho \wedge tr \upharpoonright B \in \mathcal{F}_{UD}[[Q]]\rho \\ &\quad \vee tr \upharpoonright A \in \mathcal{F}_{UD}[[P]]\rho \wedge (tr \upharpoonright B, \{\}) \in \mathcal{F}_{UF}[[Q]]\rho\} \\ \mathcal{F}_{UD}[[P \parallel Q]]\rho &\hat{=} \{tr \mid \exists tr_P \in \mathcal{F}_{UD}[[P]]\rho, (tr_Q, \{\}) \in \mathcal{F}_{UF}[[Q]]\rho \bullet \\ &\quad tr \text{ interleaves } (tr_P, tr_Q)\} \\ &\cup \\ &\{tr \mid \exists (tr_P, \{\}) \in \mathcal{F}_{UF}[[P]]\rho, tr_Q \in \mathcal{F}_{UD}[[Q]]\rho \bullet \\ &\quad tr \text{ interleaves } (tr_P, tr_Q)\} \\ \mathcal{F}_{UD}[[P \setminus A]]\rho &\hat{=} \{tr \setminus A \frown tr' \mid tr \in \mathcal{F}_{UD}[[P]]\rho\} \\ &\cup \\ &\{u \setminus A \frown tr' \mid u \in \mathcal{F}_I[[P]]\rho \wedge \#(u \setminus A) < \infty\} \end{aligned}$$

$$\begin{aligned}
\mathcal{F}_{UD}[\![f(P)]\!] \rho &\hat{=} \{f(tr) \frown tr' \mid tr \in \mathcal{F}_{UD}[\![P]\!] \rho\} \\
\mathcal{F}_{UD}[\![f^{-1}(P)]\!] \rho &\hat{=} \{tr \mid f(tr) \in \mathcal{F}_{UD}[\![P]\!] \rho\} \\
\mathcal{F}_{UD}[\![Y]\!] \rho &\hat{=} \rho(Y)
\end{aligned}$$

The semantic function \mathcal{F}_I

Define $\mathcal{E}_I = \mathcal{V} \rightarrow \mathcal{M}_I$. The semantic function

$$\mathcal{F}_I : CSP \rightarrow \mathcal{E}_I \rightarrow \mathcal{M}_I$$

is defined by the following set of equations:

$$\begin{aligned}
\mathcal{F}_I[\![Loop]\!] \rho &\hat{=} \{u \mid u \in \Sigma^* \cup \Sigma^\omega\} \\
\mathcal{F}_I[\![Stop]\!] \rho &\hat{=} \{\} \\
\mathcal{F}_I[\![Skip]\!] \rho &\hat{=} \{\} \\
\mathcal{F}_I[\![P ; Q]\!] \rho &\hat{=} \{u \mid u \in \mathcal{F}_I[\![P]\!] \rho \wedge \checkmark \notin \sigma(u)\} \\
&\cup \\
&\{tr \frown u' \mid (tr \frown \langle \checkmark \rangle, \{\}) \in \mathcal{F}_{UF}[\![P]\!] \rho \wedge \checkmark \notin \sigma(tr) \\
&\quad \wedge u' \in \mathcal{F}_I[\![Q]\!] \rho\} \\
&\cup \\
&\{tr \frown u \mid tr \in \mathcal{F}_{UD}[\![P ; Q]\!] \rho\} \\
\mathcal{F}_I[\![P \overset{t_0}{\triangleright} Q]\!] \rho &\hat{=} \mathcal{F}_I[\![P]\!] \rho \cup \mathcal{F}_I[\![Q]\!] \rho \\
\mathcal{F}_I[\![P \square Q]\!] \rho &\hat{=} \mathcal{F}_I[\![P]\!] \rho \cup \mathcal{F}_I[\![Q]\!] \rho \\
\mathcal{F}_I[\![a : A \longrightarrow P_a]\!] \rho &\hat{=} \bigcup_{a \in A} \{\langle a \rangle \frown u \mid u \in \mathcal{F}_I[\![P_a]\!] \rho\} \\
\mathcal{F}_I[\![\prod_{i \in I} P_i]\!] \rho &\hat{=} \bigcup_{i \in I} \mathcal{F}_I[\![P_i]\!] \rho \\
\mathcal{F}_I[\![P _A \parallel_B Q]\!] \rho &\hat{=} \{u \mid u \in (A \cup B)^\omega \\
&\quad \wedge u \upharpoonright A \in \mathcal{F}_I[\![P]\!] \rho \vee (u \upharpoonright A, \{\}) \in \mathcal{F}_{UF}[\![P]\!] \rho \\
&\quad \wedge u \upharpoonright B \in \mathcal{F}_I[\![Q]\!] \rho \vee (u \upharpoonright B, \{\}) \in \mathcal{F}_{UF}[\![Q]\!] \rho \\
&\quad \} \\
&\cup \\
&\{tr \frown u \mid tr \in \mathcal{F}_{UD}[\![P _A \parallel_B Q]\!] \rho\} \\
\mathcal{F}_I[\![P \parallel Q]\!] \rho &\hat{=} \{u \mid \exists u_P, u_Q \bullet (\#u_P = \infty \vee \#u_Q = \infty) \\
&\quad \wedge u \text{ interleaves}(u_P, u_Q) \\
&\quad \wedge (u_P, \{\}) \in \mathcal{F}_{UF}[\![P]\!] \rho \vee u_P \in \mathcal{F}_I[\![P]\!] \rho \\
&\quad \wedge (u_Q, \{\}) \in \mathcal{F}_{UF}[\![Q]\!] \rho \vee u_Q \in \mathcal{F}_I[\![Q]\!] \rho\} \\
&\cup \\
&\{tr \frown u \mid tr \in \mathcal{F}_{UD}[\![P \parallel Q]\!] \rho\} \\
\mathcal{F}_I[\![P \setminus A]\!] \rho &\hat{=} \{u \setminus A \mid u \in \mathcal{F}_I[\![P]\!] \rho \wedge \#(u \setminus A) = \infty\} \\
&\cup \\
&\{tr \frown u \mid tr \in \mathcal{F}_{UD}[\![P \setminus A]\!] \rho\}
\end{aligned}$$

$$\begin{aligned}
\mathcal{F}_I[f(P)]\rho &\hat{=} \{f(u)\hat{\ } \mid u \in \mathcal{F}_I[P]\rho\} \\
&\cup \\
&\{tr\hat{\ }u \mid tr \in \mathcal{F}_{UD}[f(P)]\rho\} \\
\mathcal{F}_I[f^{-1}(P)]\rho &\hat{=} \{u \mid f(u) \in \mathcal{F}_I[P]\rho\} \\
\mathcal{F}_I[Y]\rho &\hat{=} \rho(Y)
\end{aligned}$$

The semantic function \mathcal{F}_{UF}

Define $\mathcal{E}_{UF} = \mathcal{V} \rightarrow \mathcal{M}_{UF}$. The semantic function

$$\mathcal{F}_{UF} : CSP \rightarrow \mathcal{E}_{UF} \rightarrow \mathcal{M}_{UF}$$

is defined by the following set of equations:

$$\begin{aligned}
\mathcal{F}_{UF}[Loop]\rho &\hat{=} \{(tr, X) \mid tr \in \Sigma^* \wedge X \subseteq \Sigma\} \\
\mathcal{F}_{UF}[Stop]\rho &\hat{=} \{(\langle \rangle, X) \mid X \subseteq \Sigma\} \\
\mathcal{F}_{UF}[Skip]\rho &\hat{=} \{(\langle \rangle, X) \mid \surd \notin X\} \\
&\cup \\
&\{(\langle \surd \rangle, X) \mid X \subseteq \Sigma\} \\
\mathcal{F}_{UF}[P; Q]\rho &\hat{=} \{(tr, X) \mid \surd \notin \sigma(tr) \wedge \\
&\quad (tr, X \cup \{\surd\}) \in \mathcal{F}_{UF}[P]\rho\} \\
&\cup \\
&\{(tr\hat{\ }tr', X) \mid \surd \notin \sigma(tr) \\
&\quad \wedge (tr\hat{\}\langle \surd \rangle, \{\}) \in \mathcal{F}_{UF}[P]\rho \\
&\quad \wedge (tr', X) \in \mathcal{F}_{UF}[Q]\rho\} \\
&\cup \\
&\{(tr, X) \mid tr \in \mathcal{F}_{UD}[P; Q]\rho\} \\
\mathcal{F}_{UF}[P \overset{to}{\triangleright} Q]\rho &\hat{=} \mathcal{F}_{UF}[Q]\rho \cup \{(tr, X) \mid (tr, X) \in \mathcal{F}_{UF}[P]\rho \wedge tr \neq \langle \rangle\} \\
\mathcal{F}_{UF}[P \square Q]\rho &\hat{=} \{(\langle \rangle, X) \mid (\langle \rangle, X) \in \mathcal{F}_{UF}[P]\rho \cap \mathcal{F}_{UF}[Q]\rho\} \\
&\cup \\
&\{(tr, X) \mid tr \neq \langle \rangle \wedge (tr, X) \in \mathcal{F}_{UF}[P]\rho \cup \mathcal{F}_{UF}[Q]\rho\} \\
\mathcal{F}_{UF}[a : A \longrightarrow P_a]\rho &\hat{=} \{(\langle \rangle, X) \mid X \cap A = \{\}\} \\
&\cup \\
&\bigcup_{a \in A} \{(\langle a \rangle\hat{\ }tr, X) \mid (tr, X) \in \mathcal{F}_{UF}[P_a]\rho\} \\
\mathcal{F}_{UF}[\prod_{i \in I} P_i]\rho &\hat{=} \bigcup_{i \in I} \mathcal{F}_{UF}[P_i]\rho \\
\mathcal{F}_{UF}[P _A \parallel_B Q]\rho &\hat{=} \{(tr, Z) \mid \exists X, Y \bullet (tr \upharpoonright A, X \upharpoonright A) \in \mathcal{F}_{UF}[P]\rho \wedge \\
&\quad (tr \upharpoonright B, Y \upharpoonright B) \in \mathcal{F}_{UF}[Q]\rho \wedge \\
&\quad (X \upharpoonright A) \cup (Y \upharpoonright B) = Z \upharpoonright A \cup B \wedge \\
&\quad tr = tr \upharpoonright (A \cup B)\} \\
&\cup \\
&\{(tr, X) \mid tr \in \mathcal{F}_{UD}[P _A \parallel_B Q]\rho\}
\end{aligned}$$

$$\begin{aligned}
\mathcal{F}_{UF}[[P \parallel Q]]\rho &\hat{=} \{(tr, X) \mid \exists tr_P, tr_Q \bullet tr \text{ interleaves}(tr_P, tr_Q) \\
&\quad \wedge (tr_P, X) \in \mathcal{F}_{UF}[[P]]\rho \\
&\quad \wedge (tr_Q, X) \in \mathcal{F}_{UF}[[Q]]\rho\} \\
&\cup \\
&\{(tr, X) \mid tr \in \mathcal{F}_{UD}[[P \parallel Q]]\rho\} \\
\mathcal{F}_{UF}[[P \setminus A]]\rho &\hat{=} \{(tr \setminus A, X) \mid (tr, X \cup A) \in \mathcal{F}_{UF}[[P]]\rho\} \\
&\cup \\
&\{(tr, X) \mid tr \in \mathcal{F}_{UD}[[P \setminus A]]\rho\} \\
\mathcal{F}_{UF}[[f(P)]]\rho &\hat{=} \{(f(tr), X) \mid (tr, f^{-1}(X)) \in \mathcal{F}_{UF}[[P]]\rho\} \\
&\cup \\
&\{(tr, X) \mid tr \in \mathcal{F}_{UD}[[f(P)]]\rho\} \\
\mathcal{F}_{UF}[[f^{-1}(P)]]\rho &\hat{=} \{(tr, X) \mid (f(tr), f(X)) \in \mathcal{F}_{UF}[[P]]\rho\} \\
\mathcal{F}_{UF}[[Y]]\rho &\hat{=} \rho(Y)
\end{aligned}$$

The semantics $\mathcal{F}_{UI}[\mu Y \circ P]\rho$ of a recursive term is defined to be the least fixed point of the function

$$F = \lambda S \bullet \mathcal{F}_I[[P]](\rho[S/Y]) : \mathcal{M}_{UI} \rightarrow \mathcal{M}_{UI}$$

It is established in [Ros88] that this is well-defined, and that this is equal to $F^\alpha(\perp)$ for some ordinal α , where $\perp = \mathcal{F}_{UI}[[Loop]]\rho$.

Infinite Timed Failures

The information ordering on behaviours is defined as follows:

$$(s', \aleph') \preceq (s, \aleph) \Leftrightarrow \exists s'' \bullet s = s' \frown s'' \wedge \aleph' \subseteq \aleph \triangleleft \text{begin}(s'')$$

We formally define \mathcal{M}_{TI} to be those subsets S of $T\Sigma_{\leq}^\omega \times IRSET$ satisfying axioms 1–3 given below, and axiom 4 to follow.

1. $(\langle \rangle, \{\}) \in S$
2. $(s, \aleph) \in S \wedge (s', \aleph') \preceq (s, \aleph) \Rightarrow (s', \aleph') \in S$
3. $(s, \aleph) \in S \Rightarrow$
 $\exists \aleph' \in IRSET \bullet \aleph \subseteq \aleph' \wedge (s, \aleph') \in S \wedge \forall (t, a) \in \mathbf{R}_+ \times \Sigma \bullet$
 $(t, a) \notin \aleph' \Rightarrow (s \triangleleft t \frown \langle (t, a) \rangle, \aleph' \triangleleft t) \in S$
 \wedge
 $(t > 0 \wedge \neg \exists \varepsilon > 0 \bullet ((t - \varepsilon, t) \times \{a\} \subseteq \aleph'))$
 $\Rightarrow (s \triangleleft t \frown \langle (t, a) \rangle, \aleph' \triangleleft t) \in S$

Axioms 1 and 2 require that an element of \mathcal{M}_{TI} must be a non-empty downward closed set of behaviours. Axiom 3 requires that on every execution, timed events must be either possible or refusable.

A set of behaviours T is finitely variable if for every time t , the set $T \triangleleft t$ is a complete partial order under \preceq . A set of behaviours T is closed if

$$T = \overline{T} = \{(s, \aleph) \mid \forall t \bullet (s, \aleph) \triangleleft t \in T\}$$

Let \mathcal{CL} be the set of finitely variable closed sets of behaviours satisfying axioms 1–3. Then axiom 4 states that

$$4. \quad S = \bigcup \{Q \in \mathcal{CL} \mid S \supseteq Q\}$$

The semantic function \mathcal{F}_{TI}

Define $\mathcal{E}_{TI} = \mathcal{V} \rightarrow \mathcal{M}_{TI}$. The semantic function

$$\mathcal{F}_{TI} : CSP \rightarrow \mathcal{E}_{TI} \rightarrow \mathcal{M}_{TI}$$

is defined by the following set of equations:

$$\begin{aligned} \mathcal{F}_{TI}[\text{Loop}] \rho &\hat{=} \{(s, \aleph) \mid s \in T\Sigma_{\leq}^{\omega} \wedge \aleph \in IRSET\} \\ \mathcal{F}_{TI}[\text{Stop}] \rho &\hat{=} \{(\langle \rangle, \aleph) \mid \aleph \in IRSET\} \\ \mathcal{F}_{TI}[\text{Skip}] \rho &\hat{=} \{(\langle \rangle, \aleph) \mid \checkmark \notin \sigma(\aleph)\} \\ &\quad \cup \\ &\quad \{((t, \checkmark), \aleph) \mid t \geq 0 \wedge \checkmark \notin \sigma(\aleph \uparrow [0, t])\} \\ \mathcal{F}_{TI}[P ; Q] \rho &\hat{=} \{(s, \aleph) \mid \checkmark \notin \sigma(s) \wedge \\ &\quad (s, \aleph \cup ([0, \text{end}(s, \aleph)) \times \{\checkmark\})) \in \mathcal{F}_{TI}[P] \rho \\ &\quad \vee \\ &\quad s = s_P \frown s_Q \wedge \checkmark \notin \sigma(s_P) \wedge \\ &\quad (s_Q, \aleph) - t \in \mathcal{F}_{TI}[Q] \rho \wedge \\ &\quad (s_P \frown ((t, \checkmark), \aleph \triangleleft t \cup ([0, t) \times \{\checkmark\}))) \in \mathcal{F}_{TI}[P] \rho\} \\ \mathcal{F}_{TI}[P \triangleright^{t_0} Q] \rho &\hat{=} \{(s, \aleph) \mid \text{begin}(s) \leq t_0 \wedge (s, \aleph) \in \mathcal{F}_{TI}[P] \rho\} \\ &\quad \cup \\ &\quad \{(s, \aleph) \mid \text{begin}(s) \geq t_0 \wedge (\langle \rangle, \aleph \triangleleft t_0) \in \mathcal{F}_{TI}[P] \rho \\ &\quad \wedge \\ &\quad (s, \aleph) - t_0 \in \mathcal{F}_{TI}[Q] \rho\} \\ \mathcal{F}_{TI}[P \square Q] \rho &\hat{=} \{(\langle \rangle, \aleph) \mid (\langle \rangle, \aleph) \in \mathcal{F}_{TI}[P] \rho \cap \mathcal{F}_{TI}[Q] \rho\} \\ &\quad \cup \\ &\quad \{(s, \aleph) \mid s \neq \langle \rangle \wedge (s, \aleph) \in \mathcal{F}_{TI}[P] \rho \cup \mathcal{F}_{TI}[Q] \rho \\ &\quad \wedge \\ &\quad (\langle \rangle, \aleph \triangleleft \text{begin}(s)) \in \mathcal{F}_{TI}[P] \rho \cap \mathcal{F}_{TI}[Q] \rho\} \\ \mathcal{F}_{TI}[a : A \longrightarrow P_a] \rho &= \{(\langle \rangle, \aleph) \mid A \cap \sigma(\aleph) = \{\}\} \\ &\quad \cup \\ &\quad \{((t, a) \frown (s + t), \aleph) \mid \\ &\quad \quad a \in A \wedge t \geq 0 \wedge A \cap \sigma(\aleph \triangleleft t) = \{\} \\ &\quad \quad \wedge (s, \aleph - t) \in \mathcal{F}_{TI}[P_a] \rho\} \\ \mathcal{F}_{TI}[\prod_{i \in I} P_i] \rho &\hat{=} \bigcup_{i \in I} \mathcal{F}_{TI}[P_i] \rho \\ \mathcal{F}_{TI}[P \parallel_B Q] \rho &\hat{=} \{(s, \aleph) \mid \exists \aleph_P, \aleph_Q \bullet \\ &\quad \aleph \uparrow (A \cup B) = (\aleph_P \uparrow A) \cup (\aleph_Q \uparrow B) \\ &\quad \wedge s = s \uparrow (A \cup B) \\ &\quad \wedge (s \uparrow A, \aleph_P) \in \mathcal{F}_{TI}[P] \rho \\ &\quad \wedge (s \uparrow B, \aleph_Q) \in \mathcal{F}_{TI}[Q] \rho\} \end{aligned}$$

$$\begin{aligned}
\mathcal{F}_{TI}[[P \parallel Q]]\rho &\hat{=} \{(s, \aleph) \mid \exists s_P, s_Q \bullet \forall t \bullet \text{strip}(s \uparrow t) \text{interleaves}(\text{strip}(s_P \uparrow t), \text{strip}(s_Q \uparrow t)) \wedge \\
&\quad (s_P, \aleph) \in \mathcal{F}_{TI}[[P]]\rho \wedge \\
&\quad (s_Q, \aleph) \in \mathcal{F}_{TI}[[Q]]\rho\} \\
\mathcal{F}_{TI}[[P \setminus A]]\rho &\hat{=} \{(s \setminus A, \aleph) \mid (s, \aleph \cup ([0, \infty) \times A) \in \mathcal{F}_{TI}[[P]]\rho\} \\
\mathcal{F}_{TI}[[f(P)]]\rho &\hat{=} \{(f(s), \aleph) \mid (s, f^{-1}(\aleph)) \in \mathcal{F}_{TI}[[P]]\rho\} \\
\mathcal{F}_{TI}[[f^{-1}(P)]]\rho &\hat{=} \{(s, \aleph) \mid (f(s), f(\aleph)) \in \mathcal{F}_{TI}[[P]]\rho\} \\
\mathcal{F}_{TI}[[Y]]\rho &\hat{=} \rho(Y)
\end{aligned}$$

The semantics $\mathcal{F}_{TI}[[\mu Y \circ P]]\rho$ for a recursive term is defined to be the least fixed point of the function

$$\lambda S \bullet \mathcal{F}_{TI}[[P]](\rho[S/Y]) : \mathcal{M}_{TI} \rightarrow \mathcal{M}_{TI}$$

It is established in [MRS95] that this is well-defined.