

Process Algebra for Discrete Event Simulation

P.G. Harrison* **B. Strulo†**

Department of Computing
Imperial College
LONDON SW7 2BZ

Abstract

We present a process algebra or programming language, based on CCS, which may be used to describe discrete event simulations with parallelism. It has extensions to describe the passing of time and probabilistic choice, either discrete, between a countable number of processes, or continuous to choose a random amount of time to wait. It has a clear operational semantics and we give approaches to denotational semantics given in terms of an algebra of equivalences over processes. It raises questions about when two simulations are equivalent and what we mean by non-determinism in the context of the specification of a simulation. It also exemplifies some current approaches to adding time and probability to process algebras.

1 Introduction

Imagine we wish to simulate the behaviour of a complex system with computerised components, such as a telephone network. First, let us look at the implementation of such a complex system. When it is implemented, typically work will start with some type of specification of the system as a whole. Such a specification may never be completely explicit but we certainly require some underlying understanding of the whole system. This will include parts of the system (we will call them modules) which are to be implemented and other parts (we will call them workloads) that represent the demands of the environment on the modules, such as customer calls in our telephone example. From this the implementor will proceed to specifications of the actual modules, such as digital exchanges and their software. Finally executable programs (or pieces of hardware or a mixture) will be constructed for these modules.

Two directions are apparent in the evolution of these implementation processes. Firstly the development of appropriate high level languages, both for execution and specification, has been essential to manage this complexity. Secondly much work has gone into understanding the semantics of such languages formally with the intention of reasoning mathematically about the implementations.

Now consider the process of making some model of the performance of such a system. We will refer to such a model as a simulation. By this we mean an explicit algorithmic

*Supported by the CEC as part of the ESPRIT-BRA QMIPS project 7269

†Supported by a grant from the SERC

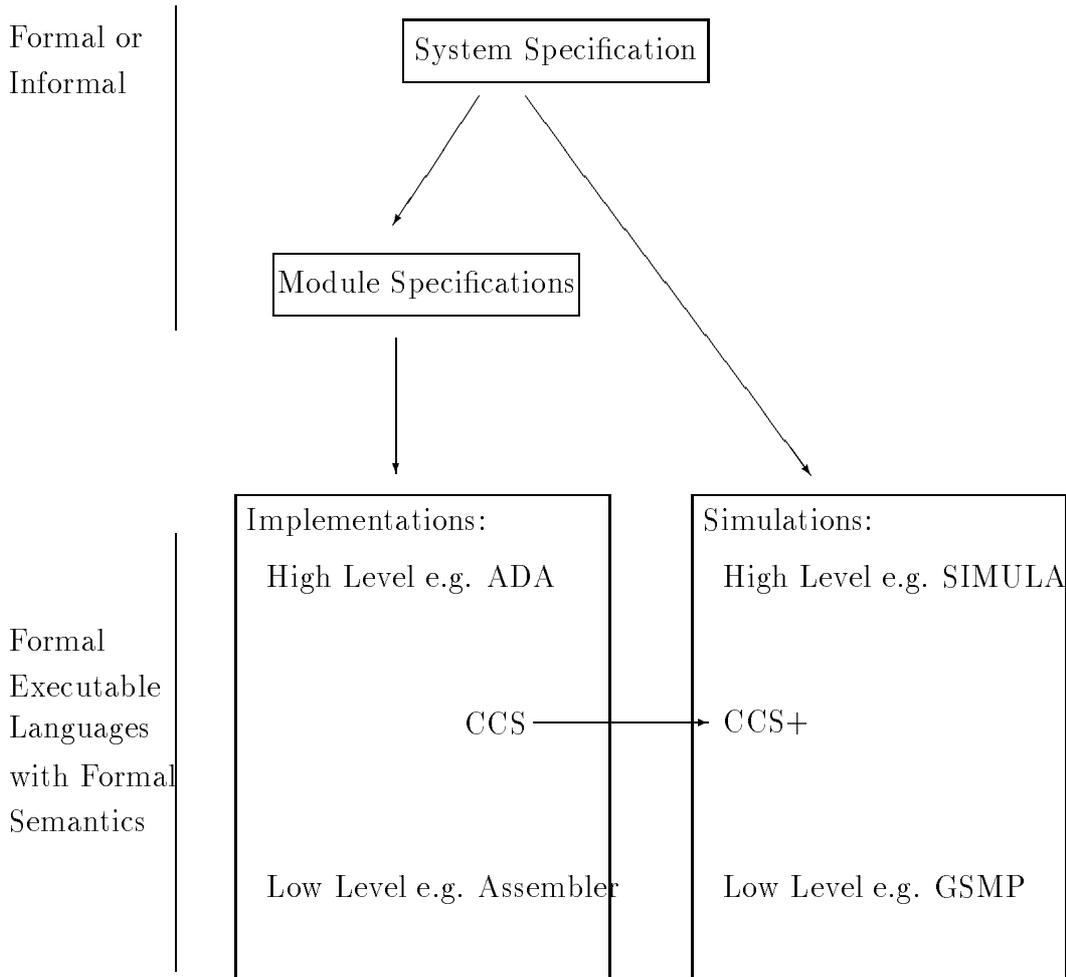


Figure 1: Relationship between implementation and simulation

model which may be used either for analysis or for execution to determine performance experimentally. It seems that even if our intention is abstract analysis we must have in mind some underlying simulation in this sense. So we look for formalisms that are both high level and have a formal semantics. In figure 1 we show schematically the kind of comparison we are considering. Here CCS+ represents the language we will describe.

We require a formal semantics because such a semantics makes the behaviour of the language explicit and rigorous. Existing languages with the facilities to describe simulations often have complex behaviour described in terms of things like priority queues and event lists. Inevitable ambiguities in such informal descriptions can lead to unexpected behaviour or worse to simulations which are not true to their specification. However, in addition, we hope to derive from a formal semantics, tools to analyze simulations. If we can assign a meaning to a simulation then we can say that two simulations are equal if their meanings are equal. Then we can perform mechanical simplification and abstraction or at least prove the validity of our manual simplifications and abstractions. This parallels the use of formal semantics in conventional programming languages, as discussed in say

[7].

We also want our language to be high level. We could certainly consider Generalised Semi-Markov Processes (GSMPs) as a semantics for simulation. But a GSMP description is very far from any current notion of system specification. Given a system of any complexity it is very hard to see that a GSMP description is a correct model of that system. We need to find a formalism nearer to specification, more high level and thus more natural for describing real complex concurrent systems. But we must still have a rigorous semantics and thus make our formalism a bridge between informal descriptions of a system and mathematical models of that system.

So by making the language high level we aim to reduce the chance of the original simulation being unfaithful to the specification and by making its semantics formal we aim to reduce the chance of any simplifications being unfaithful to the original.

What is the difference between an implementation and a simulation? Firstly, in a simulation, we need to be able to model our workloads while for an implementation they are supplied by the outside world. Further we need to model them in some detail so as to represent the demands we expect them to impose on our implementation. Similarly when we abstract from detail in our simulation by regarding sub-modules as black boxes we need to describe their behaviour in more detail than may interest an implementor since we want to represent their performance in addition to their abstract behaviour. Thus in a simulation we need to be able to describe both probabilistic behaviour with explicit distributions and also explicit behaviour in time.

Secondly, in a simulation, we are often much more concerned directly with parallelism. Single modules in an implementation are generally sequential and so the implementor may be able to use sequential languages both for implementation and for module specification. Essentially he may be able to move consideration of parallelism up to the level of the system specification. In simulation we are modelling the behaviour of the whole system. Our simulation has to be faithful to the whole system specification and so we will inevitably need to be able to describe parallelism explicitly if this is how the system is to operate.

Work on formal semantics has very clearly elucidated the behaviour of sequential systems. Formal semantics for parallelism have proved more controversial and it is only recently that some consensus about the main approaches has been reached. Surveying the literature has suggested process algebras as one of the leading contenders [8, 12]. Many approaches to extending process algebras with actual time values [3, 1, 2, 6, 13, 16, 10] or with time and probabilities [5] have also been discussed.

We use a selection from these approaches with some new techniques to obtain a process algebra which may be used to describe discrete event simulations. So we take a relatively high level language that has a formal semantics of parallelism (CCS from [8]) and add constructs to represent a real-number time and various probabilistic features. We then extend its formal semantics to cover these additions. In this way we can derive semantically sound axioms relating sub-expressions of a program. Such axioms can form the basis of mechanisable program transformations which, for example, may allow simpler or more efficient or more mathematically tractable simulations to be derived from another with a clearer specification. Again this compares to program transformation in declarative languages [4].

We present a process algebra or programming language, based on CCS, which may be used to describe concurrent processes. It has extensions to describe the passing of time and probabilistic choice, either discrete, between a countable number of processes,

or continuous to choose a random amount of time to wait. It thus has sufficient facilities to specify discrete event simulations. It has a clear operational semantics and we discuss equivalences over these operational semantics that imply equations on processes that can be used to work formally with processes.

Because it allows both non-determinism and probabilistic choice, it can be used to show the difference between these concepts and their relevance to simulation. It also exemplifies some current approaches to adding time and probability to process algebras.

2 Language

First we give the language syntax and informally present the intended meaning.

The language is based on the timed CCS of [15] with some features of the temporal CCS of [9] though their calculus does not generalise directly to real number time and they have more power in that they have a weak $+$ operator that we do not. It ends up very close to [2] except we have only two options for the upper limit to the time of execution of an action. Our impatient prefixing is equivalent to Chen's $\alpha|_0^0$ while our idling prefixing is equivalent to his $\alpha(s)|_0^\infty$.

Our language is defined so that CCS is a sub-calculus in the sense that the operational semantics of the pure CCS terms in our language is the same as in CCS. Our rules for evolution (as in the mentioned calculi) does not behave in exactly the way expected for the time required to compute. Thus for example it is easy to write programs which perform an infinite amount of computation in finite time and thus force time to stop for the whole system. Other processes simply deadlock in time and also prevent time passing in the whole system. We justify this by interpreting time evolution as representing simulated time rather than the time in which real calculation takes place.

As in CCS, communication and synchronisation are represented via an alphabet of atomic actions or labels. Each action has a conjugate action represented by an overline and the conjugate to the conjugate is the original action. There is also a distinguished self-conjugate invisible (or silent) action which represents internal activity of a process invisible to the outside. Thus we have an alphabet Act including names, co-names and the invisible action τ .

We use two sorts and two sets of variables, PVar of sort *Process expression*, and TVar of sort *Time*. We use an algebra of time expressions giving a set of expressions TExp . This has the operator $+$ indicating addition and \div indicating non-negative subtraction *i.e.*

$$x \div y = \begin{cases} x - y & x \geq y \\ 0 & \text{otherwise} \end{cases}$$

The other symbols are 0 and an arbitrary number of other constant symbols. We will assume that these constants are non-negative real numbers (we will write \mathbb{R}^+). Equality for this sort is defined in the usual way. We will use time expressions and their values interchangeably.

Our language is then standard (basic) CCS with this new sort of Time. We then distinguish four forms of prefixing, one which describes a prefix of a timed delay, represented (t). This uses the value of the time expression t which may have free variables. We use two separate forms of action prefixing. Simple prefixing $\alpha.P$ is distinguished from $\alpha[s \leftarrow t]$ where the action prefix binds the time variable s to the time of occurrence of the action

α . Then the final form binds time variables to the sampling of a random variable. This is represented $\mathcal{R}[s \leftarrow f]$.

Finally we add a probabilistic choice, represented by $\sum_{i \in I} [q_i] P_i$.

So we define the process expressions as

Definition 1 *The set of process expressions, \mathbf{PExp} , is the least set including*

$$\begin{array}{cccccccc} X & \text{nil} & (t).P & \alpha.P & \alpha[s \leftarrow t].P & \mathcal{R}[s \leftarrow f].P & & \\ & & & & & & \sum_{i \in I} [q_i] P_i & P + Q & P \mid Q & P[S] & P \setminus A & \text{rec}(X = P) \end{array}$$

where $X \in \mathbf{PVar}$, $P, Q \in \mathbf{PExp}$, $s \in \mathbf{TVar}$, $t \in \mathbf{TExp}$, $\alpha \in \mathbf{Act}$, I is a finite indexing set with the $P_i \in \mathbf{PExp}$, S is a relabelling function $S : \mathbf{Act} \rightarrow \mathbf{Act}$, $A \subset \mathbf{Act}$ is a set of actions with $\tau \notin A$, q_i are probabilities in $[0, 1]$ with $\sum q_i = 1$, and f is a distribution function taking values in \mathbb{R}^+ .

Here nil is a process constant representing a deadlocked or terminated process which can perform no action and must simply idle from now on. $(t).P$ represents a process which must delay for t time units and then becomes process P .

$\alpha.P$ and $\alpha[s \leftarrow t].P$ both represent processes which can immediately perform an α action. If they do they become P but in the latter case the free variable s is replaced by the value t and so becomes $P\{t/s\}$. The essential difference is that the former prefixing is “impatient” and cannot delay. The latter expression can delay for any time t' but records this delay and becomes $\alpha[s \leftarrow t + t'].P$. This latter prefixing binds the variable s in P . Of course if s is not free in P then we need not concern ourselves with s or t , and for convenience we may write $\alpha[\].P$ in this case (the brackets distinguishing it from impatient prefixing).

We view this delaying behaviour as analogous to that of the value-passing part of CCS. We imagine a global clock, broadcasting the passage of time, to which all processes waiting for a synchronisation listen. Then the performance of the action causes the time value variable to become instantiated to the passing of time it has heard from that clock. Then the process may act on that value later, perhaps by delaying less later if it has already been delayed.

$P + Q$ represents a process which can evolve until P or Q are ready to participate in a labelled transition. Then it will choose non-deterministically between them. Once we have shown that no ambiguity arises we will use $\sum_{i \in I} P_i$ to represent finite sums.

$\sum_{i \in I} [q_i] P_i$ represents a process which will immediately and randomly become one of the P_i on the basis of the q_i probabilities. This is a straight-forward finite probabilistic choice.

$\mathcal{R}[s \leftarrow f].P$ represents a process which can immediately sample a random variable with the specified distribution and become $P\{t/s\}$ for some t . Although this is a form of uncountable probabilistic branching the resulting processes are highly restricted. They are identical except in the values of certain time expressions and hence certain delays.

The other constructs have the same meaning as basic CCS. We define free variables of either sort in the standard way.

3 Semantics

3.1 Transition system

The operational semantics of these processes are described in terms of three sorts of transitions between processes. Thus we envisage every state in the system as represented by a process and the transitions representing state changes. We define probabilistic transitions which we assume to be resolved first. The idea is that any probabilistic part may be assumed under any circumstances to have already been resolved before any other actions are considered. We will show that our transition system has this property (proposition 1).

Then we have the standard CCS type of *labelled* transitions which are labelled from Act and happen non-deterministically following the usual behaviour of a CCS-type language. Transitions labelled with τ must happen immediately *i.e.* before any evolution can take place. We show this property as proposition 4. Note that our impatient actions must also happen before any evolution. We will call impatient actions and τ actions *immediate*.

When no more immediate transitions are possible the system becomes *stable*. Then, all parts of the process synchronously undergo a timed *evolution* transition via their (t) actions. This represents the passing of time and continues until further probabilistic or immediate transitions are possible. To be sure of this we need to show that the process cannot evolve past the potential of performing a sampling or immediate transition. This is shown in proposition 3.

This assumption that immediate transitions take priority over evolution is sometimes called the Maximal Progress assumption and is discussed in [15]. In fact our language is more expressive than his since we can also write impatient actions that are not τ *e.g.* in processes such as $Clock = tick.(1).Clock$ which ticks once a time unit. Here each tick cannot be made to idle and so the clock cannot be made to run slow.

While this is similar to some ideas of timed calculi we envisage it rather with two separate notions of time. We might imagine that immediate transitions must take some real time to process. But this is not the simulated time that passes during evolution transitions. Thus we imagine the system performing calculations to decide its next step and resolving any exposed probabilistic choices. It calculates until eventually stability is reached. Then all processes co-operatively allow simulated time to pass until at least one process is ready to perform more calculations.

The possible transitions are described via inference systems in the usual way. We say the transitions are the least relation between processes that satisfies the given inference rules.

3.2 Probabilistic transitions

We start with the probabilistic transitions in Figure 2. We name the rules for use in proofs. Each transition has an associated probability measure, shown above it. It also has an associated index, shown as a subscript. We give a syntax for indices as

Definition 2 *The set of index atoms is defined by*

$$\text{Atom} = \{(i) : i \in \mathbb{N}\} \cup \{[t] : t \in \mathbb{R}^+\}$$

The set of basic indices, Index_0 , is the least set including

$$a.\langle \rangle \quad \langle \rangle \quad l \mid l' \quad l + l'$$

where $l, l' \in \text{Index}_0$, $a \in \text{Atom}$.

Note the way there is a single a value carried throughout the inference and on top of it an index “skeleton” full of $\langle \rangle$. Note also the strong way we wish probabilistic transitions to happen first - the Delay rule. Only a labelled prefix can hide a probabilistic transition. However this is just enough to prevent the inference reaching down to a recursion and an infinite chain of inference.

This relation captures our fundamental idea of how we wish probabilistic transitions to operate but it is not clear how it represents probabilities. It is even non-deterministic since a sum or parallel combination with two probabilistic sub-terms is given two possible transitions. To clarify the position we combine these multiple transitions into our actual relation with another inference system, figure 3. Here we use a syntax for indices as

Definition 3 *The set of general indices, Index , is the least set including*

$$\langle \rangle \quad a.l \quad l \mid l' \quad l + l'$$

where $l, l' \in \text{Index}$, $a \in \text{Atom}$, $i \in \mathbb{N}$, $t \in \mathbb{R}^+$.

Note that $\text{Index}_0 \subset \text{Index}$.

Here we use the If function to generate to extend the index as we generate the final transition.

Definition 4 *The label extension function $\text{If} : \text{Index}_0 \times \text{Index} \rightarrow \text{Index}$ is defined by*

$$\begin{aligned} \text{If}(\langle \rangle, m) &= m \\ \text{If}(a.\langle \rangle, m) &= a.m \\ \text{If}(l + l', m + m') &= \text{If}(l, m) + \text{If}(l', m') \\ \text{If}(l \mid l', m \mid m') &= \text{If}(l, m) \mid \text{If}(l', m') \end{aligned}$$

Note that this definition covers sufficient cases since if $P \xrightarrow[\circ]{p}_{l+l'} Q$ then Q is of the form $Q' + Q''$ and similarly for \mid .

By our guardedness restriction we know that for every process there is some maximum length of inference both for $\xrightarrow[\circ]$ and for \leftrightarrow . We will call these numbers $\text{maxp}_0(P)$ and $\text{maxp}(P)$. We will also regularly perform induction by the lexicographic ordering on $(\text{maxp}(P), \text{maxp}_0(P))$. We will call this “induction on maxp ”.

This is the usual sort of presentation for operational semantics of process algebras. We must demonstrate the sense in which it provides a well defined probabilistic semantics. In other words we wish to show we have defined some sort of probability distribution over the transitions from P .

Our first problem is that this transition relation gives derivations which are prefixes of each other. For example we have both

$$\begin{aligned} \mathcal{R}[s \leftarrow f]. \mathcal{R}[s' \leftarrow f']. P \xrightarrow[\circ]{f(t)}_{[t]} \mathcal{R}[s' \leftarrow f']. P\{t/s\} \\ \mathcal{R}[s \leftarrow f]. \mathcal{R}[s' \leftarrow f']. P \xrightarrow[\circ]{f(t), f(t')}_{[t], [t']} P\{t/s\}\{t'/s'\} \end{aligned}$$

We cannot easily interpret these as probabilities since here for example the probability of transiting to the set $\{\mathcal{R}[s' \leftarrow f']. P\{t/s\} : t \in \mathbb{R}^+\} \cup \{P\{t/s\}\{t'/s'\} : t, t' \in \mathbb{R}^+\}$ could

$\frac{\sum_{i \in I} [q_i] P_i \xrightarrow{q_i}_{\circ} [i, \langle \rangle] P_i}{}$	Resolve
$\frac{\mathcal{R}[s \leftarrow f]. P \xrightarrow{f(t)}_{\circ} (t, \langle \rangle) P \{t/s\}}{}$	Sample
$\frac{P \xrightarrow{q}_{\circ} P'}{(t). P \xrightarrow{q}_{\circ} (t). P'}$	Delay
$\frac{P \xrightarrow{q}_{\circ} P'}{P + Q \xrightarrow{q}_{\circ} (l+\langle \rangle, a) P' + Q}$	Sum1
$\frac{Q \xrightarrow{q}_{\circ} Q'}{P + Q \xrightarrow{q}_{\circ} \langle \rangle + l P + Q'}$	Sum2
$\frac{P \xrightarrow{q}_{\circ} P'}{P \mid Q \xrightarrow{q}_{\circ} l \mid \langle \rangle P' \mid Q}$	Parallel1
$\frac{Q \xrightarrow{q}_{\circ} Q'}{P \mid Q \xrightarrow{q}_{\circ} \langle \rangle \mid l P \mid Q'}$	Parallel2
$\frac{P \xrightarrow{q}_{\circ} P'}{P[S] \xrightarrow{q}_{\circ} P'[S]}$	Relabel
$\frac{P \xrightarrow{q}_{\circ} P'}{P \setminus A \xrightarrow{q}_{\circ} P' \setminus A}$	Restrict
$\frac{P \{ \text{rec}(X = P) / X \} \xrightarrow{q}_{\circ} P'}{\text{rec}(X = P) \xrightarrow{q}_{\circ} P'}$	Fix

Figure 2: Basic probabilistic transitions

$\frac{P \xrightarrow{p}_{\circ} P'}{P \xrightarrow{p}_{\circ} P'}$	Any-Samples
$\frac{P \xrightarrow{p}_{\circ} P' \quad P' \xrightarrow{q}_{\circ} P''}{P \xrightarrow{p, q}_{\circ} P''}$	Multi-Samples

Figure 3: Actual probabilistic transitions

naïvely be evaluated to 2. The simplest way round this is to only look at descendants which cannot \hookrightarrow .

Thus we define

Definition 5 For a given process, P ,

$$\text{Index}_P = \{m \in \text{Index} : \exists P', q \quad P \xrightarrow[q]{m} P' \quad \neg P' \hookrightarrow\}$$

$p_P : \text{Index}_P \rightarrow [0, 1]$ is

$$p_P(m) = q \quad \text{iff} \quad \exists P' \quad P \xrightarrow[q]{m} P'$$

and $q_P : \text{Index}_P \rightarrow \text{Proc}$ is

$$q_P(m) = P' \quad \text{iff} \quad \exists q \quad P \xrightarrow[q]{m} P'$$

We need to show these definitions are consistent so we prove

Lemma 1 If $P \xrightarrow[q]{\circ} Q$ and $P \xrightarrow[q']{\circ} Q'$ then $q = q'$ and $Q \equiv Q'$.

This shows we have some sort of probability “distribution” function. But over what can we integrate this distribution? We start by treating it as a distribution over our index space. The set of indices associated with all the transitions from the given process form the sample space, *i.e.* it is this index which distinguishes the separate transitions.

Now we prove

Lemma 2 p_P is measurable.

Indices are trees of real and natural numbers and so it is straightforward to define measurable sets. We define:

Definition 6 Borel is the σ -field of measurable sets from Index.

Then it is straightforward to define integration and the probability associated with any such set of indices:

Definition 7 The index probability measure, $\mu^* : \text{Proc} \times \text{Borel} \rightarrow [0, 1]$ is

$$\mu^*(P, I) = \int_I p_P(l) dl$$

We then wish to move this structure over to Proc. But what is the appropriate σ -field of measurable sets on Proc. Using the usual definition of $q_P^{-1}(S) = \{l \in \text{Index}_P : q_P(l) \in S\}$, we define:

Definition 8 The measurable sets w.r.t. P are

$$\overline{B}_P = q_P^{-1}(\text{Borel}) = \{S : \exists I \in \text{Borel} \quad S = q_P^{-1}(I)\}$$

The σ -field of measurable sets is \overline{B} where

$$\overline{B} = \bigcap_{P \in \text{Proc}} \overline{B}_P$$

$\frac{}{\alpha.P \xrightarrow{\alpha} P}$	Prefix1
$\frac{}{\alpha[s \leftarrow t].P \xrightarrow{\alpha} P\{t/s\}}$	Prefix2
$\frac{P \xrightarrow{\alpha} P'}{(0).P \xrightarrow{\alpha} P'}$	No-Delay
$\frac{P \xrightarrow{\alpha} P' \quad \neg P \hookrightarrow \neg Q \hookrightarrow \quad Q \xrightarrow{\alpha} Q' \quad \neg P \hookrightarrow \neg Q \hookrightarrow}{P + Q \xrightarrow{\alpha} P' \quad P + Q \xrightarrow{\alpha} Q'}$	Sum
$\frac{P \xrightarrow{\alpha} P' \quad \neg Q \xrightarrow{q} \iota \quad Q \xrightarrow{\alpha} Q' \quad \neg P \xrightarrow{q} \iota}{P Q \xrightarrow{\alpha} P' Q \quad P Q \xrightarrow{\alpha} P Q'}$	Parallel
$\frac{P \xrightarrow{\alpha} P' \quad Q \xrightarrow{\bar{\alpha}} Q' \quad \neg P Q \xrightarrow{q} \iota}{P Q \xrightarrow{\tau} P' Q'}$	Communicate
$\frac{P \xrightarrow{\alpha} P'}{P[S] \xrightarrow{S(\alpha)} P'[S]}$	Relabel
$\frac{P \xrightarrow{\alpha} P'}{P \setminus A \xrightarrow{\alpha} P' \setminus A} \quad \alpha \notin A \text{ and } \bar{\alpha} \notin A$	Restrict
$\frac{P\{\text{rec}(X = P)/X\} \xrightarrow{\alpha} P'}{\text{rec}(X = P) \xrightarrow{\alpha} P'}$	Fix

Figure 4: Labelled transitions

Now we may define our probability distribution over this σ -field in the usual way as:

Definition 9 *The process probability measure, $\mu : \text{Proc} \times \overline{B} \rightarrow [0, 1]$ is*

$$\mu(P, S) = \mu^*(P, q_P^{-1}(S))$$

Thus $\mu(P, S)$ is the total probability that P may become a member of S by a probabilistic sampling transition. Finally, to prove that $\mu(P, \cdot)$ is a probability we need

Lemma 3 $\mu(P, \text{Proc}) = 1$ for all P with $P \hookrightarrow$.

Thus, for each process P , these probabilistic transitions induce a measure $\mu(P, \cdot)$ over the σ -field of sets of processes \overline{B} .

3.3 Labelled transitions

These are defined in Figure 4. These are very similar to the standard CCS rules. There are preconditions to force all probabilistic choices to be resolved first. The Prefix2 rule shows the way in which time variables are bound by the execution of an action. The No-Delay rule allows zero delays to be disregarded when any transition is to be inferred. We will thus later be able to deduce that $(0).P = P$.

3.4 Evolution transitions

Eventually no more immediate or probabilistic transitions will be possible. The only thing the process can then do is to evolve in time.

As in [15] a central aim of this design is to give the Maximal Progress property that τ actions always happen before evolution. The problem here is that it is difficult to prevent a parallel construct from evolving past a potential τ operation. Thus we would like to have a rule

$$\frac{P \xrightarrow{t} P' \quad Q \xrightarrow{t} Q'}{P \mid Q \xrightarrow{t} P' \mid Q'}$$

But consider $P \equiv \alpha$, $Q \equiv \bar{\alpha}$. Both of these can delay arbitrarily but in parallel we wish them to be forced to act immediately without delay. In [15] this is done by “Timed Sorts”. We use a slightly different approach.

First we define a simple evolution relation \xrightarrow{t}_\circ . Note that we will not distinguish between time expressions and time values. The definitions are in Figure 5.

Note that these rules do not allow the inference of a zero time evolution (trivial induction). Also we have no rules to allow $\alpha.P$ to evolve.

These rules capture our essential ideas about evolution. However they only allow evolutions between the simplest changes of syntactic form of processes. Thus for example we do have $(1).(1).P \xrightarrow{1}_\circ (1).P \xrightarrow{1}_\circ P$ but not $(1).(1).P \xrightarrow{2}_\circ P$. This is necessary to ensure that parallel constructs are checked for immediate transitions at every stage at which they are possible. But we do wish to combine allowable transitions so we add a new level of transition. Our actual evolution transitions are thus defined in Figure 6. Thus we now have $(1).(1).P \xrightarrow{1}_\circ (1).P$ and so by Any-Delay we have $(1).(1).P \xrightarrow{1} (1).P$. But we also have $(1).P \xrightarrow{1}_\circ P$ and so by Add-Delay we can infer $(1).(1).P \xrightarrow{2} P$ as required.

3.5 Properties

We have stated informally that if our processes can sample they must do so first. We prove:

Proposition 1 *If $P \xrightarrow{a}_l$ then $\neg P \xrightarrow{a} \neg P \xrightarrow{t}$.*

Thus if probabilistic transitions are possible they are the only transitions possible and happen before any others.

We list some other properties of our system and compare them to those noted by other authors for their temporal process algebra. In particular [11] gives a concise collection of important properties of various systems and we will adopt their terminology for these.

3.5.1 Time determinism

Most systems of timed process algebras have the property that time evolution introduces no non-determinism - this arises only from labelled actions. Our calculus has this property:

Proposition 2 *If $P \xrightarrow{t} P'$ and $P \xrightarrow{t} P''$ then $P' \equiv P''$.*

$\frac{\neg P \hookrightarrow}{(t+u).P \rightsquigarrow_{\circ}^t (u).P} \quad t > 0$	Reduce-Delay
$\frac{P \rightsquigarrow_{\circ}^t P'}{(0).P \rightsquigarrow_{\circ}^t P'}$	No-Delay
$\frac{\neg P \hookrightarrow}{(t).P \rightsquigarrow_{\circ}^t P} \quad t > 0$	End-Delay
$\frac{}{\text{nil} \rightsquigarrow_{\circ}^t \text{nil}} \quad t > 0$	Nil
$\frac{}{\alpha[s \leftarrow t].P \rightsquigarrow_{\circ}^{t'} \alpha[s \leftarrow t+t'].P} \quad \alpha \neq \tau \text{ and } t' > 0$	Idle
$\frac{P \rightsquigarrow_{\circ}^t P' \quad Q \rightsquigarrow_{\circ}^t Q'}{P+Q \rightsquigarrow_{\circ}^t P'+Q'}$	Sum
$\frac{\neg P \mid Q \xrightarrow{\tau} \quad P \rightsquigarrow_{\circ}^t P' \quad Q \rightsquigarrow_{\circ}^t Q'}{P \mid Q \rightsquigarrow_{\circ}^t P' \mid Q'}$	Parallel
$\frac{P \rightsquigarrow_{\circ}^t P'}{P[S] \rightsquigarrow_{\circ}^t P'[S]}$	Relabel
$\frac{P \rightsquigarrow_{\circ}^t P'}{P \setminus A \rightsquigarrow_{\circ}^t P' \setminus A}$	Restrict
$\frac{P\{\text{rec}(X=P)/X\} \rightsquigarrow_{\circ}^t P'}{\text{rec}(X=P) \rightsquigarrow_{\circ}^t P'}$	Fix

Figure 5: Simple Evolution transitions

$\frac{P \rightsquigarrow_{\circ}^t P'}{P \rightsquigarrow^t P'} \quad \text{Any-Delay}$
$\frac{P \rightsquigarrow^t P' \quad P' \rightsquigarrow_{\circ}^{t'} P''}{P \rightsquigarrow^{t+t'} P''} \quad \text{Add-Delay}$

Figure 6: Extended Evolution transitions

3.5.2 Time additivity

Most other systems also require that if a process can idle for some period it can also idle for any smaller period ([16] - “time continuity”). Our calculus does have this property but our requirement that τ is immediate makes this rather more difficult than it might seem. We must show that a process that can evolve $P \xrightarrow{t+t'} P'$ does not have the lurking capability of a $P \xrightarrow{t} P'' \xrightarrow{\tau}$ which would imply $\neg P'' \xrightarrow{t'} P'$.

We prove two intermediate lemmas.

Lemma 4 *If $P \xrightarrow{t}_\circ P'$ and $P \xrightarrow{t'}_\circ P''$ with $t' > t$ and $P' \xrightarrow{\alpha}$ then $P \xrightarrow{\alpha}$.*

Proof

Induction on length of inference of $P \xrightarrow{t'}_\circ$.

- Reduce-Delay: $P \equiv (t' + u).Q$ and so $P' \equiv (t' - t + u).Q$ and so $\neg P' \xrightarrow{\alpha}$
- Nil & End-Delay similar
- No-Delay: $P \equiv (0).Q$
So $Q \xrightarrow{t}_\circ Q'$ and $Q \xrightarrow{t'}_\circ$ and $Q' \xrightarrow{\alpha}$
So $Q \xrightarrow{\alpha}$ by ind. hyp. and so $P \xrightarrow{\alpha}$ as required.
- Idle: $P \equiv \alpha[s \leftarrow t].Q$ and so $P \equiv P'$
- Sum: $P \equiv P_1 + P_2$
So $P_1 \xrightarrow{t'}_\circ$ and $P_1 \xrightarrow{t}_\circ P'_i$ and similarly for P_2 .
Also w.l.o.g. $P'_i \xrightarrow{\alpha}$. But since $P_1 \xrightarrow{t'}_\circ$ is a shorter inference $P_1 \xrightarrow{\alpha}$ and hence $P \xrightarrow{\alpha}$ as required.
- Parallel: $P \equiv R \mid Q$
So $R \xrightarrow{t'}$ and $Q \xrightarrow{t'}$ and $R \xrightarrow{t} R'$ and $Q \xrightarrow{t} Q'$
Also $R' \mid Q' \xrightarrow{\alpha}$ so either
 - $R' \xrightarrow{\alpha}$ and so $R \xrightarrow{\alpha}$ by ind. hyp. and so $P \xrightarrow{\alpha}$ as required.
 - $Q' \xrightarrow{\alpha}$ similarly.
 - $R' \xrightarrow{\beta}$ and $Q' \xrightarrow{\bar{\beta}}$. Then by ind. hyp $R \mid Q \xrightarrow{\tau}$ which contradicts $P \xrightarrow{t'}_\circ$.
- Other inferences trivial inductive steps.

□

This lemma shows that although a process may be able to obtain new capabilities through evolution it can do so only at the end of its simple evolution period. Note however that this is not true for the extended evolution relation since *e.g.* $(1).\alpha[] . P \xrightarrow{2} \alpha[] . P$.

Lemma 5 *If $P \xrightarrow{t+t'}_\circ P''$ then $\exists P'$ with $P \xrightarrow{t}_\circ P'$ and $P' \xrightarrow{t'}_\circ P''$.*

Proof

Induction on length of inference of $P \xrightarrow{t+t'}_\circ P'$. The Parallel step requires the use of lemma 4 to show that $\neg P' \xrightarrow{\tau}$. □

Proposition 3 *If $P \xrightarrow{t+t'} P''$ then $\exists P'$ with $P \xrightarrow{t} P'$ and $P' \xrightarrow{t'} P''$.*

Proof

Induction on length of inference of $P \xrightarrow{t+t'} P''$.

- Any-Delay: $P \xrightarrow{t+t'} P''$ and so $\exists P'$ with $P \xrightarrow{t}_\circ P'$ and $P' \xrightarrow{t'}_\circ P''$. Thus $\exists P'$ with $P \xrightarrow{t} P'$ and $P' \xrightarrow{t'} P''$ as required.
- Add-Delay: We have $P \xrightarrow{t+t'} P''$ because $P \xrightarrow{u} Q$ and $Q \xrightarrow{u'} P''$ with $u + u' = t + t'$. If $t = u$ then the result is trivial. Otherwise, if $t < u$ then we have the result by considering $P \xrightarrow{t+(u-t)} Q$ for which the result is true by the inductive hypothesis. Alternatively, if $u < t$ then we may consider $Q \xrightarrow{(t-u)+(u'-(t-u))} P''$ for which we know the result to be true.

□

Thus evolution does behave as expected for a real number clock *i.e.* it reaches every intermediate time. This is also needed to ensure our naïve interpretation of the maximal progress assumption. We now know that an evolving process may evolve for a continuous interval either forever or to some maximum t . Once this t is reached it becomes a new process with either immediate or probabilistic transitions.

3.5.3 Deadlock-freeness

In our calculus we do have states (such as $(\alpha.\text{nil}) \setminus \alpha$) which can do nothing - not even an evolution transition - since their impatient action prevents them from evolving. Thus our calculus does not have the property of “deadlock-freeness” from [11]. Such time-stop processes are intuitively difficult to understand. However time evolution in our calculus is intended to represent simulated time passing and so a process which prevents this passing is simply part of a simulation which fails to reach its time-advance step and so is not a well formed simulation but is a understandable program. We will define a symbol to represent such programs

Definition 10 *The time stop program is*

$$\mathbf{0} = (\alpha.\text{nil}) \setminus \alpha$$

Alternatively we could have simply added $\mathbf{0}$ to our language with no transition rules.

3.5.4 Action urgency

As already discussed, we have both impatient and patient actions (the notion of “urgency” from [11]). However our calculus is unusual in that it also has the maximal progress (“minimal delay” or “tau-urgency”) property.

Proposition 4 *If $P \xrightarrow{\tau}$ then $\neg P \xrightarrow{t}$.*

This, together with proposition 1, means that expressions can be divided into three disjoint sets. Each state is either *sampling* or *immediate* or *evolving*.

3.5.5 Persistency

We do have the important and interesting “persistency” property that the passage of time cannot (alone) remove the capability of performing actions.

Proposition 5 *If $P \xrightarrow{t} Q$ and $P \xrightarrow{\alpha}$ then $Q \xrightarrow{\alpha}$.*

Thus time evolution does not remove the capability to perform actions. It can, however, allow new capabilities as in

$$a[].P + (1).b[].Q \xrightarrow{1} a[].P + b[].Q$$

Capabilities disappear when some action (e.g. a τ action) makes the process change state. This is the mechanism for time-outs as in

$$a[].P + (1).\tau.Q \xrightarrow{1} a[].P + \tau.Q \xrightarrow{\tau} Q$$

3.5.6 Finite-variability

In our calculus we can easily define processes that perform infinite computation in finite time. Thus we do not have the “finite-variability” (“non-Zenoness, well-timedness”) property and also not the “bounded control” property. Questions of whether this is realistic or satisfactory are less relevant in this simulation context.

3.5.7 Free variables

Finally we have an obligation to show that free variables are not going to cause a problem for our interpretation of processes as programs.

Lemma 6 *If P is a process and $P \xrightarrow{\alpha} P'$ then P' is a process.*

If P is a process and $P \xrightarrow{a}_1 P'$ then P' is a process.

If P is a process and $P \xrightarrow{t} P'$ then P' is a process.

3.6 The meaning of Processes

We have given processes an operational semantics as transition systems. Thus they are an (uncountable) set of states, each state with one of three kinds of transitions from them. Either they are *immediate* states with a countable set of immediate transitions to other states, or they are *sampling* states which sample probabilistic distributions to choose their next state, or they are *evolving* states which evolve deterministically on the basis of a set of clocks (exposed (t) expressions) to some new state.

To relate this to discrete event simulation we consider the immediate transitions as representing internal computation and synchronization of the simulation. So a process (simulation) makes probabilistic choices and internal computations until all its component agents become stable (*i.e.* unable to act immediately). Then all such agents bid a time-advance real number. The smallest of these is chosen and simulated time is advanced by that number. Then further computations take place.

4 Strong bisimulation

We now wish to define an equivalence over these process expressions. For the immediate and evolution transitions the standard notion of bisimulation seems appropriate while for the probabilistic transitions the notion is similar but slightly more complex.

We wish to make processes equivalent only if they define equal measures over their descendants. But we do not wish to distinguish processes if their measures differ only on sets which separate equivalent descendants. Thus we do not wish to distinguish $\sum_{i \in I} [q_i] P_i$ from $\sum_{i \in I} [q_i] Q_i$ if we are considering P_i and Q_i equivalent. Thus we must disallow comparison of the measure for sets including P_i and not Q_i for example. So we define the sub σ -field over which we require the measures to be equal.

Definition 11 For an equivalence relation $\mathcal{R} \subset \text{Proc} \times \text{Proc}$ we define $\overline{B}(\mathcal{R}) \subset \overline{B}$ as

$$\{S \in \overline{B} : S \text{ is a union of } \mathcal{R}\text{-equivalence classes} \}$$

In other words $\overline{B}(\mathcal{R})$ are all the measurable sets which don't separate \mathcal{R} -equivalent processes.

Then we may define our fundamental equivalence:

Definition 12 An equivalence relation $\mathcal{R} \subset \text{Proc} \times \text{Proc}$ is a strong bisimulation iff $P \mathcal{R} Q$ implies for all $\alpha \in \text{Act}$, $t \in \mathbb{R}^+$, $S \in \overline{B}(\mathcal{R})$:

1. if $P \xrightarrow{\alpha} P'$, then $\exists Q'$ such that $Q \xrightarrow{\alpha} Q'$ and $P' \mathcal{R} Q'$.
2. if $Q \xrightarrow{\alpha} Q'$, then $\exists P'$ such that $P \xrightarrow{\alpha} P'$ and $P' \mathcal{R} Q'$.
3. if $P \xrightarrow{t} P'$, then $\exists Q'$ such that $Q \xrightarrow{t} Q'$ and $P' \mathcal{R} Q'$.
4. if $Q \xrightarrow{t} Q'$, then $\exists P'$ such that $P \xrightarrow{t} P'$ and $P' \mathcal{R} Q'$.
5. If P and Q are probabilistic then $\mu(P, S) = \mu(Q, S)$.

Part 5 of this definition states that the equivalence is “good” if the total probability of moving from P to any set of states which isn't “too fine” is the same for any related process Q .

Definition 13 We say two processes P and Q are in strong bisimulation, denoted by $P \sim Q$, iff there is any strong bisimulation with $P \mathcal{R} Q$.

Thus

$$\sim = \cup \{ \mathcal{R} : \mathcal{R} \text{ is a bisimulation} \}$$

We regularly need to convert an arbitrary relation into an equivalence relation. Thus we write \mathcal{R}^* to indicate the reflexive, symmetric, transitive closure of \mathcal{R} .

Proposition 6 \sim is the largest strong bisimulation

We show $+$ and $|$ are commutative and associative and have **nil** as a unit. We show $\setminus A$ and $[S]$ pass through the other operators and **rec** satisfies the usual law.

Theorem 1

1. $P + \text{nil} \sim P$
2. $P + Q \sim Q + P$
3. $P + (Q + R) \sim (P + Q) + R$
4. $P | \text{nil} \sim P$
5. $P | Q \sim Q | P$
6. $P | (Q | R) \sim (P | Q) | R$
7. $\text{nil} \setminus A \sim \text{nil}$
8. $(\alpha.P) \setminus A \sim \alpha.(P \setminus A)$ if $\alpha \notin A$ and $\bar{\alpha} \notin A$
9. $(\alpha.P) \setminus A \sim \mathbf{0}$ if $\alpha \in A$ or $\bar{\alpha} \in A$
10. $(\alpha[s \leftarrow t].P) \setminus A \sim \alpha[s \leftarrow t].(P \setminus A)$ if $\alpha \notin A$ and $\bar{\alpha} \notin A$
11. $(\alpha[s \leftarrow t].P) \setminus A \sim \text{nil}$ if $\alpha \in A$ or $\bar{\alpha} \in A$
12. $P + Q \setminus A \sim (P \setminus A) + (Q \setminus A)$
13. $(t).P \setminus A \sim (t).(P \setminus A)$
14. $\mathcal{R}[s \leftarrow f].P \setminus A \sim \mathcal{R}[s \leftarrow f].(P \setminus A)$
15. $(\sum [p_i] P_i) \setminus A \sim \sum [p_i] (P_i \setminus A)$
16. $\text{nil}[S] \sim \text{nil}$
17. $(\alpha.P)[S] \sim S(\alpha).(P[S])$
18. $(\alpha[s \leftarrow t].P)[S] \sim S(\alpha)[s \leftarrow t].(P[S])$
19. $(P + Q)[S] \sim P[S] + Q[S]$
20. $(t).P[S] \sim (t).(P[S])$
21. $\mathcal{R}[s \leftarrow f].P[S] \sim \mathcal{R}[s \leftarrow f].(P[S])$
22. $(\sum [p_i] P_i)[S] \sim \sum [p_i] (P_i[S])$
23. $\text{rec}(X = P) \sim P\{\text{rec}(X = P)/X\}$

Since we are only interested in processes modulo bisimulation and for this equivalence we have associativity (and commutativity) for $+$ and $|$ we will eliminate brackets and abbreviate multiple sums by $\sum_{i \in I} P_i$ and multiple parallels by $\prod_{i \in I} P_i$ (for I finite).

We might have expected $P + P \sim P$ but this is not true for probabilistic P where the two P on the LHS can choose independently (without resolving the $+$) and become different. It is of course true for P in the form $\sum (t_i).\alpha_i[s_i \leftarrow t_i].P_i$.

These last are not very remarkable. The more interesting laws are those that follow:

Theorem 2