First-Order Axioms for Asynchrony

Peter Selinger*
Department of Mathematics
University of Pennsylvania
Philadelphia, PA 19104-6395

Abstract. We study properties of asynchronous communication independently of any concrete concurrent process paradigm. We give a general-purpose, mathematically rigorous definition of several notions of asynchrony in a natural setting where an agent is asynchronous if its input and/or output is filtered through a buffer or a queue, possibly with feedback. In a series of theorems, we give necessary and sufficient conditions for each of these notions in the form of simple first-order or second-order axioms. We illustrate the formalism by applying it to asynchronous CCS and the core join calculus.

Introduction

The distinction between *synchronous* and *asynchronous* communication is a relevant issue in the design and analysis of distributed and concurrent networks. Intuitively, communication is said to be synchronous if messages are sent and received simultaneously, via a 'handshake' or 'rendez-vous' of sender and receiver. It is asynchronous if messages travel through a communication medium with possible delay, such that the sender cannot be certain if or when a message has been received.

Asynchronous communication is often studied in the framework of concurrent process paradigms such as the asynchronous π-calculus, which was originally introduced by Honda and Tokoro [9], and was independently discovered by Boudol [6] as a result of his work with Berry on chemical abstract machines [5]. Another such asynchronous paradigm is the join calculus, which was recently proposed by Fournet and Gonthier as a calculus of mobile agents in distributed networks with locality and failure [7, 8].

In this paper, we study properties of asynchronous communication in general, not with regard to any particular process calculus. We give a general-purpose, mathematically rigorous definition of asynchrony, and we then show that this notion can be axiomatized. We model processes by labeled transition systems with input and output, a framework that is sufficiently general to fit concurrent process paradigms such as the π-calculus or the join calculus, as well as data flow models and other such formalisms. These transition systems are similar to Lynch and Stark's input/output automata [10], but our treatment is more category-theoretical and close in spirit to Abramsky's interaction categories [1, 2].

Various properties of asynchrony have been exploited in different contexts by many authors. For instance, Lynch and Stark [10] postulate a form of *input receptivity* for

*This research was supported by an Alfred P. Sloan Doctoral Dissertation Fellowship.

their automata. Palamidessi [13] makes use of a certain *confluence* property to prove that the expressive power of the asynchronous π-calculus is strictly less than that of the synchronous π-calculus. Axioms similar to ours have been postulated by [4] and by [14] for a notion of asynchronous labeled transition systems, but without the input/output distinction which is central to the our approach.

The main novelty of this paper is that our axioms are not postulated *a priori*, but derived from more primitive notions. We define asynchrony in elementary terms: an agent is asynchronous if its input and/or output is filtered through a communication medium, such as a buffer or a queue, possibly with feedback. We then show that our first- and second-order axioms precisely capture each of these notions. This characterization justifies the axioms *a posteriori*. As a testbed and for illustration, we apply these axioms to an asynchronous version of Milner's CCS, and to the core join calculus.

Due to limitations of space, most proofs are omitted in this abbreviated version of the paper. Only the proof of Theorem 2.1 is included as a typical example of the type of reasoning that is employed here. A full version of the paper is available from the author and will also appear as part of his Ph.D. Thesis.

Acknowledgments. I would like to thank Catuscia Palamidessi, Davide Sangiorgi, Benjamin Pierce, Dale Miller, Steve Brookes, Ian Stark, and Glynn Winskel for discussions and helpful comments on this work.

1 An Elementary Definition of Asynchrony

If R is a binary relation, we write R^{-1} for the inverse relation and R^* for the reflexive, transitive closure of R. We also write $\leftarrow$ for $\rightarrow^{-1}$, etc. The binary identity relation on a set is denoted Δ. The composition of two binary relations R and Q is written $R \circ Q$ or simply RQ, *i.e.* $xRQz$ if there exists y such that $xRyQz$. The disjoint union of two sets X and Y is denoted by $X + Y$.

1.1 Labeled Transition Systems and Bisimulation

To keep this paper self-contained, we summarize the standard definitions for labeled transition systems and weak and strong bisimulation.

Definition. A ***labeled transition system (LTS)*** is a tuple $\mathbf{S} = \langle S, A, \rightarrow_{\mathbf{S}}, s_0 \rangle$, where S is a set of ***states***, A is a set of ***actions***, $\rightarrow_{\mathbf{S}} \subseteq S \times A \times S$ is a ***transition relation*** and $s_0 \in S$ is an ***initial state***. We call A the ***type*** of $\mathbf{S}$, and we write $\mathbf{S} : A$.

We often omit the subscript on $\rightarrow_{\mathbf{S}}$, and we write $|\mathbf{S}|$ for the set of states S. For $\alpha \in A$, we regard $\xrightarrow{\alpha}$ as a binary relation on $|\mathbf{S}|$ via $s \xrightarrow{\alpha} s'$ iff $\langle s, \alpha, s' \rangle \in \rightarrow$.

Definition. Let $\mathbf{S}$ and $\mathbf{T}$ be LTSs of type A. A binary relation $R \subseteq |\mathbf{S}| \times |\mathbf{T}|$ is a ***strong bisimulation*** if for all $\alpha \in A$, $R\xrightarrow{\alpha} \subseteq \xrightarrow{\alpha} R$ and $R^{-1}\xrightarrow{\alpha} \subseteq \xrightarrow{\alpha} R^{-1}$. In diagrams:

$$\begin{array}{ccc} s & R & t \\ \downarrow{\scriptstyle\alpha} & & \\ t' & & \end{array} \Rightarrow \exists s'.\ \begin{array}{ccc} s & R & t \\ \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\alpha} \\ s' & R & t' \end{array} \qquad \text{and} \qquad \begin{array}{ccc} s & R & t \\ \downarrow{\scriptstyle\alpha} & & \\ s' & & \end{array} \Rightarrow \exists t'.\ \begin{array}{ccc} s & R & t \\ \downarrow{\scriptstyle\alpha} & & \downarrow{\scriptstyle\alpha} \\ s' & R & t' \end{array}$$

Next, we consider LTSs with a distinguished action $\tau \in A$, called the ***silent*** or the ***unobservable*** action. Let $\stackrel{\tau}{\Rightarrow}$ be the relation $\stackrel{\tau}{\rightarrow}^*$. For $a \in A \setminus \tau$, let $\stackrel{a}{\Rightarrow}$ be the relation $\stackrel{\tau}{\rightarrow}^* \stackrel{a}{\rightarrow} \stackrel{\tau}{\rightarrow}^*$. A binary relation $R \subseteq |\mathbf{S}| \times |\mathbf{T}|$ is a ***weak bisimulation*** if for all $\alpha \in A$, $R \stackrel{\alpha}{\rightarrow} \subseteq \stackrel{\alpha}{\Rightarrow} R$ and $R^{-1} \stackrel{\alpha}{\rightarrow} \subseteq \stackrel{\alpha}{\Rightarrow} R^{-1}$. In diagrams:

$$\begin{array}{ccc} s & R & t \\ \downarrow{\scriptstyle \alpha} & & \\ t' & & \end{array} \Rightarrow \exists s'.\ \begin{array}{ccc} s & R & t \\ \alpha\Downarrow & & \downarrow\alpha \\ s' & R & t' \end{array} \qquad \text{and} \qquad \begin{array}{ccc} & s & R\ t \\ \alpha\downarrow & & \\ & s' & \end{array} \Rightarrow \exists t'.\ \begin{array}{ccc} s & R & t \\ \downarrow\alpha & & \Downarrow\alpha \\ s' & R & t' \end{array}$$

It is well-known that there is a maximal strong bisimulation, which we denote by $\sim$, and a maximal weak bisimulation, which we denote by $\approx$. We say that $s \in |\mathbf{S}|$ and $t \in |\mathbf{T}|$ are ***strongly (weakly) bisimilar*** if $s \sim t$ ($s \approx t$). Finally, $\mathbf{S}$ and $\mathbf{T}$ are said to be strongly (weakly) bisimilar if $s_0 \sim t_0$ ($s_0 \approx t_0$).

The relations $\sim$ and $\approx$, as binary relations on an LTS $\mathbf{S}$, are equivalence relations. We denote the respective equivalence classes of a state s by $[s]_\sim$ and $[s]_\approx$. On the quotient $\mathbf{S}/\!\sim$, we define transitions $[s]_\sim \stackrel{a}{\rightarrow} [t]_\sim$ iff $s \stackrel{a}{\rightarrow} \sim t$, making it into a well-defined transition system. Similarly, on $\mathbf{S}/\!\approx$, we define $[s]_\approx \stackrel{a}{\rightarrow} [t]_\approx$ iff $s \stackrel{a}{\rightarrow} \approx t$. For all $s \in \mathbf{S}$, one has $s \sim [s]_\sim$ and $s \approx [s]_\approx$, and hence $\mathbf{S} \sim (\mathbf{S}/\!\sim)$ and $\mathbf{S} \approx (\mathbf{S}/\!\approx)$. We say that $\mathbf{S}$ is ***$\sim$-reduced*** if $\mathbf{S} = \mathbf{S}/\!\sim$, and ***$\approx$-reduced*** if $\mathbf{S} = \mathbf{S}/\!\approx$.

1.2 Input, Output and Sequential Composition

So far we have distinguished only one action: the silent action τ. We will now add further structure to the set of actions by distinguishing input and output actions. Let *in* and *out* be constants. For any sets X and Y, define a set of ***input actions*** $\mathit{In}\, X := \{\mathit{in}\} \times X$, and a set of ***output actions*** $\mathit{Out}\, Y := \{\mathit{out}\} \times Y$. Note that $\mathit{In}\, X$ and $\mathit{Out}\, Y$ are disjoint. We will write input and output actions as $\mathit{in}\, x$ and $\mathit{out}\, x$ instead of $\langle \mathit{in}, x\rangle$ and $\langle \mathit{out}, x\rangle$, respectively. Let B be a set whose elements are not of the form $\mathit{in}\, x$, $\mathit{out}\, y$ or τ. The elements of $\mathbf{B} + \{\tau\}$ are called ***internal actions***.

Definition. We define $X \rightarrow_B Y$ to be the set $\mathit{In}\, X + \mathit{Out}\, Y + B + \{\tau\}$. A labeled transition system $\mathbf{S}$ of type $X \rightarrow_B Y$ is called an ***LTS with input and output***, or simply an ***agent***. If B is empty, we will omit the subscript in $X \rightarrow_B Y$.

Our labeled transition systems with input and output are similar to the input/output automata of Lynch and Stark [10]. However, we consider a notion of sequential composition that is more in the spirit of Abramsky's interaction categories [1, 2]. Given two agents $\mathbf{S}\colon X \rightarrow_B Y$ and $\mathbf{T}\colon Y \rightarrow_B Z$, we define $\mathbf{S};\mathbf{T}\colon X \rightarrow_B Z$ by feeding the output of $\mathbf{S}$ into the input of $\mathbf{T}$. This is a special case of parallel composition and hiding. Notice that this notion of sequential composition is different from the one of CSP or ACP, where $\mathbf{T}$ cannot start execution until $\mathbf{S}$ is finished.

Definition 1.1. Let $\mathbf{S}\colon X \rightarrow_B Y$ and $\mathbf{T}\colon Y \rightarrow_B Z$ be agents with respective initial states s_0 and t_0. The ***sequential composition*** $\mathbf{S};\mathbf{T}$ is of type $X \rightarrow_B Z$. It has states $|\mathbf{S}| \times |\mathbf{T}|$ and initial state $\langle s_0, t_0\rangle$. The transitions are given by the following rules:

$$\frac{s \stackrel{\alpha}{\rightarrow}_{\mathbf{S}} s' \qquad \alpha \text{ not output}}{\langle s,t\rangle \stackrel{\alpha}{\rightarrow}_{\mathbf{S};\mathbf{T}} \langle s',t\rangle} \qquad \frac{t \stackrel{\alpha}{\rightarrow}_{\mathbf{T}} t' \qquad \alpha \text{ not input}}{\langle s,t\rangle \stackrel{\alpha}{\rightarrow}_{\mathbf{S};\mathbf{T}} \langle s,t'\rangle} \qquad \frac{s \stackrel{\mathit{out}\, y}{\longrightarrow}_{\mathbf{S}} s' \qquad t \stackrel{\mathit{in}\, y}{\longrightarrow}_{\mathbf{T}} t'}{\langle s,t\rangle \stackrel{\tau}{\rightarrow}_{\mathbf{S};\mathbf{T}} \langle s',t'\rangle}$$

Example 1.2. For any set X, define an agent $\mathcal{I}_X$ of type $X \to X$ with states $X + \{\bot\}$, initial state $\bot$ and transitions $\bot \xrightarrow{in\ x} x$ and $x \xrightarrow{out\ x} \bot$, for all $x \in X$. $\mathcal{I}_X$ acts as a buffer of capacity one: A possible sequence of transitions is

$$\bot \xrightarrow{in\ x} x \xrightarrow{out\ x} \bot \xrightarrow{in\ y} y \xrightarrow{out\ y} \bot \xrightarrow{in\ z} z \xrightarrow{out\ z} \bot \ldots$$

Let $X = \{x\}$. Then $\mathcal{I}_X$ and $\mathcal{I}_X ; \mathcal{I}_X$ are the following agents:

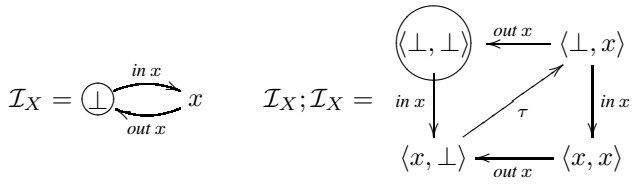

Here the initial state of each agent is circled. When representing agents in diagrams like these, it is often convenient to omit the names of the states, and to identify weakly bisimilar states. With that convention, we write:

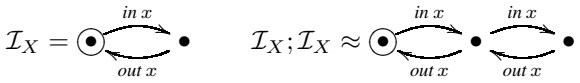

Note that $\mathcal{I}_X ; \mathcal{I}_X$ is a queue of capacity 2. In general, for any set Y, $\mathcal{I}_Y ; \mathcal{I}_Y$ is a first-in, first-out queue of capacity 2.

Two LTSs $\mathbf{S}$ and $\mathbf{T}$ of type A are ***isomorphic*** if there is a bijection between $|\mathbf{S}|$ and $|\mathbf{T}|$ preserving $\rightarrow$ and initial states.

Lemma 1.3. *1. Sequential Composition of labeled transition systems is associative up to isomorphism.*

2. Sequential Composition of agents respects both weak and strong bisimulation, i.e.

$$\frac{\mathbf{S}_1 \approx \mathbf{S}_2 \qquad \mathbf{T}_1 \approx \mathbf{T}_2}{\mathbf{S}_1 ; \mathbf{T}_1 \approx \mathbf{S}_2 ; \mathbf{T}_2} \qquad \text{and} \qquad \frac{\mathbf{S}_1 \sim \mathbf{S}_2 \qquad \mathbf{T}_1 \sim \mathbf{T}_2}{\mathbf{S}_1 ; \mathbf{T}_1 \sim \mathbf{S}_2 ; \mathbf{T}_2}$$

Unfortunately, agents do not form a category under sequential composition: there are no identity morphisms. In Section 1.4, we will introduce two categories of agents, one of which has unbounded buffers as its identity morphisms, and the other one queues.

1.3 Buffers and Queues

For any set X, let X^* be the free monoid and X^{**} the free commutative monoid generated by X. The elements of X^* are finite sequences. The empty sequence is denoted by ϵ. The elements of X^{**} are finite multisets. The empty multiset is denoted by $\emptyset$. We define the following agents of type $X \to_B X$:

1. The ***buffer*** $\mathcal{B}_X$ has states X^{**}, initial state $\emptyset$, and transitions $w \xrightarrow{in\ x} wx$ and $xw \xrightarrow{out\ x} w$, for all $w \in X^{**}$ and $x \in X$.

2. The ***queue*** $\mathcal{Q}_X$ has states X^*, initial state ϵ, and transitions $w \xrightarrow{in\ x} wx$ and $xw \xrightarrow{out\ x} w$, for all $w \in X^*$ and $x \in X$.

The only difference between the definitions of $\mathcal{B}_X$ and $\mathcal{Q}_X$ is whether the states are considered as sequences or multisets. We will write $\mathcal{B}$ and $\mathcal{Q}$ without subscript if X is clear from the context. $\mathcal{B}$ acts as an infinite capacity buffer which does not preserve the order of messages. For example, one possible sequence of transitions is

$$\emptyset \xrightarrow{in\ x} x \xrightarrow{in\ y} xy \xrightarrow{in\ z} xyz \xrightarrow{out\ y} xz \xrightarrow{out\ x} z \xrightarrow{in\ w} wz \ldots$$

$\mathcal{Q}$ acts as an infinite capacity first-in, first-out queue. A possible sequence of transitions is

$$\epsilon \xrightarrow{in\ x} x \xrightarrow{in\ y} xy \xrightarrow{out\ x} y \xrightarrow{in\ z} yz \xrightarrow{in\ w} yzw \xrightarrow{out\ y} zw \ldots$$

Lemma 1.4. *1. $\mathcal{B};\mathcal{B} \approx \mathcal{B}$ and $\mathcal{B};\mathcal{B} \not\sim \mathcal{B}$.*

2. $\mathcal{Q};\mathcal{Q} \approx \mathcal{Q}$ and $\mathcal{Q};\mathcal{Q} \not\sim \mathcal{Q}$.

3. $\mathcal{Q};\mathcal{B} \approx \mathcal{B}$ and $\mathcal{Q};\mathcal{B} \not\sim \mathcal{B}$.

4. If $|X| \geq 2$, then $\mathcal{B};\mathcal{Q} \not\approx \mathcal{B}$ and $\mathcal{B};\mathcal{Q} \not\approx \mathcal{Q}$.

The remainder of this paper is devoted to examining the effect of composing arbitrary agents with buffers and queues.

1.4 Notions of Asynchrony

In the asynchronous model of communication, messages are assumed to travel through a communication medium or *ether*. Sometimes, the medium is assumed to be first-in, first-out (a queue); sometimes, as in the asynchronous π-calculus, messages might be received in any order (a buffer).

Our approach is simple: we model the medium explicitly. An asynchronous agent is one whose output and/or input behaves as if filtered through either a buffer $\mathcal{B}$ or a queue $\mathcal{Q}$.

Definition 1.5. An agent $\mathbf{S}\colon X \to_B Y$ is

out-buffered	if	$\mathbf{S} \approx \mathbf{S};\mathcal{B}$	***out-queued***	if	$\mathbf{S} \approx \mathbf{S};\mathcal{Q}$
in-buffered	if	$\mathbf{S} \approx \mathcal{B};\mathbf{S}$	***in-queued***	if	$\mathbf{S} \approx \mathcal{Q};\mathbf{S}$
buffered	if	$\mathbf{S} \approx \mathcal{B};\mathbf{S};\mathcal{B}$	***queued***	if	$\mathbf{S} \approx \mathcal{Q};\mathbf{S};\mathcal{Q}$

We use the word ***asynchrony*** as a generic term to stand for any such property. Distinguishing these six different notions will allow us to study them separately. Yet another notion of asynchrony, incorporating feedback, will be defined in Section 3.2.

Remark. Because of Lemma 1.4, the operation of pre- or post-composing an agent with $\mathcal{B}$ or $\mathcal{Q}$ is idempotent up to $\approx$. Consequently, any agent of the form $\mathbf{S};\mathcal{B}$ is out-buffered, any agent of the form $\mathcal{B};\mathbf{S}$ is in-buffered, an agent is buffered iff it is in- and out-buffered, and so on. Also, each of the six properties is invariant under weak bisimulation.

Let B be a set. Buffered agents $\mathbf{S}\colon X \to_B Y$ form the morphisms of a category $\mathbf{Buf}_B$, whose objects are sets X, Y, etc.; the identity morphism on X is given by the buffer $\mathcal{B}_X$. Similarly, queued agents form a category $\mathbf{Que}_B$. These categories have a symmetric monoidal structure, which will be described in Section 3.1.

1.5 Examples

Example 1.6. The first example shows the effect of post-composing different agents with the buffer $\mathcal{B}$. Notice that although $\mathcal{B}$ has infinitely many states, $\mathbf{S};\mathcal{B}$ may have only finitely many states up to weak bisimulation.

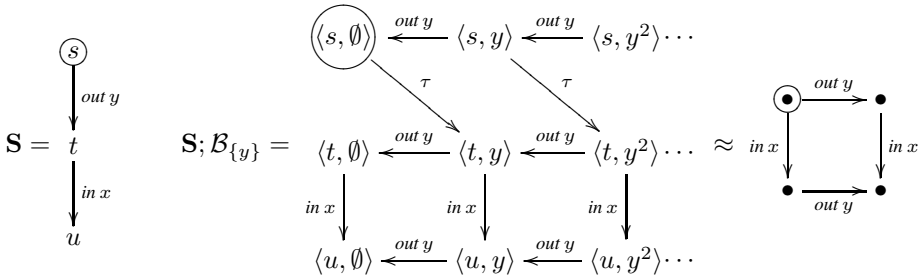

Example 1.7.

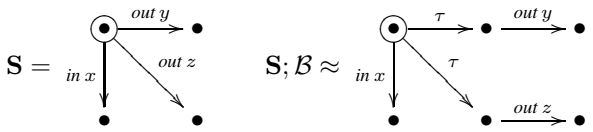

Example 1.8. Here is an example on in-bufferedness. Notice that an input action is possible at every state of $\mathcal{B};\mathbf{S}$.

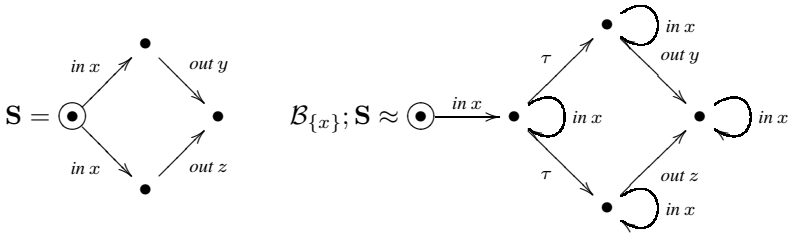

2 First-Order Axioms for Asynchrony

In this section, we will give necessary and sufficient conditions for each of the notions of asynchrony from Definition 1.5. These conditions take the form of *first-order axioms*, by which we mean axioms that use quantification only over states and actions, but not over subsets of states or actions. The axioms, which are shown in Tables 1 through 2, characterize each of our notions of asynchrony *up to weak bisimulation*; this means, an LTS is asynchronous iff it is weakly bisimilar to one satisfying the axioms. It is possible to lift the condition "up to weak bisimulation" at the cost of introducing second-order axioms; this is the subject of Section 6.

Table 1: First-order axioms for out-buffered agents

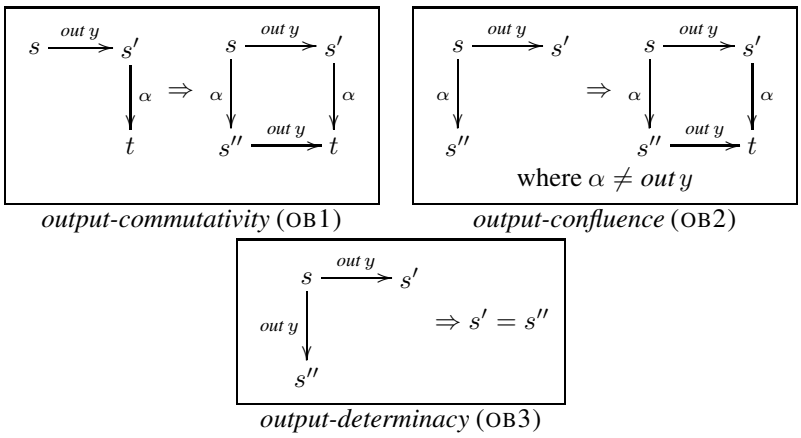

output-commutativity (OB1) *output-confluence* (OB2)

output-determinacy (OB3)

2.1 Out-Buffered Agents

Table 1 lists three axioms for out-buffered agents. We use the convention that variables are implicitly existentially quantified if they occur only on the right-hand-side of an implication, and all other variables are implicitly universally quantified. Thus the axioms are:

(OB1) *Output-commutativity*: output actions can always be delayed.

(OB2) *Output-confluence*: when an output action and some other action are possible, then they can be performed in either order with the same result. In particular, neither action precludes the other.

(OB3) *Output-determinacy*: from any state s, there is at most one transition $\mathit{out}\, y$ for each $y \in Y$.

Each of these axioms is plausible for the behavior of a buffer. Output-determinacy is maybe the least intuitive of the three properties; the idea is that once an output action is stored in a buffer, there is only one way of retrieving it. Together, these axioms characterize out-bufferedness up to weak bisimulation:

Theorem 2.1 (Characterization of out-buffered agents). *An agent* $\mathbf{S}$ *is out-buffered if and only if* $\mathbf{S} \approx \mathbf{T}$ *for some* $\mathbf{T}$ *satisfying* (OB1)–(OB3).

This is a direct consequence of the following proposition:

Proposition 2.2. *1. Every agent of the form* $\mathbf{S}; \mathcal{B}$ *satisfies* (OB1)–(OB3).

2. If $\mathbf{S}$ *satisfies* (OB1)–(OB3)*, then* $\mathbf{S} \approx \mathbf{S}; \mathcal{B}$.

Proof. 1. Clearly, the buffer $\mathcal{B}$ satisfies (OB1)–(OB3). Moreover, these conditions are preserved by arbitrary sequential composition from the left.

Table 2: First-order axioms for in-buffered agents

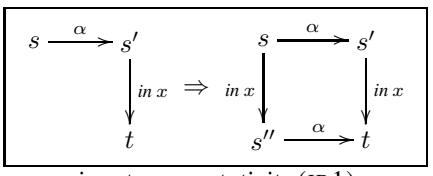

input-commutativity (IB1)

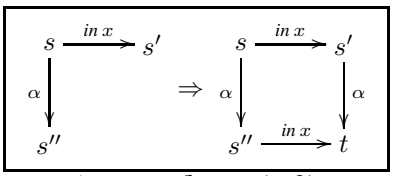

input-confluence (IB2)

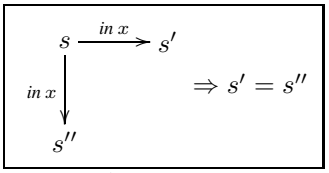

input-determinacy (IB3)

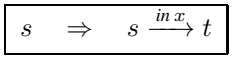

input-receptivity (IB4)

2. Suppose $\mathbf{S}\colon X \to_B Y$ satisfies (OB1)–(OB3). For a sequence $w = y_1 y_2 \cdots y_n \in Y^*$, we write $s \xrightarrow{out\ w} t$ if $s \xrightarrow{out\ y_1} \xrightarrow{out\ y_2} \cdots \xrightarrow{out\ y_n} t$ ($n \geq 0$). Note that if $w' \in Y^*$ is a permutation of w, then $s \xrightarrow{out\ w'} t$ iff $s \xrightarrow{out\ w} t$ by (OB1). Consider the relation $R \subseteq |\mathbf{S}| \times |\mathbf{S};\mathcal{B}|$ given by $sR\langle t, w\rangle$ iff $s \xrightarrow{out\ w} t$. Clearly, R relates initial states. We show that R is a weak bisimulation. In one direction, suppose

$$\begin{array}{ccc} s & R & \langle t, w\rangle \\ {\scriptstyle\alpha}\downarrow & & \\ s'. & & \end{array}$$

Two cases arise:

Case 1: $\alpha = out\, y$ for some $y \in w$. By the definition of R, $s \xrightarrow{out\ y} s'' \xrightarrow{out\ w'} t$, where $w = yw'$. By (OB3), we have $s' = s''$. Therefore $s'R\langle t, w'\rangle$, and also $\langle t, w\rangle \xrightarrow{\alpha} \langle t, w'\rangle$.✓

Case 2: $\alpha \neq out\, y$ for all $y \in w$. From $s \xrightarrow{out\ w} t$ and $s \xrightarrow{\alpha} s'$, we get $s' \xrightarrow{out\ w} t'$ and $t \xrightarrow{\alpha} t'$ by repeated application of (OB2). Therefore $s'R\langle t', w\rangle$ and $\langle t, w\rangle \stackrel{\alpha}{\Rightarrow} \langle t', w\rangle$ (notice the use of $\Rightarrow$ here, which is necessary in case α is an output action).✓

In the other direction, suppose

$$\begin{array}{ccc} s & R & \langle t, w\rangle \\ & & \downarrow{\scriptstyle\alpha} \\ & & \langle t', w'\rangle. \end{array}$$

We distinguish three cases for $\langle t, w\rangle \xrightarrow{\alpha} \langle t', w'\rangle$, depending on which rule in Definition 1.1 was used.

Case 1: $t \xrightarrow{\alpha} t'$, $w = w'$ and α not output. Then $s \xrightarrow{out\ w} t \xrightarrow{\alpha} t'$, which implies $s \xrightarrow{\alpha} s' \xrightarrow{out\ w} t'$ by repeated application of (OB1), *i.e.* $s \xrightarrow{\alpha} s'R\langle t', w\rangle$.✓

Case 2: $t = t'$, $w \xrightarrow{\alpha} w'$ and α not input. Since $\mathcal{B}$ has only input and output transitions, α must be *out* y for some $y \in Y$ with $w = yw'$. Then $s \xrightarrow{\textit{out}\, y} s' \xrightarrow{\textit{out}\, w'} t$, *i.e.* $s \xrightarrow{\alpha} s' R \langle t, w' \rangle$.✓

Case 3: $t \xrightarrow{\textit{out}\, y} t'$, $w \xrightarrow{\textit{in}\, y} w'$ and $\alpha = \tau$. In this case, $w' = wy$ and $s \xrightarrow{\textit{out}\, w} t \xrightarrow{\textit{out}\, y} t'$, hence $sR\langle t', w' \rangle$.✓ □

Remark 2.3. Theorem 2.1 generalizes to other notions of equivalence of processes, as long as they are coarser than weak bisimulation. Indeed, if $\cong$ is an equivalence of processes such that $\approx \subseteq \cong$, then for any agent $\mathbf{S}$, there exists some out-buffered $\mathbf{T}$ with $\mathbf{S} \cong \mathbf{T}$ iff there exists $\mathbf{T}'$ satisfying (OB1)–(OB3) and $\mathbf{S} \cong \mathbf{T}'$. This is a trivial consequence of Theorem 2.1. Similar remarks apply to the other results in this section and in Section 3.

2.2 In-Buffered Agents and Queues

The axioms for in-buffered agents are listed in Table 2. The main difference to the out-buffered case is the property *input-receptivity*: an in-buffered agent can perform any input action at any time. This was illustrated in Example 1.8. The input/output automata of Lynch and Stark [10] have this property, and so does Honda and Tokoro's original version of the asynchronous π-calculus [9].

Remark. Somewhat surprisingly, the axioms in Table 2 are not independent. In fact, (IB1) and (IB2) are equivalent in the presence of (IB3) and (IB4). We present all four axioms in order to highlight the analogy to the output case.

Theorem 2.4 (Characterization of in-buffered agents). *An agent* $\mathbf{S}$ *is in-buffered if and only if* $\mathbf{S} \approx \mathbf{T}$ *for some* $\mathbf{T}$ *satisfying* (IB1)–(IB4).

The axioms can be adjusted to accommodate queues rather than buffers: In (OB1) and (OB2), change the side conditions to "α not output". Change (OB3) to "if $s \xrightarrow{\textit{out}\, y} s'$ and $s \xrightarrow{\textit{out}\, z} s''$ then $y = z$ and $s' = s''$". In (IB1) and (IB2), change the side conditions to "α not input". Then the analogs of Theorems 2.1 and 2.4 hold.

3 More Constructors and Asynchrony with Feedback

3.1 Agent Constructors

In this section, we will introduce some operations on agents, such as renaming and hiding of actions, parallel composition and feedback.

1. *Domain extension.* If $\mathbf{S}$ is an LTS of type A, and if $A \subseteq A'$, then $\mathbf{S}$ can also be regarded as an LTS of type A'.

2. *Domain restriction (hiding).* If $\mathbf{S}$ is an LTS of type A, and if $\tau \in A' \subseteq A$, then $\mathbf{S}|_{A'}$ is defined to be the LTS of type A' which has the same states as $\mathbf{S}$, and whose transitions are those of $\mathbf{S}$ restricted to $|\mathbf{S}| \times A' \times |\mathbf{S}|$.

Domain extension and domain restriction are special cases of the following, general renaming construct:

3. *General renaming and hiding.* Let $\mathbf{S}$ be an LTS of type A and let $r \subseteq A \times A'$ be a relation such that $\tau r \alpha'$ iff $\tau = \alpha'$. Define $\mathbf{S}_r$ to be the LTS of type A' that has the same states and initial state as $\mathbf{S}$ and transitions $s \xrightarrow{\alpha}_{\mathbf{S}_r} t$ iff $s \xrightarrow{\alpha'}_{\mathbf{S}} t$ for some $\alpha r \alpha'$.

Let us now turn to various forms of parallel composition.

4. *Parallel composition without interaction.* Let $\mathbf{S}$ and $\mathbf{T}$ be LTSs of type A. Then $\mathbf{S}\|\mathbf{T}$ is the LTS of type A with states $|\mathbf{S}| \times |\mathbf{T}|$ and initial state $\langle s_0, t_0 \rangle$, and whose transitions are given by the rules

$$\frac{s \xrightarrow{\alpha}_{\mathbf{S}} s'}{\langle s,t \rangle \xrightarrow{\alpha}_{\mathbf{S}\|\mathbf{T}} \langle s',t \rangle} \qquad \frac{t \xrightarrow{\alpha}_{\mathbf{T}} t'}{\langle s,t \rangle \xrightarrow{\alpha}_{\mathbf{S}\|\mathbf{T}} \langle s,t' \rangle}.$$

5. *Symmetric monoidal structure.* Let $X \oplus X'$ be the disjoint union of sets. For $\mathbf{S}\colon X \to_B Y$ and $\mathbf{T}\colon X' \to_B Y'$, define $\mathbf{S} \oplus \mathbf{T}\colon X \oplus X' \to_B Y \oplus Y'$ to be the agent $\mathbf{S}_r \| \mathbf{T}_q$, where r and q are the inclusions of $X \to_B Y$, respectively $X' \to_B Y'$ into $X \oplus X' \to_B Y \oplus Y'$. Then $\oplus$ defines a symmetric monoidal structure on the categories **Buf** and **Que**. The tensor unit is given by the agent $\mathbf{I}$ of type $\emptyset \to \emptyset$ with one state and no transitions.

The constructors we have considered so far, including sequential composition, are not sufficient to build arbitrary networks. What is missing is the ability to construct loops. The next constructor allows the output of an agent to be connected to its own input:

6. *Self-composition (feedback).* Let $\mathbf{S}\colon X \to_B Y$. Let $O \subseteq Y \times X$ be a set of pairs. Define $\mathbf{S} \circlearrowleft O$, the self-composition of $\mathbf{S}$ along O, to be the LTS of type $X \to_B Y$ whose states are identical with those of $\mathbf{S}$, and whose transitions are given by the rules

$$\frac{s \xrightarrow{\alpha}_{\mathbf{S}} t}{s \xrightarrow{\alpha}_{\mathbf{S} \circlearrowleft O} t} \qquad \frac{s \xrightarrow{\textit{out } y} \xrightarrow{\tau} \xrightarrow{\textit{in } x}_{\mathbf{S}} t \qquad \langle y, x \rangle \in O}{s \xrightarrow{\tau}_{\mathbf{S} \circlearrowleft O} t}.$$

   In the common case where $\mathbf{S}\colon X \to_B X$ and $O = \{\langle x, x \rangle \mid x \in X\}$, we will write $\mathbf{S}^\circ$ instead of $\mathbf{S} \circlearrowleft O$.

We can use self-composition to define both sequential and parallel composition.

7. *Sequential composition.* The sequential composition of agents was defined in Definition 1.1. Alternatively, one can define it from the more primitive notions of direct sum, feedback and hiding: Let $\mathbf{S}\colon X \to_B Y$ and $\mathbf{T}\colon Y \to_B Z$. Then $\mathbf{S} \oplus \mathbf{T}\colon X \oplus Y \to_B Y \oplus Z$, and with $\Delta Y = \{\langle y, y \rangle \mid y \in Y\}$, one gets $\mathbf{S}; \mathbf{T} \approx ((\mathbf{S} \oplus \mathbf{T}) \circlearrowleft \Delta Y)|_{X \to_B Z}$.

8. *Parallel composition (with interaction).* Let $\mathbf{S}, \mathbf{T}\colon X \to_B X$. The parallel composition $\mathbf{S}|\mathbf{T}$ is defined to be the agent $(\mathbf{S}\|\mathbf{T})^\circ$.

Proposition 3.1. *All of the agent constructors in this section respect weak bisimulation. For instance, if* $\mathbf{S} \approx \mathbf{S}'$ *and* $\mathbf{T} \approx \mathbf{T}'$*, then* $\mathbf{S}_r \approx \mathbf{S}'_r$ *and* $\mathbf{S}\|\mathbf{T} \approx \mathbf{S}'\|\mathbf{T}'$*, etc.*

Table 3: First-order axioms for out-buffered agents with feedback

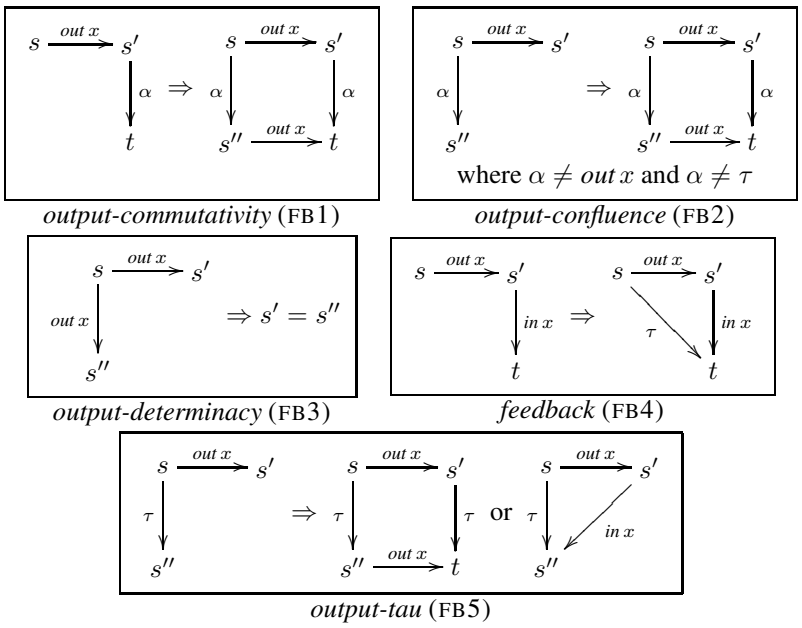

output-commutativity (FB1) *output-confluence* (FB2)

output-determinacy (FB3) *feedback* (FB4)

output-tau (FB5)

3.2 Asynchrony with Feedback

In concurrent process calculi such as CCS or the π-calculus, messages that are emitted from a process are immediately available as input to all processes, including the sending process itself. In our setting, this is best modeled by requiring that all processes are of type $X \to X$ for one fixed set X, and by using self-composition to feed the output back to the input.

In the presence of feedback, out-bufferedness takes a slightly different form, which is expressed in the following definition.

Definition. An agent $\mathbf{S}\colon X \to_B X$ is ***out-buffered with feedback*** if $\mathbf{S} \approx \mathbf{R}^\circ$ for some out-buffered agent $\mathbf{R}$.

Example 3.2. The following agent $\mathbf{S}$ is out-buffered with feedback but not out-buffered:

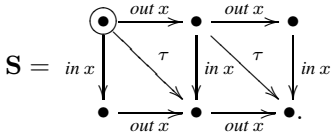

Remark. Recently, Amadio, Castellani and Sangiorgi [3] have given a definition of asynchronous bisimulation, which accounts for the fact that an agent of type $X \to X$

Table 4: Transitions for asynchronous CCS

$$(act)\ \frac{}{\alpha.P \xrightarrow{\alpha} P} \qquad (sum)\ \frac{G \xrightarrow{\alpha} P}{G + G' \xrightarrow{\alpha} P} \qquad (sum')\ \frac{G' \xrightarrow{\alpha} P}{G + G' \xrightarrow{\alpha} P}$$

$$(comp)\ \frac{P \xrightarrow{\alpha} P'}{P|Q \xrightarrow{\alpha} P'|Q} \qquad (comp')\ \frac{Q \xrightarrow{\alpha} Q'}{P|Q \xrightarrow{\alpha} P|Q'}$$

$$(synch)\ \frac{P \xrightarrow{\alpha} P' \qquad Q \xrightarrow{\bar{\alpha}} Q'}{P|Q \xrightarrow{\tau} P'|Q'} \qquad (res)\ \frac{P \xrightarrow{\alpha} P' \qquad \alpha \notin L \cup \bar{L}}{P \setminus L \xrightarrow{\alpha} P' \setminus L}$$

$$(rel)\ \frac{P \xrightarrow{\alpha} P'}{P[f] \xrightarrow{f\alpha} P'[f]} \qquad (rec)\ \frac{P \xrightarrow{\alpha} P' \qquad A \stackrel{\text{def}}{=} P}{A \xrightarrow{\alpha} P'}$$

might receive a message, and then immediately send it again, without this interaction being observable on the outside. Feedback is concerned with the dual phenomenon, namely a process that sends a message and then immediately receives it again.

Out-bufferedness with feedback is characterized up to weak bisimulation by the first-order axioms that are listed in Table 3.

Theorem 3.3 (Characterization of out-buffered agents with feedback).
An agent $\mathbf{S}\colon X \to_B X$ *is out-buffered with feedback if and only if* $\mathbf{S} \approx \mathbf{T}$ *for some agent* $\mathbf{T}$ *satisfying* (FB1)–(FB5).

4 Example: Asynchronous CCS

In this section, we will show that an asynchronous version of Milner's Calculus of Communicating Systems (CCS) [11, 12] fits into the framework outlined in the previous section of out-buffered labeled transition systems with feedback.

Let $X = \{a, b, c, \dots\}$ be an infinite set of ***names***, and let $\bar{X} = \{\bar{a}, \bar{b}, \bar{c}, \dots\}$ be a corresponding set of ***co-names***, such that X and $\bar{X}$ are disjoint and in one-to-one correspondence via $\bar{}$. We also write $\bar{\bar{a}} = a$. Names correspond to input-actions, and co-names to output-actions. Let $\tau \notin X + \bar{X}$, and let $Act = X + \bar{X} + \{\tau\}$ be the set of ***actions***, ranged over by the letters $\alpha, \beta, \dots$; Let the letter L range over subsets of X, and write $\bar{L}$ for $\{\bar{a} \mid a \in L\}$. Let the letter f range over ***relabeling functions***, which are functions $f : X \to X$. Any relabeling function extends to $f : Act \to Act$ by letting $f\bar{a} = \overline{fa}$ and $f\tau = \tau$.

Let $A, B, C, \dots$ range over a fixed set of ***process constants***. Asynchronous CCS ***processes*** $P, Q, \dots$ and ***guards*** $G, H, \dots$ are given by the following grammars:

$$P ::= \bar{a}.0 \mid P|P \mid P \setminus L \mid P[f] \mid A \mid G$$

$$G ::= a.P \mid \tau.P \mid G + G \mid \mathbf{0}$$

Assume a set of ***defining equations*** $A \stackrel{\text{def}}{=} P$, one for each process constant A. The operational semantics of asynchronous CCS is given in terms of a labeled transition system $\mathbf{S}_{\text{CCS}} = \langle S, Act, \rightarrow \rangle$, which is defined in Table 4. The states are CCS processes. Notice that we have not specified a distinguished initial state; this is more convenient in this context, and no harm is done. Also notice that there is no rule for $\mathbf{0}$. This is because the process $\mathbf{0}$ is inert, *i.e.* there are no transitions $\mathbf{0} \stackrel{\alpha}{\rightarrow} P$.

Theorem 4.1. *The labeled transition system* $\mathbf{S}_{\text{CCS}}$ *is out-buffered with feedback.*

5 Example: The Core Join Calculus

The join calculus was introduced by Fournet and Gonthier in [7] and further developed in [8]. It is a concurrent, message passing calculus like the π-calculus. However, the reaction rule is simpler and closer to the semantics of a chemical abstract machine. Here, we will only be concerned with the *core* join calculus.

Let $x, y, \ldots$ range over a countable set $\mathcal{N}$ of ***names***. Let $\tilde{x}, \tilde{y}, \ldots$ range over sequences of names. Core join calculus ***processes*** $P, Q, \ldots$ and ***rules*** $R, S, \ldots$ are given by the following grammars:

$$P ::= x\langle \tilde{y} \rangle \mid P|P \mid \mathbf{def}\ R_1 \wedge \ldots \wedge R_m\ \mathbf{in}\ P \qquad R ::= x_1(\tilde{v}_1)| \ldots |x_n(\tilde{v}_n) \rhd P$$

A process of the form $x\langle \tilde{v} \rangle$ is called a ***message***. In the rule $R = x_1(\tilde{v}_1)| \ldots |x_n(\tilde{v}_n) \rhd P$, the names $\tilde{v}_1 \ldots \tilde{v}_n$ are bound, and they are assumed to be distinct. The names $x_1 \ldots x_n$ are called the ***defined names*** of R, denoted $dn(R)$. Finally, all of the defined names of $R_1, \ldots, R_m$ are bound in the process $\mathbf{def}\ R_1 \wedge \ldots \wedge R_m\ \mathbf{in}\ P$. For a more comprehensive treatment, see [7, 8].

The semantics of the core join calculus is given in the style of a chemical abstract machine. A ***state*** $\Delta \vdash_N \Pi$ is a multiset Δ of rules together with a multiset Π of processes. N is a set of names, such that $fn(\Delta, \Pi) \subseteq N$. We identify states up to α-equivalence, *i.e.* up to renaming of bound variables. The transitions of this machine follow a simple idea: the processes on the right hand side evolve according to the rules on the left-hand side. There are two kinds of transitions: ***structural*** transitions, denoted $\rightharpoonup$, and ***reactions***, denoted $\mapsto$:

$$(str1) \qquad \Delta \vdash_N \Pi, P|Q \quad \rightharpoonup \quad \Delta \vdash_N \Pi, P, Q$$

$$(str2) \qquad \Delta \vdash_N \Pi, \mathbf{def}\ R_1 \wedge \ldots \wedge R_m\ \mathbf{in}\ P \quad \rightharpoonup \quad \Delta, R_1, \ldots, R_m \vdash_{N'} \Pi, P$$
$$\text{where } N' = N + dn(R_1, \ldots, R_m)$$

$$(join) \qquad \Delta \vdash_N \Pi, x_1\langle \tilde{y_1} \rangle, \ldots, x_n\langle \tilde{y_n} \rangle \quad \mapsto \quad \Delta \vdash_N \Pi, [\tilde{y_1}/\tilde{v}_1, \ldots, \tilde{y_n}/\tilde{v}_n]P$$
$$\text{where } (x_1(\tilde{v}_1)| \ldots |x_n(\tilde{v}_n) \rhd P) \in \Delta$$

The rule $(join)$ is of course only applicable is the length of $\tilde{y_i}$ and $\tilde{v}_i$ are the same, for all i. Note that in the rule $(str2)$, the sets N and $dn(R_1, \ldots, R_m)$ must be disjoint; this may necessitate renaming some bound variables in $\mathbf{def}\ R_1 \wedge \ldots \wedge R_m\ \mathbf{in}\ P$.

Table 5: Second-order axioms for out-buffered agents

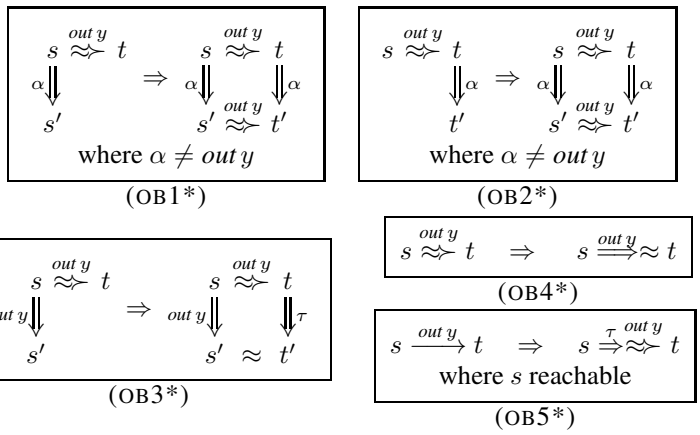

Remark. In the original formulation of the join calculus [7, 8], the structural rules are assumed to be reversible. We adopt a different convention here.

To fit the join calculus into our framework, we make it into a labeled transition system with input and output. Let $X = \{x\langle\tilde{y}\rangle \mid x \in \mathcal{N}, \tilde{y} \in \mathcal{N}^*\}$ be the set of messages. We add input and output transitions:

$$(in) \qquad \Delta \vdash_N \Pi \quad \xrightarrow{in\, x\langle\tilde{y}\rangle} \quad \Delta \vdash_{N \cup \{x,\tilde{y}\}} \Pi, x\langle\tilde{y}\rangle$$

$$(out) \quad \Delta \vdash_N \Pi, x\langle\tilde{y}\rangle \quad \xrightarrow{out\, x\langle\tilde{y}\rangle} \quad \Delta \vdash_N \Pi$$

Further, we let $\xrightarrow{\tau} = \rightharpoonup \cup \mapsto$. With these definitions, the join calculus defines a labeled transition system $\mathbf{S}_{\text{join}} \colon X \to X$.

Theorem 5.1. *The labeled transition system* $\mathbf{S}_{\text{join}}$ *defined by the core join calculus is out-buffered with feedback.*

6 Other Characterizations of Asynchrony

In Sections 2 and 3, we have characterized notions of asynchrony by first-order axioms *up to weak bisimulation*. It is possible to remove the words "up to weak bisimulation", *i.e.* to characterize asynchrony directly. This happens at the cost of introducing second-order axioms. The shift to second-order seems to be inevitable, since weak bisimulation itself is a second-order notion.

The axioms for out-buffered agents are given in Tables 5. It is possible to give corresponding axioms for in-bufferedness, out-queuedness and in-queuedness.

Theorem 6.1. *An agent* $\mathbf{S} \colon X \to_B Y$ *is out-buffered if and only if for each* $y \in Y$ *there exists a binary relation* $\overset{out\, y}{\approx\!\!>} \subseteq |\mathbf{S}| \times |\mathbf{S}|$ *satisfying* (OB1*)–(OB5*).

7 Conclusions and Future Work

We have shown how to abstractly characterize various notions of asynchrony in a general-purpose framework, independently of any particular process paradigm. This can be done by first-order axioms up to weak bisimulation, or by higher-order axioms "on the nose". The present framework of labeled transition systems with input and output can be used to model asynchronous communication in CCS, as well as the join calculus. To give an adequate treatment of calculi with explicit, dynamic scoping operators, such as the π-calculus, one should equip these labeled transition systems with the ability to handle dynamically created names. Work is in progress on a notion of fibered labeled transition system that can be used to model this more general situation.

References

[1] S. Abramsky. Interaction categories and communicating sequential processes. In A. W. Roscoe, editor, *A Classical Mind: Essays in honour of C. A. R. Hoare*, pages 1–16. Prentice Hall International, 1994.

[2] S. Abramsky, S. Gay, and R. Nagarajan. Interaction categories and typed concurrent programming. In *Proceedings of the 1994 Marktoberdorf Summer School*. Springer, 1994.

[3] R. M. Amadio, I. Castellani, and D. Sangiorgi. On bisimulations for the asynchronous π-calculus. In *CONCUR '96*, Springer LNCS 1119, pages 147–162, 1996.

[4] M. A. Bednarczyk. *Categories of asynchronous systems*. PhD thesis, University of Sussex, 1988.

[5] G. Berry and G. Boudol. The chemical abstract machine. *Theoretical Computer Science*, 96:217–248, 1992.

[6] G. Boudol. Asynchrony and the π-calculus. Technical Report 1702, INRIA, Sophia-Antipolis, 1992.

[7] C. Fournet and G. Gonthier. The reflexive cham and the join-calculus. In *POPL '96*, 1996.

[8] C. Fournet, G. Gonthier, J.-J. Levy, L. Maranget, and D. Remy. A calculus of mobile agents. In *CONCUR '96*, Springer LNCS 1119, pages 406–421, 1996.

[9] K. Honda and M. Tokoro. An object calculus for asynchronous communication. In *Proc. ECOOP 91, Geneve*, 1991.

[10] N. A. Lynch and E. W. Stark. A proof of the Kahn principle for input/output automata. *Information and Computation*, 82:81–92, 1989.

[11] R. Milner. *A Calculus of Communicationg Systems*. Springer LNCS 92. 1980.

[12] R. Milner. Operational and algebraic semantics of concurrent processes. Technical report, University of Edinburgh, Nov. 1987. Chapter for the Handbook of Theoretical Computer Science.

[13] C. Palamidessi. Comparing the expressive power of the synchronous and the asynchronous π-calculus. In *POPL '97 (Paris)*, 1997.

[14] M. W. Shields. Concurrent machines. *Theoretical Computer Science*, 28:449–465, 1985.