

Distributed Processes and Location Failures

(Extended Abstract)

James Riely and Matthew Hennessy*

Abstract

Site failure is an essential aspect of distributed systems; nonetheless its effect on programming language semantics remains poorly understood. To model such systems, we define a process calculus in which processes are run at distributed *locations*. The language provides operators to kill locations, to test the status (dead or alive) of locations, and to spawn processes at remote locations. Using a variation of bisimulation, we provide alternative characterizations of strong and weak barbed congruence for this language, based on an operational semantics that uses *configurations* to record the status of locations. We then derive a second, symbolic characterization in which configurations are replaced by logical formulae. In the strong case the formulae come from a standard propositional logic, while in the weak case a temporal logic with past time modalities is required. The symbolic characterization establishes that, in principle, barbed congruence for such languages can be checked efficiently using existing techniques.

1 Introduction

Many semantic theories have been proposed for concurrent processes [18, 16, 6]. Although these theories have been fruitfully applied to the analysis of some distributed systems, for the most part they ignore an essential feature of such systems, namely their *distribution*.

As a simple example consider two implementations of a client-server application in which the client can demand an interactive service provided by the server, such as previewing or updating a document. In one implementation (System A) the server spawns a process to handle the document at its own site, the remote location, and the client previews the document remotely. In the other (System B) the server sends a process, including the document, to the client site, and the client previews the document locally. Using the semantic theories mentioned above it would be difficult to distinguish between these implementations, as the only difference between them is the location at which activity occurs. We aim to develop a useful *extensional* theory of systems which would take this type of property into account.

*Research funded by EPSRC project GR/K60701. Authors' address: School of Cognitive and Computing Sciences, Univ. of Sussex, Falmer, Brighton, BN1 9QH, UK, {jamesri,matthewh}@cogs.susx.ac.uk
Acknowledgement: We thank Flavio Corradini and Alan Jeffrey; both made important comments and suggestions in the early stages of this work.

In [8, 20, 10] such theories have been proposed. All of these theories, however, are based on a very strong assumption: that an observer, or user, can determine the location at which every action is performed. Here we start from a weaker premise: that in distributed systems sites are liable to *failure*. The model of failure we have adopted is a *fail stop* model in which failures are independent of each other and the number of failures that can occur is unbounded. Assuming that sites can fail, it is easy to see that Systems A and B, outlined above, are indeed different: if, after the client has begun interaction with the document, a failure occurs at the remote site, then in System A the client deadlocks, while in System B it can continue operation unaffected.

Our work is motivated by the papers [2, 12]. In these papers, distributed languages with location failures are defined and shown to be very expressive. In both of these papers, the semantics is based on *barbed equivalence*, which requires quantification over all program contexts and thus is difficult to use directly. In each of the cited works, the authors provide a translation from their language into a simpler (non-distributed) language and prove that the translations are adequate or fully abstract in some sense. While these translations provide theoretical results about the relative expressiveness of distributed and interleaving calculi, they are sufficiently complicated to make reasoning about examples, even simple ones, very difficult.

By restricting attention to an asynchronous language, Amadio [4] has recently improved on the results of [2], providing simpler translations. Although our work developed independently of [4], the language we study has much in common with the language developed there. The main difference is that our language has no value-passing, allowing us to concentrate on the effects of location failure and simplifying the statement of many of our results. Since the issues raised by failures and value passing are largely independent, this paper may be seen as providing two extensional views of a language similar to Amadio's; the first of these is concrete, as is his translation, the second is more abstract.

In [Section 2](#), we consider a simple language for *located processes* based on pure CCS [18], with which we assume familiarity. For example $(a.p)_\ell$ is a process located at ℓ which, if ℓ is alive, may perform the action a and then behave as $(p)_\ell$. In addition to the usual operators of CCS we have the following new operators: $\text{spawn}(\ell, p)$ which starts process p running at location ℓ ; $\text{kill } \ell.p$ which, if location ℓ is alive, kills ℓ (with the result that any process located at ℓ is deactivated) and then behaves as p ; and $\text{if } \ell \text{ then } p \text{ else } q$ which silently evolves to either p or q , depending on whether ℓ is alive or dead when the test is performed.

We give an operational semantics for this language in terms of a labelled transition system. The judgments depend on a set L , of *live* locations, and are of the form $L \triangleright P \xrightarrow{\alpha} L' \triangleright P'$, where P and P' are located processes and α is either a visible action, which permits synchronization, or the internal action τ . To decide on an appropriate equivalence between process terms we follow the approach advocated in [22]. We define both strong and weak barbed equivalence between processes, $\dot{\sim}$ and $\dot{\approx}$. We then dictate that the required equivalence, which we refer to as *barbed bisimulation equivalence*, is defined (for example in the weak case) as: $P \approx Q$ if and only if for every suitable context $\mathbb{C}[\]$, $\mathbb{C}[P] \dot{\approx} \mathbb{C}[Q]$. Although this may be reasonable, it is not a very useful definition; the reader is invited to determine whether the following pairs of processes should be equivalent or distinguished.

$$\begin{array}{ll}
P_1 = [(\alpha)_\ell | (\bar{\alpha} + \tau.a)_k] \setminus \alpha & P_2 = [(if\ k\ then\ \alpha\ else\ nil)_\ell | (\bar{\alpha}.a)_k] \setminus \alpha \\
Q_1 = [(\alpha + \tau)_\ell | (\bar{\alpha}.a)_k] \setminus \alpha & Q_2 = (spawn(k, a))_\ell
\end{array}$$

In [Section 3](#) we define two bisimulation-based relations, strong and weak *Located-Failure equivalence* (*LF-equivalence*) and show that these coincide with the indirectly defined barbed congruences. Since LF-equivalence is defined using bisimulations, the problem of deciding that two systems are semantically congruent can, in principle, be solved using standard proof techniques associated with bisimulation [\[18\]](#). However, constructing an LF-bisimulation requires that one consider the behavior of the systems under all possible sequences of kills, by both the systems themselves and the environment. The number of states that must be explored may be exponentially larger than the number needed to construct a CCS bisimulation.

In [Section 4](#) we use the ideas of [\[15\]](#) to give alternative *symbolic* characterizations of LF-equivalence that can be decided using a much smaller state space. The idea is to replace the operational judgments $L \triangleright P \xrightarrow{\alpha} L' \triangleright P'$ with judgments of the form $P \xrightarrow[\varphi]{\alpha} P'$, where φ is a logical formula that describes the circumstances under which the action α can be performed. In the strong case the required logic is straightforward: a propositional logic that describes the state (dead or alive) of the sites in the system. In the weak case, however, we require a more complicated logic that can express statements of the form *site ℓ was alive at some point in the past*. Using these symbolic transitions, the standard definition of *symbolic bisimulation* [\[15\]](#) requires only minor modification to capture \simeq and \cong ; hence the symbolic proof techniques and tools of [\[15\]](#) may be used to check the new semantic equivalences proposed in this paper.

In this extended abstract we have omitted several formal definitions and all proofs. The full version [\[21\]](#) includes additional results and examples, including a discussion of basic processes and comparisons with other equivalences.

2 The Language

The syntax of processes is parameterized with respect to several syntactic sets. We assume a set *Loc* of *locations* k, ℓ, m , a set *PConst* of *process constants* A used to define recursive processes, and a set *Act* of *communication actions* a, b, c , such that every action $a \in Act$ has a complement $\bar{a} \in Act$ ($\bar{\cdot}$ is a bijection on *Act*). The set $Act_\tau = Act \cup \{\tau\}$ of *actions* α includes also the distinguished *silent action* τ . The formal syntax is as follows. Most of the operators should be familiar from CCS; all of the new constructs have been described in the introduction.

$$\begin{array}{l}
p, q (\in BProc) ::= \alpha.p \mid spawn(\ell, p) \mid kill\ \ell.p \mid if\ \ell\ then\ p\ else\ q \mid A \mid \sum_{i \in I} P_i \\
\quad \quad \quad \mid p|q \mid p \setminus a \mid p[f] \\
P, Q (\in LProc) ::= P|Q \mid P \setminus a \mid P[f] \mid (p)_\ell
\end{array}$$

We have adopted a two-level syntax which distinguishes between *basic* processes p and *located* processes P . Intuitively, a basic process corresponds to what one normally thinks of as a *process*: a collection of threads of computation that must be run at a single site. A located process, instead, corresponds to a *distribution* of basic processes over several sites.

Note that many basic processes may be located at a single site, and a basic process may share a private channel (unknown to other basic processes running at the same site) with a remote process.

The ability of a process to perform an action is dependent on the set of live locations, and consequently the transition relation determining the operational semantics is defined between *configurations*. A *liveset* L is any subset of Loc . A *configuration* $(L \triangleright P)$ is a pair comprising a liveset L and a located process term P . The set of all configurations is $Config$, ranged over by C and D .

In giving the intensional semantics of processes, it will be convenient for later development if we distinguish executions of the operator $kill\ell.p$ depending upon whether ℓ is alive or dead at the time of execution. To capture this distinction, we extend the set of actions to the set $KAct = Act \cup \{kill\ell \mid \ell \in Loc\}$, which includes the *kill actions* $kill\ell$. Unless otherwise specified, μ ranges over $KAct_\tau = KAct \cup \{\tau\}$. In [Table 1](#) we define the transition relation $(\xrightarrow{\mu}) \subseteq Config \times Config$. The definition uses the following simple structural equivalence on processes:

$$(p|q)_\ell \equiv (p)_\ell | (q)_\ell \quad (p \setminus a)_\ell \equiv (p)_\ell \setminus a \quad (p[f])_\ell \equiv (p)_\ell [f]$$

While the transition relation \longrightarrow distinguishes effective kill actions from those that have no effect, a basic tenet of our study is that the precise moment of location failure should be unobservable. Thus we extract from \longrightarrow a transition relation \mapsto in which all kill actions have been replaced with silent actions. It is this derived relation \mapsto that we take to be fundamental.

Definition 1 (\mapsto). $C \xrightarrow{a} C'$ iff $C \xrightarrow{a} C'$ □
 $C \xrightarrow{\tau} C'$ iff $C \xrightarrow{\tau} C'$ or $\exists k: C \xrightarrow{killk} C'$

Most of the rules in [Table 1](#) are straightforward, being inherited directly from CCS, modulo the constraint that the process $(p)_\ell$ can only move if ℓ is alive. Note that the three new operators — kill, spawn and the conditional — are modeled as τ -transitions; this reflects the fact that in a distributed system the implementation of these operators would involve some computation and thus the passage of some time.

We now discuss the problem of defining an appropriate semantic equivalence for located processes, based on the transition relation \mapsto . An obvious possibility is to adapt the bisimulation equivalences of CCS [18]. (Strong) CCS-*bisimulation* is the largest symmetric relation \sim^{ccs} on configurations such that whenever $C \sim^{ccs} D$ and $C \xrightarrow{\alpha} C'$ there exists a D' such that $D \xrightarrow{\alpha} D'$ and $C' \sim^{ccs} D'$. A weak version of this relation, $\dot{\sim}^{ccs}$, can be obtained by adapting this definition to the *weak* transition relation \Longrightarrow , defined as usual. To see that CCS-bisimulation is not suitable for our language, for example is not a congruence, consider the processes $P_3 = [(\alpha.a)_\ell | (\bar{\alpha})_k] \setminus \alpha$ and $Q_3 = [(\alpha)_\ell | (\bar{\alpha}.a)_k] \setminus \alpha$. $P_3 \sim^{ccs} Q_3$, but P_3 and Q_3 can be distinguished by a context that kills location ℓ , if this kill action is performed after the initial communication on α .

The use of $\dot{\sim}^{ccs}$ for CCS has been justified in [22] by the fact that it coincides with the congruence obtained from a simple notion of observation called *barbed bisimulation*. Similar results have been obtained for lazy and eager functional languages [1, 14, 7], giving further evidence for the reasonableness of this approach. Roughly, two processes are barbed

Table 1 Transition system with configurations (symmetric rules for | omitted)

Act _c)	$L \triangleright (\alpha.p)_\ell \xrightarrow{\alpha} L \triangleright (p)_\ell$	if	$\ell \in L$
Spawn _c)	$L \triangleright (\text{spawn}(k, p))_\ell \xrightarrow{\tau} L \triangleright (p)_k$	if	$\ell \in L$
Kill1 _c)	$L \triangleright (\text{kill } m.p)_\ell \xrightarrow{\text{kill } m} L \setminus \{m\} \triangleright (p)_\ell$	if	$\ell \in L, m \in L$
Kill2 _c)	$L \triangleright (\text{kill } m.p)_\ell \xrightarrow{\tau} L \triangleright (p)_\ell$	if	$\ell \in L, m \notin L$
Cond1 _c)	$L \triangleright (\text{if } m \text{ then } p \text{ else } q)_\ell \xrightarrow{\tau} L \triangleright (p)_\ell$	if	$\ell \in L, m \in L$
Cond2 _c)	$L \triangleright (\text{if } m \text{ then } p \text{ else } q)_\ell \xrightarrow{\tau} L \triangleright (q)_\ell$	if	$\ell \in L, m \notin L$
Sum _c)	$L \triangleright (\sum_{i \in I} p_i)_\ell \xrightarrow{\mu} L' \triangleright (p'_j)_k$	if	$L \triangleright (p_j)_\ell \xrightarrow{\mu} L' \triangleright (p'_j)_k, j \in I$
Def _c)	$L \triangleright (A)_\ell \xrightarrow{\mu} L' \triangleright (p'_k)$	if	$L \triangleright (p)_\ell \xrightarrow{\mu} L' \triangleright (p'_k), A \stackrel{\text{def}}{=} p$
Str _c)	$L \triangleright P \xrightarrow{\mu} L' \triangleright Q$	if	$P \equiv P', L \triangleright P' \xrightarrow{\mu} L' \triangleright Q', Q' \equiv Q$
Par _c)	$L \triangleright P Q \xrightarrow{\mu} L' \triangleright P' Q$	if	$L \triangleright P \xrightarrow{\mu} L' \triangleright P'$
Comm _c)	$L \triangleright P Q \xrightarrow{\tau} L' \triangleright P' Q'$	if	$L \triangleright P \xrightarrow{a} L' \triangleright P', L \triangleright Q \xrightarrow{\bar{a}} L' \triangleright Q'$
Restr _c)	$L \triangleright P \setminus a \xrightarrow{\mu} L' \triangleright P' \setminus a$	if	$L \triangleright P \xrightarrow{\mu} L' \triangleright P', \mu \notin \{a, \bar{a}\}$
Ren _c)	$L \triangleright P[f] \xrightarrow{f(\mu)} L' \triangleright P'[f]$	if	$L \triangleright P \xrightarrow{\mu} L' \triangleright P'$

bisimilar if every silent transition of one can be matched by a silent transition of the other in such a way that the derived states are capable of exactly the same observable actions; in addition, the derived states must also be barbed bisimilar. For our language, the formal definition is as follows.

Definition 2 (Barbed bisimulation). Weak *barbed bisimilarity* (\approx) is the largest symmetric relation over configurations such that whenever $C \approx D$: (a) $C \xrightarrow{\tau} C'$ implies that for some $D', D \xrightarrow{\varepsilon} D'$ and $C' \approx D'$; and (b) for every $a, C \xrightarrow{a}$ implies $D \xrightarrow{a}$. Strong barbed bisimilarity (\approx) is obtained by replacing $\xRightarrow{\varepsilon}$ by $\xrightarrow{\varepsilon}$ everywhere in the definition. \square

Barbed bisimulation is a very weak relation; for example, it is not preserved by parallel composition. However, by closing over all contexts we arrive at a reasonable semantic equivalence that by definition enjoys an important property, namely that it is a congruence.

Definition 3 (Barbed equivalence). Located processes P and Q are (weak) *barbed equivalent* ($P \approx Q$) if for every context $\mathbb{C}[\]$ such that $\mathbb{C}[P]$ and $\mathbb{C}[Q]$ are configurations, $\mathbb{C}[P] \approx \mathbb{C}[Q]$. Strong barbed equivalence (\approx) is obtained in the same manner from \approx . \square

Because it requires quantification over all contexts, barbed equivalence is difficult to use directly. For example the processes P_1 and Q_1 , given in the introduction, are distinguished by \approx whereas P_2 and Q_2 are identified; it is far from obvious why. Even worse, processes P_5 and Q_5 (given in [Section 3](#)) are related, although establishing this fact requires that one prove that P_1 and Q_1 are *related* under the assumption that ℓ is alive at the time P_1 and Q_1 are compared, that is, ℓ is *initially alive*.

We end this section with some additional, simpler examples. The processes $(a)_\ell | (b)_k$ and $(b)_\ell | (a)_k$ can be distinguished by a context that kills ℓ . The same context can be used

to distinguish the basic processes $\text{spawn}(\ell, a)$ and $\text{spawn}(k, a)$, regardless of where they are located. These examples indicate that although the location of an action is not reflected directly in the operational semantics they do impinge on the behavior of processes. The order in which kill actions are executed is also significant. For example $\text{kill } \ell. \text{kill } k$ can be distinguished from $\text{kill } k. \text{kill } \ell$ using the process $(a)_\ell | (b)_k$.

3 Located-Failures Equivalence

In this section and the next, we provide alternative characterizations of barbed equivalence for our language. Note that if $L \triangleright P \xrightarrow{\mu} L' \triangleright P'$, then L' is determined by L and μ . To emphasize this, we adopt the following notation. For each action μ , we define a function “ iafter_μ ” which reflects the immediate effect of action μ on a liveset. We also define the relations $\xrightarrow[L]{\mu}$ and $\xRightarrow[L]{\mu}$ on process terms, which capture the capability of action μ under liveset L .

$$\text{iafter}_\mu(L) \stackrel{\text{def}}{=} \begin{cases} L \setminus \{k\}, & \text{if } \mu = \text{kill } k \\ L, & \text{if } \mu \in \text{Act} \cup \{\tau, \varepsilon\} \end{cases} \quad \left| \begin{array}{l} P \xrightarrow[L]{\mu} P' \stackrel{\text{def}}{\Leftrightarrow} L \triangleright P \xrightarrow{\mu} \text{iafter}_\mu(L) \triangleright P' \\ P \xRightarrow[L]{\mu} P' \stackrel{\text{def}}{\Leftrightarrow} L \triangleright P \xRightarrow{\mu} \text{iafter}_\mu(L) \triangleright P' \end{array} \right.$$

For example, $\text{iafter}_\alpha(L) = L$ for any α , and $\text{iafter}_{\text{kill } \ell}(\{\ell, k\}) = \{k\}$. If $P = (\alpha.a)_\ell | (\bar{\alpha})_k$, then $P \xrightarrow[\{\ell, k\}]{a} \text{nil}$, but P has no a -transition under the liveset $\{k\}$.

We first present the strong case.

Definition 4 (Strong LF-equivalence). Let $\mathcal{S} = \{\mathcal{S}_L\}_{L \subseteq \text{Loc}}$ be an indexed family of relations on $L\text{Proc}$. \mathcal{S} is a *strong LF-bisimulation* if for every L , \mathcal{S}_L is symmetric and whenever $P \mathcal{S}_L Q$:

- (a) $P \xrightarrow[L]{\mu} P'$ implies $\exists Q' : Q \xrightarrow[L]{\mu} Q'$ and $P' \mathcal{S}_{\text{iafter}_\mu(L)} Q'$
- (b) for every $k \in L$ $P \mathcal{S}_{L \setminus \{k\}} Q$

P and Q are *strong LF-equivalent under L* ($P \simeq_L Q$) if there exists a strong LF-bisimulation \mathcal{S} with $P \mathcal{S}_L Q$.

P and Q are *strong LF-equivalent* ($P \simeq Q$), if $P \simeq_L Q$ for every subset L of Loc . \square

In the full paper, we prove that \simeq and \sim coincide. The alternative characterization of weak barbed equivalence is more complicated: it is not sufficient to change the strong arrows in [Definition 4](#) to weak arrows. To see this, consider the following processes:

$$\begin{aligned} P_5 &= [(b.\beta.\alpha + b.(\alpha + \tau))_\ell | (\bar{\beta}.(\bar{\alpha} + \tau.a) + \bar{\alpha}.a)_k] \setminus \alpha \setminus \beta \\ Q_5 &= [(b.(\alpha + \tau))_\ell | (\bar{\alpha}.a)_k] \setminus \alpha \end{aligned}$$

If ℓ is initially dead, P_5 and Q_5 are clearly equivalent: both are strong equivalent to nil . If ℓ is initially alive, however, the situation is not so clear. The questionable move is P_5 's b -transition to $P_1 \simeq [(\alpha)_\ell | (\bar{\alpha} + \tau.a)_k] \setminus \alpha$. To match this move Q_5 must perform a weak b -transition to $Q_1 \simeq [(\alpha + \tau)_\ell | (\bar{\alpha}.a)_k] \setminus \alpha$. But P_1 and Q_1 are not barbed equivalent: if ℓ

is dead, then Q_1 is capable of a a transition that P_1 cannot match. This would lead one to believe that P and Q are *not* barbed equivalent; however, they are.

Intuitively this is true because when P_5 reaches P_1 , ℓ must be alive; thus P_1 and Q_1 need only be compared under the constraint that ℓ is initially alive. Once this comparison has begun, the environment can distinguish Q_1 from P_1 only by killing ℓ , but it cannot control internal activity on the part of P_1 before ℓ is dead.

Definition 5 (Weak LF-equivalence). For $\mu \in Act_\tau$, define $\hat{\mu}$ such that $\hat{a} = a$ and $\hat{\tau} = \varepsilon$. The definition of \cong is similar to that for \simeq , except that when $P \mathcal{S}_L Q$, we require:

- (a) $P \xrightarrow[\!L]{\mu} P'$ implies $\exists Q' : Q \xrightarrow[\!L]{\hat{\mu}} Q'$ and $P' \mathcal{S}_{\text{iafter}_\mu(L)} Q'$
- (b) for every $k \in L$ $\exists Q' : Q \xrightarrow[\!L]{\varepsilon} \cdot \xrightarrow[\!L \setminus \{k\}]{\varepsilon} Q'$ and $P \mathcal{S}_{L \setminus \{k\}} Q'$ □

Whereas the first clause in the definition of weak LF-bisimulation is as one would expect, the second clause is somewhat surprising. It says, in effect, that if the environment kills a location k , then Q must be able to (silently) evolve to a process Q' that matches P ; but in reaching Q' , Q may exploit the intermediate states of the system (that is, k alive, then k dead).

Theorem 6. For all located processes $P \cong Q$ if and only if $P \approx Q$. □

4 Symbolic characterizations

While the LF-equivalences provide a great deal of insight into the meaning of barbed equivalence in distributed process description languages such as ours, they are unwieldy to use in practice. For the most part, this is due to the use of configurations in the operational semantics. In this section, we improve this situation by defining a *symbolic* transition system directly on located process terms, then giving characterizations of strong and weak LF-equivalence using these symbolic transitions. As one should expect, the weak case is quite a bit more subtle than the strong.

We begin by giving the symbolic operational semantics. The symbolic transition relation makes use of propositional formulae π, ρ , which are given a semantics in terms of livesets. Intuitively, a formula indicates a set of constraints on the status of locations (dead or alive) at the time that the transition is enabled. If $P \xrightarrow[\!0 \wedge 1]{\mu} P'$ then if location 0 is dead and 1 is alive, P is capable of making an μ -transition to P' ; that is, if $0 \notin L$ and $1 \in L$ then $P \xrightarrow[\!L]{\mu} P'$. In [Table 2](#) we define the transition relation $\xrightarrow[\!L]{\mu} \subseteq LProc \times LProc$. The two transition systems are related by the fact that $P \xrightarrow[\!L]{\mu} P'$ if and only if there exists a π such that $P \xrightarrow[\!L]{\mu} P'$ and $L \models \pi$.

The standard definition of symbolic bisimulation [[15](#)] requires that we define entailment between formulae, which we do in the standard way:

$$\pi \Vdash \rho \text{ iff } \forall L : L \models \pi \text{ implies } L \models \rho$$

Note that entailment is a preorder on formulae. If $\pi \Vdash \rho$ we say that π is *stronger* than ρ . ff is the strongest formula under \Vdash , tt the weakest.

Table 2 Symbolic transition system (symmetric rules for | omitted)

Act _s)	$(\alpha.p)_\ell \xrightarrow{\alpha} (p)_\ell$		
Spawn _s)	$(\text{spawn}(k, p))_\ell \xrightarrow{\tau} (p)_k$		
Kill1 _s)	$(\text{kill } m.p)_\ell \xrightarrow{\frac{\text{kill } m}{\ell \wedge m}} (p)_\ell$		
Kill2 _s)	$(\text{kill } m.p)_\ell \xrightarrow{\frac{\tau}{\ell \wedge \bar{m}}} (p)_\ell$		
Cond1 _s)	$(\text{if } m \text{ then } p \text{ else } q)_\ell \xrightarrow{\frac{\tau}{\ell \wedge m}} (p)_\ell$		
Cond2 _s)	$(\text{if } m \text{ then } p \text{ else } q)_\ell \xrightarrow{\frac{\tau}{\ell \wedge \bar{m}}} (q)_\ell$		
Sum _s)	$(\sum_{i \in I} P_i)_\ell \xrightarrow{\mu} (p'_j)_\ell$	if	$(p_j)_\ell \xrightarrow{\mu} (p'_j)_\ell, j \in I$
Def _s)	$(A)_\ell \xrightarrow{\mu} (p')_\ell$	if	$(p)_\ell \xrightarrow{\mu} (p')_\ell, A \stackrel{\text{def}}{=} p$
Str _s)	$P \xrightarrow{\mu} Q$	if	$P \equiv P', P' \xrightarrow{\mu} Q', Q' \equiv Q$
Par _s)	$P Q \xrightarrow{\mu} P' Q$	if	$P \xrightarrow{\mu} P'$
Comm _s)	$P Q \xrightarrow{\frac{\tau}{\pi \wedge \bar{\rho}}} P' Q'$	if	$P \xrightarrow{a} P', Q \xrightarrow{\bar{a}} Q'$
Restr _s)	$P \setminus a \xrightarrow{\mu} P' \setminus a$	if	$P \xrightarrow{\mu} P', \mu \notin \{a, \bar{a}\}$
Ren _s)	$P[f] \xrightarrow{\frac{f(\mu)}{\pi}} P'[f]$	if	$P \xrightarrow{\mu} P'$

We must also identify a set of formulae suitable as parameters in the recursive definition of symbolic equivalence, that is, the analogs of the parameters L in the definition of LF-equivalence. Intuitively, when we say that P and Q are LF-equivalent under L , we are limiting attention to a single possible world, namely that in which exactly the sites in L are alive. The idea of symbolic equivalences, instead, is to treat many possible worlds simultaneously (via entailment). In the case of strong LF-bisimulation, where $P \simeq_L Q$ and $M \subseteq L$ imply $P \simeq_M Q$, this is achieved by restricting attention to *negative formulae* — formulae which contain no positive atoms — in the recursive definition of symbolic equivalence. Finally, we identify a transformation on formulae (indexed by actions) which specifies the conditions under which residual processes are to be compared:

$$M \models \text{after}_\alpha(\rho) \text{ iff } \exists L: L \models \rho \text{ and } M \subseteq L$$

$$M \models \text{after}_{\text{kill } k}(\rho) \text{ iff } \exists L: L \models \rho \text{ and } M \subseteq L \setminus \{k\}$$

Definition 7 (Strong symbolic bisimulation). Let \mathcal{S} be a family of relations on $LProc$ indexed by negative formulae ϑ . \mathcal{S} is a *strong symbolic bisimulation* if for every ϑ , \mathcal{S}_ϑ is symmetric and whenever $P \mathcal{S}_\vartheta Q$ and $P \xrightarrow{\mu} P'$ then for some π_i, ρ_i , and Q_i :


- (a) $\vartheta \wedge \pi \Vdash \bigvee_i \rho_i$,
- (b) $\rho_i \Vdash \pi_i$,
- (c) $Q \xrightarrow{\mu/\pi_i} Q_i$, and
- (d) $P' \mathcal{S}_{\text{after}_\mu(\rho_i)} Q_i$

We write $P \simeq_\vartheta^s Q$ to indicate that there exists a symbolic bisimulation \mathcal{S} with $P \mathcal{S}_\vartheta Q$. \square

Theorem 8. $P \simeq_L Q$ iff $\exists \vartheta: P \simeq_\vartheta^s Q$ and $L \models \vartheta$. In addition, $(\simeq) = (\simeq_{\text{tt}}^s)$. \square

As a first attempt to define weak symbolic bisimulation, let us try simply replacing the strong transitions in [Definition 7](#) with weak edges defined by conjoining formulae. For

example, we would have $P \xrightarrow[\text{tt}]{\varepsilon} P$ and $P \xrightarrow[\pi \wedge \rho]{a} P'$ if $P \xrightarrow{\tau} \cdot \xrightarrow{\rho} P'$. Unfortunately, this definition does not suffice. Consider the processes P_5 and Q_5 , previously defined; these have the following symbolic transition graphs (where we have write $\xrightarrow{\mu}$ as $\xrightarrow{\mu\pi}$):



As noted in [Section 3](#), in order to prove these processes equivalent we must compare the processes P_1 and Q_1 under the assumption that ℓ is initially alive, but using our provisional definition we would end up comparing P_1 and Q_1 under the assumption $\text{tt} = \text{neg}(\ell \wedge k)$, which is not strong enough to prove that they are related.

As a second attempt, we might simply allow all positive information to carry over into the recursive formula ϑ_i , that is, change the last clause of [Definition 7](#) to $\vartheta_i = \rho_i$. Whereas our first attempt produced an equivalence that was too strong, the revised definition is too weak. For example, the following processes would be identified even though they are not barbed equivalent.

$$\begin{array}{ccc}
 P_6 = [(\alpha.a)_\ell | (\bar{\alpha})_k] \setminus \alpha & & Q_6 = [(\alpha)_\ell | (\bar{\alpha}.a)_k] \setminus \alpha \\
 \downarrow \tau_{\ell \wedge k} & \approx & \downarrow \tau_{\ell \wedge k} \\
 P'_6 & & Q'_6 \\
 \downarrow a_\ell & & \downarrow a_k \\
 \bullet & & \bullet
 \end{array}$$

Here P'_6 and Q'_6 would be compared under the formula $\ell \wedge k$. This formula, however, says something more than we would like, namely that ℓ and k remain alive until P'_6 and Q'_6 execute their first action. More complicated examples can be constructed to show that we must be able to express properties such as “ ℓ and k must have been alive, then ℓ must have died, and after that k must have died.”

Our solution is to define weak symbolic edges using a *past-time temporal logic* [17], interpreted over sequences of livesets. A *live sequence* \mathcal{L} is a finite nonempty sequence of livesets $\langle L_1, \dots, L_n \rangle$, such that for every i between 1 and $n - 1$ there exists a location k such that $L_{i+1} = L_i \setminus \{k\}$. For example, $\langle \{\ell\}, \emptyset \rangle$ is a live sequence, but $\langle \{\ell\}, \{\ell\} \rangle$ and $\langle \{\ell, k\}, \emptyset \rangle$ are not. We write $\mathcal{L}_{(i)}$ for the i^{th} element of \mathcal{L} and, where clear from context, use n to refer to the length of \mathcal{L} . Thus, for example, \mathcal{L} models $\bar{\ell}$ if $\ell \notin \mathcal{L}_{(n)}$ and \mathcal{L} models $\diamond\varphi$ if \mathcal{L} or some prefix of \mathcal{L} models φ . Because live sequences must be strictly decreasing, $\ell \wedge \diamond\bar{\ell}$ is unsatisfiable; however $\langle \{\ell\}, \emptyset \rangle \models \bar{\ell} \wedge \diamond\ell$. Weak symbolic transitions are defined as follows:

$$\begin{array}{l}
 P \xrightarrow[\text{tt}]{\varepsilon} P \\
 P \xrightarrow[\diamond(\varphi \wedge \pi)]{\varepsilon} P' \text{ if } P \xrightarrow[\varphi]{\varepsilon} \cdot \xrightarrow{\tau} P' \\
 P \xrightarrow[\varphi \wedge \pi]{\mu} P' \text{ if } P \xrightarrow[\varphi]{\mu} \cdot \xrightarrow{\tau} P' \\
 P \xrightarrow[\varphi \wedge \pi]{\alpha} P' \text{ if } P \xrightarrow[\varphi]{\varepsilon} \cdot \xrightarrow{\alpha} P' \\
 P \xrightarrow[\bar{k} \wedge \Theta(\varphi \wedge \pi)]{\text{kill}k} P' \text{ if } P \xrightarrow[\varphi]{\varepsilon} \cdot \xrightarrow{\text{kill}k} P'
 \end{array}$$

Intuitively $P \xrightarrow[\varphi]{\mu} P'$ means that P can perform the action μ to become P' in an environment where the change in live sets satisfies the formula φ . For example if $\varphi_1 = (\ell \wedge k) \S \ell$ and $\varphi_2 = (\ell \wedge k) \S k$ then P_6 has the symbolic transition $\xrightarrow[\varphi_1]{a}$ but not $\xrightarrow[\varphi_2]{a}$, whereas for Q_6 it is the opposite.

As parameters to the weak relation we simply take Boolean formulae, but now interpreted on the initial liveset of a live sequence. Rather than use two logics in the definition or introduce additional operators, we define the function “initially” which converts Boolean formulae into temporal formulae with this interpretation in mind. The transformation function for generating formulae after an action is performed, which we call “finally”, must then transform temporal formula into Boolean ones. The definitions are as follows: (In the full paper, we show how to calculate these functions.)

$$\begin{aligned} \mathcal{L} \models \text{initially}(\pi) & \text{ iff } \mathcal{L}_{(1)} \models \pi \\ M \models \text{finally}(\varphi) & \text{ iff } \exists \mathcal{L} : \mathcal{L} \models \varphi \text{ and } M = \mathcal{L}_{(n)} \end{aligned}$$

Definition 9 (Weak symbolic bisimulation). Similar to [Definition 7](#), except that when $P \mathcal{S}_\pi Q$ and $P \xrightarrow[\varphi]{\hat{\mu}} P'$ we require:

- | | | |
|---|--|---|
| (a) $\text{initially}(\pi) \wedge \varphi \Vdash \bigvee_i \psi_i,$
(b) $\psi_i \Vdash \varphi_i,$ | (c) $Q \xrightarrow[\varphi_i]{\hat{\mu}} Q_i,$ and
(d) $P' \mathcal{S}_{\text{finally}(\psi_i)} Q_i$ | □ |
|---|--|---|

Theorem 10. $P \approx_L Q$ iff $\exists \pi : L \models \pi$ and $P \approx_\pi^s Q$. In addition, $(\approx) = (\approx_{\text{tt}}^s)$. □

5 Conclusions

In this paper we have proposed a new semantic theory for distributed systems which takes into account the possibility of failures at sites. This theory is an adaptation of standard bisimulation-based theories [\[18\]](#) using an operational semantics for *located processes*. The new semantic equivalences are justified in terms of *barbed bisimulations* [\[22\]](#). We also give *symbolic* characterizations of the new equivalences, which means that they can be investigated using the symbolic methods of [\[15\]](#).

Site failure has also played a role in languages studied in [\[2, 4, 12\]](#). In these papers abstract languages based on Facile [\[13\]](#) or the pi-calculus [\[19, 5\]](#) are studied. The original motivation for this paper was to provide an alternative characterization of barbed equivalence for languages such as these. Although we have not treated value passing or references, we postulate that our results can be extended in a straightforward way to value-passing languages which retain the assumption that all failures are independent, such as the languages in [\[2, 4\]](#). More delicate is the extension to languages such as the distributed join-calculus [\[12\]](#) in which the independence assumption is dropped. In this case the logical language used for symbolic bisimulations must be extended to allow statements about the interdependence of locations; we leave this to future work.

A number of location-based equivalences already exist in the literature [\[8, 9, 20, 10\]](#); however, none of these theories addresses the possible failure of sites. Their emphasis, rather, is to define a measure of the concurrency or distribution of a process: two processes

are deemed equivalent only if, informally, they have the same degree of concurrency. In the full paper we give a series of counter-examples which show that \approx is incomparable with all of the equivalences proposed in these papers; we also discuss variations on the language and model of failure.

References

- [1] Samson Abramsky. The lazy lambda calculus. In *Research Topics in Functional Programming*, pages 65–117. Addison-Wesley, 1990.
- [2] Roberto Amadio and Sanjiva Prasad. Localities and failures. In *FST-TCS*, volume 880 of *LNCS*. Springer, 1994.
- [3] Roberto Amadio. From a concurrent λ -calculus to the π -calculus. In *Foundations of Computation Theory*, volume 965 of *LNCS*. Springer, 1995.
- [4] Roberto Amadio. An asynchronous model of locality, failure, and process mobility. Technical report, Laboratoire d'Informatique de Marseille, 1997.
- [5] Roberto Amadio, Ilaria Castellani, and Davide Sangiorgi. On bisimulations for the asynchronous π -calculus. In *CONCUR96*, volume 1119 of *LNCS*, pages 147–162. Springer, 1996.
- [6] J. C. M. Baeten and W. P. Weijland. *Process Algebra*. Cambridge University Press, 1990.
- [7] G. Boudol. A lambda calculus for (strict) parallel functions. *Information and Control*, 108:51–127, 1994.
- [8] G. Boudol, I. Castellani, M. Hennessy, and A. Kiehn. A theory of processes with localities. *Formal Aspects of Computing*, 6:165–200, 1994.
- [9] I. Castellani. Observing distribution in processes: static and dynamic localities. *International Journal of Foundations of Computer Science*, 6(6):353–393, 1995.
- [10] Flavio Corradini. *Space, Time and Nondeterminism in Process Algebras*. PhD thesis, Università Degli Studi di Roma “La Sapienza”, 1996.
- [11] C. Fournet and G. Gonthier. The reflexive CHAM and the join-calculus. In *POPL94*. ACM Press, 1994.
- [12] C. Fournet, G. Gonthier, J.J. Levy, L. Marganet, and D. Remy. A calculus of mobile agents. In *CONCUR96*, volume 1119 of *LNCS*, pages 406–421. Springer, 1996.
- [13] A. Giacalone, P. Mishra, and S. Prasad. A symmetric integration of concurrent and functional programming. *International Journal of Parallel Programming*, 18(2):121–160, 1989.
- [14] Andrew D. Gordon. Bisimilarity as a theory of functional programming. In *MFPS*, volume 1 of *ENTCS*, <http://pigeon.elsevier.nl/mcs/tcs/pc/Menu.html>. Elsevier, 1995.
- [15] M. C. B. Hennessy and H. Lin. Symbolic bisimulations. *Theoretical Computer Science*, 138:353–389, 1995.
- [16] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
- [17] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent System: Specification*. Springer, 1992.
- [18] Robin Milner. *Communication and concurrency*. Prentice-Hall, 1989.
- [19] Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes. *Information and Computation*, 100(1), September 1992.

- [20] Ugo Montanari and Daniel Yankelovich. Partial order localities. In *ICALP92*, volume 623 of *LNCS*, pages 617–628. Springer, 1992.
- [21] James Riely and Matthew Hennessy. Distributed processes and location failures. Technical Report 2/97, University of Sussex, Department of Computer Science, <http://www.cogs.susx.ac.uk>, 1997.
- [22] Davide Sangiorgi. *Expressing Mobility in Process Algebras: First-Order and Higher-Order Paradigms*. PhD thesis, University of Edinburgh, 1992.