

Priority in Process Algebras

Rance Cleaveland

*Department of Computer Science, State University of New York at Stony Brook,
Stony Brook, NY 11794, USA, rance@eos.ncsu.edu*

Gerald Lüttgen

*Institute for Computer Applications in Science and Engineering (ICASE),
NASA Langley Research Center, Hampton, VA 23681, USA, luttgen@icase.edu*

V. Natarajan

*Networking Hardware Division, IBM Corporation,
Research Triangle Park, NC 27709, USA, nataraj@raleigh.ibm.com*

Abstract

This chapter surveys the semantic ramifications of extending traditional process algebras with notions of *priority* that allow for some transitions to be given precedence over others. The need for these enriched formalisms arises when one wishes to model system features such as *interrupts*, *prioritized choice*, or *real-time behavior*.

Approaches to priority in process algebras can be classified according to whether the induced notion of *pre-emption* on transitions is *global* or *local* and whether priorities are *static* or *dynamic*. Early work in the area concentrated on global pre-emption and static priorities and led to formalisms for modeling interrupts and aspects of real-time, such as *maximal progress*, in *centralized* computing environments. More recent research has investigated localized notions of pre-emption in which the *distribution* of systems is taken into account, as well as dynamic priority approaches, i.e., those where priority values may change as systems evolve. The latter allows one to model behavioral phenomena such as scheduling algorithms and also enables the efficient encoding of real-time semantics.

Technically, this chapter studies the different models of priorities by presenting extensions of Milner's *Calculus of Communicating Systems* (CCS) with static and dynamic priority as well as with notions of global and local pre-emption. In each case the operational semantics of CCS is modified appropriately, behavioral theories based on *strong* and *weak bisimulation* are given, and related approaches for different process-algebraic settings are discussed.

Key words: priority, pre-emption, bisimulation.

Contents

1	Introduction	1
1.1	Classification of Approaches to Priority	2
1.2	Summary of this Chapter	3
1.3	Organization of this Chapter	4
2	Basic Language and Notation	4
3	Static Priority and Global Pre-emption	6
3.1	Operational Semantics	7
3.2	Semantic Theory Based on Strong Bisimulation	7
3.3	Semantic Theory Based on Weak Bisimulation	9
3.4	Example	14
3.5	Prioritization and Deprioritization Operators	15
3.6	Extension to Multi-Level Priority Schemes	18
3.7	Concluding Remarks and Related Work	20
4	Static Priority and Local Pre-emption	20
4.1	Locations	22
4.2	Operational Semantics	23
4.3	Semantic Theory Based on Strong Bisimulation	25
4.4	Semantic Theory Based on Weak Bisimulation	28
4.5	Example	31
4.6	Discussion on the Alteration of Some Design Decisions	32
4.7	Extension to Multi-Level Priority Schemes	34
4.8	Camilleri and Winskel's Approach	37
4.9	Relating Both Priority Approaches	39
4.10	Concluding Remarks and Related Work	40
5	Dynamic Priority and Global Pre-emption	41
5.1	Real-Time Semantics	42
5.2	Dynamic Priority Semantics	44
5.3	Relating Dynamic Priority and Real-Time Semantics	46
5.4	Concluding Remarks and Related Work	47
6	Priority in Other Process-Algebraic Frameworks	48
7	Conclusions and Directions for Future Work	50
8	Sources and Acknowledgments	53
	References	53

1 Introduction

Traditional *process algebras* [7,39,42,56] provide a framework for reasoning about the *communication potential* of *concurrent* and *distributed systems*. Such theories typically consist of a simple *calculus* with a well-defined *operational semantics* [1,67] given as *labeled transition systems*; a *behavioral equivalence* is then used to relate implementations and specifications, which are both formalized as terms in the calculus. In order to facilitate *compositional reasoning*, in which systems are verified on the basis of the behavior of their components, researchers have devoted great attention to the definition of *behavioral congruences*, which allow the substitution of “equals for equals” inside larger systems.

Although many case studies (see, e.g., [3]) prove the utility of the process-algebraic approach to system modeling and verification, many systems in practice cannot be modeled accurately within this framework. One reason is that traditional process algebras focus exclusively on expressing the potential *non-determinism* that the interplay of concurrent processes may exhibit; they do not provide any means for encoding differing levels of *urgency* among transitions that might be enabled from a given system state. Typical examples of urgency include:

- **interrupts**, where non-urgent transitions at a state are *pre-empted* whenever an interrupt is raised;
- **programming language constructs**, such as the PRIORITY construct in *occam* [43], that impose an order on transitions;
- **real-time behavior** that is semantically founded on the well-known *synchrony hypothesis* [13] or *maximal progress assumption* [78]; and
- **scheduling algorithms** which also rely on the concept of pre-emption.

In each of these cases urgency provides a means for *restricting nondeterminism*. This mechanism is simply ignored in traditional process algebras. As a consequence, the resulting system models are often not faithful since they contain spurious paths that cannot be traversed by the real-world systems themselves [16,30].

As a simple example of the need for integrating concepts of urgency in process algebra consider the interrupt-based system depicted in Figure 1. It consists of two processes, *A* that flips back and forth between two states and *B* that checks if *A* is running properly. Whenever *B* receives a *check* message it requests status information from *A* via *interrupt port* \underline{i} , which in turn responds by $\overline{\text{ok}}$. In the absence of an indication that a communication on \underline{i} is more urgent than one on *back and forth*, the process *A* can ignore a *check* request indefinitely.




Fig. 1. Back-and-forth example

1.1 Classification of Approaches to Priority

A number of approaches have been proposed for introducing sensitivity to aspects of priority [5,12,16,22–24,27,29,30,33,35,37,45,46,52,53,62,63,69,72,73]. One may classify these approaches according to the following two criteria.

Static vs. dynamic priority:

In *static* approaches transitions are assigned priority values that do not change as the system under consideration evolves. These schemes find application in the modeling of interrupts or prioritized choice constructs. In the former case, interrupts have a fixed urgency level associated with them; in the latter priorities of transitions are fixed by the static program syntax. Almost all priority approaches to process algebra published so far deal with static priorities. The exceptions are [16,22], which present models that allow priority values of transitions to change as systems evolve. Such *dynamic* behavior is useful in modeling scheduling algorithms and real-time semantics.

Global vs. local pre-emption:

This criterion refers to the *scope* of the priority values. In the case of centralized systems, priorities generally have a *global* scope in the sense that transitions in one process may pre-empt transitions in another. We refer to this kind of pre-emption, which has been advocated by Baeten, Bergstra, and Klop [5] and by Cleaveland and Hennessy [27] in the late eighties, as *global pre-emption*. In contrast, in a distributed system containing several sites, transitions should only be allowed to pre-empt those at the same site or *location*. This kind of pre-emption, which was first studied by Camilleri and Winskel [24] in the early nineties, is called *local pre-emption*.

Based on this classification scheme, the body of this chapter investigates the following different semantics for a prototypical process-algebraic language:

static/ global, static/local, and dynamic/global. Only the combination of dynamic priority and local pre-emption is omitted, since research has not yet been carried out on this topic.

Some caveats about terminology are in order here. Other process algebra researchers have used the term “pre-emption” in a setting without priorities [17]; in their usage pre-emption occurs when the *execution* of one transition removes the possibility of another. In our priority-oriented framework, we say that pre-emption occurs when the *presence* of one transition disables other transitions. Berry [12] refers to this latter notion as *must pre-emption* and to the former as *may pre-emption*. In this chapter, whenever we speak of “pre-emption,” we mean “must pre-emption.” It should also be noted that our concept of global pre-emption and dynamic priority differs from the notion of *globally dynamic* priority found in [72]; as the distinction is somewhat technical we defer further discussion on this point to later in the chapter.

1.2 Summary of this Chapter

This chapter surveys existing work on priority in process algebras. In order to focus on some of the technical issues involved with priority and pre-emption, we introduce a simple framework for their illustration. This framework extends Milner’s *Calculus of Communicating Systems* (CCS) [56] and its *bisimulation*-based semantic theory by attaching priority values to actions. Although familiarity with CCS is not a prerequisite for reading this chapter, some knowledge of it would be advantageous since not all standard definitions and notations are re-stated here. For our language three different semantics are given: one reflecting static priorities and global pre-emption, one for static priorities and local pre-emption, and one capturing dynamic priorities. The common language allows for a detailed comparison of the semantic concepts; in addition, the classification scheme presented above helps us to categorize most published approaches to priority. These have been proposed for a variety of well-known process algebras, such as the already mentioned CCS, the *Algebra of Communicating Processes* (ACP) [9], *Communicating Sequential Processes* (CSP) [42], the *Calculus of Broadcasting Systems with Priorities* (PCBS) [69], *Synchronous CCS* (SCCS) [56], and *Asynchronous Communicating Shared Resources* (ACSR) [23].

Technically, for the process algebras with static priority to be presented in this chapter we develop a semantic theory based on the notion of bisimulation [56,65]. Our aim is to carry over the standard algebraic results from CCS [56], including abstractness results as well as axiomatic, logical, and algorithmic characterizations. More precisely, we investigate both *strong* and *weak bisimulations* that are based on naive adaptations of the standard defi-

nitions as given by Milner; we especially characterize the *largest congruences* contained in these relations. These abstractness results indicate that the behavioral relations are semantically adequate and useful for formally reasoning about concurrent and distributed systems. Moreover, we present sound and complete *axiomatizations* for the obtained strong bisimulations with respect to finite processes, i.e., those which do not contain recursion. These axiomatizations testify to the mathematical tractability of the semantic theories presented here. We also characterize the attendant notions of prioritized strong and weak bisimulations as standard bisimulations on alternative transition relations so that well-known *partition-refinement algorithms* [26,48,64] for their computation become applicable. This also allows for establishing logical characterizations of the behavioral relations by adapting *Hennessey-Milner logic* [19,56]. In the case of the dynamic priority semantics, we prove a one-to-one correspondence with traditional *real-time semantics* [10,59,78] in terms of strong bisimulation. Because of this close relationship, semantic theories developed for real-time process algebras can be carried over to the dynamic priority setting.

1.3 Organization of this Chapter

The remainder of this chapter is organized as follows. The next section introduces our language, defines some formal notation used throughout the chapter, and discusses some basic design decisions we have taken. Section 3 presents a semantics of the language based on static priority and global pre-emption; Section 4 then develops a semantics based on static priority and local pre-emption. A dynamic priority approach is illustrated in Section 5. Related work is referred to in each of the last three sections, while Section 6 surveys several priority approaches adopted for different process-algebraic frameworks. Section 7 contains our conclusions and suggestions for future work. The final section points to the sources of the research compiled in this chapter.

2 Basic Language and Notation

As mentioned above, the language considered here is an extension of Milner’s CCS [56], a process algebra characterized by *handshake communication* and *interleaving semantics* for parallel composition. Syntactically, CCS includes notations for *visible actions*, which are either sends or receives on ports, and a distinguished *invisible*, or *internal* action. The semantics of CCS is then given via a transition relation that labels execution steps with actions. When a sender and receiver synchronize, the resulting action is internal. Consequently, transitions labeled by visible actions can be seen as representing only “po-

tential” computation steps, since in order for them to occur they require a contribution from the environment. Transitions labeled by internal actions describe complete synchronizations and, therefore, should be viewed as “real” computation steps.

In order to capture priorities, the syntax of our language differs from CCS in that the *port* set exhibits a priority scheme, i.e., priorities are attached to ports. Our notion of pre-emption then stipulates that a system cannot engage in transitions labeled by actions with a given priority, whenever it is able to perform a transition labeled by an *internal* action of higher priority. In this case we say that the lower-priority transition is pre-empted by the higher-priority internal transition. In accordance with the above discussion visible actions *never* have pre-emptive power over actions of lower priority, because visible actions only indicate the potential for execution. An algebraic justification of this design decision can be found in Section 3.5.

Technically, priority values are taken from some finite domain together with a strict order. For the sake of simplicity we use finite initial intervals \mathcal{N} of the natural numbers \mathbb{N} in what follows. We adopt the convention that smaller numbers mean higher priorities; so 0 is the highest priority. Intuitively, visible actions represent potential communications that a process may be willing to engage in with its environment. Formally, let $\{\Lambda_k \mid k \in \mathcal{N}\}$ denote an \mathcal{N} -indexed family of countably infinite, disjoint sets of ports. Intuitively, Λ_k contains the ports with priority k that processes may synchronize over. Then the set of *actions* \mathcal{A}_k with priority k may be defined by $\mathcal{A}_k =_{\text{df}} \Lambda_k \cup \bar{\Lambda}_k \cup \{\tau_k\}$, where $\bar{\Lambda}_k =_{\text{df}} \{\bar{\lambda}_k \mid \lambda_k \in \Lambda_k\}$ and $\tau_k \notin \Lambda_k$. An action $\lambda_k \in \Lambda_k$ may be thought of as representing the receipt of an input on port λ that has priority k , while $\bar{\lambda}_k \in \bar{\Lambda}_k$ constitutes the deposit of an output on λ . The invisible action τ_k represent internal computation steps with priority k . For better readability we write $\lambda:k$, if $\lambda_k \in \Lambda_k$, $\bar{\lambda}:k$, if $\bar{\lambda}_k \in \bar{\Lambda}_k$, and $\tau:k$ for τ_k . The sets of all ports Λ , all complementary ports $\bar{\Lambda}$, and all actions \mathcal{A} are defined by $\cup\{\Lambda_k \mid k \in \mathcal{N}\}$, $\cup\{\bar{\Lambda}_k \mid k \in \mathcal{N}\}$, and $\cup\{\mathcal{A}_k \mid k \in \mathcal{N}\}$, respectively. In what follows, we use $\alpha:k, \beta:k, \dots$ to range over \mathcal{A} and $a:k, b:k, \dots$ to range over $\Lambda \cup \bar{\Lambda}$. We also extend $\bar{\cdot}$ to all visible actions $a:k$ by $\bar{a}:k =_{\text{df}} a:k$. Finally, if $L \subseteq \Lambda \cup \bar{\Lambda}$ then $\bar{L} = \{\bar{a}:k \mid a:k \in L\}$. The *syntax* of our language is defined by the following BNF.

$$\begin{array}{l}
 P \quad ::= \quad \mathbf{0} \quad | \quad x \quad | \quad \alpha:k.P \quad | \quad P + P \quad | \quad P|P \quad | \\
 \quad \quad \quad P[f] \quad | \quad P \setminus L \quad | \quad \mu x.P
 \end{array}$$

Here, f is a *finite relabeling*, i.e., a mapping on \mathcal{A} with $f(\tau:k) = \tau:k$ for all $k \in \mathcal{N}$, $f(\bar{a}:k) = f(a:k)$ for all $a:k \in \Lambda \cup \bar{\Lambda}$, and $|\{\alpha:k \mid f(\alpha:k) \neq \alpha:k\}| < \infty$. Moreover, a relabeling preserves priority values, i.e., for all $a:k \in \Lambda \cup \bar{\Lambda}$ we have $f(a:k) = b:k$ for some $b:k \in \Lambda_k \cup \bar{\Lambda}_k$. Furthermore, the *restriction set*

L is a subset of $\Lambda \cup \overline{\Lambda}$, and x is a *variable* taken from a set \mathcal{V} . Sometimes it is convenient to write $C \stackrel{\text{def}}{=} P$ for $\mu C.P$, where identifier C is interpreted as variable. We adopt the standard definitions for *sort* of a process, *free* and *bound variables*, *open* and *closed terms*, *guarded recursion*, and *contexts* [56]. We refer to closed and guarded terms as *processes* and use P, Q, R, \dots to range over the set \mathcal{P} of processes. Finally, we denote syntactic equality by \equiv .

Although our framework allows for multi-level *priority schemes* we often restrict ourselves to a two-level priority framework, i.e. we choose $\mathcal{N} = \{0, 1\}$. The reason is that even in this simple setting most semantic and technical issues regarding the introduction of priority to process algebra can be illustrated. However, we also discuss how the obtained results can be carried over to multi-level priority schemes. In order to improve readability within the two-level priority scheme we often write $\underline{\alpha}$ for the “prioritized” action $\alpha:0$, α for the “unprioritized” action $\alpha:1$, \underline{A} for \mathcal{A}_0 , and A for \mathcal{A}_1 . Moreover, we let δ and γ represent elements taken from $\underline{A} \cup A$. Finally, we want to emphasize again that α and $\underline{\alpha}$ are considered to be different ports; i.e., the priority value is part of a port. Thus, in a CCS-based framework only complementary actions having the *same* priority value can engage in a communication. We discuss the consequences of lifting this restriction in Section 3.7 for frameworks involving global pre-emption and in Section 4.6 for those involving local pre-emption. It should be noted that the dynamic priority approach presented in Section 5 also differs in its interpretation of ports, actions, and priority values. Finally, our language does not provide any means for changing priority values of actions. However, we will discuss in Section 3.5 the effect of introducing additional operators, called *prioritization* and *deprioritization*, which increase and decrease the priority values of actions, respectively.

3 Static Priority and Global Pre-emption

In this section we introduce a semantics of our language, restricted to a two-level priority scheme, based on static priority and global pre-emption. We refer to this language as CCS^{sg} (CCS with static priority and global pre-emption) and develop its semantic theory along the lines mentioned in Section 1.2. The organization of this section is as follows. Section 3.1 formally introduces the operational semantics for CCS^{sg} . The following two sections show how to adapt the notions of strong bisimulation and observational congruence to CCS^{sg} , respectively. Section 3.4 applies the semantic theory to our introductory back-and-forth example. The consequences of adding prioritization and deprioritization operators to CCS^{sg} are discussed in Section 3.5. Finally, Section 3.6 comments on the extension of CCS^{sg} to multi-level priority schemes, whereas Section 3.7 presents our concluding remarks and related work.

3.1 Operational Semantics

The *semantics* of a process $P \in \mathcal{P}$ is given by a labeled transition system $\langle \mathcal{P}, \mathcal{A}, \longrightarrow, P \rangle$, where \mathcal{P} is the set of *states*, \mathcal{A} is the *alphabet*, $\longrightarrow \subseteq \mathcal{P} \times \mathcal{A} \times \mathcal{P}$ is the *transition relation*, and P is the *start state*. We write $P \xrightarrow{\gamma} P'$ instead of $\langle P, \gamma, P' \rangle \in \longrightarrow$ and say that P may engage in action γ and thereafter behave like process P' . Moreover, we let $P \xrightarrow{\gamma}$ stand for $\exists P' \in \mathcal{P}. P \xrightarrow{\gamma} P'$. The presentation of the operational rules requires *prioritized initial action sets* $\underline{\mathcal{I}}(P)$ which are defined as the least sets satisfying the equations in Table 1. Intuitively, $\underline{\mathcal{I}}(P)$ denotes the set of all prioritized actions in which P can initially engage. For convenience, we also write $\underline{\underline{\mathcal{I}}}(P)$ for $\underline{\mathcal{I}}(P) \setminus \{\underline{\tau}\}$.

Table 1
Prioritized initial action sets for CCS^{sg}

$\underline{\mathcal{I}}(\alpha.P)$	$= \{\alpha\}$	$\underline{\mathcal{I}}(\mu x.P)$	$= \underline{\mathcal{I}}(P[\mu x.P/x])$
$\underline{\mathcal{I}}(P + Q)$	$= \underline{\mathcal{I}}(P) \cup \underline{\mathcal{I}}(Q)$	$\underline{\mathcal{I}}(P Q)$	$= \underline{\mathcal{I}}(P) \cup \underline{\mathcal{I}}(Q) \cup \{\underline{\tau} \mid \underline{\mathcal{I}}(P) \cap \overline{\underline{\mathcal{I}}}(Q) \neq \emptyset\}$
$\underline{\mathcal{I}}(P[f])$	$= \{f(\alpha) \mid \alpha \in \underline{\mathcal{I}}(P)\}$	$\underline{\mathcal{I}}(P \setminus L)$	$= \underline{\mathcal{I}}(P) \setminus (L \cup \overline{L})$

The rules in Plotkin-style notation [67] in Table 2 formally define the transition relation and capture the following operational behavior. Process $\gamma.P$ may engage in action γ and then behave like P . The *summation operator* $+$ denotes *nondeterministic choice*. Process $P + Q$ may behave like process P (Q) if Q (P) does not pre-empt an unprioritized transition by performing a prioritized internal transition. $P[f]$ behaves exactly as process P with the actions renamed with respect to *relabeling* f . Process $P | Q$ stands for the *parallel composition* of P and Q , according to an *interleaving semantics* with *synchronized communication* on complementary actions on the same priority level, resulting in the internal action $\underline{\tau}$ or τ . However, if Q (P) is capable of engaging in a prioritized internal transition, then unprioritized transitions of P (Q) are pre-empted. The *restriction operator* $\setminus L$ prohibits in $P \setminus L$ the execution of transitions labeled by actions in $L \cup \overline{L}$ and, thus, permits the *scoping* of actions. Finally, $\mu x.P$ denotes a *recursively defined* process that is a distinguished solution to the equation $x = P$.

3.2 Semantic Theory Based on Strong Bisimulation

The semantic theory for CCS^{sg} is based on the notion of *bisimulation* [56,65]. First, *strong bisimulation* [56] is adapted from CCS to our setting as follows; we refer to this relation as *prioritized strong bisimulation*.

Table 2
Operational semantics for CCS^{sg}

<u>Act</u> $\frac{\overline{\quad}}{\alpha.P \xrightarrow{\alpha} P}$	Act $\frac{\overline{\quad}}{\alpha.P \xrightarrow{\alpha} P}$
<u>Sum1</u> $\frac{P \xrightarrow{\alpha} P'}{P + Q \xrightarrow{\alpha} P'}$	Sum1 $\frac{P \xrightarrow{\alpha} P'}{P + Q \xrightarrow{\alpha} P'} \quad \perp \notin \underline{\mathcal{I}}(Q)$
<u>Sum2</u> $\frac{Q \xrightarrow{\alpha} Q'}{P + Q \xrightarrow{\alpha} Q'}$	Sum2 $\frac{Q \xrightarrow{\alpha} Q'}{P + Q \xrightarrow{\alpha} Q'} \quad \perp \notin \underline{\mathcal{I}}(P)$
<u>Com1</u> $\frac{P \xrightarrow{\alpha} P'}{P Q \xrightarrow{\alpha} P' Q}$	Com1 $\frac{P \xrightarrow{\alpha} P'}{P Q \xrightarrow{\alpha} P' Q} \quad \perp \notin \underline{\mathcal{I}}(P Q)$
<u>Com2</u> $\frac{Q \xrightarrow{\alpha} Q'}{P Q \xrightarrow{\alpha} P Q'}$	Com2 $\frac{Q \xrightarrow{\alpha} Q'}{P Q \xrightarrow{\alpha} P Q'} \quad \perp \notin \underline{\mathcal{I}}(P Q)$
<u>Com3</u> $\frac{P \xrightarrow{a} P' \quad Q \xrightarrow{\bar{a}} Q'}{P Q \xrightarrow{\tau} P' Q'}$	Com3 $\frac{P \xrightarrow{a} P' \quad Q \xrightarrow{\bar{a}} Q'}{P Q \xrightarrow{\tau} P' Q'} \quad \perp \notin \underline{\mathcal{I}}(P Q)$
<u>Rel</u> $\frac{P \xrightarrow{\alpha} P'}{P[f] \xrightarrow{f(\alpha)} P'[f]}$	Rel $\frac{P \xrightarrow{\alpha} P'}{P[f] \xrightarrow{f(\alpha)} P'[f]}$
<u>Res</u> $\frac{P \xrightarrow{\alpha} P'}{P \setminus L \xrightarrow{\alpha} P' \setminus L} \quad \alpha \notin L \cup \overline{L}$	Res $\frac{P \xrightarrow{\alpha} P'}{P \setminus L \xrightarrow{\alpha} P' \setminus L} \quad \alpha \notin L \cup \overline{L}$
<u>Rec</u> $\frac{P[\mu x.P/x] \xrightarrow{\alpha} P'}{\mu x.P \xrightarrow{\alpha} P'}$	Rec $\frac{P[\mu x.P/x] \xrightarrow{\alpha} P'}{\mu x.P \xrightarrow{\alpha} P'}$

Definition 3.1 (Prioritized Strong Bisimulation)

A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is called a prioritized strong bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$ and $\gamma \in \mathcal{A}$ the following condition holds.

$$P \xrightarrow{\gamma} P' \text{ implies } \exists Q'. Q \xrightarrow{\gamma} Q' \text{ and } \langle P', Q' \rangle \in \mathcal{R} .$$

We write $P \simeq Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some prioritized strong bisimulation \mathcal{R} .

Table 3
Axiomatization of \simeq

<p>(A1) $t + u = u + t$</p> <p>(A3) $t + t = t$</p> <p>(E) Let $t = \sum_i \gamma_i.t_i$ and $u = \sum_j \delta_j.u_j$. Then $t u = \sum_i \gamma_i.(t_i u) + \sum_j \delta_j.(t u_j) +$ $\sum_{\gamma_i \equiv \delta_j} \{\underline{\tau}.(t_i u_j) \mid \gamma_i \in \underline{A}\} + \sum_{\gamma_i \equiv \delta_j} \{\tau.(t_i u_j) \mid \gamma_i \in A\}$</p>	<p>(A2) $t + (u + v) = (t + u) + v$</p> <p>(A4) $t + \mathbf{0} = t$</p> <p>(Res1) $\mathbf{0} \setminus L = \mathbf{0}$</p> <p>(Res2) $(\gamma.t) \setminus L = \mathbf{0}$</p> <p>(Res3) $(\gamma.t) \setminus L = \gamma.(t \setminus L)$</p> <p>(Res4) $(t + u) \setminus L = (t \setminus L) + (u \setminus L)$</p>
	<p>(Rel1) $\mathbf{0}[f] = \mathbf{0}$</p> <p>(Rel2) $(\gamma.t)[f] = f(\gamma).(t[f])$</p> <p>(Rel3) $(t + u)[f] = t[f] + u[f]$</p> <p>(P) $\underline{\tau}.t + \alpha.u = \underline{\tau}.t$</p>

It is easy to see that \simeq is an equivalence and that it is the *largest* prioritized strong bisimulation. The following result, which enables compositional reasoning, can be proved straightforwardly using standard techniques [1,27,76].

Theorem 3.2 *\simeq is a congruence.*

An *axiomatization* of \simeq for *finite* processes, i.e., guarded and closed CCS^{sg} terms not containing recursion, can be developed closely along the lines of [27]. We write $\vdash t = u$ if process term t can be rewritten to u using the axioms in Table 3. These axioms correspond to the ones presented in [56], except that Axiom (P) dealing with global pre-emption is added. In *Expansion Axiom* (E) the symbol \sum stands for the indexed version of $+$, where the empty sum denotes the inaction process $\mathbf{0}$. The next theorem states that our equations characterize prioritized strong bisimulation for finite CCS^{sg} processes. Its proof can be found in [27]; it uses techniques described in [56].

Theorem 3.3 *Let t, u be finite processes. Then $t \simeq u$ if and only if $\vdash t = u$.*

3.3 Semantic Theory Based on Weak Bisimulation

The behavioral congruence developed in the previous section is too strong for verifying systems in practice, as it requires that two equivalent terms match each other's transitions exactly, even those labeled by internal actions. In process algebra one remedies this problem by developing a semantic congruence

that abstracts from internal transitions. We start off with the definition of a *naive prioritized weak bisimulation* which is an adaptation of Milner's *observational equivalence* [56].

Definition 3.4 (Naive Prioritized Weak Transitions)

- (1) $\hat{\tau} =_{df} \hat{\tau} =_{df} \epsilon$, $\hat{a} =_{df} \underline{a}$, and $\hat{a} =_{df} a$
- (2) $\xrightarrow{\epsilon}_\times =_{df} (\xrightarrow{\tau} \cup \xrightarrow{\tau})^*$
- (3) $\xrightarrow{\gamma}_\times =_{df} \xrightarrow{\epsilon}_\times \circ \xrightarrow{\gamma} \circ \xrightarrow{\epsilon}_\times$

Observe that this transition relation ignores priority levels for $\xrightarrow{\epsilon}_\times$. This is in accordance with the fact that a priority value is part of an action and, thus, unobservable for internal actions.

Definition 3.5 (Naive Prioritized Weak Bisimulation)

A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is a naive prioritized weak bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$, and $\gamma \in \mathcal{A}$ the following condition holds.

$$P \xrightarrow{\gamma} P' \text{ implies } \exists Q'. Q \xrightarrow{\hat{\gamma}}_\times Q' \text{ and } \langle P', Q' \rangle \in \mathcal{R} .$$

We write $P \approx_\times Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for a naive prioritized weak bisimulation \mathcal{R} .

Naive prioritized weak bisimulation can be shown to be an equivalence. Unfortunately, \approx_\times is not a congruence for CCS^{sg} with respect to parallel composition, summation, and recursion. Whereas the compositionality defect for summation and recursion is similar to the one for CCS [56], the defect with respect to parallel composition is due to pre-emption. As an example consider the processes $P \stackrel{\text{def}}{=} a.\mathbf{0} + \underline{b}.\mathbf{0}$ and $Q \stackrel{\text{def}}{=} a.\mathbf{0} + \tau.(a.\mathbf{0} + \underline{b}.\mathbf{0})$. It is easy to see that $P \approx_\times Q$. However, when composing these processes in parallel with process $\bar{b}.\mathbf{0}$ then $Q \mid \bar{b}.\mathbf{0} \xrightarrow{a} \mathbf{0} \mid \bar{b}.\mathbf{0}$ whereas $P \mid \bar{b}.\mathbf{0} \not\xrightarrow{a}_\times$, i.e., $P \mid \bar{b}.\mathbf{0} \not\approx_\times Q \mid \bar{b}.\mathbf{0}$. Therefore, one has to be more careful when defining the prioritized weak transition relation. Transitions labeled by visible actions may turn into internal transitions when composed with an environment and, thus, may gain pre-emptive power. An adequate notion of weak transitions must take the processes' potential of engaging in visible prioritized transitions into account.

3.3.1 *Prioritized Weak Bisimulation*

Despite the lack of compositionality, the definition of \approx_\times reflects an intuitive approach to abstracting from internal computation. For handling the congruence problem one may consider the following fact from universal algebra.

Proposition 3.6 *For an equivalence \mathcal{R} in an algebra \mathfrak{R} the largest congruence $\mathcal{R}^+ \subseteq \mathcal{R}$ exists and $\mathcal{R}^+ = \{ \langle P, Q \rangle \mid \forall \mathfrak{R}\text{-contexts } C[X]. \langle C[P], C[Q] \rangle \in \mathcal{R} \}$. Here, a \mathfrak{R} -context $C[X]$ is a term in \mathfrak{R} with one free occurrence of variable X .*

Consequently, we know that \approx_\times contains a largest congruence \approx_\times^+ for CCS^{sg} and devote the rest of this section to characterizing this congruence. We first define a new weak transition relation which takes pre-emption into account.

Definition 3.7 (Prioritized Weak Transitions) *Let $L \subseteq \underline{A} \setminus \{\perp\}$.*

- (i) $\hat{\tau} =_{df} \epsilon$, $\hat{\tau} =_{df} \epsilon$, $\hat{a} =_{df} \underline{a}$, and $\hat{a} =_{df} a$
- (ii) $P \xrightarrow[L]{\alpha} P'$ if $P \xrightarrow{\alpha} P'$ and $\underline{\mathcal{I}}(P) \subseteq L$
- (iii) $\xRightarrow{\epsilon} =_{df} (\xrightarrow{\tau})^*$
- (iv) $\xrightarrow[L]{\epsilon} =_{df} (\xrightarrow{\tau} \cup \xrightarrow[L]{\tau})^*$
- (v) $\xRightarrow{\alpha} =_{df} \xRightarrow{\epsilon} \circ \xrightarrow{\alpha} \circ \xRightarrow{\epsilon}$
- (vi) $\xrightarrow[L]{\alpha} =_{df} \xrightarrow[L]{\epsilon} \circ \xrightarrow[L]{\alpha} \circ \xRightarrow{\epsilon}$

Intuitively, we have made the transition relation sensitive to pre-emption by introducing conditions involving prioritized initial action sets and by preserving priority levels of internal actions. In the remainder, we show that prioritized initial action sets are an adequate means for measuring pre-emption potentials. In this light, $P \xrightarrow[L]{\alpha} P'$ states that P can evolve to P' by performing the unprioritized action α , if the environment does not offer any prioritized communication on some port in L .

Definition 3.8 (Prioritized Weak Bisimulation)

A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is a prioritized weak bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$, $\underline{a} \in \underline{A}$, and $\alpha \in A$ the following conditions hold.

- (1) $\perp \notin \underline{\mathcal{I}}(P)$ implies $\exists Q'. Q \xrightarrow[L]{\epsilon} Q'$, $\underline{\mathcal{I}}(Q') \subseteq L$ where $L = \underline{\mathcal{I}}(P)$, $\tau \notin \underline{\mathcal{I}}(Q')$, and $\langle P, Q' \rangle \in \mathcal{R}$.
- (2) $P \xrightarrow{\alpha} P'$ implies $\exists Q'. Q \xrightarrow[L]{\hat{\alpha}} Q'$, and $\langle P, Q' \rangle \in \mathcal{R}$.
- (3) $P \xrightarrow{\alpha} P'$ implies $\exists Q'. Q \xrightarrow[L]{\hat{\alpha}} Q'$, where $L = \underline{\mathcal{I}}(P)$, and $\langle P', Q' \rangle \in \mathcal{R}$.

We write $P \cong Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some prioritized weak bisimulation \mathcal{R} .

This new version of weak bisimulation is algebraically more robust than the naive one. Condition (1) of Definition 3.8 is necessary for achieving compositionality with respect to parallel composition.

Proposition 3.9 *The equivalence \cong is a congruence with respect to prefixing, parallel composition, relabeling, and restriction. Moreover, \cong is characterized as the largest congruence contained in \approx_\times , in the sub-algebra of CCS^{sg} induced by these operators and recursion.*

Although \cong is itself not a congruence, this relation provides the basis for obtaining a congruence, as is made precise in the next section.

3.3.2 Prioritized Observational Congruence

The compositionality defect of \cong with respect to summation is handled by the following notion of *prioritized observational congruence*. Unfortunately, the summation fix presented in [56], which requires an initial internal transition to be matched by a nontrivial internal weak transition, is not sufficient in order to achieve a congruence based on prioritized weak bisimulation. To see why, let $D \stackrel{\text{def}}{=} \tau.E$ and $E \stackrel{\text{def}}{=} \tau.D$. Now define $P \stackrel{\text{def}}{=} \tau.D$ and $Q \stackrel{\text{def}}{=} \tau.E$. By Definition 3.8 we may observe $P \cong Q$, but $P + a.\mathbf{0} \not\cong Q + a.\mathbf{0}$ since the former can perform an a -transition whereas the latter cannot. It turns out that we have to require that observationally congruent processes must possess the same initial action sets; a requirement which is stronger than Condition (1) of Definition 3.8.

Definition 3.10 *Define $P \cong^1 Q$ if for all $\underline{a} \in \underline{A}$ and $a \in A$ the following conditions and their symmetric counterparts hold.*

- (1) $\mathcal{I}(P) \supseteq \mathcal{I}(Q)$
- (2) $P \xrightarrow{\underline{a}} P'$ implies $\exists Q'. Q \xrightarrow{\underline{a}} Q'$ and $P' \cong Q'$.
- (3) $P \xrightarrow{a} P'$ implies $\exists Q'. Q \xrightarrow[L]{a} Q'$, where $L = \underline{\mathcal{I}}(P)$, and $P' \cong Q'$.

The following theorem states the desired algebraic result for \cong^1 .

Theorem 3.11 \cong^1 is the largest congruence contained in \approx_\times , i.e., $\cong^1 = \approx_\times^+$.

Whereas the proof of the congruence property of \cong^1 is standard (cf. [56]), the “largest” part is proved by using the following fact from universal algebra.

Proposition 3.12 *Let \mathcal{R}_1 and \mathcal{R}_2 be equivalences in an algebra \mathfrak{R} such that $\mathcal{R}_1^+ \subseteq \mathcal{R}_2 \subseteq \mathcal{R}_1$. Then $\mathcal{R}_1^+ = \mathcal{R}_2^+$.*

For the purposes of this section, we choose $\mathcal{R}_1 = \approx_\times$ and $\mathcal{R}_2 = \cong$. The following theorem, which establishes $\mathcal{R}_2^+ = \cong^1$, can be proved along the lines of a corresponding theorem in [56]; for details see [53].

Theorem 3.13 \cong^1 is the largest congruence contained in \cong .

In order to apply Proposition 3.12, $\approx_\times^+ \subseteq \cong \subseteq \approx_\times$ needs to be established. The inclusion $\cong \subseteq \approx_\times$ immediately follows from the definition of the naive prioritized weak and the prioritized weak transition relations. Thus, one is left with $\approx_\times^+ \subseteq \cong$. This inclusion turns out to be difficult to show directly. Therefore, the auxiliary relation $\cong_a =_{\text{df}} \{\langle P, Q \rangle \mid C_{PQ}[P] \approx_\times C_{PQ}[Q]\}$ is defined, which lies in between. Using the abbreviation $\underline{\mathcal{S}}$ for the union of the

finite prioritized sorts of P and Q , we define $C_{PQ}[X] \stackrel{\text{def}}{=} X \mid H_{PQ}$ and

$$H_{PQ} \stackrel{\text{def}}{=} \underline{c}.0 + \sum_{\substack{L \subseteq \overline{\underline{S}}, \\ \underline{b} \in \underline{S}}} \underline{\tau}. \left(\begin{array}{c} \underline{d}_{L,\underline{b}}.H_{PQ} + \\ D_L + \underline{e}.H_{PQ} + \\ \underline{b}.H_{PQ} \end{array} \right).$$

Moreover, D_L is defined as $\sum_{\underline{a} \in L} \underline{a}.0$. The actions $\underline{c}, \underline{d}_{L,\underline{b}}, \underline{e}$ for all $L \subseteq \overline{\underline{S}}$ and $\underline{b} \in \underline{S}$, and their complements, are supposed to be “fresh” actions, i.e., not in $\underline{S} \cup \overline{\underline{S}}$. By Proposition 3.6 we may conclude $\approx_x^+ \subseteq \approx_a$. The other necessary inclusion $\approx_a \subseteq \approx$ is established by showing that \approx_a is a prioritized weak bisimulation. Summarizing, Theorem 3.11 is a consequence of Proposition 3.12 when assembling our results as illustrated in Figure 2, where an arrow from relation R_1 to relation R_2 stands for $R_1 \subseteq R_2$.




Fig. 2. Proof sketch of Theorem 3.11

3.3.3 Operational Characterization

The aim of this section is to show how prioritized weak bisimulation can be computed by adapting standard partition-refinement algorithms [26,48,64] developed for strong bisimulation [56]. To this end, we provide an operational characterization of prioritized weak bisimulation as strong bisimulation by introducing an *alternative prioritized weak transition relation*.

Definition 3.14 Let $P, P' \in \mathcal{P}$, $\alpha \in A$, and $\underline{\alpha} \in \underline{A}$.

- (1) $\xrightarrow{\hat{\alpha}}_* =_{df} \xrightarrow{\hat{\alpha}}$ and
- (2) $P \xrightarrow{\hat{\alpha}}_* P'$ if $\exists P'' . \underline{\tau} \notin \underline{\mathcal{I}}(P'')$ and $P \xrightarrow{\underline{\epsilon}} P'' \xrightarrow[\underline{L}]{\hat{\alpha}} P'$ for $L = \underline{\mathcal{I}}(P'')$.

Observe that the alternative prioritized weak transition relation is not parameterized by prioritized initial action sets. It can be computed efficiently by using *dynamic programming* techniques.

Definition 3.15 A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is called an alternative prioritized weak bisimulation if for all $\langle P, Q \rangle \in \mathcal{R}$ and $\gamma \in \mathcal{A}$ the following condition holds.

$$P \xRightarrow{\hat{\gamma}}_* P' \text{ implies } \exists Q'. Q \xRightarrow{\hat{\gamma}}_* Q' \text{ and } \langle P', Q' \rangle \in \mathcal{R} .$$

We write $P \approx_* Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some alternative prioritized weak bisimulation \mathcal{R} .

Theorem 3.16 (Operational Characterization) $\approx = \approx_*$.

The proof of this characterization result is straightforward [53]. It should be mentioned that the presented characterization may also serve as a basis for defining a Hennessy-Milner logic for prioritized weak bisimulation along the lines of [56].

3.4 Example

As a simple example, we return to the back-and-forth system introduced in Section 1, which can be formalized in CCS^{sg} as follows: $\text{Sys} \stackrel{\text{def}}{=} (A \mid B) \setminus \{\underline{i}\}$ where $A \stackrel{\text{def}}{=} \text{back}.A' + \underline{i}.\tau.\overline{\text{ok}}.\underline{i}.A$, $A' \stackrel{\text{def}}{=} \text{forth}.A + \underline{i}.\tau.\overline{\text{ok}}.\underline{i}.A'$, and $B \stackrel{\text{def}}{=} \text{check}.\underline{i}.\underline{i}.B$. Intuitively, \underline{i} is an internal *interrupt*, and thus prioritized and restricted (via $\setminus \{\underline{i}\}$), which is invoked whenever **check** is executed. Hence, in such a state process A cannot engage in a transition labeled by **back** or **forth** according to our pre-emptive operational semantics, but must accept the communication on prioritized port \underline{i} . One can think of the τ -action in the definition of process A as representing some internal activities determining the current status of the system. The CCS^{sg} semantics of **Sys** is shown in Figure 3.




Fig. 3. Semantics of **Sys**

In the sequel, we prove that **Sys** meets its specification **Spec**, which is given by

$$\begin{aligned} \text{Spec} &\stackrel{\text{def}}{=} \text{back}.\text{Spec}' + \text{check}.\overline{\text{ok}}.\text{Spec} \\ \text{Spec}' &\stackrel{\text{def}}{=} \text{forth}.\text{Spec} + \text{check}.\overline{\text{ok}}.\text{Spec}' . \end{aligned}$$

First, the validity of $\text{Sys} \cong \text{Spec}$ is proved by the relation presented in Table 4, whose symmetric closure is a prioritized weak bisimulation that contains $\langle \text{Sys}, \text{Spec} \rangle$.

Table 4

A relation whose symmetric closure is a prioritized weak bisimulation

$\langle \text{Sys} \quad \quad \quad , \text{Spec} \quad \rangle$,	$\langle (A' B) \setminus \{\underline{i}\} \quad \quad \quad , \text{Spec}' \quad \rangle$,
$\langle (A \bar{\underline{i}}.\underline{i}.B) \setminus \{\underline{i}\} \quad \quad \quad , \overline{\text{ok}}.\text{Spec} \quad \rangle$,	$\langle (\tau.\overline{\text{ok}}.\bar{\underline{i}}.A \underline{i}.B) \setminus \{\underline{i}\} \quad \quad \quad , \overline{\text{ok}}.\text{Spec} \quad \rangle$,
$\langle (\overline{\text{ok}}.\bar{\underline{i}}.A \underline{i}.B) \setminus \{\underline{i}\} \quad \quad \quad , \overline{\text{ok}}.\text{Spec} \quad \rangle$,	$\langle (\bar{\underline{i}}.A \underline{i}.B) \setminus \{\underline{i}\} \quad \quad \quad , \text{Spec} \quad \rangle$,
$\langle (A' \bar{\underline{i}}.\underline{i}.B) \setminus \{\underline{i}\} \quad \quad \quad , \overline{\text{ok}}.\text{Spec}' \quad \rangle$,	$\langle (\tau.\overline{\text{ok}}.\bar{\underline{i}}.A' \underline{i}.B) \setminus \{\underline{i}\} \quad \quad \quad , \overline{\text{ok}}.\text{Spec} \quad \rangle$,
$\langle (\overline{\text{ok}}.\bar{\underline{i}}.A' \underline{i}.B) \setminus \{\underline{i}\} \quad \quad \quad , \overline{\text{ok}}.\text{Spec}' \quad \rangle$,	$\langle (\bar{\underline{i}}.A' \underline{i}.B) \setminus \{\underline{i}\} \quad \quad \quad , \text{Spec}' \quad \rangle$ }

In addition, both processes only possess visible initial actions, and their prioritized initial action sets are identical. Hence, we may conclude $\text{Sys} \cong^1 \text{Spec}$.

3.5 Prioritization and Deprioritization Operators

There are several other language constructs worth considering when dealing with priority. Of particular interest are the unary operators introduced by Cleaveland and Hennessy in [27] which correspond to the *prioritization* of a visible unprioritized action, written $\lceil a$ for $a \neq \tau$, and to the *deprioritization* of a visible prioritized action, written $\lfloor \underline{a}$ for $\underline{a} \neq \underline{\tau}$. The operational semantics of these operators is formally defined in Table 5. Their presence requires that (i) every prioritized port \underline{a} corresponds one-to-one to an unprioritized port a , (ii) every relabeling f satisfies $f(\underline{a}) = \underline{f(a)}$, and (iii) every restriction set L obeys the property “ $a \in L$ if and only if $\underline{a} \in L$.” Intuitively, $P \lceil a$ prioritizes all a -transition which P can perform, while $P \lfloor \underline{a}$ deprioritizes all \underline{a} -transitions in which P can engage, provided the newly deprioritized transition is also available to P . Note that the notion of priority is still static and not dynamic since the prioritization and deprioritization operators are static operators. Thus, the change of priority values affects a process in its whole and is not limited to its initial behavior.

Including prioritization and deprioritization operators with CCS^{sg} does not have any consequences for prioritized strong bisimulation since it is compositional with respect to these operators [27]. The axiomatization of prioritized strong bisimulation for finite processes can also be extended to cover the new operators. The necessary additional axioms are stated in Table 6.

The presence of the prioritization and the deprioritization operator allows us to formally justify the design decision that only prioritized internal actions have pre-emptive power over unprioritized actions. For this purpose assume

Table 5
Semantics for the prioritization and the deprioritization operator

Prio1 $\frac{P \xrightarrow{a} P'}{P[a \xrightarrow{a} P']_a} \tau \notin \underline{\mathcal{I}}(P)$	Deprio1 $\frac{P \xrightarrow{\underline{a}} P'}{P[\underline{a} \xrightarrow{a} P']_{\underline{a}}} \tau \notin \underline{\mathcal{I}}(P)$
Prio2 $\frac{P \xrightarrow{a} P'}{P[a \xrightarrow{a} P']_a} \tau \in \underline{\mathcal{I}}(P)$	Deprio2 $\frac{P \xrightarrow{\underline{a}} P'}{P[\underline{a} \xrightarrow{a} P']_{\underline{a}}} \tau \in \underline{\mathcal{I}}(P)$
Prio3 $\frac{P \xrightarrow{\gamma} P'}{P[a \xrightarrow{\gamma} P']_a} \gamma \neq a$	Deprio3 $\frac{P \xrightarrow{\gamma} P'}{P[\underline{a} \xrightarrow{\gamma} P']_{\underline{a}}} \gamma \neq \underline{a}$

Table 6
Axioms for the prioritization and the deprioritization operator

(Prio1)	$\mathbf{0}[a = \mathbf{0}$	
(Prio2)	$(a.t)[a = \underline{a}.(t[a)$	
(Prio3)	$(\gamma.t)[a = \gamma.(t[a)$	$\gamma \neq a$
(Prio4)	$(t + \tau.u + \underline{b}.v)[a = (t + \tau.u)[a + \underline{b}.(v[a)$	
(Prio5)	$(t + \delta.u + \gamma.v)[a = (t + \delta.u)[a + (t + \gamma.v)[a$	$\delta, \gamma \in \mathcal{A} \setminus \{\tau\}$
<hr/>		
(Deprio1)	$\mathbf{0}[\underline{a} = \mathbf{0}$	
(Deprio2)	$(\underline{a}.t)[\underline{a} = a.(t[\underline{a}$	
(Deprio3)	$(\gamma.t)[\underline{a} = \gamma.(t[\underline{a}$	$\gamma \neq \underline{a}$
(Deprio4)	$(t + \tau.u + \underline{b}.v)[\underline{a} = (t + \tau.u)[\underline{a} + \underline{b}.(v[\underline{a}$	
(Deprio5)	$(t + \delta.u + \gamma.v)[\underline{a} = (t + \delta.u)[\underline{a} + (t + \gamma.v)[\underline{a}$	$\delta, \gamma \in \mathcal{A} \setminus \{\tau\}$

that (i) pre-emption is not encoded in the side conditions of the operational rules but, equivalently, in the notion of bisimulation [27] and that (ii) the naive view of pre-emption gives all prioritized actions pre-emptive power. Thus, a naive bisimulation \sim_n demands the following condition for equivalent processes $P \sim_n Q$ and unprioritized actions $\alpha \in A$:

$$(P \xrightarrow{\alpha} P' \wedge \not\exists \underline{\beta}. P \xrightarrow{\underline{\beta}}) \text{ implies } (\exists Q'. Q \xrightarrow{\alpha} Q', \not\exists \underline{\beta}. Q \xrightarrow{\underline{\beta}} \wedge P' \sim_n Q')$$

and vice versa. The condition for prioritized actions can be adopted from standard strong bisimulation. Regarding the properties of \sim_n , it turns out that it

is not a congruence; e.g., $a.\mathbf{0} + \underline{b}.\mathbf{0} \sim_n \underline{b}.\mathbf{0}$ but $(a.\mathbf{0} + \underline{b}.\mathbf{0}) \setminus \{b\} \not\sim_n (\underline{b}.\mathbf{0}) \setminus \{b\}$ since the former process can engage in an a -transition while the latter is deadlocked. Thus, the question arises what the largest congruence contained in \sim_n is; it turns out that it is prioritized strong congruence as defined before (see [27] for a formal treatment). This shows that in a pre-emptive semantics only prioritized internal actions may pre-empt unprioritized actions. However, the above algebraic result is only correct if we include the deprioritization operator in our language. A non-trivial characterization of the largest congruence contained in \sim_n with respect to our original language is still an open problem.

For the language extended by the prioritization and the deprioritization operator, an observational congruence together with an axiomatic characterization with respect to finite processes has been developed in [62,63], which is briefly reviewed here. For this purpose, we need to refine the prioritized weak transition relation. First, we re-define \xrightarrow{a} to $\xrightarrow{\epsilon} \circ \xrightarrow{a} \circ \xrightarrow{\epsilon}$, i.e., a weak unprioritized a -transition consists of an a -transition that is preceded and trailed by *prioritized* internal transitions only. Moreover, we replace $\underline{\mathcal{I}}(P)$ by $\underline{\mathcal{I}}(P) \cup \mathcal{I}(P)$ in the definition of $\xrightarrow{\epsilon}$ since one has to take into account that unprioritized actions may turn into prioritized ones when applying the prioritization operator. Finally, we write $P \xrightarrow[L]{\tau} P'$ whenever $P \xrightarrow[L]{\epsilon} P'$ and $P \not\equiv P'$. Consequently, visible weak unprioritized transitions only abstract from prioritized internal actions. The reason for this restriction is that, otherwise, prioritized weak bisimulation would not be compositional with respect to the prioritization and the deprioritization operator. In contrast, the original prioritized weak transition relation allows an α -transition to be preceded by any sequence of $\underline{\tau}$ - and τ -transitions (satisfying a condition on prioritized initial action sets) and only to be trailed by $\underline{\tau}$ -transitions.

The notions of prioritized weak bisimulation and prioritized observational congruence are defined in [62,63] as follows, where $P \Downarrow$ stands for $\exists P'. P \xrightarrow{\epsilon} P'$ and $P' \not\xrightarrow{\tau}$.

Definition 3.17 *A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is an extended prioritized weak bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$, $\gamma \in \mathcal{A} \setminus \{\tau\}$ the following conditions hold.*

- (1) $P \Downarrow$ implies $Q \Downarrow$.
- (2) $P \xrightarrow{\gamma} P'$ implies $\exists Q'. Q \xrightarrow{\hat{\gamma}} Q'$, and $\langle P, Q' \rangle \in \mathcal{R}$.
- (3) $P \xrightarrow{\tau} P'$ implies $\exists Q'. Q \xrightarrow[L]{\epsilon} Q'$, $L = \underline{\mathcal{I}}(P) \cup \mathcal{I}(P)$, and $\langle P', Q' \rangle \in \mathcal{R}$.

We write $P \approx_{\text{pd}} Q$ if there exists an extended prioritized weak bisimulation \mathcal{R} such that $\langle P, Q \rangle \in \mathcal{R}$.

Definition 3.18 *We define $P \approx_{\text{pd}}^1 Q$ if for all $\gamma \in \mathcal{A} \setminus \{\tau\}$ the following conditions and their symmetric counterparts hold.*

Table 7

Axioms for the τ -laws

($\underline{\tau}1$)	$\gamma.(1.t + t) = \gamma.t$	$1 \in \{\tau, \underline{\tau}\}$
($\underline{\tau}2$)	$\underline{\tau}.t = \underline{\tau}.t + t$	
($\underline{\tau}3$)	$\gamma.(t + \underline{\tau}.u) = \gamma.(t + \underline{\tau}.u) + \gamma.u$	
($\tau1$)	$t + \tau.(u + \tau.v) = t + \tau.(u + \tau.v) + \tau.v$	$\vdash_I t \sqsubseteq_i v$

Table 8

Axiomatization of \sqsubseteq_i (Axioms I)

(iC1) $\underline{\alpha}.t \sqsubseteq_i \underline{\alpha}.u$	(iC2) $\mathbf{0} \sqsubseteq_i \gamma.t \quad \gamma \in \mathcal{A} \setminus \{\underline{\tau}\}$	(iC3) $\alpha.t \sqsubseteq_i \mathbf{0}$
---	---	---

- (1) $P \xrightarrow{\gamma} P'$ implies $\exists Q' Q \xrightarrow{\gamma} Q'$ and $P' \approx_{\text{pd}} Q'$.
(2) $P \xrightarrow{\tau} P'$ implies $\exists Q'. Q \xrightarrow[L]{\tau} Q'$, where $L = \underline{\mathcal{I}}(P) \cup \underline{\mathcal{I}}(P)$, and $P' \approx_{\text{pd}} Q'$.

The observational congruence \approx_{pd}^1 possesses nice algebraic properties for our language extended by the prioritization and the deprioritization operator, including a largest congruence result similar to Theorem 3.11 and a sound and complete axiomatization for finite processes. For the latter, the axiomatization for prioritized strong bisimulation is augmented with suitable τ -laws as shown in Table 7 (cf. [56]). The relation \sqsubseteq_i , occurring in the side condition of Axiom ($\tau1$), is the pre-congruence on finite processes generated by the axioms presented in Table 8 using the laws of inequational reasoning. We write $\vdash_I t \sqsubseteq_i u$ if t can be related to u by Axioms (iC1), (iC2), and (iC3). Intuitively, $\vdash_I t \sqsubseteq_i u$ holds, whenever (i) $\tau \in \underline{\mathcal{I}}(t)$ if and only if $\tau \in \underline{\mathcal{I}}(u)$ and (ii) $\underline{\mathcal{I}}(t) \subseteq \underline{\mathcal{I}}(u)$.

Finally, it should be noted that applications underline the importance of the additional freedom of abstracting from internal transitions, which is gained by leaving out the prioritization and the deprioritization operator. In fact, the observational congruence \approx_{pd}^1 does not relate the processes **Sys** and **Spec** of our back-and-forth example. This is due to the presence of the unprioritized internal action in **Sys**.

3.6 Extension to Multi-Level Priority Schemes

We now remark on the extension of CCS^{sg} to a multi-level priority scheme. To do so, we first alter the definition of prioritized initial action sets to capture the priority-level of actions. We define sets $I^k(P)$ for process P with respect to priority value k as shown in Table 9.

Table 9
Potential initial action sets for CCS^{sg}

$$\begin{array}{l}
\overline{I^k(\alpha:l.P)} = \{\alpha:l \mid l = k\} \quad \overline{I^k(P[f])} = \{f(\alpha:l) \mid \alpha:l \in I^k(P)\} \\
\overline{I^k(\mu x.P)} = \overline{I^k(P[\mu x.P/x])} \quad \overline{I^k(P+Q)} = \overline{I^k(P)} \cup \overline{I^k(Q)} \\
\overline{I^k(P \setminus L)} = \overline{I^k(P) \setminus (L \cup \overline{L})} \quad \overline{I^k(P|Q)} = \overline{I^k(P)} \cup \overline{I^k(Q)} \cup \{\tau:k \mid \overline{I^k(P)} \cap \overline{I^k(Q)} \neq \emptyset\}
\end{array}$$

Using this definition of initial action sets and the convention “ $\tau \notin I^{<k}(P)$ if $\exists l < k. \tau:l \in I^l(P)$,” the operational semantics depicted in Table 2 can be re-stated as follows, as exemplarily shown for Rule (Com3).

$$\text{Com3} \quad \frac{P \xrightarrow{a:k} P' \quad Q \xrightarrow{\overline{a}:k} Q'}{P|Q \xrightarrow{\tau:k} P'|Q'} \quad \tau \notin I^{<k}(P|Q)$$

Observe that the sets $I^k(P)$ may contain actions in which P cannot initially engage, since their definition does not consider pre-emption. In fact, the set $\mathcal{I}^k(P)$ of actions with priority value k , in which P can indeed initially engage is given by $\{\alpha:k \in I^k(P) \mid \tau:l \notin I^l(P) \text{ for all } l < k\}$. However, it is easy to show that $\tau \notin I^{<k}(P)$ if and only if $\tau \notin \mathcal{I}^{<k}(P)$ [53]. Thus, the side condition of Rule (Com3) captures our intuition that $P|Q$ cannot engage in a more urgent internal transition.

Table 10
Prioritized weak transition relation

$$\begin{array}{l}
\overline{\xrightarrow{\epsilon:0}} =_{\text{df}} (\overline{\xrightarrow{\tau:0}})^* \quad \overline{P \xrightarrow[\underline{L}]{\alpha:k} P'} \text{ if } P \xrightarrow{\alpha:k} P' \text{ and } \overline{\mathcal{I}^l(P)} \subseteq \underline{L} \text{ for all } l < k \\
\overline{\xrightarrow[\underline{L}]{\epsilon:k}} =_{\text{df}} (\{\overline{\xrightarrow[\underline{L}]{\tau:l}} \mid l \leq k\})^* \quad \overline{\xrightarrow[\underline{L}]{\alpha:k}} =_{\text{df}} \overline{\xrightarrow[\underline{L}]{\epsilon:k}} \circ \overline{\xrightarrow[\underline{L}]{\alpha:k}} \circ \overline{\xrightarrow{\epsilon:0}}
\end{array}$$

The re-development of the bisimulation-based semantic theory proceeds along the lines of the above sections and does not raise any new semantic issues. For example, the notion of prioritized observational congruence is defined as follows, where (i) the prioritized weak transition relation is given by the rules in Table 10, (ii) $\overline{\mathcal{I}^k(P)} =_{\text{df}} \overline{\mathcal{I}^k(P)} \setminus \{\tau:k\}$, (iii) the relation \cong_{ml} is the adaption of prioritized weak bisimulation to the multi-level priority scheme, (iv) $\overline{\mathcal{I}(P)} =_{\text{df}} \bigcup \{\overline{\mathcal{I}^k(P)} \mid k \in \mathcal{N}\}$, and (v) $\overline{\mathcal{I}^{<k}(P)} =_{\text{df}} \overline{\mathcal{I}^{<k}(P)} \setminus \{\tau:l \mid l < k\}$.

Definition 3.19 *Processes P and Q are prioritized observational congruent if for all actions $\alpha:k$ the following conditions and their symmetric counterparts hold.*

- (1) $\overline{\mathcal{I}(P)} \supseteq \overline{\mathcal{I}(Q)}$
- (2) $P \xrightarrow{\alpha:k} P'$ implies $\exists Q'. Q \xrightarrow[\underline{L}]{\alpha:k} Q'$, where $\underline{L} = \overline{\mathcal{I}^{<k}(P)}$, and $P' \cong_{\text{ml}} Q'$

Details of this extension of CCS^{sg} can be found in [53].

3.7 Concluding Remarks and Related Work

We conclude by first commenting on the design decision that priority values are considered to be part of ports, which implies that only complementary actions having the same priority can synchronize. Lifting this design decision by allowing $a:k$ and $\bar{a}:l$, where $k \neq l$, to synchronize leads to the question of which priority value to assign to the resulting τ . One can imagine several obvious choices for this function, e.g., maximum or minimum. In addition, [35,37] recommend to use the sum of the priority values of the actions involved. Unfortunately, while a specific function may be suitable for certain examples, it is difficult to motivate for general applications. In the next section we will see that such a function is superfluous when dealing with *local pre-emption*.

Regarding related work, Gerber and Lee developed a real-time process algebra, the *Calculus of Communicating Shared Resources* (CCSR) [34], that explicitly takes the availability of system resources into account. Semantically, synchronizations between processes are modeled in an interleaving fashion using instantaneous transitions, whereas the access of resources is truly concurrent and consumes time. In CCSR a priority structure may be defined over resources in order to indicate their importance, e.g., for ensuring that deadlines are met. The underlying concept of priority is similar to that of CCS^{sg} in that priorities are static and pre-emption is global. In [35] a resource-based prioritized (strong) bisimulation for CCSR, together with axiomatizations with respect to several classes of processes [21] is presented.

Prasad also extended his *Calculus of Broadcasting Systems* (CBS) [68] for dealing with a notion of static priority [69]. He refers to the priority calculus as PCBS. For PCBS nice semantic theories based on Milner's strong and weak bisimulation [56] have been developed. Remarkably, these theories do not suffer from the technical subtleties which have been encountered for CCS^{sg} , although the concept of pre-emption is basically the same. The reason is that PCBS uses a much simpler model for communication which is based on the principle of *broadcasting*. In this setting, priority values are only attached to output actions, which cannot be restricted or hidden as in traditional process algebras. Finally, it should be mentioned that PCBS contains an operator, called *translate*, which allows for the prioritization and the deprioritization of actions.

4 Static Priority and Local Pre-emption

This section provides a new semantics for our language, subsequently referred to as CCS^{sl} (CCS with static priority and local pre-emption), which is distin-

guished from the one developed in the previous section by the design decision that it only allows actions to pre-empt others at the same “location” and, therefore, captures a notion of *localized precedence*. This constraint reflects an essential intuition about distributed systems, namely, that the execution of a process on one processor should not affect the behavior of a process on another processor unless the designer explicitly builds an interaction, e.g., a synchronization, between them.




Fig. 4. Example system

The following example demonstrates the necessity to consider locations when reasoning about priority within distributed systems. The example system consists of an application that manipulates data from two memory benches (cf. Figure 4). In order to improve the efficiency in the computer system, each bench, **Bench1** and **Bench2**, is connected to a direct-memory-access (DMA) controller. To overcome the low speed of most memory modules, application **Appl** works alternately with each memory bench. We model **Appl** in CCS^{sl} by $\text{Appl} \stackrel{\text{def}}{=} \underline{\text{fetch1}}.\underline{\text{fetch2}}.\text{Appl}$. Each memory bench, **Bench1** and **Bench2**, is continuously able to serve the application or to allow the external DMA controller to access the memory via channel **dma**. However, if a memory bench has to decide between both activities, then it chooses the former since the progress of the application is considered more important. Consequently, we define $\text{Bench1} \stackrel{\text{def}}{=} \underline{\text{fetch1}}.\text{Bench1} + \text{dma}.\text{Bench1}$ and $\text{Bench2} \stackrel{\text{def}}{=} \underline{\text{fetch2}}.\text{Bench2} + \text{dma}.\text{Bench2}$. The overall system **Sys** is given by

$$\text{Sys} \stackrel{\text{def}}{=} (\text{Appl} \mid \text{Bench1} \mid \text{Bench2}) \setminus \{\underline{\text{fetch1}}, \underline{\text{fetch2}}\} .$$

Since the application uses the memory cells alternately, the DMA is expected to be allowed to access the memory bench which is currently not serving the application. However, when using the approach to priority involving global pre-emption presented in Section 3, all **dma**-transitions in the labeled transition system of **Sys** are pre-empted since the application can indefinitely engage in a prioritized communication, i.e., direct-memory-access is never granted.

Generally speaking, one would expect that priorities at different *sites* of a distributed system do not influence the behavior of each other, i.e., priorities at different sites are supposed to be incomparable. The semantics given in

Section 3 does not permit this distinction to be made; the net effect is that some computations that one would expect to find in a distributed system are improperly suppressed. It has been proposed to remedy this shortcoming regarding distributed systems by introducing *local pre-emption* [24,29].

The remainder of this section is organized as follows. The next section introduces a notion of locations that is used in Section 4.2 for the definition of the operational semantics of CCS^{sl} for a two-level priority scheme. Sections 4.3 and 4.4 develop the semantic theories based on strong and weak bisimulation, respectively, while Section 4.5 re-considers the direct-memory-access example. The consequences of lifting some design decisions in CCS^{sl} are discussed in Section 4.6. After extending CCS^{sl} to a multi-level priority scheme in Section 4.7 and presenting another approach to priority taken from [24] in Section 4.8, a formal comparison of the two approaches is given in Section 4.9. Finally, Section 4.10 concludes with some additional remarks and with comments on related work.

4.1 Locations

We now introduce the notion of *location*, which will be used in the next section in the operational semantics for CCS^{sl} as a basis for deciding when one transition pre-empts another. Intuitively, a location represents the “address(es)” of subterm(s) inside a larger term; when a system performs an action, CCS^{sl} semantics will also note the location of the subterm(s) that “generate(s)” this action. Observe that because of the potential for synchronization more than one subterm may be involved in an action. The account of locations closely follows that of [29,60].

Formally, let $\mathcal{A}_{\text{addr}} =_{\text{df}} \{L, R, l, r\}$ be the *address alphabet*, and let \bullet be a special symbol not in $\mathcal{A}_{\text{addr}}$. Then, $\mathcal{A}\text{ddr} =_{\text{df}} \{\bullet s \mid s \in \mathcal{A}_{\text{addr}}^*\}$ represents the set of (process) *addresses* ranged over by v, w . Intuitively, an element of $\mathcal{A}\text{ddr}$ represents the address of a subterm, with \bullet denoting the current term, l (r) representing the left (right) subterm of $+$, and L (R) the left (right) subterm of $|$. For example, in the process $(a.\mathbf{0} \mid b.\mathbf{0}) + c.\mathbf{0}$ the address of $a.\mathbf{0}$ is $\bullet Ll$, of $b.\mathbf{0}$ is $\bullet Rl$, and of $c.\mathbf{0}$ is $\bullet r$. If $\bullet s_1$ and $\bullet s_2$ are addresses, then we write $\bullet s_1 \cdot \bullet s_2 = \bullet s_1 s_2$ to represent address concatenation (where $s_1 s_2$ represents the usual concatenation of elements in $\mathcal{A}_{\text{addr}}^*$). Further, if $V \subseteq \mathcal{A}\text{ddr}$ and $\zeta \in \mathcal{A}_{\text{addr}}$, then we write $V \cdot \zeta$ for $\{v \cdot \zeta \mid v \in V\}$. Occasionally, we omit \bullet from addresses.

As mentioned in the previous section, we want to adopt the view that processes at different sides of the parallel composition operator are logically – not necessarily physically – executed on different processors. Thus, priorities

on different sides of the parallel composition operator are distributed and, therefore, should be incomparable. However, priorities on different sides of the summation operator should be comparable since argument processes of summation are logically scheduled on the same processor. This intuition is formalized in the following *location comparability relation* on addresses which is adapted from [37].

Definition 4.1 (Location Comparability Relation)

The location comparability relation \bowtie on addresses is the smallest reflexive and symmetric subset of $\mathit{Addr} \times \mathit{Addr}$ such that for all $v, w \in \mathit{Addr}$:

- (1) $\langle v \cdot l, w \cdot r \rangle \in \bowtie$, and
- (2) $\langle v, w \rangle \in \bowtie$ implies $\langle v \cdot \zeta, w \cdot \zeta \rangle \in \bowtie$ for $\zeta \in \mathcal{A}_{\mathit{addr}}$.

In the sequel we write $v \bowtie w$ instead of $\langle v, w \rangle \in \bowtie$. If $v \in \mathit{Addr}$ then we use $[v]$ to denote the set $\{w \in \mathit{Addr} \mid v \bowtie w\}$. Note that the location comparability relation is not transitive, e.g., we have $Ll \bowtie r$ and $r \bowtie Rl$, but $Ll \not\bowtie Rl$, since $L \not\bowtie R$.

We may now define the set $\mathcal{L}oc$ of (transition) *locations* as $\mathit{Addr} \cup (\mathit{Addr} \times \mathit{Addr})$. Intuitively, a transition location records the addresses of the components in a term that participate in the execution of a given action. In our language, transitions are performed by single processes or pairs of processes (in the case of a synchronization). We define $\langle v, w \rangle \cdot \zeta =_{\text{df}} \langle v \cdot \zeta, w \cdot \zeta \rangle$ and $[\langle v, w \rangle] =_{\text{df}} [v] \cup [w]$ where $v, w \in \mathit{Addr}$ and $\zeta \in \mathcal{A}_{\mathit{addr}}$. We use m, n, o, \dots to range over $\mathcal{L}oc$ in what follows.

4.2 Operational Semantics

The operational semantics of a CCS^{sl} process P is given by a labeled transition system $\langle \mathcal{P}, \mathcal{L}oc \times \mathcal{A}, \longrightarrow, P \rangle$. The transition relation $\longrightarrow \subseteq \mathcal{P} \times (\mathcal{L}oc \times \mathcal{A}) \times \mathcal{P}$ with respect to unprioritized actions is defined in Table 12 using Plotkin-style operational rules [67] whereas for prioritized actions the same rules as for CCS^{sg} apply (see Table 2). We write $P \xrightarrow{m, \alpha} P'$ if $\langle P, \langle m, \alpha \rangle, P' \rangle \in \longrightarrow$ and $\alpha \in \mathcal{A}$, and say that P may engage in action α offered from location m and thereafter behave like process P' .

The presentation of the operational rules requires *distributed prioritized initial action sets* which are defined as the least sets satisfying the rules in Table 11. Intuitively, $\underline{\mathcal{I}}_m(P)$ denotes the set of all prioritized initial actions of P from location m . Note that these sets are either empty or contain exactly one initial action. $\underline{\mathcal{I}}_m(P) = \emptyset$ means that either m is not a location of P or P is incapable of performing a prioritized action at location m . Additionally, let us denote the set $\bigcup \{ \underline{\mathcal{I}}_m(P) \mid m \in M \}$ of all distributed prioritized initial actions of P from

Table 11

Distributed prioritized initial action sets for CCS^{sl}

$\underline{\mathcal{I}}_m(\mu x.P) = \underline{\mathcal{I}}_m(P[\mu x.P/x])$	$\underline{\mathcal{I}}_\bullet(\underline{\alpha}.P) = \{\underline{\alpha}\}$
$\underline{\mathcal{I}}_{m.l}(P + Q) = \underline{\mathcal{I}}_m(P)$	$\underline{\mathcal{I}}_{n.r}(P + Q) = \underline{\mathcal{I}}_n(Q)$
$\underline{\mathcal{I}}_m(P[f]) = \{f(\underline{\alpha}) \mid \underline{\alpha} \in \underline{\mathcal{I}}_m(P)\}$	$\underline{\mathcal{I}}_m(P \setminus L) = \underline{\mathcal{I}}_m(P) \setminus (L \cup \bar{L})$
$\underline{\mathcal{I}}_{m.L}(P Q) = \underline{\mathcal{I}}_m(P)$	$\underline{\mathcal{I}}_{n.R}(P Q) = \underline{\mathcal{I}}_n(Q)$
$\underline{\mathcal{I}}_{(m.L, n.R)}(P Q) = \{\underline{\tau} \mid \underline{\mathcal{I}}_m(P) \cap \overline{\underline{\mathcal{I}}_n(Q)} \neq \emptyset\}$	

Table 12

Operational semantics for CCS^{sl}

Act	$\frac{-}{\alpha.P \xrightarrow{\bullet, \alpha} P}$	Sum1	$\frac{P \xrightarrow{m, \alpha} P'}{P + Q \xrightarrow{m.l, \alpha} P'} \quad \underline{\tau} \notin \underline{\mathcal{I}}(Q)$
Rel	$\frac{P \xrightarrow{m, \alpha} P'}{P[f] \xrightarrow{m, f(\alpha)} P'[f]}$	Sum2	$\frac{Q \xrightarrow{n, \alpha} Q'}{P + Q \xrightarrow{n.r, \alpha} Q'} \quad \underline{\tau} \notin \underline{\mathcal{I}}(P)$
Res	$\frac{P \xrightarrow{m, \alpha} P'}{P \setminus L \xrightarrow{m, \alpha} P' \setminus L} \quad \alpha \notin L \cup \bar{L}$	Com1	$\frac{P \xrightarrow{m, \alpha} P'}{P Q \xrightarrow{m.L, \alpha} P' Q} \quad \underline{\mathcal{I}}_{[m]}(P) \cap \overline{\underline{\mathcal{I}}(Q)} = \emptyset$
Rec	$\frac{P[\mu x.P/x] \xrightarrow{m, \alpha} P'}{\mu x.P \xrightarrow{m, \alpha} P'}$	Com2	$\frac{Q \xrightarrow{n, \alpha} Q'}{P Q \xrightarrow{n.R, \alpha} P Q'} \quad \underline{\mathcal{I}}_{[n]}(Q) \cap \overline{\underline{\mathcal{I}}(P)} = \emptyset$
		Com3	$\frac{P \xrightarrow{m, \alpha} P' \quad Q \xrightarrow{n, \bar{\alpha}} Q'}{P Q \xrightarrow{(m.L, n.R), \tau} P' Q'} \quad \underline{\mathcal{I}}_{[m]}(P) \cap \overline{\underline{\mathcal{I}}(Q)} = \emptyset \wedge \underline{\mathcal{I}}_{[n]}(Q) \cap \overline{\underline{\mathcal{I}}(P)} = \emptyset$

locations $M \subseteq \mathcal{Loc}$ by $\underline{\mathcal{I}}_M(P)$ and the set $\underline{\mathcal{I}}_{\mathcal{Loc}}(P)$ of all distributed prioritized initial actions of P by $\underline{\mathcal{I}}(P)$. We also define analogue sets restricted to visible actions: $\underline{\mathcal{I}}_M(P) =_{\text{df}} \underline{\mathcal{I}}_M(P) \setminus \{\underline{\tau}\}$ and $\underline{\mathcal{I}}(P) =_{\text{df}} \underline{\mathcal{I}}(P) \setminus \{\underline{\tau}\}$, respectively.

The side conditions of the operational rules guarantee that a process does not perform an unprioritized action if it can engage in a prioritized synchronization or in internal computation, i.e., a $\underline{\tau}$ -transition, from a comparable location. In contrast to the global notion of pre-emption defined in Section 3, the local notion here is much weaker since $\underline{\mathcal{I}}_{[m]}(P) \subseteq \underline{\mathcal{I}}(P)$ for all $m \in \mathcal{Loc}$ and $P \in \mathcal{P}$. In other words, local pre-emption does not pre-empt as many transitions as global pre-emption does. The difference between CCS^{sl} and CCS^{sg} semantics arises by the side conditions of the rules for parallel composition

with respect to unprioritized transitions. Since locations on different sides of parallel composition $P|Q$ are incomparable, τ 's arising from a location of P (Q) cannot pre-empt the execution of a transition, even an unprioritized one, of Q (P). Only if P (Q) engages in a prioritized synchronization with Q (P) can unprioritized actions from a comparable location of P (Q) be pre-empted.

4.3 Semantic Theory Based on Strong Bisimulation

Just as in Section 3, we present an equivalence relation for CCS^{sl} processes that is based on bisimulation [65]. Our aim is to characterize the largest congruence contained in the “naive” adaptation of strong bisimulation [56] to our framework obtained by ignoring location information.

Definition 4.2 (Naive Distributed Prioritized Strong Bisimulation)

A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is called naive distributed prioritized strong bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$ and $\gamma \in \mathcal{A}$ the following condition holds.

$$P \xrightarrow{\gamma} P' \text{ implies } \exists Q'. Q \xrightarrow{\gamma} Q' \text{ and } \langle P', Q' \rangle \in \mathcal{R} .$$

We write $P \simeq Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some naive distributed prioritized strong bisimulation \mathcal{R} .

Although \simeq is an equivalence, it is unfortunately – in contrast to Section 3.2 – not a congruence. The lack of compositionality is demonstrated by the following example, which embodies the traditional view that “parallelism = non-determinism.” We have $a.\underline{b}.\mathbf{0} + \underline{b}.a.\mathbf{0} \simeq a.\mathbf{0}|\underline{b}.\mathbf{0}$ but $(a.\underline{b}.\mathbf{0} + \underline{b}.a.\mathbf{0})|\bar{b}.\mathbf{0} \not\approx (a.\mathbf{0}|\underline{b}.\mathbf{0})|\bar{b}.\mathbf{0}$, since the latter process can perform an a -transition, while the corresponding a -transition of the former is pre-empted because the right process in the summation can engage in a prioritized communication. The above observation is not surprising since the distribution of processes influences the pre-emption of transitions and, consequently, the bisimulation. However, we know by Proposition 3.6 that \simeq includes a largest congruence \simeq^+ for CCS^{sl} .

4.3.1 Distributed Prioritized Strong Bisimulation

In the remainder we develop a characterization of \simeq^+ . To do so, we need to take *local* pre-emption into account.

Definition 4.3 A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is a distributed prioritized strong bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$, $\alpha \in A$, $\underline{\alpha} \in \underline{A}$, and $m \in \mathcal{L}oc$ the following conditions hold.

$$(1) P \xrightarrow{\alpha} P' \text{ implies } \exists Q'. Q \xrightarrow{\alpha} Q' \text{ and } \langle P', Q' \rangle \in \mathcal{R} .$$

- (2) $P \xrightarrow{m,\alpha} P'$ implies
 $\exists Q', n. Q \xrightarrow{n,\alpha} Q', \underline{\mathcal{I}}_{[n]}(Q) \subseteq \underline{\mathcal{I}}_{[m]}(P)$, and $\langle P', Q' \rangle \in \mathcal{R}$.

We write $P \simeq^1 Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some distributed prioritized strong bisimulation \mathcal{R} .

Intuitively, the distributed prioritized initial action set of a process with respect to some location is a measure of the pre-emptive power of the process relative to that location. Thus, the second condition of Definition 4.3 states that an unprioritized action α from some location m of process P must be matched by the same action from some location n of Q and that the pre-emptive power of Q relative to n is at most as strong as the pre-emptive power of P relative to m . The next theorem is the main result of this section.

Theorem 4.4 \simeq^1 is the largest congruence contained in \simeq .

We refer for the proof to [53]; the context needed in the largest congruence proof is similar to the one used in Section 3.3.

4.3.2 Axiomatic Characterization

In this section we present an axiomatization of \simeq^1 for *finite* processes, for which we introduce a new binary summation operator \oplus to the process algebra CCS^{sl} . This operator is called *distributed summation* and is needed for giving an *Expansion Axiom* (cf. Axiom (E) in Table 13). Its operational semantics is defined below and differs from the nondeterministic choice operator $+$ in that a location in its left argument is never comparable to one in its right argument.

$$\begin{array}{ll} \underline{\text{dSum1}} \frac{t \xrightarrow{\alpha} t'}{t \oplus u \xrightarrow{\alpha} t'} & \underline{\text{dSum1}} \frac{t \xrightarrow{m,\alpha} t'}{t \oplus u \xrightarrow{m.L,\alpha} t'} \\ \underline{\text{dSum2}} \frac{u \xrightarrow{\alpha} u'}{t \oplus u \xrightarrow{\alpha} u'} & \underline{\text{dSum2}} \frac{u \xrightarrow{n,\alpha} u'}{t \oplus u \xrightarrow{n.R,\alpha} u'} \end{array}$$

It can easily be checked that \simeq^1 is compositional with respect to \oplus .

Now, we turn to the axiom system for distributed prioritized strong bisimulation. We write $\vdash_E t = u$ if term t can be rewritten to u using the axioms in Tables 13 and 14 as well as Axioms (A1)–(A4), Axioms (Res1)–(Res4), Axioms (Rel1)–(Rel3), and Axiom (P) from Table 3. Axioms (lc1), (D1), (S2), and (S3) involve side conditions. Regarding Axiom (lc1), we introduce the unary predicate \natural over processes (of the form $\sum_{j \in J} \gamma_j.t_j$ for some nonempty index set J) together with the following proof rules: (i) $\natural \underline{\alpha}.t$ and (ii) $\natural t$ and $\natural u$ implies $\natural(t + u)$. Intuitively, $\natural(\sum_{j \in J} \gamma_j.t_j)$ if and only if $\gamma_j \in \underline{A}$ for all $j \in J$. The

Table 13

Axiomatization of \simeq^1 (Axioms E)

(iA1) $t \oplus u = u \oplus t$	(iA2) $t \oplus (u \oplus v) = (t \oplus u) \oplus v$
(iA3) $t \oplus t = t$	(iA4) $t \oplus \mathbf{0} = t$
(E) $t \equiv \bigoplus_i \sum_j \gamma_{ij}.t_{ij}$ and $u \equiv \bigoplus_k \sum_l \delta_{kl}.u_{kl}$ implies $t u =$ $\bigoplus_i \sum_j (\gamma_{ij}.(t_{ij} u) + \sum_k \sum_l \{\tau.(t_{ij} u_{kl}) \gamma_{ij} \equiv \bar{\delta}_{kl}, \gamma_{ij}, \delta_{kl} \in A\}$ $\quad + \sum_k \sum_l \{\underline{\tau}.(t_{ij} u_{kl}) \gamma_{ij} \equiv \bar{\delta}_{kl}, \gamma_{ij}, \delta_{kl} \in \underline{A}\}) \oplus$ $\bigoplus_k \sum_l (\delta_{kl}.(t u_{kl}) + \sum_i \sum_j \{\tau.(t_{ij} u_{kl}) \gamma_{ij} \equiv \bar{\delta}_{kl}, \gamma_{ij}, \delta_{kl} \in A\}$ $\quad + \sum_i \sum_j \{\underline{\tau}.(t_{ij} u_{kl}) \gamma_{ij} \equiv \bar{\delta}_{kl}, \gamma_{ij}, \delta_{kl} \in \underline{A}\})$	
(iRes4) $(t \oplus u) \setminus L = (t \setminus L) \oplus (u \setminus L)$ (iRel3) $(t \oplus u)[f] = t[f] \oplus u[f]$	

Table 14

Axioms E (continued)

(D1) $(t \oplus t') + (u \oplus u') = ((t \oplus t') + u') \oplus ((u \oplus u') + t')$	
$(\vdash_I t \sqsubseteq_i t', \vdash_I u \sqsubseteq_i u')$	
(D2) $(t \oplus u) + \underline{\alpha}.v = (t + \underline{\alpha}.v) \oplus (u + \underline{\alpha}.v)$	
(lc1) $t \oplus \underline{\alpha}.u = t + \underline{\alpha}.u$	(lt)
(lc2) $(\underline{\alpha}.t + u) = (\underline{\alpha}.t + u) \oplus \underline{\alpha}.t$	
(S1) $(t + \underline{\alpha}.u) \oplus (t' + \underline{\alpha}.u') = (t + \underline{\alpha}.u + \underline{\alpha}.u') \oplus (t' + \underline{\alpha}.u')$	
(S2) $(t + \alpha.v) \oplus (u + \alpha.v) = (t + \alpha.v) \oplus u$	($\vdash_I t \sqsubseteq_i u$)
(S3) $t \oplus u = t + u$	($\vdash_I t =_i u$)

relation \sqsubseteq_i is defined as in Section 3.5 (see Table 8). The axioms in Table 13 are basically those given in Table 3 and augmented with the corresponding axioms for the distributed summation operator. Moreover, the Expansion Axiom has been adapted for our algebra (cf. Axiom (E), where \sum is the indexed version of $+$, and \bigoplus is the indexed version of \oplus). Note that parallelism in CCS^{sl} cannot be resolved in nondeterminism by using operator $+$ only, since priorities on different sides of $|$ are incomparable, but priorities on different sides of $+$ are comparable. This is the motivation for introducing operator \oplus . The axioms in Table 14 are new and show how we may “restructure” locations. They deal with the *distributivity* of the summation operators (Axioms (D1) and (D2)), the *interchangeability* of the summation operators (Axioms (lc1) and (lc2)), and the *saturation* of locations (Axioms (S1), (S2), and (S3)), respectively.

Theorem 4.5 *Let t, u be finite processes. Then $\vdash_E t = u$ if and only if $t \simeq^1 u$.*

The proof of this theorem can be found in [29,53].

4.3.3 Operational Characterization

The following definition introduces an equivalence \simeq_* which characterizes \simeq^1 as standard strong bisimulation [53]. It uses the notation $P \xrightarrow{\alpha} P'$ for $P, P' \in \mathcal{P}$, $\alpha \in A$, and $L \subseteq \underline{A} \setminus \{\tau\}$ whenever $\exists m \in \mathcal{L}oc. P \xrightarrow{m, \alpha} P'$ and $\underline{\mathcal{L}}_{[m]}(P) \subseteq L$. Note that these enriched transitions take local pre-emption potentials into account, thereby avoiding the explicit annotation of transitions with locations.

Definition 4.6 *A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is an alternative distributed prioritized strong bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$, $\alpha \in A$, $\underline{\alpha} \in \underline{A}$, and $L \subseteq \underline{A} \setminus \{\tau\}$ the following conditions hold.*

- (1) $P \xrightarrow{\underline{\alpha}} P'$ implies $\exists Q'. Q \xrightarrow{\underline{\alpha}} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$.
- (2) $P \xrightarrow{\alpha}_L P'$ implies $\exists Q'. Q \xrightarrow{\alpha}_L Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$.

We write $P \simeq_* Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some alternative distributed prioritized strong bisimulation \mathcal{R} .

Similar to Section 3.3.3 we obtain an operational characterization of our behavioral relation.

Theorem 4.7 (Operational Characterization) $\simeq^1 = \simeq_*$.

4.4 Semantic Theory Based on Weak Bisimulation

As for CCS^{sg}, we develop a coarser bisimulation-based congruence by abstracting from internal actions. We start off with the definition of a *naive distributed prioritized weak bisimulation* which is an adaptation of observational equivalence [56].

Definition 4.8 (Naive Distributed Prioritized Weak Transitions)

- (1) $\hat{\gamma} =_{df} \epsilon$, if $\gamma \in \{\tau, \tau\}$, and $\hat{\gamma} =_{df} \gamma$, otherwise
- (2) $\xrightarrow{\epsilon}_\times =_{df} (\tau \cup \cup \{ \xrightarrow{m, \tau} \mid m \in \mathcal{L}oc \})^*$
- (3) $\xrightarrow{\underline{\alpha}}_\times =_{df} \xrightarrow{\epsilon}_\times \circ \xrightarrow{\underline{\alpha}} \circ \xrightarrow{\epsilon}_\times$
- (4) $\xrightarrow{m, \underline{\alpha}}_\times =_{df} \xrightarrow{\epsilon}_\times \circ \xrightarrow{m, \underline{\alpha}} \circ \xrightarrow{\epsilon}_\times$

In the following we write $P \xrightarrow{\alpha}_\times P'$ for $\exists m \in \mathcal{L}oc. P \xrightarrow{m, \alpha}_\times P'$.

Definition 4.9 (Naive Distributed Prioritized Weak Bisimulation)

A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is a naive distributed prioritized weak bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$ and $\gamma \in \mathcal{A}$ the following condition holds.

$$P \xrightarrow{\gamma} P' \text{ implies } \exists Q'. Q \xrightarrow{\hat{\gamma}}_{\times} Q' \text{ and } \langle P', Q' \rangle \in \mathcal{R} .$$

We write $P \approx_{\times} Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some naive distributed prioritized weak bisimulation \mathcal{R} .

It is fairly easy to see that \approx_{\times} is not a congruence for CCS^{sl} . One compositionality defect arises with respect to parallel composition and is similar to the one mentioned for naive distributed prioritized strong bisimulation. Another defect, which is carried over from CCS, is concerned with the summation operators.

4.4.1 *Distributed Prioritized Weak Bisimulation*

We devote the rest of this section to characterizing the largest congruence contained in naive distributed prioritized weak bisimulation. To do so, we first re-define the weak transition relation.

Definition 4.10 (Distributed Prioritized Weak Transitions)

For $L, M \subseteq \underline{\mathcal{A}} \setminus \{\tau\}$ we define the following notations.

- (1) $\hat{\epsilon} =_{df} \epsilon$, $\hat{a} =_{df} a$, $\hat{\tau} =_{df} \tau$, and $\hat{a} =_{df} a$
- (2) $P \xrightarrow[L]{m, \alpha} P'$ if $P \xrightarrow{m, \alpha} P'$ and $\underline{\mathcal{I}}_{[m]}(P) \subseteq L$
- (3) $\xrightarrow{\epsilon} =_{df} (\xrightarrow{\tau} \cup \cup \{ \xrightarrow[m]{m, \tau} \mid m \in \text{Loc} \})^*$
- (4) $\xrightarrow{\alpha} =_{df} \xrightarrow{\epsilon} \circ \xrightarrow{\alpha} \circ \xrightarrow{\epsilon}$
- (5) $\xrightarrow[L]{\epsilon} =_{df} (\xrightarrow{\tau} \cup \cup \{ \xrightarrow[L]{m, \tau} \mid m \in \text{Loc} \})^*$
- (6) $P \xrightarrow[L, M]{m, \alpha} P'$ if $\exists P'', P'''. P \xrightarrow[L]{\epsilon} P'' \xrightarrow[L]{m, \alpha} P''' \xrightarrow{\epsilon} P'$ and $\underline{\mathcal{I}}(P'') \subseteq M$.

Intuitively, these definitions are designed to reflect constraints that a process environment must satisfy in order for the given transition to be enabled. Thus, $P \xrightarrow[L]{m, \alpha} P'$ means that P can engage in action α at location m to P' provided that the environment does not offer a prioritized communication involving actions in L . If the environment were to offer such a communication, the result would be a τ at a comparable location to m in P , which would pre-empt the α . In a similar vein, $P \xrightarrow{\epsilon} P'$ holds if P can evolve to P' via a nonpre-emptable sequence of internal transitions, regardless of the environment's behavior. These internal transitions should therefore involve either τ , which can never be pre-empted, or τ , in which case no prioritized actions should be enabled at the same location. Likewise, $P \xrightarrow[L]{\epsilon} P'$ means that, so long as the environment does not offer to synchronize with P on a prioritized action in L , process P may engage in a sequence of internal computation steps and become P' . Finally,

the M -parameter in $\xrightarrow[L, M]{m, \alpha}$ provides a measure of the pre-emptive impact that a process has on its environment. From the definition, $P \xrightarrow[L, M]{m, \alpha} P'$ is true if P can engage in some internal computation followed by α , so long as the environment refrains from synchronizations in L , and then engage in some nonpre-emptable internal computation to arrive at P' . In addition, the state at which α is enabled should only offer prioritized communications in M .

Note that the definition of $P \xrightarrow[L]{\epsilon} P'$ is in accordance with our intuition that internal actions, and therefore their locations, are unobservable. Moreover, the environment of P is not influenced by internal actions performed by P , since priorities arising from different sides of the parallel composition operator are incomparable. Therefore, the parameter M is unnecessary in the definition of the relation $\xrightarrow[L]{\epsilon}$. Finally, for notational convenience $\xrightarrow[L, M]{m, \epsilon}$ is interpreted as $\xrightarrow[L]{\epsilon}$.

Definition 4.11 (Distributed Prioritized Weak Bisimulation)

A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is a distributed prioritized weak bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$, $\alpha \in A$, $\underline{\alpha} \in \underline{A}$, and $m \in \mathcal{L}oc$ the following conditions hold.

- (1) $\exists Q', Q''. Q \xrightarrow{\epsilon} Q'' \xrightarrow{\epsilon} Q', \underline{\mathcal{I}}(Q'') \subseteq \underline{\mathcal{I}}(P)$, and $\langle P, Q' \rangle \in \mathcal{R}$.
- (2) $P \xrightarrow{\underline{\alpha}} P'$ implies $\exists Q'. Q \xrightarrow{\hat{\alpha}} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$.
- (3) $P \xrightarrow{m, \alpha} P'$ implies $\exists Q', n. Q \xrightarrow[L, M]{n, \hat{\alpha}} Q', L = \underline{\mathcal{I}}_{[m]}(P), M = \underline{\mathcal{I}}(P)$, and $\langle P', Q' \rangle \in \mathcal{R}$.

We write $P \cong Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some distributed prioritized weak bisimulation \mathcal{R} .

Condition (1) of Definition 4.11 guarantees that distributed prioritized weak bisimulation is compositional with respect to parallel composition. Its necessity is best illustrated by the following example. The processes $P \stackrel{\text{def}}{=} \tau.\underline{a}.\mathbf{0}$ and $Q \stackrel{\text{def}}{=} \underline{a}.\mathbf{0}$ would be considered equivalent if Condition (1) were absent. However, the context $C[X] \stackrel{\text{def}}{=} X | (\underline{a}.\mathbf{0} + b.\mathbf{0})$ distinguishes them. The following proposition is the CCS^{sl} equivalent of Proposition 3.9.

Proposition 4.12 *The equivalence relation \cong is a congruence with respect to prefixing, parallel composition, relabeling, and restriction. Moreover, \cong is characterized as the largest congruence contained in \approx_{\times} , in the sub-algebra of CCS^{sl} induced by these operators and recursion.*

4.4.2 Distributed Prioritized Observational Congruence

Analogue to Section 3, the summation fix presented in [56] is not sufficient in order to achieve a congruence relation.

Definition 4.13 We define $P \cong^1 Q$ if for all $\alpha \in A$, $\underline{\alpha} \in \underline{A}$, and $m \in \mathcal{Loc}$ the following conditions and their symmetric counterparts hold.

- (1) $\underline{\mathcal{I}}(P) \supseteq \underline{\mathcal{I}}(Q)$
- (2) $P \xrightarrow{\underline{\alpha}} P'$ implies $\exists Q'. Q \xrightarrow{\underline{\alpha}} Q'$ and $P' \cong Q'$.
- (3) $P \xrightarrow{m, \underline{\alpha}} P'$ implies $\exists Q'. n. Q \xrightarrow{L, M}^{n, \underline{\alpha}} Q', L = \underline{\mathcal{I}}_{[m]}(P), M = \underline{\mathcal{I}}(P)$, and $P' \cong Q'$.

Now, we can state the main result of this section, which can be proved by using the technique already presented in Section 3.3.2; see [53] for details.

Theorem 4.14 \cong^1 is the largest congruence contained in \approx_\times .

4.4.3 Operational Characterization

We now characterize distributed prioritized weak bisimulation as standard bisimulation over an appropriately defined transition relation. To begin with, we introduce a family of relations \xrightarrow{M} on processes, where $M \subseteq \underline{A} \setminus \{\underline{\tau}\}$, by defining $P \xrightarrow{M} P'$ if $\exists P''. P \xrightarrow{\underline{\epsilon}} P'' \xrightarrow{\underline{\epsilon}} P'$ and $\underline{\mathcal{I}}(P'') \subseteq M$. Moreover, we write $P \xrightarrow{L, M}^{\hat{\alpha}} P'$ whenever there exists some $m \in \mathcal{Loc}$ such that $P \xrightarrow{L, M}^{m, \hat{\alpha}} P'$.

Definition 4.15 A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is an alternative distributed prioritized weak bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$, $\alpha \in A$, $\underline{\alpha} \in \underline{A}$, and $L, M \subseteq \underline{A} \setminus \{\underline{\tau}\}$ the following conditions hold.

- (1) $P \xrightarrow{M} P'$ implies $\exists Q'. Q \xrightarrow{M} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$.
- (2) $P \xrightarrow{\hat{\alpha}} P'$ implies $\exists Q'. Q \xrightarrow{\hat{\alpha}} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$.
- (3) $P \xrightarrow{L, M}^{\hat{\alpha}} P'$ implies $\exists Q'. Q \xrightarrow{L, M}^{\hat{\alpha}} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$.

We write $P \cong_* Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some alternative distributed prioritized weak bisimulation \mathcal{R} .

Theorem 4.16 (Operational Characterization) $\cong = \cong_*$.

The interested reader can find the proof of this theorem in [53].

4.5 Example

We now return to the direct-memory-access example system introduced in Section 4. The CCS^{sl} semantics of Sys, which corresponds to our intuition regarding distributed systems, is given in Figure 5 where we abstract away the locations.




Fig. 5. Semantics of the dma-system

As stated before, the application uses the two memory cells alternately. Thus, the DMA is expected to be allowed to access the free memory bench. Accordingly, the specification of the system can be formalized by $\text{Spec} \stackrel{\text{def}}{=} \text{dma.Spec}$. It is easy to see that the symmetric closure of

$$\{\langle \text{Spec}, \text{Sys} \rangle, \langle \text{Spec}, (\overline{\text{fetch2}}.\text{Appl} \mid \text{Bench1} \mid \text{Bench2}) \setminus \{\text{fetch1}, \text{fetch2}\} \rangle\}$$

is a distributed prioritized weak bisimulation. Therefore, $\text{Spec} \cong \text{Sys}$ as expected, i.e., system Sys meets its specification Spec .

4.6 Discussion on the Alteration of Some Design Decisions

Up to now we have restricted the number of priority levels in CCS^{sl} to two and communication to complementary actions having the same priority. In this section we study the implications of the removal of these restrictions, which leads to a new version of CCS^{sl} , called $\text{CCS}_{\text{ml}}^{\text{sl}}$ (CCS^{sl} with a multi-level priority scheme), that is formally defined in the next section.

Allowing communication between unprioritized actions and complementary prioritized actions raises the question of whether the resulting internal action should be τ or $\underline{\tau}$. When dealing with local pre-emption, this decision has no important consequences for sequential communicating processes, i.e., those in *standard concurrent form* [56]; however, it is of obvious importance for processes like $(\underline{a}.\mathbf{0} \mid \overline{a}.\mathbf{0}) + b.\mathbf{0}$ in which one has to decide if the b -transition is enabled. One reasonable view is that a communication should be pre-empted whenever one communication partner is pre-empted, i.e., it cannot engage in a communication. This implies that the *minimal* priority of the complementary actions ought to be assigned to the internal action. To reflect this in the operational semantics, one could replace Rules (Com1), (Com2), and (Com3) for parallel composition by the ones presented in Table 15 plus their symmetric counterparts. The side conditions involve sets $\mathcal{I}(P)$ that include all unprioritized visible actions in which P can initially engage.

It turns out that the largest congruence results concerning distributed prioritized strong bisimulation and distributed prioritized observational congruence can be carried over to the new calculus. However, the new semantics has algebraic shortcomings since parallel composition is *not associative*, as illustrated

Table 15
Modified operational rules

$$\begin{array}{l}
\text{Com1} \frac{P \xrightarrow{m, \alpha} P'}{P|Q \xrightarrow{m \cdot L, \alpha} P'|Q} \underline{\mathcal{I}}_{[m]}(P) \cap (\overline{\underline{\mathcal{I}}(Q)} \cup \overline{\mathcal{I}(Q)}) = \emptyset \\
\text{Com3a} \frac{P \xrightarrow{m, \alpha} P' \quad Q \xrightarrow{n, \bar{\alpha}} Q'}{P|Q \xrightarrow{\langle m \cdot L, n \cdot R \rangle, \tau} P'|Q'} \underline{\mathcal{I}}_{[m]}(P) \cap (\overline{\underline{\mathcal{I}}(Q)} \cup \overline{\mathcal{I}(Q)}) = \emptyset \wedge \\
\underline{\mathcal{I}}_{[n]}(Q) \cap (\overline{\underline{\mathcal{I}}(P)} \cup \overline{\mathcal{I}(P)}) = \emptyset \\
\text{Com3b} \frac{P \xrightarrow{m, \alpha} P' \quad Q \xrightarrow{n, \bar{\alpha}} Q'}{P|Q \xrightarrow{\langle m \cdot L, n \cdot R \rangle, \tau} P'|Q'} \underline{\mathcal{I}}_{[n]}(Q) \cap (\overline{\underline{\mathcal{I}}(P)} \cup \overline{\mathcal{I}(P)}) = \emptyset
\end{array}$$

by the following example. Consider process $(b \cdot \mathbf{0} + \underline{a} \cdot \mathbf{0}) | (\bar{a} \cdot \mathbf{0} + \underline{c} \cdot \mathbf{0}) | \bar{c} \cdot \mathbf{0}$. When computing the semantics in a left-associative manner, the initial b -transition is pre-empted according to Rule (Com1) since \underline{a} may potentially communicate with \bar{a} . However, when first composing the second and third parallel components, the \bar{a} -transition is pre-empted, and consequently the b -transition is enabled by Rule (Com1). The reason for this problem is that transitions are pre-empted because the considered process can *potentially* engage in a higher prioritized communication from a comparable location. However, this potential communication cannot take place if the communication partner is itself pre-empted. The same problem also arises when extending CCS^{sl} to multiple priority levels, even if communication is only allowed on complementary actions of the same priority, as can be observed by using a slight adaptation of the previous example: $(b : 2 \cdot \mathbf{0} + a : 1 \cdot \mathbf{0}) | (\bar{a} : 1 \cdot \mathbf{0} + c : 0 \cdot \mathbf{0}) | \bar{c} : 0 \cdot \mathbf{0}$.

One can imagine two approaches to fixing the problems with the first (and second) alteration to the theory. One is to change the operational semantics; in particular, the side conditions could be weakened such that an unprioritized transition is only pre-empted when a prioritized action from a comparable location can *actually* engage in a communication. This approach has not been investigated in the literature, yet. The second solution follows an approach developed in [24] for a different setting and involves the use of a syntax restriction on processes prohibiting output actions, i.e., actions in $\bar{\Lambda}$, from occurring as initial actions that are in the scope of $+$. Hence, all potential communication partners are also actual ones, and the standard side conditions for parallel composition are sufficient to encode the desired notion of pre-emption. It is important to mention that the proposed syntax restriction still allows one to specify many practically relevant examples within the calculus. Indeed, a similar restriction may be found in the programming language `occam` [43].

4.7 Extension to Multi-Level Priority Schemes

For $\text{CCS}_{\text{ml}}^{\text{sl}}$ we allow a multi-level priority scheme and communication between complementary actions having potentially different priorities. As seen in the previous section, both of these relaxations yield a semantics for which parallel composition is not associative. However, we have also argued that this problem vanishes if the syntax is restricted such that output actions can never be pre-empted. We adapt the syntax restriction proposed by Camilleri and Winskel [24], which states that initial actions in the scope of a comparable summation operator must be input actions. Therefore, input and output actions are explicitly distinguished in $\text{CCS}_{\text{ml}}^{\text{sl}}$; the internal action τ is also treated as input action. In the following, we let a, b, \dots range over the set Λ of input ports and \bar{a}, \bar{b}, \dots over the set $\bar{\Lambda}$ of output ports. Moreover, we let γ stand for any input action and let α range over $\mathcal{A} =_{\text{df}} \Lambda \cup \bar{\Lambda} \cup \{\tau\}$. Since in the restricted syntax priority values of output actions need never be compared with other priority values, there are no priority values associated with output actions at all. The syntax of $\text{CCS}_{\text{ml}}^{\text{sl}}$ is formally defined by the following BNF for P .

$$\begin{aligned} I & ::= \mathbf{0} \mid x \mid \gamma:k.I \mid I + I \mid I \oplus I \mid I \mid I \mid I[f] \mid I \setminus L \mid \mu x.I \\ P & ::= \mathbf{0} \mid x \mid \alpha:k.P \mid I + I \mid P \oplus P \mid P \mid P \mid P[f] \mid P \setminus L \mid \mu x.P \end{aligned}$$

Here, f is an *injective*, finite relabeling, $L \subseteq \Lambda \cup \bar{\Lambda}$ is a restriction set, and x is a variable taken from a countable domain \mathcal{V} . A relabeling satisfies the properties $f(\Lambda) \subseteq \Lambda$, $f(\bar{\Lambda}) \subseteq \bar{\Lambda}$, $f(\tau) = \tau$, and $f(\bar{a}) = \overline{f(a)}$. Thus, in addition to the requirements of a finite relabeling in CCS, relabelings in $\text{CCS}_{\text{ml}}^{\text{sl}}$ may only map input ports to input ports and output ports to output ports. Since actions attached with different priority values do not represent different ports here, relabelings and restriction sets do not deal with priority values. Especially, the priority value of a relabeled transition remains the same, i.e., there is no explicit or implicit mechanism for prioritization or deprioritization (cf. Section 3.5). In the remainder, we let $\mathcal{P}_{\text{ml}}^{\text{sl}}$ denote the set of all $\text{CCS}_{\text{ml}}^{\text{sl}}$ processes.

Table 16

Initial output action sets for $\text{CCS}_{\text{ml}}^{\text{sl}}$

$\overline{\mathcal{I}}(\mu x.P) = \overline{\mathcal{I}}(P[\mu x.P/x])$	$\overline{\mathcal{I}}(\bar{a}.P) = \{a\}$
$\overline{\mathcal{I}}(P \mid Q) = \overline{\mathcal{I}}(P) \cup \overline{\mathcal{I}}(Q)$	$\overline{\mathcal{I}}(P \oplus Q) = \overline{\mathcal{I}}(P) \cup \overline{\mathcal{I}}(Q)$
$\overline{\mathcal{I}}(P[f]) = \{f(a) \mid a \in \overline{\mathcal{I}}(P)\}$	$\overline{\mathcal{I}}(P \setminus L) = \overline{\mathcal{I}}(P) \setminus (L \cup \bar{L})$

The semantics of $\text{CCS}_{\text{ml}}^{\text{sl}}$ processes are again labeled transition systems whose transition relations are specified by operational rules. Since transitions labeled by output actions cannot get pre-empted they do not need to take account of locations. We first present two auxiliary sets which are useful for presenting

the operational rules: (i) initial output action sets $\overline{\mathcal{I}}(P)$ of a process P and (ii) initial input action sets $I_m^k(P)$ of P with respect to a priority value k and a location m , which are defined to be the least sets satisfying the equations presented in Tables 16 and 17, respectively. For technical convenience we remove the complement of output actions in the definition of $\overline{\mathcal{I}}(\cdot)$, and we use the following four abbreviations: (i) $I_M^{<k}(P) =_{\text{df}} \bigcup \{I_m^l(P) \mid m \in M, l < k\}$, (ii) $\mathbb{I}_M^{<k}(P) =_{\text{df}} I_M^{<k}(P) \setminus \{\tau\}$, (iii) $I(P) =_{\text{df}} \bigcup \{I_m^l(P) \mid m \in \mathcal{L}oc, l \in \mathcal{N}\}$, and (iv) $\mathbb{I}(P) =_{\text{df}} I(P) \setminus \{\tau\}$.

Table 17

Initial input action sets for $\text{CCS}_{\text{ml}}^{\text{sl}}$

$I_m^k(\mu x.P)$	$= I_m^k(P[\mu x.P/x])$	$I_{\bullet}^k(\gamma:l.P)$	$= \{\gamma \mid k = l\}$
$I_{m.l}^k(P + Q)$	$= I_m^k(P)$	$I_{m.L}^k(P \oplus Q)$	$= I_m^k(P)$
$I_{n.r}^k(P + Q)$	$= I_n^k(Q)$	$I_{n.R}^k(P \oplus Q)$	$= I_n^k(Q)$
$I_m^k(P[f])$	$= \{f(\gamma) \mid \gamma \in I_m^k(P)\}$	$I_{m.L}^k(P \mid Q)$	$= I_m^k(P) \cup \{\tau \mid I_m^k(P) \cap \overline{\mathcal{I}}(Q) \neq \emptyset\}$
$I_m^k(P \setminus L)$	$= I_m^k(P) \setminus (L \cup \overline{L})$	$I_{n.R}^k(P \mid Q)$	$= I_n^k(Q) \cup \{\tau \mid I_n^k(Q) \cap \overline{\mathcal{I}}(P) \neq \emptyset\}$

Table 18

Operational semantics for $\text{CCS}_{\text{ml}}^{\text{sl}}$ wrt. output transitions

$\overline{\text{Act}}$	$\frac{-}{\overline{a}.P \xrightarrow{\overline{a}} P}$	$\overline{\text{Sum1}}$	$\frac{P \xrightarrow{\overline{a}} P'}{P \oplus Q \xrightarrow{\overline{a}} P'}$	$\overline{\text{Com1}}$	$\frac{P \xrightarrow{\overline{a}} P'}{P \mid Q \xrightarrow{\overline{a}} P' \mid Q}$
$\overline{\text{Rel}}$	$\frac{P \xrightarrow{\overline{a}} P'}{P[f] \xrightarrow{f(\overline{a})} P'[f]}$	$\overline{\text{Sum2}}$	$\frac{Q \xrightarrow{\overline{a}} Q'}{P \oplus Q \xrightarrow{\overline{a}} Q'}$	$\overline{\text{Com2}}$	$\frac{Q \xrightarrow{\overline{a}} Q'}{P \mid Q \xrightarrow{\overline{a}} P \mid Q'}$
$\overline{\text{Res}}$	$\frac{P[\mu x.P/x] \xrightarrow{\overline{a}} P'}{\mu x.P \xrightarrow{\overline{a}} P'}$	$\overline{\text{Res}}$	$\frac{P \xrightarrow{\overline{a}} P'}{P \setminus L \xrightarrow{\overline{a}} P' \setminus L} \quad \overline{a} \notin L \cup \overline{L}$		

The operational rules for $\text{CCS}_{\text{ml}}^{\text{sl}}$ semantics are formally stated in Table 18 for *output transitions* and in Table 19 for *input transitions*. As expected, the rules for output transitions coincide with the ones for plain CCS [56], whereas the rules for input transitions take local pre-emption into account, thereby using location and priority value information in their side conditions. It is worth taking a closer look at the side conditions of Rules (Sum1) and (Sum2) which differ in principle from the corresponding ones in CCS^{sl} . They guarantee that an initial $\gamma:l$ -transition of a process P is also pre-empted whenever there exists a higher prioritized initial $\gamma:k$ -transition of P , i.e., if $k < l$. This additional kind of pre-emption reflects that output transitions can communicate with a complementary input transition regardless of its priority value; i.e., if more

Table 19
Operational semantics for $\text{CCS}_{\text{m}}^{\text{sl}}$ wrt. input transitions

Act	$\frac{-}{\gamma:k.P \xrightarrow{\bullet,\gamma:k} P}$	Sum1	$\frac{P \xrightarrow{m,\gamma:k} P'}{P + Q \xrightarrow{m,l,\gamma:k} P'} \quad \tau, \gamma \notin I^{<k}(Q)$
iSum1	$\frac{P \xrightarrow{m,\gamma:k} P'}{P \oplus Q \xrightarrow{m,L,\gamma:k} P'}$	Sum2	$\frac{Q \xrightarrow{n,\gamma:k} Q'}{P + Q \xrightarrow{n,r,\gamma:k} Q'} \quad \tau, \gamma \notin I^{<k}(P)$
iSum2	$\frac{Q \xrightarrow{n,\gamma:k} Q'}{P \oplus Q \xrightarrow{n,R,\gamma:k} Q'}$	Com1	$\frac{P \xrightarrow{m,\gamma:k} P'}{P Q \xrightarrow{m,L,\gamma:k} P' Q} \quad \mathbb{I}_{[m]}^{<k}(P) \cap \overline{\mathbb{I}}(Q) = \emptyset$
Rel	$\frac{P \xrightarrow{m,\gamma:k} P'}{P[f] \xrightarrow{m,f(\gamma):k} P'[f]}$	Com2	$\frac{Q \xrightarrow{n,\gamma:k} Q'}{P Q \xrightarrow{n,R,\gamma:k} P Q'} \quad \mathbb{I}_{[n]}^{<k}(Q) \cap \overline{\mathbb{I}}(P) = \emptyset$
Rec	$\frac{P[\mu x.P/x] \xrightarrow{m,\gamma:k} P'}{\mu x.P \xrightarrow{m,\gamma:k} P'}$	Com3	$\frac{P \xrightarrow{m,a:k} P' \quad Q \xrightarrow{\bar{a}} Q'}{P Q \xrightarrow{m,L,\tau:k} P' Q'} \quad \mathbb{I}_{[m]}^{<k}(P) \cap \overline{\mathbb{I}}(Q) = \emptyset$
		Com4	$\frac{P \xrightarrow{\bar{a}} P' \quad Q \xrightarrow{n,a:k} Q'}{P Q \xrightarrow{n,R,\tau:k} P' Q'} \quad \mathbb{I}_{[n]}^{<k}(Q) \cap \overline{\mathbb{I}}(P) = \emptyset$
		Res	$\frac{P \xrightarrow{m,\gamma:k} P'}{P \setminus L \xrightarrow{m,\gamma:k} P' \setminus L} \quad \gamma \notin L \cup \bar{L}$

than one communication partner offering the matching input transition is available from a comparable location, then the one having the highest priority is taken. This kind of pre-emption requires relabelings to be restricted to injective ones, as is pointed out in [24].

The behavioral relations defined for CCS^{sl} can be adapted to $\text{CCS}_{\text{m}}^{\text{sl}}$ in a straightforward fashion, as we demonstrate by the notion of distributed prioritized strong bisimulation.

Definition 4.17 *A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is a distributed prioritized strong bisimulation for $\text{CCS}_{\text{m}}^{\text{sl}}$ if for every $\langle P, Q \rangle \in \mathcal{R}$, $\bar{a} \in \bar{\Lambda}$, $\gamma \in \Lambda \cup \{\tau\}$, $k \in \mathcal{N}$, and $m \in \mathcal{L}oc$, the following conditions hold.*

- (1) $P \xrightarrow{\bar{a}} P'$ implies $\exists Q'. Q \xrightarrow{\bar{a}} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$, and
- (2) $P \xrightarrow{m,\gamma:k} P'$ implies $\exists Q', l, n. Q \xrightarrow{n,\gamma:l} Q', \mathbb{I}_{[n]}^{<l}(Q) \subseteq \mathbb{I}_{[m]}^{<k}(P)$, and $\langle P', Q' \rangle \in \mathcal{R}$.

We write $P \simeq_{ml} Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some distributed prioritized strong bisimulation \mathcal{R} for CCS_{ml}^{sl} .

Proposition 4.18 *The relation \simeq_{ml} is compositional with respect to all operators except summation.*

The proof can be done by applying standard techniques [56] and, therefore, is omitted here. The reason for the compositionality lack with respect to summation is illustrated by the following example: $a:0.\mathbf{0} \simeq_{ml} a:1.\mathbf{0}$ holds, but $a:0.\mathbf{0} + \tau:0.\mathbf{0} \not\simeq_{ml} a:1.\mathbf{0} + \tau:0.\mathbf{0}$ since the former process can engage in a transition labeled by action a whereas the latter cannot. This defect can easily be repaired (note the analogy with weak bisimulation [56]) but is not of importance for the purposes of this chapter.

4.8 Camilleri and Winskel's Approach

In this section we briefly review the framework of Camilleri and Winskel's approach to priority [24] which we refer to as CCS^{cw} (CCS with priority due to Camilleri and Winskel). In contrast to the approaches considered so far, this process algebra with priority does not assign priority values to actions. Instead, there exists a special summation operator $\dot{+}$ in CCS^{cw} , called *prioritized choice*, which favors its left over its right argument. The syntax of CCS^{cw} terms is given by the following BNF for P .

$$\begin{aligned} I & ::= \mathbf{0} \mid x \mid \gamma.I \mid I \dot{+} I \mid I + I \mid I \mid I \mid I[f] \mid I \setminus L \mid \mu x.I \\ P & ::= \mathbf{0} \mid x \mid \alpha.P \mid I \dot{+} I \mid P + P \mid P \mid P \mid P[f] \mid P \setminus L \mid \mu x.P \end{aligned}$$

Here, the action γ , the injective, finite relabeling f , and the restriction set L satisfy the constraints mentioned in the previous section. Again, closed and guarded terms determine the set \mathcal{P}^{cw} of CCS^{cw} processes. Further, we introduce initial output and input action sets as depicted in Tables 20 and 21, respectively, and write $\mathbb{I}^{cw}(P)$ for $\mathbb{I}^{cw}(P) \setminus \{\tau\}$.

Table 20

Initial output action sets for CCS^{cw}

$\overline{\mathbb{I}}^{cw}(\bar{a}.P) = \{a\}$	$\overline{\mathbb{I}}^{cw}(\mu x.P) = \overline{\mathbb{I}}^{cw}(P[\mu x.P/x])$
$\overline{\mathbb{I}}^{cw}(P \mid Q) = \overline{\mathbb{I}}^{cw}(P) \cup \overline{\mathbb{I}}^{cw}(Q)$	$\overline{\mathbb{I}}^{cw}(P + Q) = \overline{\mathbb{I}}^{cw}(P) \cup \overline{\mathbb{I}}^{cw}(Q)$
$\overline{\mathbb{I}}^{cw}(P[f]) = \{f(a) \mid a \in \overline{\mathbb{I}}^{cw}(P)\}$	$\overline{\mathbb{I}}^{cw}(P \setminus L) = \overline{\mathbb{I}}^{cw}(P) \setminus (L \cup \bar{L})$

The semantics of a CCS^{cw} process is given by a labeled transition system whose transition relation possesses transitions of the form $\vdash_M^{cw} P \xrightarrow{\alpha} P'$, where

Table 21
Initial input action sets for CCS^{cw}

$I^{\text{cw}}(\gamma.P) = \{\gamma\}$	$I^{\text{cw}}(\mu x.P) = I^{\text{cw}}(P[\mu x.P/x])$
$I^{\text{cw}}(P \dot{+} Q) = I^{\text{cw}}(P) \cup I^{\text{cw}}(Q)$	$I^{\text{cw}}(P + Q) = I^{\text{cw}}(P) \cup I^{\text{cw}}(Q)$
$I^{\text{cw}}(P[f]) = \{f(\gamma) \mid \gamma \in I^{\text{cw}}(P)\}$	$I^{\text{cw}}(P \setminus L) = I^{\text{cw}}(P) \setminus (L \cup \bar{L})$
$I^{\text{cw}}(P \mid Q) = I^{\text{cw}}(P) \cup I^{\text{cw}}(Q) \cup \{\tau \mid I^{\text{cw}}(P) \cap \overline{I^{\text{cw}}(Q)} \neq \emptyset\}$	

Table 22
Operational semantics for CCS^{cw}

Act $\frac{}{\vdash_{\emptyset}^{\text{cw}} \alpha.P \xrightarrow{\alpha} P}$	Res $\frac{\vdash_M^{\text{cw}} P \xrightarrow{\alpha} P'}{\vdash_{M \setminus (L \cup \bar{L})}^{\text{cw}} P \setminus L \xrightarrow{\alpha} P' \setminus L} \alpha \notin L \cup \bar{L}$
Sum1 $\frac{\vdash_M^{\text{cw}} P \xrightarrow{\alpha} P'}{\vdash_M^{\text{cw}} P \dot{+} Q \xrightarrow{\alpha} P'}$	Sum2 $\frac{\vdash_N^{\text{cw}} Q \xrightarrow{\alpha} Q'}{\vdash_{N \cup I^{\text{cw}}(P)}^{\text{cw}} P \dot{+} Q \xrightarrow{\alpha} Q'} \tau, \alpha \notin I^{\text{cw}}(P)$
iSum1 $\frac{\vdash_M^{\text{cw}} P \xrightarrow{\alpha} P'}{\vdash_M^{\text{cw}} P + Q \xrightarrow{\alpha} P'}$	Com1 $\frac{\vdash_M^{\text{cw}} P \xrightarrow{\alpha} P'}{\vdash_M^{\text{cw}} P \mid Q \xrightarrow{\alpha} P' \mid Q} M \cap \overline{I^{\text{cw}}(Q)} = \emptyset$
iSum2 $\frac{\vdash_N^{\text{cw}} Q \xrightarrow{\alpha} Q'}{\vdash_N^{\text{cw}} P + Q \xrightarrow{\alpha} Q'}$	Com2 $\frac{\vdash_N^{\text{cw}} Q \xrightarrow{\alpha} Q'}{\vdash_N^{\text{cw}} P \mid Q \xrightarrow{\alpha} P \mid Q'} N \cap \overline{I^{\text{cw}}(P)} = \emptyset$
Rel $\frac{\vdash_M^{\text{cw}} P \xrightarrow{\alpha} P'}{\vdash_{f(M)}^{\text{cw}} P[f] \xrightarrow{f(\alpha)} P'[f]}$	Com3 $\frac{\vdash_M^{\text{cw}} P \xrightarrow{\alpha} P' \quad \vdash_{\emptyset}^{\text{cw}} Q \xrightarrow{\bar{\alpha}} Q'}{\vdash_M^{\text{cw}} P \mid Q \xrightarrow{\tau} P' \mid Q'} M \cap \overline{I^{\text{cw}}(Q)} = \emptyset$
Rec $\frac{\vdash_M^{\text{cw}} P[\mu x.P/x] \xrightarrow{\alpha} P'}{\vdash_M^{\text{cw}} \mu x.P \xrightarrow{\alpha} P'}$	Com4 $\frac{\vdash_{\emptyset}^{\text{cw}} P \xrightarrow{\bar{\alpha}} P' \quad \vdash_N^{\text{cw}} Q \xrightarrow{\alpha} Q'}{\vdash_N^{\text{cw}} P \mid Q \xrightarrow{\tau} P' \mid Q'} N \cap \overline{I^{\text{cw}}(P)} = \emptyset$

$M \subseteq \Lambda$. Intuitively, process P can engage in an α -transition to P' whenever the environment does not offer communications on ports in M . Despite notational differences, this is the same underlying principle as for the transition relations defined in the previous sections, which are parameterized by initial action sets. Note that $\alpha \in \bar{\Lambda}$ implies $M = \emptyset$. The CCS^{cw} transition relation is formally defined in Table 22, where $f(M)$ stands for $\{f(m) \mid m \in M\}$. Recall that the initial actions of P in $P \dot{+} Q$ are given preference over the initial actions of Q . As expected, a prioritized τ , i.e., an internal action in which the left argument

of $\dot{+}$ can initially engage, has pre-emptive power over unprioritized actions, i.e., actions in which the right argument of $\dot{+}$ can initially engage. Thus, the prioritized choice operator $\dot{+}$ in CCS^{cw} corresponds to the summation operator $+$ in $\text{CCS}_{\text{ml}}^{\text{sl}}$. In [24] the operator $+$ stands for nondeterministic choice where priorities arising from the left and the right argument are incomparable. This operator is matched by the distributed summation operator \oplus in $\text{CCS}_{\text{ml}}^{\text{sl}}$. We further investigate the correspondence of these operators in the next section.

Camilleri and Winskel also developed a bisimulation-based semantic theory for CCS^{cw} . Their notion of strong bisimulation for CCS^{cw} , as defined below, is shown to be a congruence [24].

Definition 4.19 *A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is a distributed prioritized strong bisimulation for CCS^{cw} if for every $\langle P, Q \rangle \in \mathcal{R}$, $\alpha \in \mathcal{A}$, and $M \subseteq \Lambda$ the following condition holds:*

$$\vdash_M^{\text{cw}} P \xrightarrow{\alpha} P' \text{ implies } \exists Q', N. \vdash_N^{\text{cw}} Q \xrightarrow{\alpha} Q', N \subseteq M, \text{ and } \langle P', Q' \rangle \in \mathcal{R} .$$

We write $P \simeq_{\text{cw}} Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some distributed prioritized strong bisimulation \mathcal{R} for CCS^{cw} .

4.9 Relating Both Priority Approaches

We show that the algebras $\text{CCS}_{\text{ml}}^{\text{sl}}$ and CCS^{cw} are closely related by providing an embedding of CCS^{cw} in $\text{CCS}_{\text{ml}}^{\text{sl}}$. For this purposes we define $\mathcal{N} =_{\text{df}} \{0, 1\}^*$ and the strict order $<$ on priority values to be the lexicographical order on \mathcal{N} , where 1 is less than 0.

Table 23

Translation function

$\xi^k(\mathbf{0})$	$=_{\text{df}} \mathbf{0}$	$\xi^k(P + Q)$	$=_{\text{df}} \xi^k(P) \oplus \xi^k(Q)$	$\xi^k(P \setminus L)$	$=_{\text{df}} \xi^k(P) \setminus L$
$\xi^k(x)$	$=_{\text{df}} x$	$\xi^k(P \dot{+} Q)$	$=_{\text{df}} \xi^{k0}(P) + \xi^{k1}(Q)$	$\xi^k(P[f])$	$=_{\text{df}} \xi^k(P)[f]$
$\xi^k(\gamma.P)$	$=_{\text{df}} \gamma : k. \xi^e(P)$	$\xi^k(P Q)$	$=_{\text{df}} \xi^k(P) \xi^k(Q)$	$\xi^k(\mu x.P)$	$=_{\text{df}} \mu x. \xi^k(P)$
$\xi^k(\bar{a}.P)$	$=_{\text{df}} \bar{a}. \xi^e(P)$				

We now introduce the *translation function* $\xi(\cdot) : \mathcal{P}^{\text{cw}} \longrightarrow \mathcal{P}_{\text{ml}}^{\text{sl}}$ by defining $\xi(P) =_{\text{df}} \xi^e(P)$, which maps CCS^{cw} terms to $\text{CCS}_{\text{ml}}^{\text{sl}}$ terms. The functions $\xi^k(P)$, for $k \in \mathcal{N}$, are inductively defined over the structure of CCS^{cw} processes as shown in Table 23. We note that the translation function is not *surjective*, e.g., consider process $(a:0.\mathbf{0} + b:2.\mathbf{0}) + c:1.\mathbf{0}$ on which no CCS^{cw} process is mapped. This example also shows that the notion of compositionality in CCS^{cw}

is more restrictive than the one in $\text{CCS}_{\text{ml}}^{\text{sl}}$, since a comparable summation can only be extended by summands which have a higher or a lower priority than the already considered summands. The following theorem, which is proved in [53], makes the semantic relationship between a CCS^{cw} process P and its embedding $\xi(P)$ precise.

Theorem 4.20 *Let $P, Q \in \mathcal{P}^{\text{cw}}$. Then $P \simeq_{\text{cw}} Q$ if and only if $\xi(P) \simeq_{\text{ml}} \xi(Q)$.*

As a consequence, distributed prioritized strong bisimulation for $\text{CCS}_{\text{ml}}^{\text{sl}}$ is also compositional with respect to summation in the sub-calculus of $\text{CCS}_{\text{ml}}^{\text{sl}}$ induced by CCS^{cw} .

4.10 Concluding Remarks and Related Work

A local concept of pre-emption is also considered by Hansson and Orava in [37], where CSP [42] is extended with priority by assigning natural numbers to actions. As for CCS^{sl} , they equipped their operational semantics with a notion of location and introduced a sensitivity to locations when defining pre-emption. Indeed, their work served as an inspiration for CCS^{sl} . However, the authors only conjecture that their version of strong bisimulation is a congruence, and they provide neither an axiomatization for their behavioral relation nor a theory for observational congruence. One may also criticize their semantics as not truly reflecting distributed computation. In particular, despite having a local pre-emptive semantics they compute a global priority for synchronizations.

After stressing the strong similarity of CCS^{sl} to the process algebra CCS^{cw} in the previous section, we focus on the algebraic results established in these frameworks. In [24,46] the transition relation is directly annotated with pre-emption potentials. By plugging this transition relation into the definition of standard strong bisimulation one immediately obtains a congruence. In contrast, [29] starts off by defining *naive* distributed prioritized strong bisimulation using the naive transition relation and considering pre-emption potential subsequently (by introducing the distributed prioritized initial action set condition). Then it is shown that the resulting congruence is the largest congruence in the naive equivalence. Similarly, Jensen [46] defines a naive distributed prioritized weak bisimulation based on the above-mentioned annotated transition relation. His naive weak transition relation corresponds to the distributed prioritized weak transition relation in CCS^{sl} if the parameter M is dropped. Because of the difference in the naive transition relations the abstraction result presented here is somewhat stronger than Jensen's, although the observational congruences appear to coincide.

One may wonder about the relationship between CCS^{sl} and CCS^{sg} , i.e., the

static priority & global pre-emption language in Section 3. If in CCS^{sl} the distributed summation operator is left out and pre-emption is globalized by defining $[m] =_{\text{df}} \mathcal{L}oc$ for all $m \in \mathcal{L}oc$, the operational semantics and the behavioral relations reduce to the corresponding notions presented in Section 3.

Like Camilleri and Winskel, Barrett [8] devises a semantics of *occam's* priority mechanism that is additionally concerned with fairness aspects. His framework is based on a structural operational semantics augmented with *ready-guard* sets which model possible inputs from the environment. Intuitively, these sets characterize the nature of the contexts in which a transition is enabled. Thus, they correspond to the action sets with which the CCS^{sl} and the CCS^{cw} transition relations are parameterized. Barrett is not concerned with investigating behavioral relations, but focuses on implementing *occam's* **PRIALT** and **PRIPAR** constructs on the transputer platform instead.

Other researchers have also extended Hoare's *Communicating Sequential Processes* (CSP) [42] by a concept of static priority. Inspired by the notion of priority in ADA [49], Fidge [33] has introduced new versions of the operators for external choice, parallel composition by interleaving, and parallel composition by intersection. These favor their left-hand operands similar to the operators investigated by Jensen [46]. The semantic theory in [33] is based on *failure semantics* which is made sensitive to local pre-emption. For this purpose, traces are augmented with a *preference function* which identifies the priority relation on the initial action sets of a given process. A related approach has been presented by Lowe [52]. It differs from [33] in that the underlying algebra is a timed version of CSP [31]. Additionally, Lowe aims at obtaining a fully deterministic language by employing a similar notion of priority as the one proposed by Fidge.

Finally, we remark on the notion of strong and weak bisimulations for CCS^{sl} . Since the semantic theory reflects local pre-emption, locations are implicitly occurring in our semantic equivalences. In contrast to the work on *location equivalences* [18,25,61], we do not consider locations explicitly in our relations. Our objective is not to observe locations but to capture local pre-emption.

5 Dynamic Priority and Global Pre-emption

This section develops a theory in which priorities are dynamic and pre-emption is global. The motivation for this theory originated in a desire to devise a compact model of real-time computation, and we devote significant space to establishing a tight connection between the seemingly different notions of priority and real-time [10]. For this purpose we equip our language with a dynamic priority semantics based on global pre-emption and refer to it as CCS^{dg} (CCS

with dynamic priority and global pre-emption). The connection with real-time arises when we interpret delays as priorities: the longer the delay preceding an action, the lower is its priority. This approach contrasts significantly with more traditional accounts of real-time, where the only notion of pre-emption arises in the context of the *maximal progress assumption* [78] which states that time may only pass if the system under consideration cannot engage in any further internal computation. The main result of this section is the formalization of a one-to-one correspondence between the strong-bisimulation equivalences induced by dynamic priority semantics and real-time semantics.

Unlike the process algebras with priority considered so far, actions in CCS^{dg} do have priority values that may change as systems evolve. Accordingly, we slightly alter our point of view regarding actions and priorities by separating action names from their priority values; that is, an action's priority is no longer implicit in its port name. In this vein, we take the set of actions \mathcal{A} to be $\{\alpha, \beta, \dots\}$. We also allow priority values to come from the full set \mathbb{N} of natural numbers rather than a finite set. Our syntax of processes will then require that each action is equipped with a priority value taken from \mathbb{N} .

The structure of this section is as follows. Section 5.1 briefly presents a real-time semantics for our language, whereas the dynamic priority semantics is introduced in Section 5.2. The one-to-one correspondence between dynamic priority semantics and real-time semantics is established in Section 5.3. Finally, Section 5.4 contains our concluding remarks and discusses related work.

5.1 Real-Time Semantics

We first introduce a real-time semantics for our language, referred to as CCS^{rt} semantics, which explicitly represents timing behavior. The semantics of a process is defined by a labeled transition system which contains explicit *clock transitions* – each representing a delay of one time unit – as well as *action transitions*. With respect to clock transitions, the operational semantics is set up such that processes willing to communicate with some process running in parallel are able to wait until the communication partner is ready. However, as soon as it is available, the communication has to take place, i.e., further idling is prohibited. This assumption is usually referred to as maximal progress assumption [78] or *synchrony hypothesis* [13].

Formally, the labeled transition system corresponding to a process P is a four-tuple $\langle \mathcal{P}, \mathcal{A} \cup \{1\}, \mapsto, P \rangle$, where the alphabet $\mathcal{A} \cup \{1\}$ satisfies $1 \notin \mathcal{A}$. The transition relation $\mapsto^1 \subseteq \mathcal{P} \times \mathcal{P}$ for clock transitions is defined in Table 24. With respect to action transitions, it coincides with the one for traditional CCS where Rule (Act) is replaced by axiom $\alpha:0.P \mapsto^\alpha P$. For the sake of

simplicity, we use γ as representative of $\mathcal{A} \cup \{1\}$, and write $P \xrightarrow{\gamma} P'$ instead of $\langle P, \gamma, P' \rangle \in \xrightarrow{\quad}$. If $\gamma \in \mathcal{A}$ we speak of an *action transition*, otherwise of a *clock transition*. Sometimes it is convenient to write $P \xrightarrow{\gamma}$ for $\exists P' \in \mathcal{P}. P \xrightarrow{\gamma} P'$. In order to ensure maximal progress, our operational semantics is set up in a way such that $P \not\xrightarrow{\tau}$ whenever $P \xrightarrow{\tau}$, i.e., clock transitions are pre-empted as long as P can engage in internal computation.

Table 24

Operational semantics for CCS^t (clock transitions)

$\text{tNil} \quad \frac{-}{\mathbf{0} \xrightarrow{1} \mathbf{0}}$	$\text{tRec} \quad \frac{P[\mu x.P/x] \xrightarrow{1} P'}{\mu x.P \xrightarrow{1} P'}$
$\text{tAct1} \quad \frac{-}{\alpha : k.P \xrightarrow{1} \alpha : (k-1).P} \quad k > 0$	$\text{tAct2} \quad \frac{-}{a : 0.P \xrightarrow{1} a : 0.P}$
$\text{tSum} \quad \frac{P \xrightarrow{1} P' \quad Q \xrightarrow{1} Q'}{P + Q \xrightarrow{1} P' + Q'}$	$\text{tCom} \quad \frac{P \xrightarrow{1} P' \quad Q \xrightarrow{1} Q'}{P Q \xrightarrow{1} P' Q'} \quad P Q \not\xrightarrow{\tau}$
$\text{tRel} \quad \frac{P \xrightarrow{1} P'}{P[f] \xrightarrow{1} P'[f]}$	$\text{tRes} \quad \frac{P \xrightarrow{1} P'}{P \setminus L \xrightarrow{1} P' \setminus L}$

Intuitively, the process $\alpha : k.P$, where $k > 0$, may engage in a clock transition and then behave like $\alpha : (k-1).P$. The process $\alpha : 0.P$ performs an α transition to become process P . Moreover, if $\alpha \neq \tau$, it may also idle by executing a clock transition to itself. Time has to proceed equally on both sides of summation, i.e., $P + Q$ can engage in a clock transition and, thus, delay the nondeterministic choice if and only if both P and Q can engage in a clock transition. Hence, time is a deterministic concept. Similar to summation, P and Q have to synchronize on clock transitions according to Rule (tCom). Its side condition implements maximal progress by ensuring that there is no pending communication between P and Q . Although this condition is negative, our semantics is still well-defined [76]. A semantic theory based on the notion of *bisimulation* [56] has been developed for CCS^t (cf. [59]). For the purposes of this section we restrict ourselves to (*strong*) temporal bisimulation, a congruence which is defined as follows.

Definition 5.1 (Temporal Bisimulation)

A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is called temporal bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$ and $\gamma \in \mathcal{A} \cup \{1\}$ the following holds: $P \xrightarrow{\gamma} P'$ implies $\exists Q'. Q \xrightarrow{\gamma} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$. We write $P \sim_{rt} Q$ if $\langle P, Q \rangle \in \mathcal{R}$ for some temporal bisimulation \mathcal{R} .

The reader might observe that CCS^{rt} semantics unfolds every delay value into a sequence of elementary time units. For example, the process $a:k.\mathbf{0}$ has $k+2$ states, namely $\mathbf{0}$ and $a:l.\mathbf{0}$, for $0 \leq l \leq k$ (see also Figure 6 in Section 5.3). Representing $a:k.\mathbf{0}$ by a single transition, which is labeled by $a:k$ and leads to state $\mathbf{0}$, would definitely be more efficient. This idea of compacting the state space of real-time systems can be implemented by viewing k as *priority value* assigned to action a . In other words, one may consider the delay value k as the *time-stamp* of action a [45].

5.2 Dynamic Priority Semantics

In order to make the above intuition precise, we formally introduce CCS^{dg} , i.e., a dynamic priority semantics for our language. The notion of pre-emption incorporated in CCS^{dg} is similar to CCS^{sg} ; it naturally mimics the maximal progress assumption employed in CCS^{rt} semantics. Formally, the CCS^{dg} semantics of a process P is given by a labeled transition system $\langle \mathcal{P}, \mathcal{A} \times \mathbb{N}, \longrightarrow, P \rangle$. The presentation of the operational rules for the transition relation \longrightarrow requires two auxiliary definitions.

Table 25

Potential initial action sets for CCS^{dg}

$I^k(\alpha:l.P) = \{\alpha \mid l \leq k\}$	$I^k(P \mid Q) = I^k(P) \cup I^k(Q) \cup \{\tau \mid I^k(P) \cap \overline{I^k(Q)} \neq \emptyset\}$
$I^k(P + Q) = I^k(P) \cup I^k(Q)$	$I^k(P[f]) = \{f(\alpha) \mid \alpha \in I^k(P)\}$
$I^k(\mu x.P) = I^k(P[\mu x.P/x])$	$I^k(P \setminus L) = I^k(P) \setminus (L \cup \overline{L})$

First, we introduce *potential initial action sets* as defined in Table 25, taking account of the actions in which a given process can potentially engage. Note that these sets are only supersets of the initial actions of processes because they do not take *pre-emption* into account. However, this is sufficient for our purposes regarding pre-emption, since $\tau \notin I^{<k}(P)$ if and only if $\nexists l < k. P \xrightarrow{\tau:l}$, where $I^{<k}(P) =_{\text{df}} \bigcup_{l < k} I^l(P)$, for $k > 0$, and $I^{<0}(P) =_{\text{df}} \emptyset$.

Table 26

Priority adjustment function

$[\mathbf{0}]^k =_{\text{df}} \mathbf{0},$	$[x]^k =_{\text{df}} x$	$[\mu x.P]^k =_{\text{df}} [P[\mu x.P/x]]^k$
$[\alpha:l.P]^k =_{\text{df}} \alpha:(l-k).P$	if $l > k$	$[\alpha:l.P]^k =_{\text{df}} \alpha:\mathbf{0}.P$
$[P + Q]^k =_{\text{df}} [P]^k + [Q]^k$		$[P \mid Q]^k =_{\text{df}} [P]^k \mid [Q]^k$
$[P[f]]^k =_{\text{df}} [P]^k[f]$		$[P \setminus L]^k =_{\text{df}} [P]^k \setminus L$

As second auxiliary definition for presenting the transition relation, we introduce a *priority adjustment function* as shown in Table 26. Intuitively, our semantics is set up in a way such that if one parallel component of a process engages in a transition with priority k , then the priority values of all initial actions at every other parallel component have to be decreased by k , i.e., those actions become equally “more urgent.” Thus, the semantics of parallel composition employs a kind of *fairness assumption*, and priorities have a *dynamic* character. More precisely, the priority adjustment function applied to a process P and a natural number k , denoted as $[P]^k$, returns a process term which is “identical” to P except that the priority values of the initial, top-level actions are decreased by k . Note that a priority value cannot become less than 0 and that “identical” does not mean syntactic equality but syntactic equality *up to unfolding* of recursion.

Table 27
Operational semantics for CCS^{dg}

Act1 $\frac{}{a:k.P \xrightarrow{a:l} P} l \geq k$	Act2 $\frac{}{\tau:k.P \xrightarrow{\tau:k} P}$
Sum1 $\frac{P \xrightarrow{\alpha:k} P'}{P + Q \xrightarrow{\alpha:k} P'} \tau \notin I^{<k}(Q)$	Sum2 $\frac{Q \xrightarrow{\alpha:k} Q'}{P + Q \xrightarrow{\alpha:k} Q'} \tau \notin I^{<k}(P)$
Com1 $\frac{P \xrightarrow{\alpha:k} P'}{P Q \xrightarrow{\alpha:k} P' [Q]^k} \tau \notin I^{<k}(P Q)$	Rel $\frac{P \xrightarrow{\alpha:k} P'}{P[f] \xrightarrow{f(\alpha):k} P'[f]}$
Com2 $\frac{Q \xrightarrow{\alpha:k} Q'}{P Q \xrightarrow{\alpha:k} [P]^k Q'} \tau \notin I^{<k}(P Q)$	Res $\frac{P \xrightarrow{\alpha:k} P'}{P \setminus L \xrightarrow{\alpha:k} P' \setminus L} \alpha \notin L \cup \bar{L}$
Com3 $\frac{P \xrightarrow{a:k} P' \quad Q \xrightarrow{\bar{a}:k} Q'}{P Q \xrightarrow{\tau:k} P' Q'} \tau \notin I^{<k}(P Q)$	Rec $\frac{P[\mu x.P/x] \xrightarrow{\alpha:k} P'}{\mu x.P \xrightarrow{\alpha:k} P'}$

The operational rules in Table 27 capture the following intuition. Process $a:k.P$ may engage in action a , with priority value $l \geq k$, yielding process P . The side condition $l \geq k$ reflects that k does not specify an exact priority but the *maximum* priority of the initial transition of $a:k.P$. It may also be interpreted as *lower-bound* “timing constraint.” Due to the notion of pre-emption incorporated in CCS^{dg} , $\tau:k.P$ may not perform the initial τ -transition with a lower priority than k . Process $P + Q$ may behave like P (Q) if Q (P) does not pre-empt the considered transition by being able to engage in a higher prioritized internal transition. Thus, the notion of global pre-emption reflects

implicit *upper-bound* “timing constraints.” Process $P|Q$ denotes the *parallel composition* of P and Q according to an interleaving semantics with synchronized communication on complementary actions of P and Q having the same priority value k , which results in the internal action τ attached with priority value k (cf. Rule (Com3)). The interleaving Rules (Com1) and (Com2) incorporate the dynamic behavior of priority values as explained in the previous paragraph. The side conditions of Rules (Com i) implement global pre-emption. The semantics for *relabeling*, *restriction*, and *recursion* is straightforward. As for CCS^{rt} , we may adapt a notion of strong bisimulation, referred to as *prioritized bisimulation*.

Definition 5.2 (Prioritized Bisimulation)

A symmetric relation $\mathcal{R} \subseteq \mathcal{P} \times \mathcal{P}$ is called prioritized bisimulation if for every $\langle P, Q \rangle \in \mathcal{R}$, $\alpha \in \mathcal{A}$, and $k \in \mathbb{N}$ the following holds: $P \xrightarrow{\alpha:k} P'$ implies $\exists Q'. Q \xrightarrow{\alpha:k} Q'$ and $\langle P', Q' \rangle \in \mathcal{R}$. We write $P \sim_{\text{ag}} Q$ if there exists a prioritized bisimulation \mathcal{R} such that $\langle P, Q \rangle \in \mathcal{R}$.

5.3 Relating Dynamic Priority and Real-Time Semantics

In this section we show that CCS^{dg} and CCS^{rt} semantics are closely related. The underlying intuition is best illustrated by a simple example dealing with the prefixing operator. Figure 6 depicts the dynamic priority semantics and real-time semantics of process $a:k.\mathbf{0}$. Both transition systems intuitively reflect that process $a:k.\mathbf{0}$ must at least delay k time units before it may engage in the a -transition. According to CCS^{rt} semantics, this process consecutively engages in k clock transitions passing the states $a:(k-l).\mathbf{0}$, for $0 \leq l \leq k$, before it may either continue idling in state $a:0.\mathbf{0}$ or perform the a -transition to inaction process $\mathbf{0}$. Thus, time is explicitly part of states and made visible by clock transitions, each representing a step consuming one time unit. In contrast, the dynamic priority semantics encodes the delay of at least k time units in the transitions rather than in the states. Hence, it possesses only the two states $a:k.\mathbf{0}$ and $\mathbf{0}$ connected via transitions labeled by $a:l$ for $l \geq k$. Although at first sight it seems that the price for saving intermediate states is to be forced to deal with infinite-branching, an upper bound for l can be given. In our example this upper bound is k itself, since a delay by more than k time units only results in idling and does not enable new or disable existing system behavior. Therefore, the dynamic priority transition system of $a:k.\mathbf{0}$ just consists of the two states $a:k.\mathbf{0}$ and $\mathbf{0}$ and a *symbolic* transition labeled by $a:k$, whereas the real-time transition system has $k+2$ states and $k+2$ transitions. The following proposition formally states that CCS^{dg} semantics can indeed be understood as an efficient encoding of CCS^{rt} semantics. Here, $\xrightarrow{1}^k$ denotes k consecutive clock transitions.




Fig. 6. Relating CCS^{dg} semantics and CCS^{rt} semantics

Proposition 5.3 *Let $P, P' \in \mathcal{P}$, $\alpha \in \mathcal{A}$, and $k \in \mathbb{N}$. Then $P \xrightarrow{\alpha:k} P'$ if and only if $\exists P''. P \xrightarrow{1}^k P'' \xrightarrow{\alpha} P'$.*

Proposition 5.3 is the key to prove the main result of this section.

Theorem 5.4 *Let $P, Q \in \mathcal{P}$. Then $P \sim_{\text{dg}} Q$ if and only if $P \sim_{\text{rt}} Q$.*

Consequently, prioritized and temporal bisimulation possess the same properties; especially, prioritized bisimulation is a congruence for CCS^{dg} . Again, proof details can be found in [53].

5.4 Concluding Remarks and Related Work

As shown above, real-time semantics can be encoded by dynamic priority semantics. The utility of this encoding stems from the fact that the state space of CCS^{dg} models is much smaller and the size of the transition relation is at least not worse, but in practice often better, than the one of corresponding CCS^{rt} models. This has been demonstrated by formally modeling and verifying several aspects of the widely-used *SCSI-2 bus-protocol*, for which the state space of the dynamic priority model is almost an order of magnitude smaller than the one resulting from traditional real-time semantics [16].

Regarding related work, a similar approach to the one presented above has been investigated by Jeffrey [45]. He established a formal relationship between a quantitative real-time process algebra and a process algebra with static priority which is very similar to CCS^{sg} presented in Section 3. Technically, Jeffrey also translates real-time to priority based on the idea of time-stamping. In contrast to CCS^{rt} semantics, however, a process modeled in Jeffrey’s framework may either immediately engage in an action transition or idle forever. This semantics does not allow a process to wait until a communication partner becomes available, but instead forces a “livelock” in such situations. It is only

because of this design decision that Jeffrey does not need to choose a dynamic priority framework.

In [22] a variant of CCSR [23], called CCSR92, has been introduced, which allows for modeling not only static priority but also dynamic priority. The main focus of CCSR involves the specification and verification of real-time concurrent systems, including scheduling behavior. Thus, a notion of dynamic priority, such as occurs in *priority-inheritance* and *earliest-deadline-first* scheduling algorithms, is important. In [22] dynamic priorities are given as a function of the *history* of the system under consideration. Accordingly, the operational semantics of CCSR92 is re-defined to include the historical context. The authors show that dynamic priorities do, in general, not lead to a compositional semantics and give a sufficient condition that ensures compositionality.

6 Priority in Other Process-Algebraic Frameworks

This section completes the discussion of related work by focusing on approaches to priority which (i) do either not fit in our classification scheme presented in Section 1, such as approaches for ACP [5], SCCS [72], and *stochastic* [11,41] or *probabilistic* [47,73,75] process algebras, or (ii) are concerned with process-algebraic descriptions of non-process-algebraic languages, such as *Es-terel* [12,13] and *Statecharts* [38].

Baeten, Bergstra, and Klop were the first researchers who investigated priority in process algebras [5] by developing a notion of priority in the *Algebra of Communicating Processes* (ACP) [9], a process algebra which is equipped with an axiomatic semantics. Their work is inspired by the insight that it is essential to incorporate an interrupt mechanism in process-algebraic frameworks in order to enhance their expressive power as specification and verification formalisms for concurrent systems. Therefore, a piece of syntax together with semantics defining equations is introduced in [5]. Based on a given partial order $<$ on actions a unary operator θ is defined. Intuitively, $\theta(P)$ is the context of P in which action a has precedence over action b , whenever $b < a$, i.e., non-deterministic choices between actions a and b are resolved within $\theta(P)$. Technically, the axiomatic semantics of the new language, notated as a *term rewrite system*, is shown to possess nice algebraic properties such as *confluence* and *termination*. The utility of the theory is demonstrated by simple examples dealing with interrupts, timeouts, and other aspects of system behavior. The approach in [5] differs from most other work presented in this chapter in that the partial order expressing priorities is fixed with respect to the system under consideration, i.e., the same priority relation holds at all states of the system. For example, if $a < b$ at some state of the system, then $a > b$ cannot be valid at another state, i.e., priorities in [5] are not *globally dynamic* in the

sense of [72]. It should also be mentioned that the version of ACP used in [5] does not include a designated internal action, cf. action τ in CCS; a fact which simplifies the development of algebraic theories.

Stochastic process algebras [11,41], which enhance the expressiveness of classical process algebras by integrating *performance* descriptions of concurrent systems, also define notions of priority. One example of a popular stochastic process algebra is the *Extended Markovian Process Algebra* (EMPA) [11] whose semantics is given in terms of strong bisimulation and whose static priority approach is adapted from CCS^{sg}.

Smolka and Steffen [72] have introduced static priority to the *Synchronous Calculus of Communicating Systems* (SCCS) [56]. They extended a probabilistic version of this language, known as PCCS [75], whose semantics is given in terms of *probabilistic bisimulation*. Their work shows that the concept of priority is not only related to real-time, as investigated in Section 5, but also to probability. The main idea in [72] is to allow probability guards of value 0 to be associated with alternatives of a probabilistic summation expression. Such alternatives can be chosen only if the non-zero alternatives are precluded by contextual constraints. Thus, priority may be viewed as an extreme case of probability. Most remarkably, the semantics developed in [72] does not employ a notion of pre-emption as one would expect from any priority setting. A conjecture – which if true would justify this situation – is that the very powerful hiding operator in SCCS may destroy the congruence property of bisimulation in the presence of pre-emption.

Tofts has investigated another extension of SCCS, the *Weighted Synchronous Calculus of Communicating Systems* (WSCCS) [73]. Its semantics relies upon a notion of *relative frequency* which is suitable for specifying and reasoning about aspects of priority, probability, and time in concurrent systems. In this approach, priority is encoded by means of higher ordinals; a transition has priority over another if their weights are separated at least by a factor of ω . An operator similar to the θ -operator in [5] is defined, which extracts the highest priority transitions enabled at a process state by referring to a global notion of pre-emption. In contrast to [5], Toft's operator allows for different priority structures at different states. This concept of priority yields a simpler operational semantics than the one in [72]. For WSCCS, a congruence adapted from strong bisimulation together with an equational characterization, which is sound and complete for finite processes, has been developed.

The concept of pre-emption has also been studied in other synchronous languages, most notably by Berry [12]. The technical framework investigated by him is *Esterel's zero-delay process calculus*, a theoretical version of the Esterel synchronous programming language [13]. The calculus' semantics interprets processes as deterministic mappings from input sequences to output

sequences which obey maximal progress [78]. Berry emphasizes the importance of pre-emption in control-dominated reactive and real-time programming. He suggests pre-emption operators to be considered as first-class operators which are fully orthogonal with respect to all other primitives, such as concurrency and communication. This is in contrast to the approach chosen in this chapter, in which pre-emption is implicitly encoded as side conditions of operational rules involving nondeterminism. Several examples of useful pre-emption operators are presented and axiomatized in [12], all of which are based on the ideas of *abortion* and *suspension*.

The specification language *Statecharts* [38], for which process-algebraic descriptions of Statecharts' semantics have been developed [51,54,74], extends communicating finite automata by concepts of *hierarchy* and *priority*. In Statecharts static priorities can be expressed via the absence of actions, which are called *events*, by permitting negated actions as guards, which are referred to as *triggers*. As an example, consider the following Statecharts-like term: $a : b.P + \neg b : c.Q$. This term consists of a nondeterministic choice between a b -transition with guard a to process P , and a c -transition with guard $\neg b$ to Q . Intuitively, the statechart may only engage in the latter transition if it cannot execute the former one. The reason is that the former transition produces event b which falsifies the guard of the c -transition. Thus, the b -transition is given precedence over the c -transition. In the following we argue that approaches to priority via negated events (cf. [36]) do not go well with the concept of *hiding*, which is employed in many process algebras and also in a very popular alternative to Statecharts, called ARGOS [55]. Hiding enables one to relabel a visible action into a distinguished invisible action (cf. the internal action τ in CCS). The difficulty with hiding arises when several events are hidden, i.e., all of them are relabeled to the same event and, thus, have the same implicit priority value attached to them. Hence, hiding may destroy priority schemes. In contrast, in the priority approaches considered in this chapter priorities are assigned to transitions, thereby allowing for a more fine-granular priority mechanism and avoiding the above-mentioned problem.

7 Conclusions and Directions for Future Work

This chapter has investigated various aspects of priority in process algebras. The utility of introducing priority to traditional process algebras is to enhance their expressiveness and, thereby, making them more attractive to system designers. Technically, priorities allow one to resolve potential nondeterminism, as do certain constructs in existing specification and programming languages.

Conclusions. We have illustrated the most important aspects of priority in a prototypic language which extends Milner's CCS. This language has been

equipped with several semantics according to whether priorities are static or dynamic and whether the adopted notion of pre-emption is global or local.

In practice it is easy to determine when to use static priority semantics and when to use dynamic priority semantics. For modeling interrupts and prioritized choice constructs a static notion of priority is adequate, whereas for modeling real-time or scheduling behavior dynamic priorities should be considered. However, static priority approaches may also allow for the description of a few, very simple scheduling algorithms, as shown in [46] in the presence of a prioritized parallel composition operator. In addition to the ability to express more general scheduling algorithms, the dynamic priority approach yields a more efficient verification of real-time systems, since the sizes of system models with respect to dynamic priority semantics are often several orders of magnitude smaller than the ones regarding real-time semantics [16]. If one needs to deal with both interrupt and real-time aspects at the same time, static and dynamic priority approaches must be combined. In this situation each action should be assigned two priority values, the first interpreted as a global priority value for scheduling purposes, and the second interpreted as a local priority value for modeling interrupts, where the first priority value has more weight than the second one.

Suitable guidelines supporting the decision in favor of a global or a local notion of pre-emption are the following. A semantics obeying global pre-emption is required when modeling interrupts and prioritized-choice constructs in concurrent, centralized systems or when specifying real-time and scheduling aspects. Global pre-emption also allows for making executions of action sequences atomic. This can be necessary for modeling systems accurately and, as a desired side effect, may keep system models small, thereby enhancing the efficiency of verification algorithms [30]. However, when dealing with interrupts or prioritized-choice constructs within distributed systems, the concept of global pre-emption is inadequate. Here, the use of local pre-emption does not only lead to an intuitive but also to an implementable semantics, since it does not require any knowledge about computations which are internal to other, potentially unknown sites (cf. [28]).

Technically, the three different calculi presented in Sections 3–5 are equipped with a bisimulation-based semantics. The re-development of the semantic theory of CCS for the static priority calculi included: (i) characterizations of the largest congruences contained in the naive adaptations of the standard strong and weak bisimulations, (ii) encodings of the new behavioral relations as standard strong bisimulations on enriched transition relations, and (iii) axiomatic characterizations of the prioritized strong bisimulations for finite processes. For the dynamic priority calculus, strong bisimulation served as a semantic tool for establishing a one-to-one correspondence between dynamic priority and real-time semantics. Finally, we want to point out that the semantic theo-

ries presented here show that extensions of process algebras by priority do not need to sacrifice the simplicity and the elegance that have made traditional process-algebraic approaches successful.

This chapter also surveyed related approaches to priority which are concerned with different process-algebraic calculi. We have classified them according to whether priorities are considered to be static or dynamic and whether their concept of pre-emption is global or local. The concept of priority has also been investigated in other concurrent frameworks, most notably in *Petri Nets* [14,71]. In this setting priorities are either expressed explicitly by priority relations over transitions [15] or implicitly via *inhibitor arcs* [44]. Finally, it should be mentioned that priorities can *implicitly* arise when studying causality for *mobile processes* (see, e.g., [32]). In these approaches, priorities cut off superfluous paths that only present new temporal but not causal system dependencies.

Future Work. In addition to the fact that a calculus combining dynamic priority and local pre-emption has not yet been developed, the semantic theories for CCS^{sg} and CCS^{sl} need to be completed by axiomatizing their *observational* congruences. For *finite* processes, one should be able to establish these axiomatizations using standard techniques [57]. However, for *regular* processes – i.e., the class of finite-state processes which do not contain recursion through static operators – it is not clear how to obtain a completeness result. The point is that existing methods for proving completeness of axiomatizations with respect to observational congruences rely on the possibility to remove or to insert τ -cycles in processes. In the context of pre-emption, however, this would possibly change the pre-emption potential of processes and, thus, is semantically incompatible with the prioritized observational congruences presented here. Recently, a similar problem, which was identified for the temporal process algebra PMC before [2], has been attacked in [20,40] for stochastic process calculi with priority and/or maximal progress. It remains to be seen whether these approaches and techniques can be employed for the priority settings investigated in this chapter.

Most process algebras which have been equipped with a notion of priority rely on an interleaving semantics, handshake communication, and a semantic theory based on bisimulation. Therefore, it should be investigated in which sense the presented approaches and results, especially regarding local pre-emption, can be adapted to broadcasting calculi, such as Hoare’s CSP [42]. Moreover, since for semantics based on local pre-emption the usual interleaving law is not valid, it is worth pursuing research regarding local pre-emption for non-interleaving semantic frameworks [4,77]. Preliminary considerations have been made in Jensen’s thesis [46]. However, the insights obtained by Jensen are restricted to a structural operational semantics for a CCS-based calculus, which is defined using *asynchronous transition systems* [77]. Jensen’s results

do not comprise a behavioral relation such as bisimulation (cf. [61]). Finally, we want to note that – to the best of our knowledge – extensions of higher-order process algebras [58,66] with concepts of priority do not exist, yet. Thus, it would be interesting to see if some of the approaches presented here can be carried over.

8 Sources and Acknowledgments

Major parts of this chapter have been adapted from several publications by the authors, which include two PhD theses: the results of Section 3 are taken from [27,53,62,63] and the ones of Section 4 from [29,53]; Section 5 heavily borrows from material contained in [16,53].

This work was supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-97046 while the second author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), M/S 132C, NASA Langley Research Center, Hampton, VA 23681-2199. The first author also acknowledges support by NSF grants CCR-9257963, CCR-9505662, CCR-9804091, and INT-9603441, AFOSR grant F49620-95-1-0508, and ARO grant P-38682-MA. Last but not least we would like to thank Girish Bhat, Matthew Hennessy, Michael Mendler, and Bernhard Steffen for many discussions about priority in process algebras.

References

- [1] L. Aceto, W. Fokkink, and C. Verhoef. Structural operational semantics. In this volume.
- [2] H.R. Andersen and M. Mendler. Complete axiomatization of observational congruence for PMC. Technical Report ID-TR:1993-126, Department of Computer Science, Technical University of Denmark, Lyngby, Denmark, 1993.
- [3] J.C.M. Baeten, editor. *Applications of Process Algebra*, volume 17 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1990.
- [4] J.C.M. Baeten and T. Basten. Process algebra with a partial-order semantics. In this volume.
- [5] J.C.M. Baeten, J.A. Bergstra, and J.W. Klop. Syntax and defining equations for an interrupt mechanism in process algebra. *Fundamenta Informaticae IX*, pages 127–168, 1986.
- [6] J.C.M. Baeten and J.W. Klop, editors. *CONCUR '90 (Concurrency Theory)*, volume 458 of *Lecture Notes in Computer Science*. Springer Verlag, 1990.

- [7] J.C.M. Baeten and W.P. Weijland. *Process Algebra*, volume 18 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1990.
- [8] G. Barrett. The semantics of priority and fairness in occam. In M. Main, A. Melton, M. Mislove, and D. Schmidt, editors, *Proceedings of Mathematical Foundations of Programming Semantics (MFPS '89)*, volume 442 of *Lecture Notes in Computer Science*, pages 194–208. Springer Verlag, 1989.
- [9] J.A. Bergstra and J.W. Klop. Algebra of communicating processes with abstraction. *Theoretical Computer Science*, 37(1):77–121, 1985.
- [10] J.A. Bergstra, C.A. Middelburg, and Y.S. Usenko. Discrete time process algebra and the semantics of SDL. In this volume.
- [11] M. Bernardo and R. Gorrieri. Extended Markovian process algebra. In U. Montanari and V. Sassone, editors, *CONCUR '96 (Concurrency Theory)*, volume 1119 of *Lecture Notes in Computer Science*, pages 315–330. Springer Verlag, 1996.
- [12] G. Berry. Preemption in concurrent systems. In R.K. Shyamasundar, editor, *Foundations of Software Technology and Theoretical Computer Science (FSTTCS '93)*, volume 761 of *Lecture Notes in Computer Science*, pages 72–93. Springer Verlag, 1993.
- [13] G. Berry and G. Gonthier. The ESTEREL synchronous programming language: Design, semantics, implementation. *Science of Computer Programming*, 19(2):87–152, 1992.
- [14] E. Best, R. Devillers, and M. Koutny. A consistent model for nets and process algebras. In this volume.
- [15] E. Best and M. Koutny. Petri net semantics of priority systems. *Theoretical Computer Science*, 96(1):175–215, 1992.
- [16] G. Bhat, R. Cleaveland, and G. Lüttgen. A practical approach to implementing real-time semantics. *Annals of Software Engineering*, 7, 1999. To appear.
- [17] T. Bolognesi and E. Brinksma. Introduction to the ISO specification language LOTOS. *Computer Networks and ISDN Systems*, 14:25–59, 1987.
- [18] G. Boudol, I. Castellani, M. Hennessey, and A. Kiehn. Observing localities. *Theoretical Computer Science*, 114(1):31–61, 1993.
- [19] J. Bradfield and C. Stirling. Modal process logics for processes. In this volume.
- [20] M. Bravetti and R. Gorrieri. Interactive generalized semi-markov processes. In *Process Algebra and Performance Modelling (PAPM '99)*, 1999. To appear.
- [21] P. Brémont-Grégoire, J.-Y. Choi, and I. Lee. A complete axiomatization of finite-state ACSR processes. *Information and Computation*, 138(2):124–159, 1997.

- [22] P. Brémont-Grégoire, S. Davidson, and I. Lee. CCSR92: Calculus for communicating shared resources with dynamic priorities. In Purushothaman and Zwarico [70], pages 65–85.
- [23] P. Brémont-Grégoire, I. Lee, and R. Gerber. A process algebra of communicating shared resources with dense time and priorities. *Theoretical Computer Science*, 189(1/2):179–219, 1997.
- [24] J. Camilleri and G. Winskel. CCS with priority choice. *Information and Computation*, 116(1):26–37, 1995.
- [25] I. Castellani. Locations and distributed processes. In this volume.
- [26] R. Cleaveland. Equivalence and preorder algorithms for finite-state systems. In this volume.
- [27] R. Cleaveland and M.C.B. Hennessy. Priorities in process algebras. *Information and Computation*, 87(1/2):58–77, 1990.
- [28] R. Cleaveland, G. Lüttgen, and M. Mendler. An algebraic theory of multiple clocks. In A. Mazurkiewicz and J. Winkowski, editors, *CONCUR '97 (Concurrency Theory)*, volume 1243 of *Lecture Notes in Computer Science*, pages 166–180. Springer Verlag, 1997.
- [29] R. Cleaveland, G. Lüttgen, and V. Natarajan. A process algebra with distributed priorities. *Theoretical Computer Science*, 195(2):227–258, 1998.
- [30] R. Cleaveland, V. Natarajan, S. Sims, and G. Lüttgen. Modeling and verifying distributed systems using priorities: A case study. *Software–Concepts and Tools*, 17(2):50–62, 1996.
- [31] J. Davies and S. Schneider. A brief history of Timed CSP. *Theoretical Computer Science*, 138(2):243–271, 1995.
- [32] P. Degano and C. Priami. Causality of mobile processes. In Z. Fülöp and F. Gécseg, editors, *International Conference on Automata, Languages and Programming (ICALP '95)*, volume 944 of *Lecture Notes in Computer Science*, pages 660–671. Springer Verlag, 1995.
- [33] C.J. Fidge. A formal definition of priority in CSP. *ACM Transactions on Programming Languages and Systems*, 15(4):681–705, 1993.
- [34] R. Gerber and I. Lee. CCSR: A calculus for communicating shared resources. In Baeten and Klop [6], pages 263–277.
- [35] R. Gerber and I. Lee. A resourced-based prioritized bisimulation for real-time systems. *Information and Computation*, 113(1):102–142, 1994.
- [36] S.M. German. Programming in a general model of synchronization. In R. Cleaveland, editor, *CONCUR '92 (Concurrency Theory)*, volume 630 of *Lecture Notes in Computer Science*, pages 534–549. Springer Verlag, 1992.
- [37] H. Hansson and F. Orava. A process calculus with incomparable priorities. In Purushothaman and Zwarico [70], pages 43–64.

- [38] D. Harel. Statecharts: A visual formalism for complex systems. *Science of Computer Programming*, 8:231–274, 1987.
- [39] M.C.B. Hennessy. *Algebraic Theory of Processes*. MIT Press, 1988.
- [40] H. Hermanns and M. Lohrey. Priority and maximal progress are completely axiomatisable. In D. Sangiorgi and R. de Simone, editors, *CONCUR '98 (Concurrency Theory)*, volume 1466 of *Lecture Notes in Computer Science*, pages 237–252. Springer Verlag, 1998.
- [41] H. Hermanns, M. Rettelbach, and T. Weiß. Formal characterisation of immediate actions in SPA with nondeterministic branching. *The Computer Journal*, 38(7):530–541, 1995.
- [42] C.A.R. Hoare. *Communicating Sequential Processes*. Prentice Hall, 1985.
- [43] INMOS Limited. *Occam Programming Manual*. International Series in Computer Science. Prentice Hall, 1984.
- [44] R. Janicki and M. Koutny. Semantics of inhibitor nets. *Information and Computation*, 123(1):1–17, 1995.
- [45] A. Jeffrey. Translating timed process algebra into prioritized process algebra. In J. Vytopil, editor, *Symposium on Real-Time and Fault-Tolerant Systems (FTRTFT '92)*, volume 571 of *Lecture Notes in Computer Science*, pages 493–506. Springer Verlag, 1992.
- [46] C.-T. Jensen. *Prioritized and Independent Actions in Distributed Computer Systems*. PhD thesis, Aarhus University, Denmark, 1994.
- [47] B. Jonsson, K.G. Larsen, and W. Yi. Probabilistic extensions of process algebra. In this volume.
- [48] P. Kanellakis and S.A. Smolka. CCS expressions, finite state processes, and three problems of equivalence. *Information and Computation*, 86(1):43–68, 1990.
- [49] Kempe Software Capital Enterprises. Ada95 reference manual: Language and standard libraries, 1995. Available at <http://www.adahome.com>.
- [50] L. Lamport. What it means for a concurrent program to satisfy a specification: Why no one has specified priority. In *Twelfth Annual ACM Symposium on Principles of Programming Languages (POPL '85)*, pages 78–83. IEEE Computer Society Press, 1985.
- [51] F. Levi. *Verification of Temporal and Real-Time Properties of Statecharts*. PhD thesis, University of Pisa-Genova-Udine, Pisa, Italy, 1997.
- [52] G. Lowe. Probabilistic and prioritized models of timed CSP. *Theoretical Computer Science*, 138(2):315–352, 1995.
- [53] G. Lüttgen. *Pre-emptive Modeling of Concurrent and Distributed Systems*. PhD thesis, University of Passau, Passau, Germany, 1998. Shaker Verlag.

- [54] G. Lüttgen, M. von der Beeck, and R. Cleaveland. Statecharts via process algebra. In *Concurrency Theory (CONCUR '99)*, Lecture Notes in Computer Science. Springer Verlag, 1999. To appear.
- [55] F. Maraninchi. The ARGOS language: Graphical representation of automata and description of reactive systems. In *IEEE Workshop on Visual Languages*. IEEE Computer Society Press, 1991.
- [56] R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
- [57] R. Milner. A complete axiomatisation for observational congruence of finite-state behaviours. *Information and Computation*, 81(2):227–247, 1989.
- [58] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, parts I and II. *Information and Computation*, 100(1):1–77, 1992.
- [59] F. Moller and C. Tofts. A temporal calculus of communicating systems. In Baeten and Klop [6], pages 401–415.
- [60] U. Montanari and D. Yankelevich. Location equivalence in a parametric setting. *Theoretical Computer Science*, 149(2):299–332, 1995.
- [61] M. Mukund and M. Nielsen. CCS, locations and asynchronous transition systems. In R.K. Shyamasundar, editor, *Foundations of Software Technology and Theoretical Computer Science (FSTTCS '92)*, volume 652 of *Lecture Notes in Computer Science*, pages 328–341. Springer Verlag, 1992.
- [62] V. Natarajan. *Degrees of Delay: Semantic Theories for Priority, Efficiency, Fairness, and Predictability in Process Algebras*. PhD thesis, North Carolina State University, Raleigh, NC, USA, 1996.
- [63] V. Natarajan, L. Christoff, I. Christoff, and R. Cleaveland. Priorities and abstraction in process algebra. In P.S. Thiagarajan, editor, *Foundations of Software Technology and Theoretical Computer Science (FSTTCS '94)*, volume 880 of *Lecture Notes in Computer Science*, pages 217–230. Springer Verlag, 1994.
- [64] R. Paige and R.E. Tarjan. Three partition refinement algorithms. *SIAM Journal of Computing*, 16(6):973–989, 1987.
- [65] D.M.R. Park. Concurrency and automata on infinite sequences. In P. Deussen, editor, *Proceedings of 5th GI Conference on Theoretical Computer Science*, volume 104 of *Lecture Notes in Computer Science*, pages 167–183. Springer Verlag, 1981.
- [66] J. Parrow. Mobility in process algebras. In this volume.
- [67] G.D. Plotkin. A structural approach to operational semantics. Technical Report DAIMI-FN-19, Computer Science Department, Aarhus University, Denmark, 1981.
- [68] K.V.S. Prasad. Programming with broadcasts. In E. Best, editor, *CONCUR '93 (Concurrency Theory)*, volume 715 of *Lecture Notes in Computer Science*, pages 173–187. Springer Verlag, 1993.

- [69] K.V.S. Prasad. Broadcasting with priority. In D. Sannella, editor, *Proceedings of the 5th European Symposium on Programming (ESOP '94)*, volume 788 of *Lecture Notes in Computer Science*, pages 469–484. Springer Verlag, 1994.
- [70] P. Purushothaman and A. Zwarico, editors. *First North American Process Algebra Workshop (NAPAW '92)*, Workshops in Computing. Springer Verlag, 1992.
- [71] W. Reisig. *Petri Nets: An Introduction*. Springer, 1985.
- [72] S.A. Smolka and B. Steffen. Priority as extremal probability. *Formal Aspects of Computing*, 8(5):585–606, 1996.
- [73] C. Tofts. Processes with probabilities, priority and time. *Formal Aspects of Computing*, 6(5):536–564, 1994.
- [74] A.C. Uselton and S.A. Smolka. A compositional semantics for Statecharts using labeled transition systems. In B. Jonsson and J. Parrow, editors, *CONCUR '94 (Concurrency Theory)*, volume 836 of *Lecture Notes in Computer Science*, pages 2–17. Springer Verlag, 1994.
- [75] R. van Glabbeek, S.A. Smolka, and B. Steffen. Reactive, generative, and stratified models of probabilistic processes. *Information and Computation*, 121(1):59–80, 1995.
- [76] C. Verhoef. A congruence theorem for structured operational semantics with predicates and negative premises. *Nordic Journal of Computing*, 2(2):274–302, 1995.
- [77] G. Winskel and M. Nielsen. Models for concurrency. In S. Abramsky, D.M. Gabbay, and T.S.E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 4, pages 1–148. Oxford Science Publications, 1995.
- [78] W. Yi. CCS + time = an interleaving model for real time systems. In J. Leach Albert, B. Monien, and M. Rodríguez-Artalejo, editors, *International Conference on Automata, Languages and Programming (ICALP '91)*, volume 510 of *Lecture Notes in Computer Science*, pages 217–228. Springer Verlag, 1991.