Bisimulation Can't Be Traced

Bard Bloom*
Cornell University
bard@cs.cornell.edu

Sorin Istrail†
Sandia National Laboratories
scistra@cs.sandia.gov

Albert R. Meyer‡
MIT Lab. for Computer Sci.
meyer@theory.lcs.mit.edu

July 21, 2003

Abstract

In the concurrent language CCS, two programs are considered the same if they are *bisimilar*. Several years and many researchers have demonstrated that the theory of bisimulation is mathematically appealing and useful in practice. However, bisimulation makes too many distinctions between programs. There are two straightforward programs which are not bisimilar, but nonetheless are interchangeable: one may be substituted for the other anywhere in a CCS program and no difference can be seen. Bisimulation is thus not *fully abstract*.

We consider the problem of adding operations to CCS to make bisimulation fully abstract. It is trivial to add an operation achieving full abstraction, but this operation is rather peculiar. We show that bisimulation is not fully abstract with respect to any extension of CCS by *CCS-like* operations. We give a formal description of "CCS-like," as *GSOS* and argue by proofs and counterexamples that this is indeed the right class.

In the proof of non-full-abstraction a coarser variant of bisimulation arises, a notion called *ready simulation*. We investigate the theory of ready simulation, showing that it possesses the basic properties which make bisimulation attractive. Like bisimulation, it possesses equivalent relational and logical characterizations; it has two additional equivalent characterizations as congruence with respect to all GSOS languages, and with respect to CCS extended by *process copying* and *controlled communication* operations. As a corollary, we show that bisimulation cannot be fully abstract with respect to any CCS-like languages, observing traces.

1 Introduction

One of the most basic things that a programming-language semantics should give is a notion of *program equivalence*: a statement telling when two programs do the same thing.

*Supported by an NSF Fellowship, also NSF Grant No. 8511190–DCR and ONR grant No. N00014–83–K–0125.

†Supported by NSF Grant CCR-881174

‡Supported by NSF Grant No. 8511190–DCR and by ONR grant No. N00014–83–K–0125.

Frequently, there are many choices of a notion of program equivalence, and it is not clear how to choose among them. We attempt to give some criteria for selecting one notion over another.

Two concurrent programming languages, Milner's CCS [**?**, **?**, **?**] and Hoare's CSP [**?**, **?**], share the premise that the meaning of a process is fully determined by a *synchronization tree*, namely, a rooted, unordered tree whose edges are labeled with symbols denoting basic *actions* or events. These trees are typically specified by a Structured Operational Semantics (SOS) in the style of [**?**] or by some other effective description, and so are in fact recursively enumerable trees. Both theories further agree that synchronization trees are an *over*specification of process behavior, and certain distinct trees must be regarded as equivalent processes. The notable difference in the theories is that bisimulation yields finer distinctions among synchronization trees; that is, CSP identifies trees which CCS considers different, but not conversely.

In CSP, process distinctions can be understood as based on observing *traces*, namely, *maximal* sequences of visible actions performed by a process. Two trees are *trace equivalent* iff they have the same set of traces. Given any set of operations on trees, *trace congruence* is defined to be the coarsest congruence with respect to the operations which refines trace equivalence. Thus, two CSP processes P and Q are distinguished iff there is some CSP context $C[X]$ and string s such that only one of $C[P]$ and $C[Q]$ has s as a trace. This explanation of when two synchronization trees are to be identified is thoroughly elaborated in Hennessy and DeNicola's *test equivalence* system [**?**]. On the other hand, two CCS processes are distinguished according to an "interactive" game-like protocol called *bisimulation*. Indistinguishable CCS processes are said to be *bisimilar*.

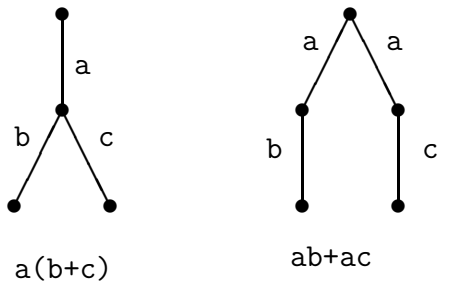

Figure 1: *trace equivalent but not trace congruent.*

A standard example is the pair of trees of Figure **??**, $a(b + c)$ and $(ab + ac)$, which are trace equivalent, but not CSP trace congruent, *viz.*, in CSP (and also CCS) they are distinct processes. Similarly, the trees of Figure **??**,

$$(abc + abd) \text{ and } a(bc + bd) \tag{1}$$

are CSP trace congruent but not bisimilar, *viz.*, equal in CSP but considered distinct in CCS

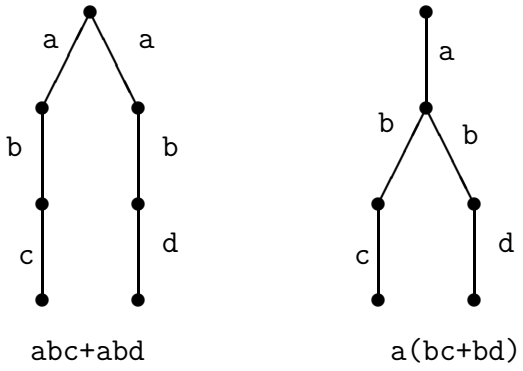

Figure 2: *CSP trace congruent but not bisimilar.*

[**?**, **?**]. The trace-based approach is developed in [**?**, **?**, **?**, **?**]. Bisimulation-based systems include [**?**, **?**, **?**, **?**, **?**, **?**, **?**].

The idea of a "silent" (aka "hidden" or "τ-") action plays an important role in both CSP and CCS theories, but creates significant technical and methodological problems; *e.g.*, bisimulation ignoring silent actions is not a congruence with respect to CCS. In this paper we assume for simplicity that *there is no silent action*. Preliminary investigations indicate that our conclusions about bisimulation generally apply to the case with silent moves; this is a matter of ongoing study.

In the absence of silent action, bisimulation is known to be a congruence with respect to all the operations of CSP/CCS, and Milner has argued extensively that in this case bisimulation yields the finest appropriate notion of the behavior of concurrent processes based on synchronization trees. Although there is some ground for refining synchronization trees further (*cf.* [**?**, **?**]), we shall accept the thesis that bisimilar trees should not be distinguished. Thus, we admit below only operations with respect to which bisimulation remains a congruence. Since bisimilar trees are easily seen to be trace equivalent, it follows in this setting that bisimulation refines any trace congruence. Our results focus on the converse question of whether further identifications should be made, *i.e.*, whether nonbisimilar processes are truly distinguishable in their observable behavior.

We noted that a pair of nonbisimilar trees P and Q can be distinguished by an "interactive" protocol. The protocol *itself* can be thought of as a new process Bisim$[P, Q]$. One might suppose that in a general concurrent programming language, it would be possible to define the new process too, and that success or failure of Bisim$[\cdot, \cdot]$ running on a pair P, Q would be easily visible to an observer who could observe traces.

However, CSP and CCS operations are very similar, and the example of Figure **??** above shows that bisimulation is a strictly finer equivalence than trace congruence with respect to CSP/CCS operations. It follows that the contexts Bisim$[\cdot, \cdot]$ distinguishing nonbisimilar processes by their traces *are not definable using the standard CSP/CCS operations*; if they

were, nonbisimilarity could be reduced to trace distinguishability. Namely, any pair of nonbisimilar trees P and Q would also be trace distinguishable by plugging them in for X in Bisim$[X, P]$ and observing the "success" trace when P is plugged in, but not when Q is plugged in.

Thus, we maintain that implicit in concurrent process theory based on bisimulation is another "interactive" kind of *meta*process, which the formalisms of CSP/CCS — languages proposed as bases for understanding interactive processes — are inadequate to define! The central question of this paper [**?**, **?**] is

> *What further operations on CCS/CSP terms are needed so that protocols reducing nonbisimilarity to trace distinguishability become definable?*

In the remainder of the paper, we argue that bisimulation *cannot* be reduced to a trace congruence with respect to any *reasonably structured* system of process constructing operations. The implications of this conclusion are discussed in the final Section **??**.

In particular, we formulate in Section **??** a general notion of a system of processes given by structured rules for transitions among terms — a GSOS system.[1] Almost all previously formulated systems of bisimulation-respecting[2] operations are definable by GSOS rule systems; the exception, in [**?**] and later papers, have been explorations of the power of structured operational semantics. Even rules with negative antecedents are allowed in GSOS systems. On the other hand, we indicate in Section **??** that any of the obvious further relaxations of the conditions defining GSOS's can result in systems which are ill-defined, countably branching, or fail to respect bisimulation. Thus, we believe that GSOS definability provides a generous and mathematically invariant constraint on what a reasonably structured system of processes might be. We therefore restrict our attention to GSOS languages.

Definition 1.1 *Two processes are* GSOS trace congruent *iff they yield the same traces in all GSOS definable contexts.*

Our first main result is that bisimulation—even restricted to finite trees—is a *strict* refinement of GSOS trace congruence. Specifically, we develop in Section **??** a characterization of GSOS trace congruence similar to the standard characterization of bisimulation and use it to prove:

Theorem 1.2 *The nonbisimilar trees $a(bc+bd)$ and $a(bc+bd)+abc$ (Figure **??**) are GSOS trace congruent.*

We remark that GSOS congruence is a strict refinement of CSP congruence. A map of the equivalences used in this paper is given in Figure **??**; the higher equivalences are finer than the lower ones. More detail on the collection of process equivalences based on synchronization trees can be found in [**?**, **?**, **?**].

[1] Originally the "G" in "GSOS" stood for "guarded recursion." We argue in Section **??** that guarded recursion is a less than essential feature for our purposes. However, the acronym has been used in too many places by too many authors to be easily changed.

[2] A language $\mathcal{L}$ *respects* an equivalence $\sim$ if $\sim$ refines $\mathcal{L}$-congruence; that is, whenever $P \sim Q$, then P and Q are $\mathcal{L}$-congruent.

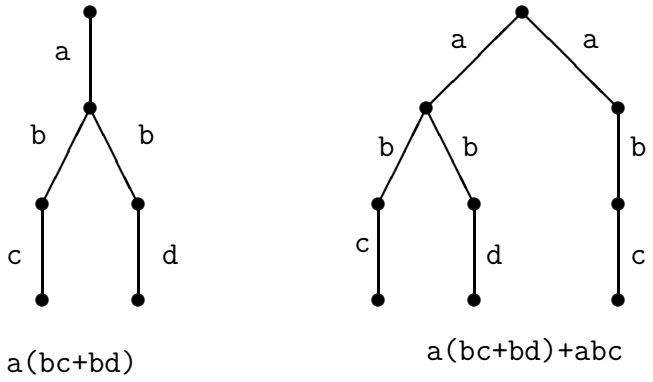

Figure 3: *Ready similar but not bisimilar.*

Theorem 1.3 *The processes $aa+ab$ and $aa+ab+a(a+b)$ (see Figure **??**) are CSP trace congruent [**?**, axiom (D5), p. 99] but not GSOS trace congruent.*

GSOS congruence has a theory which rivals that of bisimulation in its richness. It enjoys several equivalent definitions, of which the formulation as congruence with respect to all GSOS languages is perhaps the least important; the name "*ready simulation*" has been proposed. Section **??** gives a characterization of ready simulation in bisimulation-like terms. In Section **??**, we present a modal characterization of ready simulation similar to that of Hennessy and Milner in [**?**], and in Section **??** we show that it is congruence with respect to an arguably reasonable language.

Abramsky [**?**] independently raised the question of how to test distinguishability of nonbisimilar processes and formalized the operational behavior of a set of protocols which do capture bisimulation. In [**?**, **?**] we offer a similar system for the task, slightly improved in certain respects. Our thesis that no reasonably structured system — that is, defined by structured rules in the style of CCS — can capture bisimulation implies that both these systems must lack some important structured features. This paper gives a theorem demonstrating one instance of this thesis; [**?**] gives another instance.

2 Preliminaries

There are several possible ways to understand CCS; we will take a fairly denotational approach. Terms are a notation for *synchronization trees*, which are essentially infinite-state loop-free automata. We use some nonempty, finite set Act of *actions*, which are not given any further interpretation. We use a, b, c as constant actions and as metavariables ranging over actions. Unlike Milner's original CCS, [**?**], our general theory does not require any algebraic structure on the set of actions; the algebraic structure can be encoded in

Synchronization Tree Isomorphism	$P \equiv Q$	The finest acceptable process equivalence in this theory. We insist that processes with the same tree be identified, and that all languages respect tree isomorphism. This relation is *too fine*; we want the non-isomorphic processes a and $a + a$ to be identified.
Bisimulation	$P \leftrightarrows Q$	Generally accepted in this community to be the *finest* acceptable process equivalence; that is, if P and Q are bisimilar, then P and Q ought to be considered identical. Solves the a vs. $a+a$ problem of tree isomorphism. In this study we argue that bisimulation is too fine.
Ready Simulation (GSOS Congruence) ($\frac{2}{3}$-bisimulation)	$P \ \ Q$	Congruence with respect to all well-structured languages. In this paper, we present an introduction to the theory of ready simulation, suggesting that it is a mathematically appealing alternative to bisimulation, and unlike bisimulation makes only computationally meaningful distinctions.
Trace Congruence	$P \equiv_{tr}^{\mathcal{L}} Q$	Trace equivalent in all $\mathcal{L}$-contexts.
CSP Congruence (Failures Equivalence)	$P \equiv_{tr}^{\text{CSP}} Q$	Congruence with respect to CSP. This enters our discussion only incidentally.
Trace Equivalence (Automaton Equivalence)	$P \equiv_{tr} Q$	P and Q have the same finite completed traces. Not an adequate semantics for most languages.

Figure 4: Equivalences used in this study.

natural ways if desired. Note that we are trying to work in as finite a setting as possible; our action set, unlike Milner's, must be finite.[3]

Definition 2.1 *A* synchronization tree *is a rooted, unordered, finitely branching, possibly countably deep tree with edges labeled by elements of* Act. *A* countably-branching synchronization tree *is a rooted, unordered, possibly countably wide and deep tree with* Act*-labeled edges. We call synchronization trees "finitely branching" when we wish to emphasize the distinction between finitely and countably branching trees.*

Definition 2.2 *The tree P' is a* $\left\{ \begin{array}{c} \textit{child} \\ \textit{a-child} \\ \textit{descendant} \\ \textit{s-descendant} \end{array} \right\}$ *of P if there is* $\left\{ \begin{array}{c} \textit{an arc} \\ \textit{an arc labeled } a \\ \textit{a path} \\ \textit{a path labeled } s \end{array} \right\}$

[3]As discussed in more detail in Section **??**, we are interested in *distinguishing* processes rather than *expressing* them; distinctions beween processes should be observable with a finite amount of computation; in any reasonable setting, this will use only a finite number of actions.

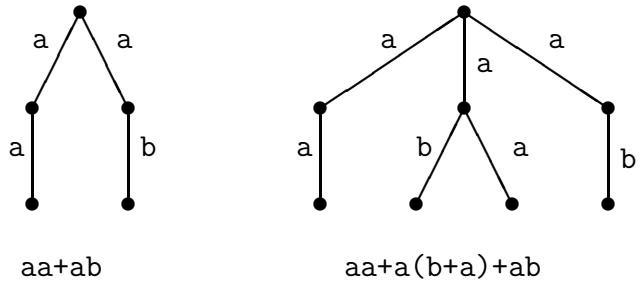

Figure 5: *CSP trace congruent but not ready similar.*

from the root of P to the root of P', where $a \in \mathsf{Act}$ and $s \in \mathsf{Act}^*$. If P' is an a-child (resp. s-descendant) of P, we write $P \xrightarrow{a} P'$ (resp. $P \xrightarrow{s} P'$).

A set Δ of trees is downward closed *if whenever $P \in \Delta$, all descendants of P are in Δ.*

It is clear how to consider a tree as an infinite-state nondeterministic automaton; the root of the tree is the start state of the automaton, and on each step it selects an edge and performs the action labeling that edge. Conversely, given such an automaton, it can be unwound into a (finitely or countably branching) synchronization tree in an obvious way.[4]

2.1 Bisimulation

Bisimulation is a pure, mathematical notion; it is independent of the language in question. Despite this, we will see that it is an adequate semantics for any well-structured language. To give a hint of the motivation of bisimulation (and because it is useful in later work) we mention an even finer adequate semantics first.

Definition 2.3 *If P or Q are synchronization trees, $P \equiv Q$ if P and Q are isomorphic as unordered edge-labeled trees.*

Although the synchronization tree semantics $\equiv$ of CCS is adequate and simple to think about, it is not the right semantics. Consider the two processes a and $a + a$ of Figure **??**. Each can only perform an a-action and then stop; $a + a$ can do that in two ways. It is generally agreed that all ways of performing a-actions and all stopped processes are the same, and so there should be no distinction between a and $a + a$. So, there is general agreement that synchronization trees are an overspecification of process behavior, and certain trees must be regarded as equivalent. The question, then, is which trees to identify. In CCS, Milner chose to identify precisely the *bisimilar* trees.

[4]Trees are used because they simplify certain definitions. Further, most core languages (without recursion) give a notation for all finite trees, but not for all finite automata.

Figure 6: a and $a + a$

As befits a good mathematical notion, there are several equivalent definitions of bisimulation.[5] The theory of bisimulation has been extensively studied [**?**]; we present only selected highlights.

Definition 2.4 *A relation $\leftrightarrows$ between synchronization trees is a* strong bisimulation relation *if, whenever $P \leftrightarrows Q$ and $a \in \mathsf{Act}$, then*

- *Whenever $P \xrightarrow{a} P'$, then for some Q', $Q \xrightarrow{a} Q'$ and $P' \leftrightarrows Q'$.*
- *Whenever $Q \xrightarrow{a} Q'$, then for some P', $P \xrightarrow{a} P'$ and $P' \leftrightarrows Q'$.*

For example, synchronization tree isomorphism and the null relation are both strong bisimulation relations. A more informative strong bisimulation relation is $\leftrightarrows_0$, where $P \leftrightarrows_0 Q$ iff P and Q are isomorphic, or if P and Q are isomorphic to a and $a + a$ respectively.

Definition 2.5 *P and Q are* bisimilar*, written $P \leftrightarrows Q$, iff there is a strong bisimulation relation $\leftrightarrows$ such that $P \leftrightarrows Q$.*

For example, $\leftrightarrows_0$ shows that $a \leftrightarrows a + a$. Indeed $P \leftrightarrows P + P$, where $P + P$ is two copies of P joined at the root. The relation $\leftrightarrows$ is itself a strong bisimulation relation, and in fact the largest one. It is an equivalence relation, and even a congruence relation with respect to all the operations of CCS.

One of the other characterizations of bisimulation will be particularly important for this study. The following logical characterization holds only for finitely branching trees. It can be extended to trees with larger branching, at the cost of introducing infinitary conjunctions and disjunctions.

Definition 2.6 *The* Hennessy-Milner formulas *over* Act *are inductively defined as:*

- tt *and* ff
- $\varphi \wedge \psi$ *and* $\varphi \vee \psi$

[5] Throughout this paper, "bisimulation" is Milner's "strong bisimulation."

- $\langle a\rangle\varphi$ *for each* $a \in \mathsf{Act}$.
- $[a]\varphi$ *for each* $a \in \mathsf{Act}$.

If φ *is a Hennessy-Milner formula and* P *is a synchronization tree, then the relation of* satisfaction, $P \models \varphi$, *is defined by:*

- $P \models \mathrm{tt}$ *always,* $P \models \mathrm{ff}$ *never;*
- $P \models (\varphi \wedge \psi)$ *iff* $P \models \varphi$ *and* $P \models \psi$, *and similarly for disjunction;*
- $P \models \langle a\rangle\varphi$ *iff for some* P', $P \xrightarrow{a} P'$ *and* $P' \models \varphi$.
- $P \models [a]\varphi$ *iff for every* P' *such that* $P \xrightarrow{a} P'$, $P' \models \varphi$.

For example, $P \models \langle a\rangle\mathrm{tt}$ iff P has an a-child. If $\varphi = \langle a\rangle[b]\langle c\rangle\mathrm{tt}$, then φ separates the processes of Figure **??**: $a(bc + bd) + abc \models \varphi$ but $a(bc + bd) \not\models \varphi$. The familiar laws of propositional logic hold for Hennessy-Milner formulas, and so we ambiguously write, *e.g.*, $\varphi_1 \wedge \varphi_2 \wedge \varphi_3$ knowing that the order of parenthesization is irrelevant. Hennessy and Milner, [**?**], show the following:

Theorem 2.7 *If* P_1 *and* P_2 *are finitely-branching synchronization trees, then* $P_1 \leftrightarrows P_2$ *iff for each Hennessy-Milner formula* φ, $P_1 \models \varphi$ *iff* $P_2 \models \varphi$.

In particular, bisimulation is fully abstract with respect to observing modal formulas. That is, if we have some way of testing $P \models \varphi$ for all P and φ, then we have a good reason to distinguish between non-bisimilar trees. However, it is hard to see how to observe modal formulas without observing too much [**?**], or to understand them as computational observations.

3 Setting

3.1 Signatures and Transition Relations

We will be studying the interaction of programming languages and semantics, and in particular we will vary the programming language. We will therefore give general definitions suitable for quantifying over languages. In general, we want to have the most powerful class of languages that can be achieved without losing the essential mathematical and aesthetic properties which characterize CCS. We propose a class of languages, the GSOS languages, and argue that it meets this goal.

A language in the style of CCS is given by a *signature* (a set of operation symbols), and an operational semantics over that signature. It is convenient to include a set of synchronization trees in each language, so that we can test two trees for equivalence in different languages without having to worry about actually defining them by terms in the two languages.[6] We will include the nodes of the trees themselves as constants in the language.

[6] We will frequently want to include *all* synchronization trees. For foundational reasons, we only include only a few representatives from each of the 2^ω isomorphism classes of synchronization trees, but we ignore this subtlety from now on.

Definition 3.1 *A* signature $\mathcal{F}$ *is a nonempty finite set of* actions $\mathsf{Act}_{\mathcal{F}}$*, a possibly empty downward-closed set* $\Delta_{\mathcal{F}}$ *of synchronization trees over* $\mathsf{Act}_{\mathcal{F}}$*, and a family of disjoint sets* $\mathcal{F}_i$ *for* $i = 0, 1, 2, \ldots$ *such that* $\bigcup_{i \in \mathbf{N}} \mathcal{F}_i$ *is finite. The elements of* $\mathcal{F}_i$ *are the* operation symbols of arity i*. We insist that* $\mathbf{0} \in \mathcal{F}_0$*,* $\mathsf{Act}_{\mathcal{F}} \subseteq \mathcal{F}_1$*, and* $+ \in \mathcal{F}_2$*.*

We fix an infinite set Var of *agent variables*; X, Y, Z range over agent variables. The set of open terms over $\mathcal{F}$ is the least set such that

- Each synchronization tree $P \in \Delta_{\mathcal{F}}$ is a term.
- Each $X \in \mathsf{Var}$ is a term
- $f(P_1, \ldots, P_k)$ is a term whenever $f \in \mathcal{F}_k$ and each P_i is a term.

A *closed* term is a term which contains no variables. Note that we have no binding operators (see Section **??**).

Aside from the required operations $\mathbf{0}$, $a(\cdot)$, and $+$, CCS has the binary operations $|$ and $\|$, and the unary operations $\backslash L$ for $L \subseteq \mathsf{Act}$ and $[\rho]$ for certain $\rho : \mathsf{Act} \to \mathsf{Act}$. Similar languages such as Meije [**?**], CSP, and ACT use fairly similar sets of operations. We will not be programming in CCS per se, so we will not give the full semantics for CCS. Standard operations are written in prefix, infix, and suffix forms following the conventions that have evolved in the field.

The operations are generally given meaning by *structured operational rules* [**?**]; a language for concurrency in the CCS/CSP style is completely defined by a signature together with a set of structured rules defining the relation $P \xrightarrow{a} Q$ on closed terms. The operational semantics are given as a labeled transition system on the set of closed terms, which we will unwind into a synchronization tree in the obvious way.

It is difficult to define "structured operational rule" in sufficient generality to cover all the ways used in the literature (*e.g.*, [**?**, **?**, **?**, **?**] as well as many places in concurrency theory) and simultaneously avoid pathologies in particular cases. The results in this paper, among other work, show that the properties of a system defined by structured rules can be quite sensitive to the form of the rules. In general, though, a structured rule has the form

$$\frac{\textit{antecedent}}{\textit{consequent}}$$

where the antecedent and consequent are facts which may have free variables. These variables are considered bound, and rules which differ only in the names of free variables are identified; *e.g.*, the following are identical:

$$\frac{x \xrightarrow{a} y}{f(x) \xrightarrow{b} g(y)} \qquad \frac{y \xrightarrow{a} x}{f(y) \xrightarrow{b} g(x)}$$

The intent of the rule is that whenever the antecedent is satisfied by some instantiation of the free variables, then so is the consequent; and conversely that whenever a fact holds, there should be some instantiation of some rule in the language with true antecedent and

that fact as consequent. Structured rules, in a variety of guises, are familiar in many areas of mathematics, computer science, and logic.

We will frequently use rule schema in a fairly informal way; *e.g.*, the following scheme describes two rules for each $a \in \mathsf{Act}$.

$$\frac{\mathit{antecedent}[a] \quad (a \in \mathsf{Act})}{\mathit{consequent\text{-}1}, \quad \mathit{consequent\text{-}2}}$$

We illustrate the informal use of rules by giving rules for the required operations. We require that all languages have the same rules for these operations. The full definitions will be given in Section **??**.

- If P is a synchronization tree and P' is an a-child of P then $P \xrightarrow{a} P'$. That is, each synchronization tree P is a process with a synchronization tree isomorphic to P.
- $\mathbf{0}$ has no rules; it denotes the null tree.
- For each $a \in \mathsf{Act}$, we have the following rule.[7] The synchronization tree of aP is that of P with a new root and an edge labeled a from the new root to the root of P.

$$aX \xrightarrow{a} X \tag{2}$$

- For each $a \in \mathsf{Act}$, we have the following two rules. The synchronization tree of $P + Q$ consists of the trees of P and Q with roots identified.

$$\frac{X \xrightarrow{a} Y}{X + X' \xrightarrow{a} Y, \qquad X' + X \xrightarrow{a} Y} \tag{3}$$

See Figure **??** for pictures of aP and $P + Q$. We write a for the process $a\mathbf{0}$ when no confusion can arise, use infix notation, and omit parentheses following usual mathematical conventions. For example, we mercifully write $a(bc + bd)$ for $a(+(b(c(\mathbf{0})), b(d(\mathbf{0}))))$. It is easy to show that $+$ is commutative and associative in the synchronization-tree semantics — that is, for all synchronization trees P and Q, $P + Q$ and $Q + P$ are isomorphic as synchronization trees. We have terms denoting all finite synchronization trees: $\mathbf{0}$ denotes the tree with no actions, and if P_i denotes the tree t_i, then $a_1P_1 + \cdots + a_nP_n$ denotes the tree with an a_i-edge to t_i for each $1 \leq i \leq n$.

Let P be a closed term of $\mathcal{F}$, and $\dot{\rightarrow}$ a transition relation over $\mathcal{F}$. The behavior of P under $\dot{\rightarrow}$ is a possibly infinite graph edge-labeled by actions, and node-labeled by terms.[8] This graph may be unwound to give a possibly infinite tree $[\![P]\!]_{\dot{\rightarrow}}$; if it is finitely branching, it is a synchronization tree giving the meaning of P in an obvious sense. Thus, given a

[7] We could make a distinction between axioms and inference rules; for our purposes, it is simplest to consider axioms such as this as inference rules with an empty set of hypotheses.

[8] There is a slight caveat here; if $P \xrightarrow{a} Q$ and there are n distinct proofs of this fact, we must add n Q-nodes to the graph, giving P and a-edge to each. With the straightforward definition, the term $a + a$ would have the same synchronization tree as a; however, the term $a(\mathbf{0} + \mathbf{0}) + a\mathbf{0}$ would have a different synchronization tree; in particular, the synchronization tree semantics would not be adequate.

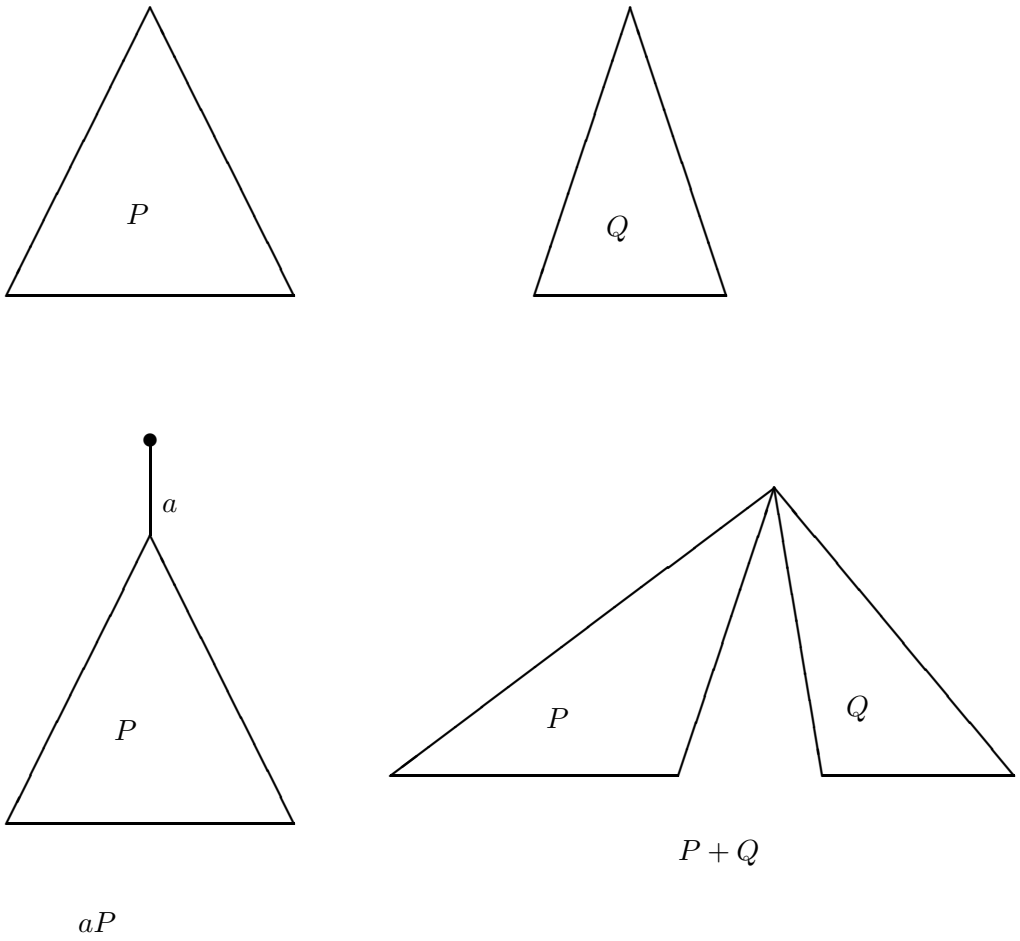

Figure 7: Effects of Basic Operations on Trees

transition relation, definitions phrased in terms of synchronization trees may be applied to processes as well; we write *e.g.* $P \leftrightarrows Q$ for $[\![P]\!]_{\rightarrow} \leftrightarrows [\![Q]\!]_{\rightarrow}$.

More generally, given an open term P of variables $X_1, \ldots, X_n$, we may interpret P as an n-ary function on synchronization trees. For example, aX denotes the function of prepending a new root and an a-branch to the original root of a tree. Examples of processes and their associated synchronization trees can be found in most of the figures in this article.

3.2 Observations

As stated in the introduction, we are mainly interested in *finite completed traces* or simply *traces*: sequences of actions that processes may take before they halt. This is formalized in the following definition.

Definition 3.2 *If $s \in \mathsf{Act}^*$, then $P \xrightarrow{s} P'$ iff there are terms $P_1, \ldots, P_{|s|}$ such that*

$$P \xrightarrow{s_1} P_1 \xrightarrow{s_2} \cdots \xrightarrow{s_{|s|}} P_{|s|} = P'$$

We write $P \xrightarrow{s}$ if for some P' we have $P \xrightarrow{s} P'$; otherwise, $P \overset{s}{\not\rightarrow}$. P is stopped *if for every $a \in \mathsf{Act}$, $P \overset{a}{\not\rightarrow}$.*

It is also possible to use *partial traces*, strings s such that $P \xrightarrow{s}$, or *infinite traces*, infinite strings s such that there exist P_i's such that $P_0 = P$ and $P_i \xrightarrow{s_{i+1}} P_{i+1}$ for each i. Partial traces are too weak for our purposes, and infinite traces are too long to observe in any practical sense in finite time. In this study, "trace" will always mean "finite terminated trace" unless otherwise specified.

We have used the notation $P \xrightarrow{a} Q$ in two ways for synchronization trees P and Q, in the senses of Definition **??** and Definition **??**; the two notions are equivalent on all synchronization trees.

Definition 3.3 *The* trace set *of P,* $\mathrm{tr}(P)$*, is* $\left\{ s \in \mathsf{Act}^* \mid P \xrightarrow{s} P' \text{ and } P' \text{ is stopped} \right\}$

We formalize "using a program" by the notion of a context.

Definition 3.4 *A* context of n holes $C[X_1, \ldots, X_n]$ *in a language $\mathcal{L}$ is simply an $\mathcal{L}$-term with free variables at most $\{X_1, \ldots, X_n\}$. The result of substituting P_i for X_i in $C[X_1, \ldots, X_n]$ is written $C[P_1, \ldots, P_n]$.*

For example, $C[X] = (X + a)|(X + b)$ is a CCS context. There are no variable-binding operations available in our language, so the familiar subtleties of renaming are unnecessary.

Definition 3.5 *Let P and Q be synchronization trees.*

1. *P and Q are* trace equivalent*, $P \equiv_{tr} Q$, iff $\mathrm{tr}(P) = \mathrm{tr}(Q)$.*
2. *P* trace approximates *Q, $P \sqsubseteq_{tr} Q$, iff $\mathrm{tr}(P) \subseteq \mathrm{tr}(Q)$.*

Figure 8: $Q_\star$ and $P_\star$: Ready Similar but not Bisimilar

3. *P and Q are* trace congruent with respect to the language $\mathcal{L}$, $P \equiv_{tr}^{\mathcal{L}} Q$, *iff for all $\mathcal{L}$-contexts $C[X]$ of one free variable, $C[P] \equiv_{tr} C[Q]$.*

4. *P* trace approximates *Q with respect to the language $\mathcal{L}$, $P \sqsubseteq_{tr}^{\mathcal{L}} Q$, iff for all $\mathcal{L}$-contexts $C[X]$ of one free variable, $C[P] \sqsubseteq_{tr} C[Q]$.*

In CSP, then, two programs are distinguished iff there is a good reason to distinguish them – a context $C[X]$ and a string s of actions such that only one of $C[P]$ and $C[Q]$ can perform s and then stop. In fact, there is a fully abstract mathematical semantics of CSP, the *failures* semantics of [**?**]; the meaning of a process P is the set of all pairs $\langle s, \mathcal{X} \rangle$, such that $P \xrightarrow{s} P'$ and $\mathcal{X}$ is a set of actions b such that $P' \stackrel{b}{\not\rightarrow}$.

It is well-known that bisimulation is strictly finer than CCS/CSP trace congruence. Logically, the refusals correspond to modal formulas of the form

$$\langle a_1 \rangle \langle a_2 \rangle \cdots \langle a_k \rangle ((\![b_1]\!]\mathrm{ff}) \wedge (\![b_2]\!]\mathrm{ff}) \wedge \cdots \wedge (\![b_l]\!]\mathrm{ff}));$$

that is, two processes are CCS/CSP congruent iff they agree on all such formulas. From this, it is routine to check that $P_\star \equiv_{tr}^{CCS/CSP} Q_\star$ but $P_\star \not\leftrightarrows Q_\star$ (see Figure **??**).

(B: Make sure this figure agrees with other definitions of $P_\star$ and $Q_\star$. :B)

It is not clear why $P_\star$ and $Q_\star$ should be considered different in CCS. In the spirit of [**?**, **?**], we investigate the question of what kinds of operations one must add to CCS to make bisimulation be trace congruence.

4 GSOS Languages

4.1 Generalizing CCS — A False Start

As we have seen, bisimulation is not trace congruence with respect to CCS. We consider various ways of generalizing CCS, attempting to refine the language's congruence to make it coincide with bisimulation.

There is a straightforward and trivial operation to add to CCS which makes bisimulation precisely equal trace congruence. This operation, called "not-bisim," takes two arguments; it produces a signal if they are not bisimilar, and is stopped if they are bisimilar.

$$\frac{X \not\leftrightarrows X'}{\mathsf{not\text{-}bisim}(X, X') \xrightarrow{a} \mathbf{0}}$$

In this language, if P and Q are not bisimilar, then the context $C[X] = \mathsf{not\text{-}bisim}(X, P)$ distinguishes them: $C[P] \not\xrightarrow{a}$, but $C[Q] \xrightarrow{a} \mathbf{0}$. There are several grounds for criticizing this operation.

- It begs the question. We have explained why we think bisimulation is important by saying that, if we consider it important, then it is important. It would be desirable to explain it in terms of something that looks more fundamental.
- We could explain any other relation between synchronization trees in the same way. For example, a slight variation on that rule would make synchronization tree isomorphism the fundamental relation in the language, distinguishing a from $a + a$.
- The rule is very difficult to apply. In Milner's original SCCS, the question of $P \leftrightarrows Q$ is not arithmetic [**?**], and probably Σ_1^1 complete. Even for finitely-branching trees, the question $P \not\leftrightarrows Q$ is r.e.-complete, and so the one-step transition relation $P \xrightarrow{a} Q$ is not decidable. It is semidecidable,[9] which is why we phrased the rule in terms of $\not\leftrightarrows$ rather than $\leftrightarrows$.
- The operation not-bisim seems to be useless. At the time of writing, we have found no researchers who exhibit any desire to have not-bisim included in a programming language. The operation is generally and accurately regarded as a theoretician's trick to get full abstraction. This operation is merely intended as an answer to the question "how could one add an operation to CCS making bisimulation into trace congruence?"; even the strongest proponents of bisimulation have not recommended not-bisim as a new primitive of CCS.

We would like to choose some criterion by which we may judge languages and see if they are reasonable. We will argue that bisimulation is not trace congruence with respect to any reasonable language; to make this formal, we will have to have some way of quantifying over all reasonable languages, and hence need some formal characterization of those languages.

One of the elegant features of CCS is the definition by a set of structured operational rules. By contrast, not-bisim is defined by an ad-hoc rule involving the predicate of bisimulation. From the form of the structured rules, CCS is guaranteed to exhibit several mathematical properties (*e.g.*, that all programs using only guarded recursion are finitely branching, and that the transition relation is decidable). We choose, therefore, to investigate languages defined by rules which look like the rules of CCS. We will check the soundness of our definition of "looking like the rules of CCS" by making sure that our languages all have

[9] Strictly, the question $P \not\leftrightarrows Q$ is r.e. given oracles of the tree constants appearing in P and Q. If there are no non-recursive trees, then the question is r.e..

the essential properties of CCS. We will try to catch all reasonable languages by taking the largest cleanly-defined class of languages with "well-structured" CCS-like definitions which have these essential properties. The remainder of this section is concerned with the discussion of our proposed class of "reasonable" CCS-like process languages.

4.2 The Purpose of Fixed Points

We are investigating the *discriminatory* power of the language, its ability to tell the difference between processes in finite time. The fixed point operations add to the *expressive* power, the ability of the language to define synchronization trees and functions on them. In general, the two kinds of power are related: increasing the discriminatory power necessarily increases the expressive power. However, the converse does not hold: it is common to discover that a language extension has increased expressive power and left discriminatory power unchanged. Most programmers and language designers are quite properly concerned with expressive power — the ability to write programs easily — and at most secondarily interested in discriminatory power. However, we are working on issues of full abstraction, in which discriminatory power is important.

One programming-language feature commonly included in core languages for concurrency is *recursion*. Not surprisingly, recursion is used to generate processes which run indefinitely. It is not necessary in the setting of this paper, as we allow an arbitrary set of processes to be included as constants.[10]

The fixed-point rule, used without caution, can be dangerous. It is possible to define processes (*e.g.*, $\mathrm{fix}\,[X \Leftarrow a + (X \parallel X)]$) which are *countably branching*; that is, there are a countable set of distinct terms Q_n such that $\mathrm{fix}\,[X \Leftarrow a + (X \parallel X)] \xrightarrow{a} Q_n$.

Infinitely branching trees and so-called "unguarded processes" [**?**] cause many problems in many aspects of the theory. It seems likely that bisimulation between countably branching trees cannot match trace congruence, for purely recursion-theoretic reasons. Unguarded recursion in Milner's original SCCS makes the transition relation $P \xrightarrow{a} Q$ undecidable; it is an open problem whether or not it is decidable in SCCS with a finite action set. The correspondence between bisimulation and Hennessy-Milner logic becomes harder; Theorem **??** fails unless infinitary conjunctions and disjunctions are used. Unguarded recursion is also incompatible with negative rules, although this could be construed as a criticism of negative rules.

For these reasons and others, restrictions are generally imposed on recursive definitions of processes. In CCS, "guarded" recursion is singled out as attractive, and in CSP and the test-equivalence system of [**?**], unguarded recursions are treated as though they diverged. The essence of these restrictions is to ensure that definable processes behave like *computable, finitely branching trees*: that there is a program which, given a and P, computes the finite set $\left\{Q \mid P \xrightarrow{a} Q\right\}$.

In the case of guarded recursion, suppose that P and Q are not trace congruent — that is, there is a context $C[X]$ and a string s of actions such that, say, $s \in \mathrm{tr}(C[P]) - \mathrm{tr}(C[Q])$.

[10] We are investigating the power of languages to build communication and control constructs, not their ability to define functions in the recursion-theoretic sense.

This context $C[X]$ may involve recursion. However, the guarded fixed point operators appearing in $C[X]$ may be unwound a suitably large but finite number of times and then replaced by $\mathbf{0}$ giving a new context, $C'[X]$, which contains no fixed point operators and also has the property that $s \in \mathrm{tr}(C'[P]) - \mathrm{tr}(C'[Q])$. That is, P and Q can be distinguished by a context not involving recursion at all. This informal argument can be formalized directly in our class of languages.

4.3 GSOS Rules

We present the general format of *GSOS structured transition rule*, and then argue that this format is a maximal right one. The argument will take the form of some theorems showing that any language defined by these rules has some desirable properties, and an array of counterexamples showing that the obvious classes of extensions do not.

Definition 4.1 *A* positive transition formula *is a triple of two terms and an action, written* $T \xrightarrow{a} T'$. *A* negative transition formula *is a pair of a term and an action, written* $T \not\xrightarrow{a}$.

Definition 4.2 *A* GSOS rule ρ *is a rule of the form:*

$$\frac{\bigcup_{i=1}^{l} \left\{ X_i \xrightarrow{a_{ij}} Y_{ij} \mid 1 \le j \le m_i \right\} \cup \bigcup_{i=1}^{l} \left\{ X_i \not\xrightarrow{b_{ik}} \mid 1 \le k \le n_i \right\}}{\mathsf{op}(X_1, \ldots, X_l) \xrightarrow{c} C[\vec{X}, \vec{Y}]}$$

where all the variables are distinct, $l \ge 0$ *is the arity of* op, $m_i, n_i \ge 0$, *and* $C[\vec{X}, \vec{Y}]$ *is a context with free variables including at most the* X*'s and* Y*'s. (It need not contain all these variables.) The operation symbol* op *is the* principal operator *of the rule;* $\mathsf{ante}(\rho)$ *is the set of antecedents, and* $\mathsf{cons}(\rho)$ *is the consequent.*

A rule is negative if it has any $x_i \not\xrightarrow{b_{ij}}$ antecedents; otherwise it is positive. Note that every X_i occurring in the antecedent of a GSOS rule must occur as an argument of the principal operator in the consequent, but not every argument of the principal operator need occur in the antecedent. In the future we will not write the ranges of the indices on the rule unless necessary.

Definition 4.3 *A* GSOS rule system $\mathcal{G}$ *over a signature* $\mathcal{F}$ *is a finite set of GSOS rules over the actions and operations in* $\mathcal{F}$, *such that precisely the required rules (??) and (??) are given for the required operations* $a(\cdot)$ *and* $+$.

Now, we must show that each GSOS rule system determines an operational semantics. The operational semantics will be given by a labeled transition system with the closed $\mathcal{F}$-terms as the processes and the actions in Act as the action. The presence of negative rules requires us to do some work to define the transition relation.

Definition 4.4 *A* (closed) substitution *is a partial map* σ *from variables to (closed) terms. We write* $P\sigma$ *for the result of substituting* $\sigma(X)$ *for each* X *occurring in* P; *if* $\sigma(X)$ *is undefined,* $P\sigma$ *is undefined.*

For example, let $\sigma(X) = Q$; then $(aX + bY)\sigma = aQ + bY$. Note that if the free variables in P are $X_1, \ldots, X_n$, then $P\sigma = P[\vec{X} := \vec{Q}]$. All substitutions in this study will be closed.

Definition 4.5 *If $\cdot$ is a transition relation, σ is a substitution, and t is a transition formula, then the predicate $\cdot, \sigma \models t$ is defined by*

$$\begin{array}{lll} \cdot, \sigma \models T \xrightarrow{a} T' & \iff & T\sigma \ {}^{a} \ T'\sigma \\ \cdot, \sigma \models T \ {}^{a} & \iff & Q.\, T\sigma \ {}^{a} \ Q. \end{array}$$

If $\mathcal{S}$ is a set of transition formulas, $\cdot, \sigma \models \mathcal{S}$ iff $\forall t \in \mathcal{S}. \cdot, \sigma \models t$.

For example, let $\cdot_{CCS}$ be the transition relation of CCS (which we write as $\dot{\rightarrow}$ in all sections in which the notation is not ambiguous). Suppose that $\sigma_1(X) = ab + ac$. Then we have $\cdot_{CCS}, \sigma_1 \models \left\{X \xrightarrow{a} b, X \xrightarrow{a} c, X \ {}^{b}\right\}$.

Definition 4.6 *If ρ is a GSOS rule, $\cdot, \sigma \models \rho$ holds iff*

$$\cdot, \sigma \models \mathsf{ante}(\rho) \text{ implies } \cdot, \sigma \models \mathsf{cons}(\rho)$$

For example, $\cdot_{CCS}, \sigma \models \rho$ for every substitution σ and every CCS rule ρ. However, we also have $\cdot_{\infty}, \sigma \models \rho$ for every CCS rule as well, where $\cdot_{\infty}$ is the universal relation: $P \ {}^{a}_{\infty} \ Q$ for all P, Q, and a.

Definition 4.7 $\cdot$ *is* sound *for $\mathcal{G}$ iff for every rule $\rho \in \mathcal{G}$ and every substitution σ, we have $\cdot, \sigma \models \rho$.*

In general, many transition relations will be sound for $\mathcal{G}$; for example, $\cdot_{\infty}$ is sound for every $\mathcal{G}$. One generally takes the *smallest* sound transition relation, showing that there is in fact a smallest one. This is not appropriate for GSOS rules; with negative rules, there may not be a smallest sound transition relation [**?**].

However, GSOS languages do define a (unique) operational semantics in an appropriate sense. The point of minimality in the usual case is to ensure that everything which is true is true for some reason, because there is some rule with that fact as consequent and a true antecedent. We make this concern explicit.

Definition 4.8 $\cdot$ *is* witnessing *for $\mathcal{G}$ iff, whenever $P \ {}^{a} \ P'$ there is a rule $\rho \in \mathcal{G}$ and a substitution σ such that $\cdot, \sigma \models \mathsf{ante}(\rho)$ and $\mathsf{cons}(\rho)\sigma = P \xrightarrow{a} P'$.*

A transition relation is witnessing if, whenever a transition happens, it happens because it was the consequent of (an instantiation of) some rule, and the antecedents of that rule were satisfied. For example, $\cdot_{CCS}$ is witnessing for CCS. However, $\cdot_{\infty}$ is not witnessing for CCS. There are no axioms for $\mathbf{0}$, yet $\mathbf{0} \ {}^{a}_{\infty} \ a$. Soundness and witnessing together select the right transition relation:

Lemma 4.9 *Let $\mathcal{G}$ be a GSOS rule system. There is a unique sound and witnessing transition relation $\dot{\rightarrow}_{\mathcal{G}}$ for $\mathcal{G}$.*

Proof: Straightforward structural induction. □

We call the unique transition relation the *standard* transition relation, and write it $\dot{\rightarrow}_{\mathcal{G}}$ or simply $\dot{\rightarrow}$.

5 Why GSOS Rules Are Desirable

In this section, we argue that GSOS rules are appropriate as a generalization of CCS. We give two theorems which, together with Lemma **??**, demonstrate that bisimulation is appropriate in the GSOS setting. Recall that bisimulation is best used with finitely branching trees; we will show that every GSOS language produces only finitely branching trees. We then give an indication of the additional power of GSOS-definable operations by some examples of operations on trees which can be defined in the GSOS setting but not in CCS.

5.1 Basic Properties

Theorem 5.1 *Let $\mathcal{G}$ be a GSOS rule system. Then the transition relation on $\mathcal{G}$ is computably finitely branching uniformly in the tree constants. That is, there is an algorithm which, given an action a, a term P, and oracles for all the tree constants occurring in P, produces the set of a-children of P; and this set is always finite.*

Proof: A straightforward recursion on P. □

It is generally accepted in the setting of process algebra that bisimilar processes should not be distinguished. This holds in a strong sense for all GSOS rule systems.

Theorem 5.2 *Let $\mathcal{G}$ be a GSOS rule system. Then bisimulation is a congruence with respect to the operations in $\mathcal{G}$. That is, if $P \leftrightarrows Q$ are synchronization trees and $C[X]$ is a context over $\mathcal{G}$, then $C[P] \leftrightarrows C[Q]$.*

Proof: This is similar to the proof of Lemma **??**, presented in Section **??**; it is best done using the machinery developed in that section. □

Corollary 5.3 *If $P \leftrightarrows Q$, then P and Q are trace congruent with respect to $\mathcal{G}$.*

Proof: It is clear that, if $R \leftrightarrows S$, then R and S have the same traces. Let $C[\cdot]$ be an arbitrary context; by Theorem **??**, $C[P] \leftrightarrows C[Q]$, and hence P and Q have the same traces in $C[\cdot]$. □

5.2 Expressive Power

GSOS rules are quite expressive, as witnessed by the fact that most structured transition rules proposed in the field have been GSOS rules. For example, the CCS restriction operations $P\ \ A$ for $A \subseteq \mathsf{Act}$ are defined by the family of rules, one for each $a \in A$:

$$\frac{X \xrightarrow{a} Y}{X\ \ A \xrightarrow{a} Y\ \ A}$$

The simple interleaving parallel composition, $\|$, is given by:

$$\frac{X \xrightarrow{a} Y}{X \parallel X' \xrightarrow{a} Y \parallel X'}\ ,\ \frac{X' \xrightarrow{a} Y}{X \parallel X' \xrightarrow{a} X \parallel Y} \tag{4}$$

(with one instance of each rule for each action a.) The standard parallel composition operation $|$ of CCS has these rules, and some extra rules for communication. Suppose that there is a distinguished action $\tau \in \mathsf{Act}$, and a permutation $\bar{\cdot}$ of $\mathsf{Act} - \{\tau\}$, such that $\bar{\bar{a}} = a$ for each action $a \neq \tau$. The remaining behavior of $|$ is given by the rule scheme

$$\frac{X \xrightarrow{a} Y, \quad X' \xrightarrow{\bar{a}} Y'}{X|X' \xrightarrow{\tau} Y|Y'}$$

The operational rules assigning behavior to CCS/CSP/ACP/Meije terms easily fit the GSOS framework.

In fact, GSOS rules go beyond the kind of structured rules needed for CCS in two aspects — the use of *negation* and *copying*. Negation allows us to define a pure form of sequential composition: $P;Q$ runs P until it stops, and then runs Q. As an operation on synchronization trees, the $P;Q$ glues a copy of Q at each leaf of P.[11]

$$\frac{X \xrightarrow{a} Y}{X;X' \xrightarrow{a} Y;X'} \qquad \frac{X' \xrightarrow{b} Y',\ \{X \not\xrightarrow{a} \mid a \in \mathsf{Act}\}}{X;X' \xrightarrow{b} Y'}$$

Copying allows us, not surprisingly, to make copies of processes. There can be more than one antecedent about the behavior of a single subprocess, and more than one copy of a process in the result in the consequent. For example, the following GSOS rules yield operations which cannot be defined in CCS [**?**].

$$\frac{X \xrightarrow{a} Y, X \xrightarrow{b} Y'}{\text{a-if-b}(X) \xrightarrow{a} \text{a-if-b}(Y)} \qquad \frac{X \xrightarrow{a} X'}{\text{double}(X) \xrightarrow{a} X' \parallel X'}$$

The operation a-if-b(X) will allow X to exhibit its a-behavior as long as it also has the possibility of performing b's at each step. The double operator produces two copies of one a-child of X, running in interleaved parallel.

6 Obvious Extensions Violate Basic Properties

There are many technical restrictions in our definition of a GSOS rule, and it is natural to ask if they can be relaxed. We indicate how various relaxations may break the key properties of GSOS systems. Note that some systems with non-GSOS rules enjoy the good properties of GSOS systems; however, this is not immediate from the syntactic specifications of these systems. GSOS rules therefore provide a language-design methodology: any language defined purely by GSOS is guaranteed to meet the basic criteria; other languages may or may

[11] It is possible to define some forms of sequential composition with positive rules. For example, sequencing in CSP runs P until it announces that it has finished by sending a special action, then runs Q. However, processes may finish without announcing that they have finished — called "deadlock" in CSP — or (in general) may announce that they have finished when in fact they are still able to continue; CSP sequencing is not identical with pure sequencing. It is evidently acceptable for programming. Frits Vaandrager has pointed out that in some ways, it is preferable; it allows us to consider *divergence*. However, divergence is not a part of Milner's original notion of strong bisimulation, and we do not consider it.

not. Perhaps more importantly, a GSOS language may be extended by GSOS operations and is still guaranteed to behave well; a well-behaved non-GSOS language extended by the same GSOS operations may cease to be well-behaved. The properties which non-GSOS systems most often violate are:

- The guarantee that bisimulation is a congruence. In fact, they typically do not respect synchronization tree isomorphism; *e.g.*, there are two stopped programs which can be distinguished. (Recall that all stopped programs are unable to take any actions, and hence they have the same synchronization tree; in fact, the null tree.)
- The requirement that $\dot{\rightarrow}$ be computably, finitely branching.
- The existence of some transition relation $\dot{\rightarrow}$ agreeing with all the rules.

We present a single, representative counterexample; the rest may be found in Appendix **??**.

Many possible extensions of the GSOS format give some kind of pattern-matching ability, which generally prevents bisimulation from being a congruence. For example, the consequent must be of the form $\mathsf{op}(\vec{X}) \xrightarrow{c} C[\vec{X}, \vec{Y}]$. If we allow more structured left-hand sides of the antecedent, we allow a certain kind of pattern matching. Consider a unary operation ζ defined by the axiom

$$\zeta(\mathbf{0}) \xrightarrow{a} \mathbf{0}$$

Now, $\mathbf{0}$ and $\mathbf{0} + \mathbf{0}$ are bisimilar; indeed, both have the null synchronization tree. However, $\zeta(\mathbf{0}) \xrightarrow{a} \mathbf{0}$ but $\zeta(\mathbf{0} + \mathbf{0})$ is stopped. This gives us a context which distinguishes between two bisimilar terms, showing that bisimulation is not a congruence with respect to this operation. Similar examples apply to other rules which can branch on more than simply the leading operator of the antecedent.

7 Theory of Ready Simulation

7.1 Overview

In this section, we develop the core of the theory of ready simulation and GSOS languages. We present and prove the equivalence of the main definitions of ready simulation, and in particular we show that ready simulation is precisely congruence with respect to all GSOS languages. We also develop a modal logic which matches ready simulation in the same way that Hennessy-Milner logic matches bisimulation. In Section **??**, we use this logic to build a GSOS language in which ready simulation is precisely congruence. As a corollary of this work, we show that bisimulation is not congruence with respect to any GSOS language.

The three main characterizations presented in this chapter are ready simulation (RS), denial logic (DL), and GSOS congruence (GC). We prove the equivalences in the order

$$\begin{array}{ccc} RS & \leftrightarrow & DL \\ \downarrow & \nearrow & \\ GC & & \end{array}$$

It is simpler to talk about synchronization trees (which are absolute) rather than process terms (which change their meaning depending on the language.)

Definition 7.1 *Let P and Q be synchronization trees.*

1. *$P \sqsubseteq_{tr}^{GSOS} Q$ iff, for all GSOS languages $\mathcal{G}$ including P and Q as trees, $P \sqsubseteq_{tr}^{\mathcal{G}} Q$.*
2. *$P \equiv_{tr}^{GSOS} Q$, P and Q are* GSOS congruent*, iff for all GSOS languages $\mathcal{G}$ including P and Q as trees, $P \equiv_{tr}^{\mathcal{G}} Q$.*

Two processes in the language $\mathcal{G}$ are GSOS congruent iff their synchronization trees with respect to $\mathcal{G}$ are.

We give several equivalent characterizations of GSOS congruence, in terms more like those defining bisimulation.

7.2 Ready Simulation and GSOS Congruence

The following characterization was discovered (after the modal characterization of Section **??**) by Larsen and Skou [**?**], and independently by R. van Glabbeek.

Definition 7.2

1. *A relation $\_\rightharpoonup'$ between synchronization trees is a* ready simulation *relation iff, whenever $P \_\rightharpoonup' Q$,*
   - *Whenever $P \xrightarrow{a} P'$ then there is a Q' such that $Q \xrightarrow{a} Q'$ and $P' \_\rightharpoonup' Q'$.*
   - *Whenever $P\ a$, then $Q\ a$.*
2. *$P \_\rightharpoonup Q$ if there is some ready simulation relation $\_\rightharpoonup'$ such that $P \_\rightharpoonup' Q$.*
3. *$P \ \ Q$ iff $P \_\rightharpoonup Q$ and $Q \_\rightharpoonup P$. In this case P and Q are said to be* ready similar*.*

A useful fact follows immediately from the definition. Let the *ready set* of P be defined by

$$\text{readies}(P) = \left\{a : P \xrightarrow{a}\right\}. \tag{5}$$

Then $P \_\rightharpoonup Q$ implies $\text{readies}(P) = \text{readies}(Q)$. In the presence of the first clause in the definition of ready simulation, $\text{readies}(P) = \text{readies}(Q)$ is equivalent to the second clause. The name "ready simulation" comes from the use of the set of actions that the process is ready to perform.

The relation $\_\rightharpoonup$ is a ready simulation relation, and in fact the largest such relation. The main result of this section is that $P \ \ Q$ iff P and Q are GSOS congruent. Proving this will take the rest of the section. Before proving it, we give some examples.

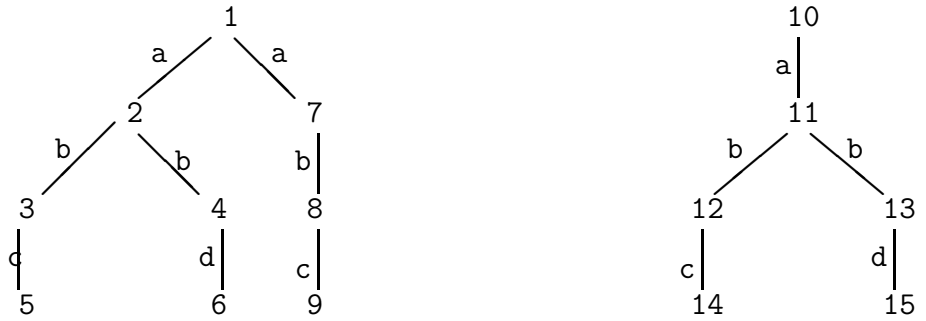

1	$\rightharpoonup_1$	10	10	$\rightharpoonup_2$	1
2	$\rightharpoonup_1$	11	11	$\rightharpoonup_2$	2
3	$\rightharpoonup_1$	12	12	$\rightharpoonup_2$	3
4	$\rightharpoonup_1$	13	13	$\rightharpoonup_2$	4
5	$\rightharpoonup_1$	14	14	$\rightharpoonup_2$	5
6	$\rightharpoonup_1$	15	15	$\rightharpoonup_2$	6
7	$\rightharpoonup_1$	11			
8	$\rightharpoonup_1$	12			
9	$\rightharpoonup_1$	14			

Figure 9: Ready Similar Processes

7.3 Examples of Ready Simulation

For any process P, we have $P \mathrel{\_\!\!\rightharpoonup} P$. Furthermore, for any P and Q we have $aP \mathrel{\_\!\!\rightharpoonup} aP + aQ$, using the relation $\mathrel{\_\!\!\rightharpoonup}$ itself:[12] the only possible transition of aP is $aP \xrightarrow{a} P$, and $aP+aQ \xrightarrow{a} P$ and $P \mathrel{\_\!\!\rightharpoonup} P$ as desired. For example, $bc \mathrel{\_\!\!\rightharpoonup} bc+bd$. The canonical example of processes which are ready similar but not bisimilar are $P_\star = abc + a(bc + bd) + abd$ and $Q_\star = a(bc + bd)$ of Figure **??**; the ready simulation relation between them is given in Figure **??**.

Note that a bisimulation relation is a ready simulation relation in each direction, and so bisimilar processes are ready similar. We therefore have:

Theorem 7.3 *Bisimulation is a strict refinement of ready simulation and hence of GSOS congruence.*

A final example of two processes which are ready similar but not bisimilar is a *lossy delay link*. A lossy two-stage link repeatedly accepts an input value v, waits one time unit, and produces as output either the value v or a signal saying that v was lost in transit. We present two ways to specify the lossy delay link in CCS. The first always receives its input correctly, but may lose it during the delay; the second may lose it either on the step that it

[12]This inequality, together with the axioms of bisimulation, gives a complete inequational axiom system for ready simulation of finite trees.

receives it or during the delay. We use the action w for waiting, and i_v and o_v for the input and output of the value v.

$$\begin{aligned} LL_1 &= \sum_v i_v.(w.o_v.LL_1 + w..LL_1) \\ LL_2 &= \sum_v (i_v.(w.o_v.LL_2 + w..LL_2) + i_v.w..LL_2) \end{aligned}$$

Lossy delay links are probably not particularly useful entities to build. However, it may be desirable to show that a system behaves correctly even if the communication channels are lossy (a safety proof, at least), and so it may be desirable to model lossy delay links.

7.4 Ready Simulation Implies GSOS Congruence

In this section, we show the following theorem:

Theorem 7.4 *Let P and Q be synchronization trees.*

1. *If $P \_\rightharpoonup Q$ then $P \sqsubseteq_{tr}^{GSOS} Q$.*
2. *If $P \ Q$ then $P \equiv_{tr}^{GSOS} Q$.*

The proof of (1) occupies the rest of the section; (2) follows immediately from (1). Recall that the operations $\mathsf{op}(\vec{X})$ were defined by rules of the form

$$\frac{X_i \stackrel{a_{ij}}{\rightarrow} Y_{ij}, \qquad X_i \ b_{ik}}{\mathsf{op}(\vec{X}) \stackrel{c}{\rightarrow} C[\vec{X}, \vec{Y}]}$$

In fact, the same is true of each context $D[\vec{X}]$ as well as the simple contexts $\mathsf{op}(\vec{X})$. That is, the behavior of $D[\vec{X}]$ can be completely captured by a set of derived rules of the form:

$$\frac{X_i \stackrel{a_{ij}}{\rightarrow} Y_{ij}, \qquad X_i \ b_{ik}}{D[\vec{X}] \stackrel{c}{\rightarrow} C[\vec{X}, \vec{Y}]}$$

We will call these constructs "ruloids" rather than "rules" because they are not the rules used to define the language and because they violate our definition of a GSOS rule (the source of the consequent has the wrong form).

Definition 7.5 *A set of ruloids R is* specifically witnessing *for a context $D[\vec{X}]$ and action c iff all the consequents of ruloids in R are of the form $D[\vec{X}] \stackrel{c}{\rightarrow} C[\vec{X}, \vec{Y}]$, and whenever $D[\vec{P}] \stackrel{c}{\rightarrow} Q$, there is a ruloid $\rho \in R$ and substitution σ such that $\mathsf{cons}(\rho)\sigma = D[\vec{P}] \stackrel{c}{\rightarrow} Q$ and $\dot{\rightarrow}, \sigma \models \mathsf{ante}(\rho)$.*

In the course of the following proof, we will use the notions of "sound and specifically witnessing" at a variety of types – e.g., concerning sets of ruloids, or functions returning ruloids. We sketch the definitions where appropriate; but they are essentially the same as Definition **??**.

Theorem 7.6 *Let $\mathcal{G}$ be a GSOS language. For each $\mathcal{G}$-context $D[\vec{X}]$ and action a, there exists a finite set $\mathcal{R}(D,a)$ of ruloids of the form*

$$\frac{X_i \stackrel{a_{ij}}{\rightarrow} Y_{ij}, \qquad X_i \; {}^{b_{ik}}}{D[\vec{X}] \stackrel{a}{\rightarrow} C[\vec{X},\vec{Y}]}$$

such that the rules in $\mathcal{R}(D,a)$ are sound and the set $\mathcal{R}(D,a)$ is specifically witnessing for $D[\vec{X}]$ and a.

Proof:

The proof is by induction on the structure of contexts. If $D[\vec{x}] = x_i$, then the set $\mathcal{R}(D,a) = \left\{ \frac{x_i \stackrel{a}{\rightarrow} y}{x_i \stackrel{a}{\rightarrow} y} \right\}$ clearly suffices.

Suppose, inductively, that $D[\vec{x}] = \mathsf{op}(D_1[\vec{x}], \ldots, D_n[\vec{x}])$. We will construct $\mathcal{R}(D,a)$, using the ruloid sets $\mathcal{R}(D_i,a)$ which have been constructed at earlier stages of induction.

Consider an arbitrary rule ρ with conclusion $\mathsf{op}(\vec{z}) \stackrel{a}{\rightarrow} E[\vec{z},\vec{w}]$. We will build a set Δ_ρ, a set of ruloids which precisely capture when $D[\vec{P}] \stackrel{a}{\rightarrow} Q$ via rule ρ.

Let PA_ρ be the set of positive antecedents of ρ, and NA_ρ the set of negative antecedents.

First, we will build the set of antecedents corresponding to PA_ρ. We'd like to replace the antecedent $t = z_i \stackrel{a_{ij}}{\rightarrow} w_{ij} \in PA_\rho$ by the antecedent $D_i[\vec{x}] \stackrel{a_{ij}}{\rightarrow} w_{ij}$; however, this is not GSOS format. We instead will replace t by the antecedents of a ruloid ρ' giving $D_i[\vec{x}]$ an a_{ij}-transition to $D'_{ij}[\vec{x},\vec{y}]$, and use $D'_{ij}[\vec{x},\vec{y}]$ in the place of w_{ij}. We will, of course, have to choose ρ''s in all possible ways for each t.

Let FP_ρ be the set of functions from PA_ρ to ruloids, such that for each $f \in FP_\rho$ and $t = \left(z_i \stackrel{a_{ij}}{\rightarrow} w_{ij}\right) \in PA_\rho$, $f(t) \in \mathcal{R}(D_i, a_{ij})$, where we rename the target variables y_{ij} if necessary in the $\mathcal{R}(D_i,a_{ij})$ to ensure that they are distinct. Note that if there are no ruloids in $\mathcal{R}(D_i,a_{ij})$ (*e.g.*, if $D_i[\vec{x}] = x_1 \backslash a_{ij}$) then there are no such f's, and our construction will leave Δ_ρ empty. Let

$$A^+_{\rho,f} = \bigcup_{t \in PA_\rho} \mathsf{ante}(f(t)) \tag{6}$$

Let σ range over substitutions, and let $P_i = \sigma(x_i)$. Then FP_ρ is sound and specifically witnessing for the set of transition formulas $D_i[\vec{x}] \stackrel{a_{ij}}{\rightarrow} w_{ij}$, in the sense that:

1. If for each i, j, $D_i[\vec{P}] \stackrel{a_{ij}}{\rightarrow} Q_{ij}$, then for some f and σ, we have $\sigma \models A^+_{\rho,f}$, and $\mathsf{cons}(f(t))\sigma = D_i[\vec{P}] \stackrel{a_{ij}}{\rightarrow} Q_{ij}$ for each t.

2. If $\sigma \models A^+_{\rho,f}$, then for each i,j, $D_i[\vec{P}] \stackrel{a_{ij}}{\rightarrow} D'_{ij}[\vec{P}, \sigma(y_{ij})]$.

The negative antecedents require a bit more work. To translate the antecedent $z_i \; {}^{b_{ik}}$, we must provide evidence that no ruloid for D_i and b_{ik} can apply to $\vec{P}$. We do this by choosing an antecedent for each ruloid in $\mathcal{R}(D_i,b_{ik})$, and asserting its opposite. (If some ruloid has no antecedents, then $D_i[\vec{P}] \stackrel{b_{ik}}{\rightarrow}$ always, so rule ρ cannot fire with $D_i[\vec{P}]$ as the i'th argument; our construction will leave Δ_ρ empty.)

The opposite of the transition formula $x_i \xrightarrow{a_{ij}} y_{ij}$ of course $x_i \ {}^{a_{ij}}$. The opposite of $x_i \ {}^{b_{ik}}$ clearly must have the form $x_i \xrightarrow{b_{ik}} y_{ik}$; however, when we are doing this, we must take care to avoid duplicate uses of variables. The details of the renaming are straightforward but tedious, and we omit them and pretend that $\mathsf{opp}(\cdot)$ is simply a function on formulas.

Let $x_i \ {}^{b_{ik}}$ be a negative hypothesis of ρ. Let $G_{\rho,i,k}$ be the set of functions g mapping $\mathcal{R}(D_i[\vec{x}], b_{ik})$ to transition formulas, such that for each ruloid $\rho' \in \mathcal{R}(D_i[\vec{x}], b_{ik})$, $g(\rho') \in \mathsf{ante}(\rho')$.

Let

$$O_{\rho,g} = \{\mathsf{opp}(g(\rho'))|\rho' \in \mathsf{dom}(g)\} \tag{7}$$

By the inductive hypothesis,

1. If $\sigma \models O_{\rho,g}$, then $D_i[\vec{x}]\sigma \ {}^{b_{ik}}$, and
2. If $D_i[\vec{x}]\sigma \ {}^{b_{ik}}$, then each rule ρ' must fail because some antecedent $g_{\rho'}$ is not satisfied. Let $g(\rho') = g_{\rho'}$ for all ρ'; then $\sigma \models O_{\rho,g}$.

Let H_ρ be the set of functions h from indices $\langle i, k\rangle$ of negative antecedents, such that $h(i,k) \in G_{\rho,i,k}$. Let

$$A^-_{\rho,h} = \bigcup_{i,k} O_{\rho,h(i,j)} \tag{8}$$

So, $A^-_{\rho,h}$ gives sufficient conditions for each negative antecedent of ρ to be satisfied; and varying h gives all ways for them to be satisfied.

We finally define

$$\Delta_\rho = \bigcup_{f \in FP_\rho, h \in H_\rho} \frac{A^+_{\rho,f} \cup A^-_{\rho,h}}{D[\vec{x}] \xrightarrow{a} E_f[\vec{x}, \vec{y}]} \tag{9}$$

where $E_f[\vec{x}, \vec{y}] = E\Big[D_i[\vec{x}], D'_{ij}[\vec{x}, \vec{y}]\Big]$ where D'_{ij} is the target of $f(z_i \xrightarrow{a_{ij}} w_{ij})$. ¿From the remarks about the $A^+_{\rho,f}$ and $A^-_{\rho,h}$ constructions, we see that Δ_ρ is sound and specifically witnessing for transitions from $D[\vec{x}]$ by rule ρ.

Finally,

$$\mathcal{R}(D, a) = \bigcup_\rho \Delta_\rho \tag{10}$$

□

Definition 7.7 *The* ruloid set *of a GSOS language $\mathcal{G}$ is the union of the sets $\mathcal{R}(D, c)$ given by Theorem* **??**.

Clearly, $D[\vec{P}] \xrightarrow{c} P'$ iff there is a ruloid in the ruloid set of $\mathcal{G}$ with consequent of the form $D[\vec{X}] \xrightarrow{c} C[\vec{X}, \vec{Y}]$ specifically witnessing this transition. Now it is fairly straightforward to prove Theorem **??**.

Lemma 7.8 *Let $\mathcal{G}$ be a GSOS language, and trees P and Q in $\mathcal{G}$. If $P \ {}_{\backsim} \ Q$, then $C[P] \ {}_{\backsim} \ C[Q]$ for each $\mathcal{G}$-context $C[X]$.*

Proof:

To show that $C[P] \_\rightharpoonup C[Q]$, it suffices to exhibit a ready simulation relation $\_\rightharpoonup^\star$ such that $C[P] \_\rightharpoonup^\star C[Q]$. The obvious candidate is the congruence extension of $\_\rightharpoonup$ itself, defined by $D[\vec{R}] \_\rightharpoonup^\star D[\vec{S}]$ whenever $\vec{R}$ and $\vec{S}$ are vectors of trees of the right length such that $R_i \_\rightharpoonup S_i$ for each i. It remains to show that $\_\rightharpoonup^\star$ is a ready simulation relation.

Suppose that $D[\vec{R}] \stackrel{c}{\rightarrow} R'$ for some R'. By Theorem **??**, this is true precisely if there is some ruloid ρ in the ruloid set of $\mathcal{G}$

$$\frac{X_i \stackrel{a_{ij}}{\rightarrow} Y_{ij}, \qquad X_i \; b_{ik}}{D[\vec{X}] \stackrel{c}{\rightarrow} C[\vec{X}, \vec{Y}]}$$

and trees R'_{ij} such that $R_i \stackrel{a_{ij}}{\rightarrow} R'_{ij}$, $R_i \; b_{ik}$, and $R' = C[\vec{R}, \vec{R'}]$.

As each $R_i \_\rightharpoonup S_i$, we know that (1) there are S'_{ij} such that $S_i \stackrel{a_{ij}}{\rightarrow} S'_{ij}$ and $R'_{ij} \_\rightharpoonup S'_{ij}$ for each i, j and (2) $S_i \; b_{ik}$ for each i, k. So, by ρ, we know that $D[\vec{S}] \stackrel{c}{\rightarrow} C[\vec{S}, \vec{S'}] = S'$. By definition of $\_\rightharpoonup^\star$, we know that

$$R' = C[\vec{R}, \vec{R'}] \;\; \_\rightharpoonup^\star \;\; C[\vec{S}, \vec{S'}] = S'$$

and so we have verified the first half of the definition of a ready simulation relation.

The second half is similar; if $D[\vec{R}]$ is unable to take a c-step, then some hypothesis of each ruloid which could allow it to take a c-step must fail. From the fact that $R_i \_\rightharpoonup S_i$, we discover that the corresponding hypothesis of each ruloid fails for $D[\vec{S}]$ as well, and so $D[\vec{S}] \; c$ as desired. □

This completes the proof of Lemma **??**. To finish Theorem **??**, we must show:

Lemma 7.9 *For all synchronization trees P and Q, if $P \_\rightharpoonup Q$, then $P \sqsubseteq_{tr} Q$.*

Proof:

Suppose that $P \_\rightharpoonup Q$ and $P \stackrel{s}{\rightarrow} P'$ with P' stopped. There is a sequence of processes $P = P_0, P_1, \ldots, P_n = P'$ such that

$$P_0 \stackrel{s_1}{\rightarrow} P_1 \stackrel{s_2}{\rightarrow} \cdots \stackrel{s_n}{\rightarrow} P_n \text{ stopped}$$

By definition, we have $P \_\rightharpoonup Q$. From the definition of $\_\rightharpoonup$, we know that there are processes $Q = Q_0, Q_1, \ldots, Q_n$ such that $Q_i \stackrel{s_{i+1}}{\rightarrow} Q_{i+1}$ and $P_i \_\rightharpoonup Q_i$ for each i. We have

$$Q_0 \stackrel{s_1}{\rightarrow} Q_1 \stackrel{s_2}{\rightarrow} \cdots \stackrel{s_n}{\rightarrow} Q_n.$$

It remains to show that Q_n is stopped. We have $P_n \_\rightharpoonup Q_n$, and so readies(P_n) = readies(Q_n). However, P_n is stopped, and so readies$(P) = \;$. Therefore Q_n is stopped, and so s is a completed trace of Q as desired. □

Theorem **??** now follows routinely. Suppose that $P \_\rightharpoonup Q$, and $C[X]$ is a context in a GSOS language $\mathcal{G}$. We have $C[P] \_\rightharpoonup C[Q]$ by Lemma **??**, and then $C[P] \sqsubseteq_{tr} C[Q]$ by Lemma **??**. Hence $P \sqsubseteq_{tr}^{\mathcal{G}} Q$. As this holds for all $\mathcal{G}$, we have $P \sqsubseteq_{tr}^{GSOS} Q$.

8 A Modal Characterization of Ready Simulation

Recall from Theorem **??** that bisimulation of finitely branching processes coincides with equivalence with respect to Hennessy-Milner formulas. A similar fact holds for ready simulation. The modal logic is useful for some purposes; *e.g.*, it characterizes the properties preserved by ready simulation. Also, the modal characterization is mathematically useful; in Section **??**, we use the modal characterization to show that ready simulation is precisely GSOS congruence.

The class of *denial formulas* is

$$\varphi ::== \mathrm{tt} \mid \mathrm{ff} \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle a \rangle \varphi \mid \neg a \tag{11}$$

The notion of satisfaction is the same for HML formulas and denial formulas, Definition **??**, with the additional clause $P \models \neg a$ iff $P \stackrel{a}{\not\rightarrow}$. Notice that $\neg a$ is equivalent to a restricted use of the $[a]$ modality, *viz.* $[a]\mathrm{ff}$, and we do not allow the full use of this modality. Denial logic is not closed under negation, either syntactically or semantically; for example, neither kind of negation of the formula $\langle a \rangle \langle a \rangle \mathrm{tt}$ is a denial formula.

Definition 8.1

1. *$P \sqsubseteq_{DL} Q$ iff for all denial formulas φ, $P \models \varphi$ implies $Q \models \varphi$.*

2. *$P \equiv_{DL} Q$ iff $P \sqsubseteq_{DL} Q$ and $Q \sqsubseteq_{DL} P$.*

Equivalently, $P \equiv_{DL} Q$ iff for all denial formulas φ, $P \models \varphi$ iff $Q \models \varphi$.

Theorem 8.2 [**?**] *If P and Q are finitely-branching synchronization trees, then*

- *$P \mathrel{\_\!\!\rightharpoonup} Q$ iff $P \sqsubseteq_{DL} Q$.*
- *$P \;\; Q$ iff $P \equiv_{DL} Q$.*

Proof: The second half follows from the first. Suppose that $P \mathrel{\_\!\!\rightharpoonup} Q$. We show that $P \models \varphi$ implies $Q \models \varphi$ by induction on φ simultaneously for all P and Q.

1. tt and ff are trivial.

2. Suppose $P \models \varphi \wedge \psi$. Then $P \models \varphi$ and $P \models \psi$, and by induction we have $Q \models \varphi$ and $Q \models \psi$ and hence $Q \models \varphi \wedge \psi$ as desired. Disjunctions are similar.

3. Suppose that $P \models \langle a \rangle \varphi$. Then there is a P' such that $P \stackrel{a}{\rightarrow} P'$ and $P' \models \varphi$. As $P \mathrel{\_\!\!\rightharpoonup} Q$, there is a Q' such that $P' \mathrel{\_\!\!\rightharpoonup} Q'$ and $Q \stackrel{a}{\rightarrow} Q'$. By induction, $Q' \models \varphi$; hence $Q \models \langle a \rangle \varphi$.

4. Suppose that $P \models \neg a$. Then $P \stackrel{a}{\not\rightarrow}$, and so $Q \stackrel{a}{\not\rightarrow}$; which is to say $Q \models \neg a$.

To prove the converse, we show that $\sqsubseteq_{DL}$ is a ready simulation relation. Suppose that $P \sqsubseteq_{DL} Q$.

- Suppose that $P \xrightarrow{a} P'$. We must show that there is some Q' such that $Q \xrightarrow{a} Q'$ and $P' \sqsubseteq_{DL} Q'$. Suppose for contradiction that there is no a-child Q' of Q such that $P' \sqsubseteq_{DL} Q'$. Q has a finite number of a-children, $Q_1, \ldots, Q_n$. For each child Q_i, there is a formula ψ_i such that $P' \models \psi_i$ but $Q_i \not\models \psi_i$. Let $\psi = \psi_1 \wedge \cdots \wedge \psi_n$; if there are no children, then let $\psi = \mathtt{tt}$. Then $P' \models \psi$ and so $P \models \langle a \rangle \psi$. However, $Q \not\models \langle a \rangle \psi$, which violates the assumption that $P \sqsubseteq_{DL} Q$

- Suppose that $P \ ^a$. Then $P \models \neg a$, and so $Q \models \neg a$. This is equivalent to $Q \ ^a$ as desired.

We have shown that $\sqsubseteq_{DL}$ is a ready simulation relation, and so $P \sqsubseteq_{DL} Q$ implies $P \_\!\!\sim Q$. $\Box$

In fact, the full syntax of denial formulas is not required; disjunctions and ff are not necessary. In particular, the formulas $\langle a \rangle(\varphi \vee \psi)$ and $(\langle a \rangle \varphi) \vee (\langle a \rangle \psi)$ are logically equivalent. The *essential denial formulas* are given by the syntax:

$$\varphi ::== \mathtt{tt} \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid \neg a.$$

Lemma 8.3 *$P \equiv_{DL} Q$ iff P and Q agree on all essential denial formulas.*

Proof: Use the fact that $\langle a \rangle(\varphi \vee \psi)$ and $(\langle a \rangle \varphi) \vee (\langle a \rangle \psi)$ are logically equivalent, and the other rules of modal logic. $\Box$

9 Ready Simulation Can Be Traced

In this section we introduce an extension CCSSS of CCS whose congruence is just ready simulation; that is, $P \equiv_{DL} Q$ iff $P \equiv_{tr}^{CCSSS} Q$. We add two operations. P is a copying operator: when P signals that it wants to fork, P forks. $S \triangleright P$ is a sort of controlled communication: S runs alone, except that it occasionally allows P the ability to take a step and communicate with it. These operations correspond to the copying and button-pushing operations in the testing scenarios of [**?**, **?**].

Using these operations, we will code denial formulas into contexts and traces, and so understand ready simulation in CCSSS. $C_\varphi[P]$ tests the process P to see if it satisfies φ, producing a characteristic kind of trace if it does and not if it does not.

Formally, we fix several distinct actions. We will use o as a sort of "visible silent action;" processes will emit o's while they are operating. The actions c_1 and c_2 are used by processes to signal to the operator that they wish to fork. In $S \triangleright P$, S uses the d action to signal that it wishes to communicate with P. There is an auxiliary operator $\triangleright$ used by $\triangleright$.

(P) usually does just what P does. However, when P signals that it wants to be forked (by the c_1 and c_2 actions),(P) forks it.

$$\frac{X \xrightarrow{a} X' \quad (a \notin \{c_1, c_2\})}{X \xrightarrow{a} X'} \qquad \frac{X \xrightarrow{c_1} X_1, X \xrightarrow{c_2} X_2}{X \xrightarrow{o} (X_1) \parallel (X_2)}$$

$S \triangleright P$ usually does just what S does; P is frozen. However, when S signals that it wants to communicate with P (by performing a d-step), $S \triangleright P$ unfreezes P and lets it take a step

in cooperation with S. This operation needs a bit of control state, telling whether or not P is frozen; we use the $\triangleright$ operator when P is frozen, and the $\triangleright$ operator when P is unfrozen.

$$\frac{S \xrightarrow{a} S' \quad (a \neq d)}{S \triangleright P \xrightarrow{a} S' \triangleright P} \qquad \frac{S \xrightarrow{d} S'}{S \triangleright P \xrightarrow{o} S' \triangleright P}$$

The operation $\triangleright$ does one step of communication and then behaves like $\triangleright$.

$$\frac{S \xrightarrow{a} S', \quad P \xrightarrow{a} P'}{S \triangleright P \xrightarrow{o} S' \triangleright P'}$$

We now define the coding of essential formulas. To make strings of actions easier to read, we write prefixing with a dot: "$d.a.t.S$" instead of "$datS$." Fix two actions t and f, distinct from the previously-mentioned actions, which we use for partial success and total failure.

$$\begin{aligned} S_{tt} &= \mathbf{0} \\ S_{\neg a} &= d.a.f \\ S_{\varphi \wedge \psi} &= c_1 S_\varphi + c_2 S_\psi \\ S_{\langle a \rangle \varphi} &= d.a.t.S_\varphi \end{aligned}$$

The context $C_\varphi[X]$ is defined to be $(S_\varphi \triangleright X)$. $C_\varphi[P]$ will compute, emitting o's while it is working. Each time it processes an $\langle a \rangle$ correctly, it will emit a t. Each time it fails to perform a $\neg a$ correctly, it will emit an f. We will show that $P \models \varphi$ iff $C_\varphi[P]$ produces a trace with enough t's and no f's.

For example,

$$\begin{aligned} C_{\langle a \rangle tt \wedge \langle b \rangle tt}[a+b] &= \big((c_1.d.a.t + c_2.d.b.t) \triangleright (a+b)\big) \\ &\xrightarrow{o} \big(d.a.t \triangleright (a+b)\big) \parallel \big(d.b.t \triangleright (a+b)\big) \\ &\xrightarrow{o} \big(a.t \triangleright (a+b)\big) \parallel \big(d.b.t \triangleright (a+b)\big) \\ &\xrightarrow{o} (t \triangleright \mathbf{0}) \parallel \big(d.b.t \triangleright (a+b)\big) \\ &\xrightarrow{t} (\mathbf{0} \triangleright \mathbf{0}) \parallel \big(d.b.t. \triangleright (a+b)\big) \\ &\xrightarrow{o} (\mathbf{0} \triangleright \mathbf{0}) \parallel \big(b.t \triangleright (a+b)\big) \\ &\xrightarrow{o} (\mathbf{0} \triangleright \mathbf{0}) \parallel (t \triangleright \mathbf{0}) \\ &\xrightarrow{t} (\mathbf{0} \triangleright \mathbf{0}) \parallel (\mathbf{0} \triangleright \mathbf{0}) \end{aligned}$$

To illustrate how the testing for $\neg a$ works, consider:

$$\begin{aligned} C_{\neg a}[(a+b)] &= \big(d.a.f \triangleright (a+b)\big) \\ &\xrightarrow{o} \big(a.f \triangleright (a+b)\big) \\ &\xrightarrow{o} (f \triangleright \mathbf{0}) \\ &\xrightarrow{f} (\mathbf{0} \triangleright \mathbf{0}) \end{aligned}$$

So, the only trace of $C_{\neg a}[a+b]$ contains an f. However, the only computation of

$$C_{\neg a}[b] = (d.a.f \triangleright b) \xrightarrow{o} (a.f \triangleright b)$$

gets stuck after performing an o.

Define $\lfloor \varphi \rfloor$ to be the number of $\langle a \rangle$'s occurring in φ; that is:

$$\begin{aligned} \lfloor tt \rfloor = \lfloor \neg a \rfloor &= 0 \\ \lfloor \varphi \wedge \psi \rfloor &= \lfloor \varphi \rfloor + \lfloor \psi \rfloor \\ \lfloor \langle a \rangle \varphi \rfloor &= 1 + \lfloor \varphi \rfloor \end{aligned}$$

We say that a trace s is φ-happy if it contains exactly $\lfloor \varphi \rfloor$ t's and no f's. A trace is φ-sad if it it contains fewer than $\lfloor \varphi \rfloor$ t's, or at least one f.

Lemma 9.1 *If $P \models \varphi$ then $C_\varphi[P]$ has a φ-happy trace. If $P \not\models \varphi$, then all traces of $C_\varphi[P]$ are φ-sad. Furthermore, no trace of $C_\varphi[P]$ has more than $\lfloor \varphi \rfloor$ t's.*

Proof:

The proof is by induction on φ.

$\varphi = tt$**:** $C_{tt}[P]$ is stopped for all P, as desired.

$\varphi = \psi \wedge \theta$**:** $C_{\psi\wedge\theta}[P] = ((c_1 S_\psi + c_2 S_\theta) \triangleright P)$. As $((c_1 S_\psi + c_2 S_\theta) \triangleright P)$ can make both c_1 and c_2 transitions, the forks the process:

$$C_{\psi\wedge\theta}[P] = ((c_1 S_\psi + c_2 S_\theta) \triangleright P) \xrightarrow{o} (S_\psi \triangleright P) \parallel (S_\theta \triangleright P) = C_\psi[P] \parallel C_\theta[P]$$

The lemma follows from the ordinary properties of sequences and interleaving.

$\varphi = \neg a$

$$C_{\neg a}[P] = (d.a.f \triangleright P) \xrightarrow{o} (a.f \triangleright P)$$

If $P \not\xrightarrow{a}$, then $(a.f \triangleright P)$ cannot move and the trace is simply the φ-happy trace o.

If $P \xrightarrow{a} P'$, then

$$(a.f \triangleright P) \xrightarrow{o} (f \triangleright P') \xrightarrow{f} (\mathbf{0} \triangleright P')$$

In this case, the trace of $C_{\neg a}[P]$ is oof, which is φ-sad.

$\varphi = \langle a \rangle \psi$

$$C_{\langle a\rangle\psi}[P] \xrightarrow{o} (a.t.S_\psi \triangleright P)$$

Consider any P' such that $P \xrightarrow{a} P'$. (If there are no such P''s, then the process is stuck and the trace is φ-sad as required.)

$$(a.t.S_\psi \triangleright P) \xrightarrow{o} (t.S_\psi \triangleright P') \xrightarrow{t} (S_\psi \triangleright P') = C_\psi[P']$$

If $P \models \varphi$, then there is a P' such that $P \xrightarrow{a} P' \models \psi$. By the induction hypothesis $C_\psi[P']$ has a ψ-happy trace, and so we have found a φ-happy trace of $C_\varphi[P]$.

If $P \not\models \varphi$, then $P' \not\models \psi$, and so every trace of $C_\psi[P]$ is ψ-sad; thus every trace of $C_\varphi[P]$ that goes through $C_\psi[P']$ is φ-sad. Every such trace must go through some $C_\psi[P']$, and so every trace of $C_\varphi[P]$ must be φ-sad.

The requirement that no trace have more than $\lfloor\varphi\rfloor$ t's is routine.

□

10 Summary of Ready Simulation

Combining the results of the previous sections, we obtain the following set of equivalent characterizations of GSOS congruence.

Theorem 10.1 *The following are equivalent:*

1. $P \mathrel{\_\!\!\sim} Q$ *(resp.* $P \;\; Q$*). (State-correspondence definition)*
2. $P \sqsubseteq_{tr}^{GSOS} Q$ *(resp.* $P \equiv_{tr}^{GSOS} Q$*). (Approximation in all GSOS languages.)*
3. $P \sqsubseteq_{DL} Q$ *(resp.* $P \equiv_{DL} Q$*) (Approximation with respect to all denial formulas)*
4. $P \models \varphi$ *implies* $Q \models \varphi$ *(resp.* $P \models \varphi$ *iff* $Q \models \varphi$ *) for all essential denial formulas* φ*.*
5. $P \sqsubseteq_{tr}^{CCSSS} Q$ *(resp.* $P \equiv_{tr}^{CCSSS} Q$*) (Trace approximation in CCSSS)*

Corollary 10.2 *Bisimulation is a strict refinement of ready simulation, and hence of GSOS congruence. In particular, the processes* $P_\star$ *and* $Q_\star$ *are trace congruent with respect to every GSOS language, although they are not bisimilar.*

There are a few other definitions of ready simulation, but they are of less interest. For example, it is possible to define the n^{th} approximant to ready simulation in the way that [**?**, **?**] defined the n^{th} approximant to bisimulation; predictably, if $P \mathrel{\_\!\!\sim_n} Q$ for all n, then $P \mathrel{\_\!\!\sim} Q$.

11 Conclusion

Should bisimulation play a significant role in process theory? It has many nice properties, a rich theory, and a tested methodology for verifying correctness of genuine, nontrivial protocols (see, *e.g.*, [**?**, **?**, **?**]). Nevertheless, we find unconvincing the arguments for taking bisimulation as a primitive notion. We maintain that computational distinctions should be made only because of observable differences "at the terminal." Global testing systems and modal logics which reduce bisimulation to such observations do not offer what we regard as a reasonable framework for defining operations on processes. We prefer to regard bisimulation as a mathematical tool which is frequently useful in proving programs correct, rather than a

characterization of what correctness should mean. In other work, the first author has used bisimulation methods to verify a silicon compilation scheme [**?**]. The compiler was correct up to bisimulation, and the correctness proof up to bisimulation was no harder than the proof up to trace congruence, so we proved the stronger theorem.

Ready simulation seems to have the mathematical properties which make bisimulation desirable. Bisimulation has several equivalent but distinct definitions, of which the existence of a bisimulation relation and equivalence with respect to a modal logic and their variants are the most useful. Ready simulation has similar definitions. Moreover, the proofs in this paper indicate that those definitions play the same role for ready simulation as they do for bisimulation; *e.g.*, the proof of Theorem **??** given in Section **??** uses the existential definition, and the proof in Section **??** that CCSSS congruence is precisely ready simulation uses the modal characterization. These results, and similar work on other aspects of ready simulation, suggest that the theory of ready simulation is very similar to that of bisimulation in character and power. Of course, ready simulation is easily justified on observational grounds, while bisimulation (despite its other name of "observational equivalence") is harder to justify.

The larger purpose of this study is to illuminate some of the issues one might wish to consider in the choice of a notion of process equivalence. We chose bisimulation as the focus of our study precisely because it had no obvious computational definition. We have given some indications of the form that a computational justification of a notion of equivalence might take: congruence with respect to some sort of well-structured language. Other forms of justification are certainly possible. However, some justification should be considered for each new notion of program equivalence; otherwise, the notion runs the risk of being unjustifiable in computational terms despite having an elegant and powerful mathematical theory. Finally, we have demonstrated that computational justification need not be incompatible with mathematical elegance.

11.1 Related Work

We have sketched the fragment of the theory of ready simulation appropriate to our discussion. There are other questions that one might wish to answer; the answers generally seem to be fairly pleasing. For example, ready simulation of finite processes has a finite axiomatization as an inequational theory [**?**]; and there is a $O(mn + n^2)$ algorithm for computing if two n-state, m-transition automata are ready similar [**?**].

In general in computer science, two programs are considered equivalent iff they are *congruent*, that is, iff one can be substituted for the other in any context and no difference is observable. This definition has two parameters: the language over which contexts can be formed, and the differences which can be seen. In this paper, we have been varying the language over GSOS and global testing languages; however, the observations have always been finite completed traces. Elsewhere, we examine other notions of observation [**?**, **?**].

Bisimulation seems to require two kinds of knowledge, the knowledge of the *possible* and *necessary* behavior of a process, interleaved arbitrarily. The possible behavior appears quite naturally in nondeterminism; each execution of a process gives a possible behavior. The necessary behavior is harder to observe. Larsen and Skou [**?**] have proposed the use

of probability, testing a process often enough to observe all possible behaviors with high probability. There is a very strong connection between ordinary and probabilistic bisimulation, and a tantalizing and debatable failure of the use of probability as a mechanism for observing bisimulation; this is discussed in [**?**, **?**].

A variation on the theme of well-structured rules is explored in [**?**]. Vaandrager and Groote allow their languages to be countably and undecidably branching. The so-called *tyft/tyxt* rules are positive rules allowing operators in antecedents; the more paradoxical operations defined in the appendix are not allowed. Ready simulation need not be a congruence with respect to such a language. The appropriate notion is the so-called $\diamond\Box$ equivalence, *viz.* equivalence with respect to Hennessy-Milner formulas in which no $\langle a \rangle$ operator occurs within the scope of a $[a]$. This is precisely tyft/tyxt congruence; in particular, bisimulation cannot be understood as congruence with respect to any language in this format.

11.2 Open Problems

Throughout this paper, we have assumed that there is no silent (hidden or "τ") action. We expect that similar results hold when there is a silent action; this remains to be verified. Silent moves make the philosophy as well as the mathematics trickier; it is no longer clear what the right questions are.

The first tricky issue is the role of weak bisimulation (Milner's "observational equivalence"). Strong bisimulation is generally accepted as an equivalence which all languages ought to respect. However, weak bisimulation is not a congruence with respect to the $+$ operation of CCS; a and τa are weakly bisimilar, but $a + b$ and $\tau a + b$ are not. In our study, it was clear that we insist that any GSOS language respect strong bisimulation; the corresponding requirement for weak bisimulation is not clear. There are several possible ways to develop the theory. For example, in CCS, weak bisimulation congruence[13] has a simple characterization as *rooted bisimulation*, or equivalently the congruence closure of bisimulation with respect to $+$ [**?**]. We could insist that all operations respect this relation. Several variants of weak bisimulation have been proposed, such as the *branching bisimulation* of [**?**], which has better equational properties than weak bisimulation. Alternatively, we could simply let bisimulation congruence be whatever it likes, as long as it respects strong bisimulation.

Even the definition of trace congruence is no longer obvious. In the (oversimplified) theory of this study, there was only one way that a process could not perform a visible action: if it is a stopped process. Given a silent move, there are now more ways that a process can do nothing visible. Consider the processes $\mathbf{0}$ (which does nothing), τ (which computes a while and then does nothing), $\tau(\tau + \tau\tau\tau)$ (which does some more complex computation before doing nothing), τ^ω (which thinks forever about what it should do next), and $\tau(\tau + \tau^\omega)$ (which may or may not think forever). The proper identifications between these processes are not obvious; and even the criteria for choosing between the possible choices are subject to some debate.

[13] P and Q are weak bisimulation congruent iff for all contexts $C[\cdot]$, $C[P]$ and $C[Q]$ are weakly bisimilar, *i.e.*, if P and Q are congruent considering weak bisimulation to be observable.

We allowed negative rules because they gave extra programming power and because our theory could handle them. There are several possibilities for combining negative rules and silent moves. Ignoring them would leave a rich theory; leaving them interpreted as they are now might also give an appropriate theory. More elaborate combinations are probably reasonable as well; for example, the hypothesis $P\ ^a$ could be interpreted as being satisfied iff the set $\left\{P' : P \stackrel{\tau^*}{\to} P'\right\}$ is finite and $P'\ ^a$ for each such P', or $P\ ^a \wedge P\ ^\tau$ as in [**?**].

Our general program remains to be elaborated for other notions of concurrency. One of the most powerful features of bisimulation is that it can be sensibly and helpfully adapted to most settings; it remains to be seen how portable ready simulation is. We have attempted to give some ways of judging basic notions of program equivalence. It may be possible to subject other theories of concurrency to the same sort of foundational treatment, which may give some assistance to the uninitiated in choosing between various kinds of theories.

12 Acknowledgments

We would like to thank Frits Vaandrager and Jan Friso Groote for many helpful email discussions, and some of the tricks used in the proof of Section **??**. Kim Larsen and Rob van Glabbeek independently discovered the characterization of ready simulation in Definition **??** and brought it to our attention, which greatly simplified the proof of Theorem **??** from the original proof directly using the modal characterization. Rob van Glabbeek also pointed out that negative GSOS rules are not necessary to capture ready simulation. We would also like to thank Samson Abramsky, Luca Aceto, Ashvin Dsouza, Georges Lauri, Robin Milner, Jon Riecke, Arie Rudich, and Sam Weber in alphabetical order for their several contributions. The first author would also like to thank Vicki Borah for boundless emotional support.

A Counterexamples for Section ??

In this section, we list more counterexamples showing that the GSOS format cannot be extended in other obvious ways.

(i) We do not allow the use of variables to do pattern-matching. For example, the following rule has two different variables taking a-steps and then becoming the same variable:

$$\frac{X_1 \stackrel{a}{\to} Y, \quad X_2 \stackrel{a}{\to} Y}{\beta(X_1, X_2) \stackrel{a}{\to} Y}$$

The context $C[Z] = \beta(a\mathbf{0}, aZ)$ distinguishes between the bisimilar processes $\mathbf{0}$ and $\mathbf{0} + \mathbf{0}$. Copying — using a source variable X_i more than once in the antecedent — does not do any harmful pattern-matching and is acceptable. Duplicated source variables in the term $\mathsf{op}(X_1, \ldots, X_n)$, however, allow too much pattern-matching to respect bisimulation; *e.g.*, the rule $\delta(X, X) \stackrel{a}{\to} \mathbf{0}$ causes as much trouble here as its cousin does in the λ-calculus.

(ii) We have insisted that the positive hypotheses be $X \stackrel{a}{\to} Y$, but the negative ones $X\ ^b$. There is no point to hypotheses $X \stackrel{a}{\to}$; simply use one of the form $X \stackrel{a}{\to} Y$ and ignore

Y. On the other hand, there is a gain of expressive power to negative hypotheses $X \not\xrightarrow{b} Y$. The gain is all in the wrong direction. It is now easy to write some countably-branching processes, but it is hard to see how the power of these rules could be useful.

$$\frac{X \not\xrightarrow{b} Y}{\xi(X) \xrightarrow{a} Y}$$

$\xi(\mathbf{0}) \xrightarrow{a} P$ for every process P; it is thus countably branching.

(iii) We gain expressive power if we allow terms rather than simply variables to appear in antecedents. If the terms appear as targets — $X \xrightarrow{a} (Y + \mathbf{0})$ — we have too much pattern-matching, and examples similar to the previous ones will show that bisimulation need not be a congruence. Terms appearing as the sources cause a more subtle problem: bisimulation will still be a congruence, but a process may be infinitely branching.[14] As an example, consider the nullary operation ω defined by two rules

$$\omega \xrightarrow{a} \mathbf{0} \qquad \frac{\omega \xrightarrow{a} Y}{\omega \xrightarrow{a} aY}$$

Then, ω is not finitely branching, for we have:

$$\begin{array}{rcl} \omega & \xrightarrow{a} & \mathbf{0} \\ \omega & \xrightarrow{a} & a\mathbf{0} \\ \omega & \xrightarrow{a} & aa\mathbf{0} \\ & \vdots & \end{array}$$

This occurs because the operation ω is defined by what amounts to an unguarded recursion: it mentions its own one-step behavior in the definition of its one-step behavior. In some circumstances, this might be a useful sort of operation; *e.g.*, we could define $\omega'(X) \equiv \sum_{i=0}^{\infty} a^i X$, a process which dawdles for a finite but unbounded time before performing X.

We could also repair this flaw by allowing terms in antecedents, but not allow any operation to refer to itself either directly or indirectly. The resulting class of languages guarantee the essential properties of GSOS languages. In fact, they have almost exactly the same properties: every operation definable with terms in antecedents is definable without them as well. On the other hand, if terms may appear in antecedents, structural induction is no longer a viable proof method.

(iv) Another possible extension is allowing multi-step antecedents [**?**]. It is in general inconsistent to have both negative and multi-step antecedents. For example, consider the operators $\cdot/a$, α, and π defined as follows:

$$\frac{X \xrightarrow{a} Y_1 \xrightarrow{a} Y_2}{X/a \xrightarrow{a} Y_2} \tag{12}$$

[14] In fact the resulting congruence is coarser than bisimulation; see [**?**] for details.

$$\frac{X \not\xrightarrow{a}}{\alpha(X) \xrightarrow{a} \mathbf{0}}$$

$$\pi \xrightarrow{a} \alpha(\pi/a) \tag{13}$$

It is not hard to show that there is no arrow relation which agrees with these rules: π/a can move iff it cannot move, for:

$$\begin{array}{ll} & \pi/a \xrightarrow{a} Z \\ \iff & \pi \xrightarrow{a} Y \xrightarrow{a} Z \\ \iff & Y = \alpha(\pi/a) \xrightarrow{a} Z \\ \iff & \pi/a \not\xrightarrow{a} \text{ and } Z = \mathbf{0} \end{array}$$

This is similar to the fact that unguarded recursion and negation do not mix. The program $\mathrm{fix}\,[X \Leftarrow \alpha(X)]$ is a term involving an unguarded recursion which can take an a-step iff it cannot take an a-step.

It is possible to have some syntactic conditions which guarantee finite branching and existence of a transition relation, but they are necessarily global. The operation π amounts to an unguarded recursion; however, there is nothing about the form of rule (**??**) indicating that it is unguarded. This rule would be perfectly acceptable in a GSOS language, for example; it is turned into an unguarded recursion by the presence of the $\cdot/a$ operator. It is nontrivial to define just what a guarded operator is in the presence of multi-step rules; the operator $\cdot/a$ "unguards" X, and an adequate calculus of guardedness remains to be developed. CSP, which has something like the $\cdot/a$ operator, forbids its use in recursions.

(v) A system with multi-step positive antecedents alone is quite possible; such systems respect bisimulation, and always have a minimal transition relation. [**?**] Two problems arise for such systems. First, sound and witnessing transition relations other than the unique minimal one may exist; second and far more important, such systems admit countably branching processes.

Consider a GSOS language containing the $/a$ operation above, the single action a, and a constant J with the rule

$$J \xrightarrow{a} J/a \tag{14}$$

Intuitively, J is a process which performs an a-move and then does whatever it does after it performs an a-move. This intuitive definition doesn't uniquely determine what J does after its a-move; neither does the formalism of sound and witnessing transition relations.

Consider a proof (used loosely) that $J/a \xrightarrow{a} K$ for some K. It must start with the rule for $/a$ used with $x = J$, and thus have hypotheses $J \xrightarrow{a} J_1$ and $J_1 \xrightarrow{a} K$. J has only one transition, from (**??**), and hence $J_1 = J/a$. That is, any proof that $J/a \xrightarrow{a} K$ is infinitely long, consisting essentially of repeated proofs of $J/a \xrightarrow{a} K$.

Now, we define two transition relations $\dot{\rightarrow}_f$ and $\dot{\rightarrow}_i$, with $P \dot{\rightarrow}_f Q$ there is a finite proof of this, and $P \dot{\rightarrow}_i Q$ if there is an infinite proof. In particular, $J/a \not\xrightarrow{a}_f$ but $J/a \xrightarrow{a}_i K$ for all processes K. Both $\dot{\rightarrow}_f$ and $\dot{\rightarrow}_i$ are easily seen to be sound and witnessing; as they disagree on J/a, they are not equal.

We know of no circumstance in which the non-uniqueness of the sound witnessing transition relation causes any difficulties; more likely, this simply reflects the fact that the definitions of soundness and witnessing are tuned for GSOS languages. There is a least transition relation in all cases, *viz.* $\dot{\rightarrow}_f$, and so a "correct" operational semantics.

(vi) More importantly, even the minimal transition relation is not necessarily finitely branching. Consider the operations $\cdot/a$ of (**??**), and κ defined by:

$$\kappa(X) \stackrel{a}{\rightarrow} X + \kappa(X)/a/a$$

Let M_0 be a synchronization tree with an infinite derivation of distinct terms

$$M_0 \stackrel{a}{\rightarrow} M_1 \stackrel{a}{\rightarrow} M_2 \stackrel{a}{\rightarrow} \cdots$$

Let $N = M_0 + \kappa(M_0)/a/a$; so

$$\kappa(M_0) \stackrel{a}{\rightarrow} N.$$

We will show that N has an infinite number of a-children. Clearly $N \stackrel{a}{\rightarrow} M_1$.

Suppose that

$$N \stackrel{a}{\rightarrow} M_i.$$

We have

$$\kappa(M_0) \stackrel{a}{\rightarrow} N \stackrel{a}{\rightarrow} M_i \stackrel{a}{\rightarrow} M_{i+1}$$

and so

$$\kappa(M_0)/a/a \stackrel{a}{\rightarrow} M_{i+1}.$$

As $N = M_0 + \kappa(M_0)/a/a$, we therefore have

$$N \stackrel{a}{\rightarrow} M_{i+1}.$$

In particular, $N \stackrel{a}{\rightarrow} M_j$ for each $j \geq 1$. As the M_j are all distinct, N has an infinite number of a-children.

(vii) However, the effects of multi-step antecedents are worse than simply infinite branching. The transition relation $\dot{\rightarrow}$ ceases to be decidable. It is straightforward to program Turing machines in a suitable GSOS language, so that $M^{(k)} \stackrel{a^{f(k)}b}{\rightarrow} \mathbf{0}$ if the k^{th} Turing machine halts on blank tape (where $f(k)$ is approximately the number of steps the Turning machine takes), and $M^{(k)}$ takes a-steps forever otherwise. Then $\kappa(M^{(k)})/a/a \stackrel{b}{\rightarrow} \mathbf{0}$ precisely if the k^{th} Turing machine halts on blank tape, which is undecidable. The same example shows that the question $P \stackrel{a}{\rightarrow}$ is undecidable.