

CCS with Priority Guards

Iain Phillips*

Department of Computing, Imperial College, London

iccp@doc.ic.ac.uk

Abstract. It has long been recognised that standard process algebra has difficulty dealing with actions of different priority, such as for instance an interrupt action of high priority. Various solutions have been proposed. We introduce a new approach, involving the addition of “priority guards” to Milner’s process calculus CCS. In our approach, priority is *unstratified*, meaning that actions are not assigned fixed levels, so that the same action can have different priority depending where it appears in a program. Unlike in other unstratified accounts of priority in CCS (such as that of Camilleri and Winskel), we treat inputs and outputs symmetrically. We introduce the new calculus, give examples, develop its theory (including bisimulation and equational laws), and compare it with existing approaches. We show that priority adds expressiveness to both CCS and the π -calculus.

1 Introduction

It has long been recognised that standard process algebra [13,9,2] has difficulty dealing with actions of different priority, such as for instance an interrupt action of high priority. Various authors have suggested how to add priority to process languages such as ACP [1,10], CCS [4,3] and CSP [8,7]. We introduce a new approach, involving the addition of “priority guards” to the summation operator of Milner’s process calculus CCS. In our approach, priority is *unstratified*, meaning that actions are not assigned fixed levels, so that the same action can have different priority depending where it appears in a program. We shall see that existing accounts of priority in CCS are either stratified [4], or else they impose a distinction between outputs and inputs, whereby prioritised choice is only made on inputs [3,5]. This goes against the spirit of CCS, where inputs and outputs are treated symmetrically, and we contend that it is unnecessary. We introduce the new calculus, give examples, develop its theory (including bisimulation and equational laws), and compare it with existing approaches. We show that priority adds expressiveness to both CCS and the π -calculus.

We start with the idea of priority. We assume some familiarity with CCS notation [13]. Consider the CCS process $a + b$. The actions a and b have equal status. Which of them engages in communication depends on whether the environment is offering the complementary actions \bar{a} or \bar{b} . By “environment” we

* Partially funded by EPSRC grant GR/K54663

mean whatever processes may be placed in parallel with $a + b$. We would like some means to favour a over b , say, so that if the environment offers both, then only a can happen. This would be useful if, for instance, a was an interrupt action. We need something more sophisticated than simply removing b altogether, since, if a cannot communicate, it should not stop b from doing so. This brief analysis points to two features of priority: (1) Priority removes (“preempts”) certain possibilities that would have existed without priority. Thus if a can communicate then b is preempted. (2) Reasoning about priority in CCS has to be parametrised by the environment.

We now explain the basic idea of priority guards. Let P be a process, let a be an action, and let U be some set of actions. Then we can form a new process $U:a.P$, which behaves like $a.P$, except that the initial action a is conditional on the environment not offering actions in \bar{U} , the CCS “complement” of U . We call U a *priority guard* in $U:a.P$. All actions in U have priority over a at this point in the computation. We call our calculus CPG (for CCS with Priority Guards).

As a simple example, if we have a CCS process $a.P + b.Q$ and we wish to give a priority over b in the choice, we add a priority guard to get $a.P + a:b.Q$ (we omit the set braces around a). Priority is specific to this choice, since the guard affects only the initial b , and not any further occurrences of b there may be in Q .

Let us see how this example is handled in two existing approaches to priority. Cleaveland and Hennessy [4] add new higher priority actions to CCS. They would write our example as $\underline{a}.P + b.Q$ (high priority actions are underlined). In their stratified calculus, actions have fixed priority levels, and only actions at the same priority level can communicate. In this paper we are interested in an unstratified approach, and so our starting point of reference is Camilleri and Winskel’s priority choice operator [3]. In their notation the example becomes $a.P \rightarrow b.Q$. They make the priority of a over b specific to the particular choice, so that b might have priority over a elsewhere in the same program. We shall compare our approach with these two existing ones in Sect. 2.

A striking by-product of adding priority guards to CCS is that we can encode mixed input and output guarded summation using priority guards and restricted parallel composition. As a simple example, $a.P + \bar{b}.Q$ can be encoded with priority guards as $\text{new } c(c:a.(P|\bar{c})|c:\bar{b}.(Q|\bar{c}))$ (where c is a fresh action). This expresses in a natural way the preemptive nature of $+$, whereby pursuing one option precludes the others. The same effect can be achieved in Camilleri and Winskel’s calculus (but only for input guards): $a.P + b.Q$ can be encoded as

$$\text{new } c((c+)a.(P|\bar{c})|(c+)b.(Q|\bar{c}))$$

but of course here we are exchanging one form of choice for another. We shall return to this encoding of summation in Sect. 6.

To end this section, we give an example, involving handling of hidden actions and the scoping of priority. We wish to program a simple interrupt. Let P be a system which consists of two processes A, B in parallel which perform actions a, b respectively, while communicating internally to keep in step with each other.

P also has an interrupt process I which shuts down A and B when the interrupt signal int is received.

$$\begin{aligned} P &\stackrel{\text{df}}{=} \text{new } \text{mid}, \text{int}_A, \text{int}_B (A|B|I) & A &\stackrel{\text{df}}{=} \text{int}_A : a.\overline{\text{mid}}.A + \text{int}_A \\ I &\stackrel{\text{df}}{=} \text{int}.\overline{(\text{int}_A.\overline{\text{int}}_B + \overline{\text{int}}_B.\overline{\text{int}}_A)} & B &\stackrel{\text{df}}{=} \text{int}_B : b.\text{mid}.B + \text{int}_B \end{aligned}$$

Without the priority guards in A and B , P could receive an int and yet A and B could continue with a and b . Actions int_A , int_B have priority over a , b , respectively. This only applies within the scope of the restriction. We can apply the usual techniques of CCS (including removing τ actions) and get

$$P = a.P_1 + b.P_2 + \text{int} \quad P_1 = b.P + \text{int} \quad P_2 = a.P + \text{int}$$

which is what we wanted. We consider this example more precisely in Sect. 8.

Notice that in the system as a whole, once the high-priority interrupt action is restricted, we have regained a standard CCS process without priority. Thus priority can be encapsulated locally, which we regard as an important feature when programming larger systems, where different priority frameworks may be in use in different subsystems.

The rest of the paper is organised as follows: First we compare our approach with related work (Sect. 2). Next we define the language of processes (Sect. 3). Then we look at reactions (Sect. 4) and labelled transitions (Sect. 5). We then look at bisimulation and equational theories for both the strong (Sect. 6) and weak cases (Sect. 7). We then return to our interrupt example (Sect. 8), and look at the extra expressiveness afforded by priority guards (Sect. 9). The paper is completed with some brief conclusions.

2 Comparison with Related Work

We refer the reader to [5] for discussion of the many approaches taken by other authors to priority. Here we restrict ourselves to comparison of our work with that of Camilleri and Winskel [3] (referred to as CW for short) and Cleaveland, Lüttgen and Natarajan [5] (CLN for short).

2.1 Camilleri and Winskel (CW)

As we have seen, CW's CCS with a prioritised choice operator $P \dot{+} Q$ allows priority to be decided in a way which is specific to each choice in a system. The idea of a priority choice between processes is interesting and natural. The authors present an operational semantics via a labelled transition relation, and define a bisimulation-based equivalence. They also give an axiomatisation of this equivalence which is complete for finite processes (i.e. those not using recursion). They do not show how to hide the τ -actions resulting from communications (though this is treated in [11]).

As we saw in the Introduction, reasoning about priority has to be parametrised on the environment. The CW transition relation is parametrised

on a set of output actions R . Thus $\vdash_R P \xrightarrow{\alpha} P'$ means that, in an environment which is ready to perform precisely the actions R , the process P can perform an action α to become P' . For example, $\vdash_R a \rightarrow b \xrightarrow{a} 0$ (any R), while $\vdash_R a \rightarrow b \xrightarrow{b} 0$ provided $\bar{a} \notin R$.

We have borrowed the idea of parametrisation on the environment for our labelled transition system for CPG. For us $P \xrightarrow{\alpha_U} P'$ means that, in an environment which offers no action in the set \bar{U} , process P can perform α to become P' . Our most basic rule is essentially $U : a.P \xrightarrow{\alpha_U} P$, provided $a \notin U$.

Note that the CW syntax shows the environment only in the prioritised choice $a.P \rightarrow b.Q$, and does this implicitly, in that $b.Q$'s "environment" is $a.P$, while $a.P$ says nothing about the actual environment. In CPG the environment is represented in the syntax directly.

There is a difference in expressiveness between CPG and CW's calculus, in that the latter cannot express cycles of priority, whereas we can in CPG. CW consider the paradoxical example $\text{new } a, b ((a \rightarrow \bar{b}) | (b \rightarrow \bar{a}))$. The problem is that there is a circularity, with a having priority over b , as well as vice versa. Can the system act? They decide to sidestep this question by breaking the symmetry in CCS between inputs and outputs, and only allowing prioritised choice on input actions. We feel that this complicates the syntax and operational semantics, and should not be necessary. There seems to be no essential reason for CW not to allow the system with circular priorities, since their environmental transition relation should be able to handle it. In our approach the example is admitted, and results in a deadlock, which would seem to be in keeping with CW's approach. We consider this example again at the end of Sect. 5.

Another reason why CW disallow priority choice on output actions is to assist in obtaining the normal form they use for proving the completeness of their equational laws for finite processes. However this normal form is still quite complicated (consisting of a sum of priority sums of sums). In our calculus CPG we have only one form of choice, and so completeness is technically simpler.

2.2 Cleaveland, Lüttgen, and Natarajan (CLN)

In CLN's basic approach [5], which is derived from earlier work of Cleaveland and Hennessy [4], actions have priority levels. Mostly they consider just two levels—ordinary actions and higher priority, underlined actions. Only actions at the same level of priority can communicate, which is really quite restrictive when one considers that two actions which are intended to communicate may have quite different priorities within their respective subsystems. Silent actions resulting from communication have preemptive power over all actions of lower priority. The authors present both strong and weak bisimulation-based equivalences (drawing on [15]), and axiomatise these for finite processes.

In our unstratified calculus CPG, by contrast, actions do not have priority levels—each priority guard operates independently, in the spirit of [3].

We referred in the Introduction to the desirability of encapsulating priorities locally. This encapsulation is present in Camilleri and Winskel's calculus (and in

our own), but not in Cleaveland and Hennessy's, since a high priority τ is treated differently from a standard τ . However, the development in [5] goes beyond the basic Cleaveland and Hennessy calculus to consider distributed priorities, where preemption is decided locally rather than globally. Consideration is also given to extending the distributed priority calculus to allow communication between actions at different levels. The authors identify a problem with associativity of parallel composition. Consider the system

$$(a + b) | (\bar{b} + \underline{c}) | \bar{c}$$

where communication is allowed between complementary actions at different levels. If this associates to the left, then a is preempted by b ; however if it associates to the right then b is preempted by c , and so a is not preempted. A similar problem is encountered when extending the distributed calculus to allow more than two levels. CLN's proposed solution is to follow CW by only allowing priorities to be resolved between *input* actions, while treating all output actions as having equal priority. We have already mentioned our reservations about this. Nevertheless the distinction between inputs and outputs gives a workable "mixed-level" calculus (distributed, multi-level, with communication between different levels). It is particularly nice that CLN show that the CW calculus can be translated faithfully and naturally into this mixed-level calculus.

It is striking that both CW and the mixed-level calculus of CLN adopt the same syntactic restriction on inputs and outputs, and also that only *strong* equivalence (τ actions not hidden) is presented for the mixed-level calculus. We shall present a weak equivalence for CPG.

3 The Language CPG

We shall denote our augmentation of CCS with priority guards by *CPG* (CCS with Priority Guards). First we define the actions of CPG. In standard CCS [13, Part I] there is a set of *names* \mathcal{N} and a disjoint set of *co-names* $\bar{\mathcal{N}}$, together with a single silent action τ . To these standard names \mathcal{N} we shall add a new disjoint set of names \mathcal{U} and a dual set $\bar{\mathcal{U}}$. These are the actions which can be used in priority guards; they can also be used in the standard way. They need to be kept separate from standard actions, since we have to be careful with them in reasoning compositionally about processes.

To see why we take this approach, consider the law $P = \tau.P$, which is valid for CCS processes.¹ In CPG, if a can be a priority guard then $a \neq \tau.a$ since there is a context in which the two sides behave differently. Indeed, $a | \bar{a} : b$ cannot perform b (since, as we shall see, b is preempted by the offer of a), whereas $\tau.a | \bar{a} : b$ can perform b initially, as a is not offered until τ has occurred. However if we know that a is a standard name then we do have $a = \tau.a$. So we can retain CCS reasoning when processes only involve standard names.

¹ We are following the formulation of CCS in [13] rather than that of [12]. Processes such as $P + (Q|R)$ are not allowed, only guarded choices $\sum \alpha_i.P_i$.

We define $\text{Std} = \mathcal{N} \cup \bar{\mathcal{N}}$, $\text{Pri} = \mathcal{U} \cup \bar{\mathcal{U}}$, $\text{Vis} = \text{Std} \cup \text{Pri}$ and $\text{Act} = \text{Vis} \cup \{\tau\}$. We let u, v, \dots range over Pri , a, b, \dots over Vis and α, β, \dots over Act . Also S, T, \dots range over finite subsets of Vis , and U, V, \dots over finite subsets of Pri . If $S \subseteq \text{Vis}$, let \bar{S} denote $\{\bar{a} : a \in S\}$, where if $\bar{a} \in \bar{\mathcal{N}} \cup \bar{\mathcal{U}}$ then $\bar{a} = a$.

Now we define processes:

Definition 1. (cf [13, Definition 4.1]) \mathcal{P} is the smallest set such that whenever P, P_i are processes then \mathcal{P} contains

1. $\sum_{i \in I} S_i : \alpha_i . P_i$ (guarded summation: I finite, each S_i finite)
2. $P_1 | P_2$ (parallel composition)
3. $\text{new } a P$ (restriction)
4. $A\langle a_1, \dots, a_n \rangle$ (identifier)

\mathcal{P} is ranged over by P, Q, R, \dots . We let M, N, \dots range over (guarded) summations. We assume that each identifier $A\langle b_1, \dots, b_n \rangle$ comes with a defining equation $A\langle a_1, \dots, a_n \rangle \stackrel{\text{df}}{=} P$, where P is a process whose free names are drawn from a_1, \dots, a_n . We will tend to abbreviate a_1, \dots, a_n by \vec{a} . We write the empty guarded summation as 0 and abbreviate $S : \alpha . 0$ by $S : \alpha$. It is assumed that the order in a summation is immaterial. We abbreviate $\emptyset : \alpha$ by α . Definition 1 is much as in standard CCS except for the priority guards S_i . The meaning of the priority guard $S : \alpha$ is that α can only be performed if the environment does not offer any action in $\bar{S} \cap \text{Pri}$. Clearly, any names in $S - \text{Pri}$ have no effect as guards, and can be eliminated without changing the behaviour of a process. We allow them to occur in the syntax, since otherwise we could not freely instantiate the parameters in an identifier. We write $u : \alpha$ instead of $\{u\} : \alpha$. Restriction is a variable-binding operator, and we write $\text{fn}(P)$ for the free names of P .

Two sublanguages of CPG are of interest:

Definition 2. Let \mathcal{P}_{Std} be the sublanguage of standard processes generated as in Definition 1 except that all names are drawn from Std (i.e. we effectively take $\mathcal{U} = \emptyset$ and $S_i = \emptyset$ in clause (1)). Let \mathcal{P}_{Ug} be the sublanguage of unguarded processes generated as in Definition 1 except that all priority guards are empty (i.e. $S_i = \emptyset$ in clause (1)).

Clearly $\mathcal{P}_{\text{Std}} \subseteq \mathcal{P}_{\text{Ug}} \subseteq \mathcal{P}$. Note that \mathcal{P}_{Std} is effectively standard CCS. The unguarded processes \mathcal{P}_{Ug} differ from \mathcal{P}_{Std} in that they may contain names in Pri . Such processes cause no problems for strong equivalence (Proposition 4), but care is needed with weak equivalence (Sect. 7), since e.g. u and $\tau.u$ ($u \in \text{Pri}$) are not weakly equivalent, as remarked above.

4 Offers and Reaction

Structural congruence is the most basic equivalence on processes, which facilitates reaction by bringing the subprocesses which are to react with each other into juxtaposition. It is defined as for CCS:

Definition 3. (cf [13, Definition 4.7]) Structural congruence, written \equiv , is the congruence on \mathcal{P} generated by the following equations:

1. Change of bound names (alpha-conversion)
2. $P|0 \equiv P$, $P|Q \equiv Q|P$, $P|(Q|R) \equiv (P|Q)|R$
3. $\text{new } a (P|Q) \equiv P|\text{new } a Q$ if $a \notin \text{fn}(P)$;
 $\text{new } a 0 \equiv 0$, $\text{new } a \text{ new } b P \equiv \text{new } b \text{ new } a P$
4. $A(\vec{b}) \equiv \{\vec{b}/\vec{a}\}P$ if $A(\vec{a}) \stackrel{\text{df}}{=} P$

Recall that a guarded action $S : a$ is conditional on other processes in the environment not offering actions in $\bar{S} \cap \text{Pri}$. Before defining reaction we define for each process P the set $\text{off}(P) \subseteq \text{Pri}$ of “higher priority” actions “offered” by P .

Definition 4. By induction on $P \in \mathcal{P}$:

1. $\text{off}(\sum_{i \in I} S_i : \alpha_i . P_i) = \{\alpha_i : i \in I, \alpha_i \in \text{Pri}, \alpha_i \notin S_i\}$
2. $\text{off}(P_1|P_2) = \text{off}(P_1) \cup \text{off}(P_2)$
3. $\text{off}(\text{new } a P) = \text{off}(P) - \{a, \bar{a}\}$
4. $\text{off}(A(\vec{b})) = \text{off}(\{\vec{b}/\vec{a}\}P)$ if $A(\vec{a}) \stackrel{\text{df}}{=} P$

In item 1 the reason that we insist $\alpha_i \notin S_i$ is that we want to equate a process such as $u : u$ with 0 , since $u : u$ can never engage in a reaction. Note that if $P \in \mathcal{P}_{\text{Std}}$ then $\text{off}(P) = \emptyset$.

In CPG, a reaction can be conditional on offers from the environment. Consider $u : b|\bar{b}$. This can react by communication on b, \bar{b} . However b is guarded by u , and so the reaction is conditional on the environment not offering \bar{u} . We reflect this by letting reaction be parametrised on sets of actions $U \subseteq \text{Pri}$. The intended meaning of $P \rightarrow_U P'$ is that P can react on its own, as long as the environment does not offer \bar{u} for any $u \in U$ (in our parlance, “eschews” U).

Definition 5. Let $P \in \mathcal{P}$ and let $S \subseteq \text{Act}$ be finite. P eschews S (written P eschews S) iff $\text{off}(P) \cap S = \emptyset$.

Definition 6. (cf [13, Definition 4.13]) The reaction relation on \mathcal{P} is the smallest relation \rightarrow on $\mathcal{P} \times \wp(\text{Pri}) \times \mathcal{P}$ generated by the following rules:

$$S : \tau.P + M \rightarrow_{S \cap \text{Pri}} P$$

$$\frac{S : a.P + M \text{ eschews } T \quad T : \bar{a}.Q + N \text{ eschews } S}{(S : a.P + M)|(T : \bar{a}.Q + N) \rightarrow_{(S \cup T) \cap \text{Pri}} P|Q} \quad \frac{P \rightarrow_U P' \quad Q \text{ eschews } U}{P|Q \rightarrow_U P'|Q}$$

$$\frac{P \rightarrow_U P'}{\text{new } a P \rightarrow_{U - \{a, \bar{a}\}} \text{new } a P'} \quad \frac{P \rightarrow_U P' \quad P \equiv Q \quad P' \equiv Q'}{Q \rightarrow_U Q'}$$

We abbreviate $P \rightarrow_{\emptyset} P'$ by $P \rightarrow P'$.

The second clause of Definition 6 is the most important. In order for an action a to react with a complementary \bar{a} , the two sides must not preempt each other (i.e. they must eschew each other's guards). Furthermore the reaction remains conditional on the environment eschewing the union of their guards. The restriction rule shows how this conditionality can then be removed by scoping. Notice that if we restrict attention to the unguarded processes \mathcal{P}_{Ug} (i.e. we let $U = \emptyset$) we recover the usual CCS reaction relation. So the new reaction relation is conservative over the old.

5 Labelled Transitions

As in standard CCS, we wish to define a transition relation on processes $P \xrightarrow{\alpha} P'$ meaning that P can perform action α and become P' . As we did with reaction, we refine the transition relation so that it is parametrised on sets of actions $U \subseteq \text{Pri}$. The intended meaning of $P \xrightarrow{\alpha}_U P'$ is that P can perform α as long as the environment eschews U . Our definition is inspired by the transition relation in [3], which is parametrised on what set of output actions the environment is ready to perform.

Definition 7. (cf [13, Definition 5.1]) *The transition relation on \mathcal{P} is the smallest relation \rightarrow on $\mathcal{P} \times \text{Act} \times \wp(\text{Pri}) \times \mathcal{P}$ generated by the following rules:*

$$\begin{array}{l}
\text{(sum)} \quad M + S : \alpha.P + N \xrightarrow{\alpha}_{S \cap \text{Pri}} P \quad \text{if } \alpha \notin S \cap \text{Pri} \\
\text{(react)} \quad \frac{P_1 \xrightarrow{a}_{U_1} P'_1 \quad P_2 \xrightarrow{\bar{a}}_{U_2} P'_2 \quad P_1 \text{ eschews } U_2 \quad P_2 \text{ eschews } U_1}{P_1 | P_2 \xrightarrow{\tau}_{U_1 \cup U_2} P'_1 | P'_2} \\
\text{(par)} \quad \frac{\frac{P_1 \xrightarrow{\alpha}_U P'_1 \quad P_2 \text{ eschews } U}{P_1 | P_2 \xrightarrow{\alpha}_U P'_1 | P_2} \quad \frac{P_2 \xrightarrow{\alpha}_U P'_2 \quad P_1 \text{ eschews } U}{P_1 | P_2 \xrightarrow{\alpha}_U P_1 | P'_2}}{P_1 | P_2 \xrightarrow{\alpha}_U P'_1 | P'_2} \\
\text{(res)} \quad \frac{P \xrightarrow{\alpha}_U P'}{\text{new } a P \xrightarrow{\alpha}_{U - \{a, \bar{a}\}} \text{new } a P'} \quad \text{if } \alpha \notin \{a, \bar{a}\} \\
\text{(ident)} \quad \frac{\{\bar{u}/\bar{a}\} P \xrightarrow{\alpha}_U P'}{A(\bar{u}) \xrightarrow{\alpha}_U P'} \quad \text{if } A(\bar{a}) \stackrel{\text{df}}{=} P
\end{array}$$

We abbreviate $P \xrightarrow{\alpha}_{\emptyset} P'$ by $P \xrightarrow{\alpha} P'$ and $\exists P'. P \xrightarrow{\alpha}_U P'$ by $P \xrightarrow{\alpha}_U$.

Proposition 1. *If $P \xrightarrow{\alpha}_U P'$ then $\alpha \notin U$ and U is finite. Moreover,*

$$\{u \in \text{Pri} : \exists U. P \xrightarrow{u}_U\} \subseteq \text{off}(P) .$$

To see that $\text{off}(P)$ can be unequal to $\{u \in \text{Pri} : \exists U. P \xrightarrow{u}_U\}$, consider $u : v.0 | \bar{u}.0$. We see that $\text{off}(u : v | \bar{u}) = \{v, \bar{u}\}$, but $u : v | \bar{u}$ cannot perform v .

As with reaction, note that if we restrict attention to the unguarded processes \mathcal{P}_{Ug} (i.e. we let $U = \emptyset$) we recover the usual CCS transition relation. So the new transition relation is conservative over the old. In applications we envisage that

the standard CCS transition relation can be used most of the time. The CPG transition relation will only be needed in those subsystems which use priority.

As an illustration of the design choices embodied in our definitions, consider the circular example of Camilleri & Winskel (Subsect. 2.1):

$$P \stackrel{\text{df}}{=} u.a + u:\bar{v} \quad Q \stackrel{\text{df}}{=} v.b + v:\bar{u} \quad R \stackrel{\text{df}}{=} \text{new } u, v (P|Q)$$

In P action u has priority over \bar{v} , while in Q action v has priority over \bar{u} . We have $P \xrightarrow{u} a$, $Q \xrightarrow{\bar{u}} 0$. For a u communication to happen, by rule (react) we need $\bar{v} \notin \text{off}(P)$, but $\text{off}(P) = \{u, \bar{v}\}$, so that the u communication cannot happen. Similarly the v communication cannot happen, and so R is strongly equivalent to 0 (strong offer equivalence is defined in the next section).

6 Strong Offer Bisimulation

Similarly to standard CCS, we define process equivalences based on strong and weak bisimulation. We consider strong bisimulation in this section and weak bisimulation (i.e. with hiding of silent actions) in the next.

The intuition behind our notion of bisimulation is that for processes to be equivalent they must make the same offers, and for a process Q to simulate a process P , Q must be able to do whatever P can, though possibly constrained by fewer or smaller priority guards. For instance, we would expect the processes $a + u:a$ and a to be equivalent, since the priority guarded $u:a$ is simulated by a .

Definition 8. (cf [13]) A symmetric relation $\mathcal{S} \subseteq \mathcal{P} \times \mathcal{P}$ is a strong offer bisimulation if $\mathcal{S}(P, Q)$ implies both that $\text{off}(P) = \text{off}(Q)$ and that for all $\alpha \in \text{Act}$,

if $P \xrightarrow{\alpha}_U P'$ then for some Q' and $V \subseteq U$, we have $Q \xrightarrow{\alpha}_V Q'$ and $\mathcal{S}(P', Q')$

Definition 9. Processes P and Q are strongly offer equivalent, written $P \stackrel{\text{off}}{\sim} Q$, iff there is some strong offer bisimulation \mathcal{S} such that $\mathcal{S}(P, Q)$.

Proposition 2. (cf [13, Prop 5.2]) \equiv is a strong offer bisimulation. Hence \equiv implies $\stackrel{\text{off}}{\sim}$. □

Proposition 3. (cf [13, Theorem 5.6]) $P \xrightarrow{\tau}_U \equiv P'$ iff $P \rightarrow_U P'$ □

Theorem 1. (cf [13, Proposition 5.29]) Strong offer equivalence is a congruence, i.e. if $P \stackrel{\text{off}}{\sim} Q$ then

- | | |
|--|---|
| 1. $S:\alpha.P + M \stackrel{\text{off}}{\sim} S:\alpha.Q + M$ | 3. $P R \stackrel{\text{off}}{\sim} Q R$ |
| 2. $\text{new } a P \stackrel{\text{off}}{\sim} \text{new } a Q$ | 4. $R P \stackrel{\text{off}}{\sim} R Q$ □ |

Note that if $P, Q \in \mathcal{P}_{\text{Ug}}$ then $P \overset{\text{off}}{\sim} Q$ iff $P \sim Q$, where $P \sim Q$ denotes that P and Q are strongly equivalent in the usual sense of [13]. So $\overset{\text{off}}{\sim}$ is conservative over \sim . In fact we can say more:

Proposition 4. *Let $P, Q \in \mathcal{P}_{\text{Ug}}$. If $P \sim Q$ then $C[P] \overset{\text{off}}{\sim} C[Q]$, for any context $C[\cdot]$. \square*

So we can reuse all the known equivalences between CCS processes when working with CPG processes.

Proposition 5. *(cf [13, Proposition 5.21]) For all $P \in \mathcal{P}$,*

$$P \overset{\text{off}}{\sim} \sum \{U:\alpha.Q : P \xrightarrow{U} Q\}$$

Proposition 6. *The following laws hold:*

$$M + S:\alpha.P \overset{\text{off}}{\sim} M + (S \cap \text{Pri}):\alpha.P \quad (1)$$

$$M + U:\alpha.P \overset{\text{off}}{\sim} M \quad \text{if } \alpha \in U \subseteq \text{Pri} \quad (2)$$

$$M + U:\alpha.P + (U \cup V):\alpha.P \overset{\text{off}}{\sim} M + U:\alpha.P \quad (3)$$

$$\begin{aligned} & (\sum U_i:\alpha_i.P_i) \mid (\sum V_j:\beta_j.Q_j) \\ & \overset{\text{off}}{\sim} \sum \{U_i:\alpha_i.(P_i \mid (\sum V_j:\beta_j.Q_j)) : \forall j, \beta_j \notin \bar{U}_i\} \\ & + \sum \{V_j:\beta_j.((\sum U_i:\alpha_i.P_i) \mid Q_j) : \forall i, \alpha_i \notin \bar{V}_j\} \\ & + \sum \{(U_i \cup V_j):\tau.P_i \mid Q_j : \alpha_i = \bar{\beta}_j \in \text{Vis}, \forall i', j'. \alpha_{i'} \notin \bar{V}_j, \beta_{j'} \notin \bar{U}_i\} \end{aligned} \quad (4)$$

$$\text{new } a (\sum U_i:\alpha_i.P_i) \overset{\text{off}}{\sim} \sum ((U_i - \{a, \bar{a}\}):\alpha_i.\text{new } a P_i : \alpha_i \neq a, \bar{a}) \quad (5)$$

Definition 10. *Let \mathcal{A}_S be the following set of axioms: the axioms of structural congruence \equiv (Definition 3) together with the five laws of Proposition 6.*

Theorem 2. *The set of axioms \mathcal{A}_S is complete for $\overset{\text{off}}{\sim}$ on finite CPG processes (a CPG process is finite if it contains no identifiers). \square*

As mentioned in the Introduction, we can encode mixed input and output guarded summation using priority guards and restricted parallel composition.

Proposition 7. *Suppose that $\{\alpha_i : i \in I\}$ are actions which cannot react with each other, i.e. there do not exist $i, j \in I$ and $a \in \text{Vis}$ such that $\alpha_i = a$ and $\alpha_j = \bar{a}$. Then*

$$\sum S_i : \alpha_i . P_i \stackrel{\text{off}}{\approx} \text{new } u \left(\prod (S_i \cup \{u\} : \alpha_i (P_i(\bar{u}))) \right)$$

where $u \in \text{Pri}$ is some fresh name not occurring in $\sum S_i : \alpha_i . P_i$ and \prod denotes parallel composition. \square

The non-reaction condition in Proposition 7 is needed, since otherwise the right-hand side would have extra unwanted reactions. The condition is not unduly restrictive, since if we have a system where the same channel a is used to pass messages both to and from a process, we can simply separate a out into two separate channels, one for each direction.

7 Weak Offer Bisimulation

We now investigate weak bisimulation, where reactions are hidden.

Definition 11. $P \Rightarrow_U P'$ iff $P \stackrel{\text{id}}{=} P'$ or $\exists U_1, \dots, U_n . P \rightarrow_{U_1} \dots \rightarrow_{U_n} P'$ with $U = U_1 \cup \dots \cup U_n$ ($n \geq 1$). We abbreviate $P \Rightarrow_{\emptyset} P'$ by $P \Rightarrow P'$.

$P \stackrel{\alpha}{\Rightarrow}_U P'$ iff $\exists U', U'' . P \Rightarrow_{U'} P'' \xrightarrow{\alpha}_{U''} P''' \Rightarrow P'$ with $U = U' \cup U''$ and $\text{off}(P'') \subseteq \text{off}(P)$.

Here $P \stackrel{\text{id}}{=} P'$ means that P and P' are identically equal. So $P \Rightarrow_U P'$ allows zero or more internal transitions with guards included in U . The condition $\text{off}(P'') \subseteq \text{off}(P)$ is needed to obtain a weak equivalence which is a congruence. The reason why we allow priority guards before performing a visible action, but not after, is as follows: For Q to simulate $P \xrightarrow{\alpha}_U P'$, Q must expect an environment offering \bar{U} up to and including performing a . After this, the environment has changed, and might be offering anything. So Q can perform further reactions to reach Q' simulating P' , but these reactions must not be subject to any priority guards.

Definition 12. A symmetric relation $\mathcal{S} \subseteq \mathcal{P} \times \mathcal{P}$ is a weak offer bisimulation if $\mathcal{S}(P, Q)$ implies both that $\text{off}(P) = \text{off}(Q)$ and that:

if $P \rightarrow_U P'$ then for some Q' and $U' \subseteq U$, we have $Q \Rightarrow_{U'} Q'$ and $\mathcal{S}(P', Q')$, and for all $a \in \text{Vis}$,

if $P \xrightarrow{a}_U P'$ then for some Q' and $U' \subseteq U$, we have $Q \xrightarrow{a}_{U'} Q'$ and $\mathcal{S}(P', Q')$.

On the sublanguage \mathcal{P}_{Std} (which corresponds to CCS) weak offer bisimulation is the same as for CCS [13, Proposition 6.3].

Definition 13. Processes P and Q are weakly offer equivalent, written $P \stackrel{\text{off}}{\approx} Q$, iff there is some weak offer bisimulation \mathcal{S} such that $\mathcal{S}(P, Q)$.

Proposition 8. *For any P, Q , if $P \overset{\text{off}}{\sim} Q$ then $P \overset{\text{off}}{\approx} Q$.* □

Theorem 3. *(cf [13, Proposition 6.17]) $\overset{\text{off}}{\approx}$ is a congruence.* □

So we have a congruence which conservatively extends CCS.

We now turn to the equational theory of weak offer equivalence. In CCS we have the law $P \approx \tau.P$ [13, Theorem 6.15]. However in CPG, $u \overset{\text{off}}{\not\approx} \tau.u$. This is because $\text{off}(u) = \{u\}$ whereas $\text{off}(\tau.u) = \emptyset$. However the usual CCS equivalence laws will still hold for the standard processes \mathcal{P}_{Std} (recall that for $P \in \mathcal{P}_{\text{Std}}$, $\text{off}(P) = \emptyset$).

Proposition 9. *(cf [13, Theorem 6.15]) The following laws hold:*

$$\tau.P \overset{\text{off}}{\approx} P \quad \text{if } \text{off}(P) = \emptyset \quad (6)$$

$$M + N + \tau.N \overset{\text{off}}{\approx} M + \tau.N \quad \text{if } \text{off}(N) \subseteq \text{off}(M) \quad (7)$$

$$M + \alpha.P + \alpha.(\tau.P + N) \overset{\text{off}}{\approx} M + \alpha.(\tau.P + N) \quad (8)$$

We stated (6), (7) and (8) because in many situations it is convenient to use conventional CCS reasoning. The next result gives the “intrinsic” τ -laws of CPG:

Proposition 10. *The following four laws hold:*

$$M + U : \tau.M \overset{\text{off}}{\approx} M \quad (9)$$

$$M + U : \tau.(N + V : \tau.P) \overset{\text{off}}{\approx} M + U : \tau.(N + V : \tau.P) + (U \cup V) : \tau.P \quad (10)$$

If $\text{off}(N + V : a.P) \subseteq \text{off}(M)$:

$$M + U : \tau.(N + V : a.P) \overset{\text{off}}{\approx} M + U : \tau.(N + V : a.P) + (U \cup V) : a.P \quad (11)$$

$$M + U : \alpha.P + U : \alpha.(\tau.P + N) \overset{\text{off}}{\approx} M + U : \alpha.(\tau.P + N) \quad (12)$$

We can derive (7) from (10) and (11). Also we can derive:

$$\tau.M \overset{\text{off}}{\approx} M \quad \text{if } \text{off}(M) = \emptyset \quad (13)$$

from (9), (10), (11). Recall that every process is strongly equivalent to a summation (Proposition 5), and so (13) is effectively as strong as (6).

Definition 14. *Let \mathcal{A}_W be the axioms \mathcal{A}_S (Definition 10) together with (9), (10), (11) and (12).*

Theorem 4. *The axioms \mathcal{A}_W are complete for $\overset{\text{off}}{\approx}$ on finite processes.* \square

Proposition 11. *(cf [13, Theorem 6.19]) Unique solution of equations. Let \vec{X} be a (possibly infinite) sequence of process variables X_i . Up to $\overset{\text{off}}{\approx}$, there is a unique sequence \vec{P} of processes which satisfy the formal equations:*

$$X_i \overset{\text{off}}{\approx} \sum_j U_{ij} : a_{ij} \cdot X_{k(ij)}$$

(notice that τ s are not allowed). \square

8 Example

We now revisit the interrupt example from Sect. 1. Recall that we had:

$$\begin{aligned} P &\stackrel{\text{df}}{=} \text{new mid, int}_A, \text{int}_B (A|B|I) & A &\stackrel{\text{df}}{=} \text{int}_A : a.\overline{\text{mid}}.A + \text{int}_A \\ I &\stackrel{\text{df}}{=} \text{int}.\overline{(\text{int}_A.\overline{\text{int}}_B + \overline{\text{int}}_B.\overline{\text{int}}_A)} & B &\stackrel{\text{df}}{=} \text{int}_B : b.\text{mid}.B + \text{int}_B \end{aligned}$$

We want to show $P \overset{\text{off}}{\approx} Q$, where

$$Q \stackrel{\text{df}}{=} a.Q_1 + b.Q_2 + \text{int} \quad Q_1 \stackrel{\text{df}}{=} b.Q + \text{int} \quad Q_2 \stackrel{\text{df}}{=} a.Q + \text{int}$$

Clearly $\text{int}_A, \text{int}_B \in \text{Pri}$. We take $a, b, \text{mid}, \text{int} \in \text{Std}$. This means that $Q \in \mathcal{P}_{\text{Std}}$. We can use Laws (4), (5) to get:

$$P \overset{\text{off}}{\approx} a.P_1 + b.P_2 + \text{int}.P_3$$

$$P_1 \overset{\text{off}}{\approx} b.P_4 + \text{int}.\tau \quad P_2 \overset{\text{off}}{\approx} a.P_4 + \text{int}.\tau \quad P_3 \overset{\text{off}}{\approx} \tau.\tau + \tau.\tau \quad P_4 \overset{\text{off}}{\approx} \tau.P + \text{int}.\tau.P_3$$

where P_1, P_2, P_3, P_4 are various states of P . We can use law (6) to get:

$$P_1 \overset{\text{off}}{\approx} b.P_4 + \text{int} \quad P_2 \overset{\text{off}}{\approx} a.P_4 + \text{int} \quad P_3 \overset{\text{off}}{\approx} 0 \quad P_4 \overset{\text{off}}{\approx} \tau.P + \text{int}$$

Then we use law (7) to get $\tau.P + \text{int} \overset{\text{off}}{\approx} \tau.P$. Notice that this needs $\text{off}(P) = \emptyset$, i.e. $a, b, \text{int} \notin \text{Pri}$. Finally:

$$P \overset{\text{off}}{\approx} a.P_1 + b.P_2 + \text{int} \quad P_1 \overset{\text{off}}{\approx} b.P + \text{int} \quad P_2 \overset{\text{off}}{\approx} a.P + \text{int}$$

By Proposition 11 we get $P \overset{\text{off}}{\approx} Q$ as we wanted.

Our reasoning was presented equationally, but could equally well have been done using bisimulation. We first unfolded the behaviour of P . Since all prioritised actions were restricted, the system P had no priorities as far as the environment was concerned. We could therefore remove silent actions and simplify using standard techniques of CCS.

In the Introduction we used plain equality ($=$) when talking about equivalence between CPG processes. This is to be interpreted as $\overset{\text{off}}{\approx}$.

9 Expressiveness

In this section we show that priorities add expressive power to both CCS and the π -calculus [14,13]. As far as we are aware, this has not been previously shown for any notion of priority in process algebra. We use the work of Ene and Muntian [6], which was inspired by that of Palamidessi [16].

Definition 15. [16] An encoding $\llbracket \cdot \rrbracket$ is a compositional mapping from (the processes of) one calculus to another. It is called uniform if $\llbracket P|Q \rrbracket = \llbracket P \rrbracket \parallel \llbracket Q \rrbracket$ and, for any renaming σ , $\llbracket \sigma(P) \rrbracket = \sigma(\llbracket P \rrbracket)$. A semantics is reasonable if it distinguishes two processes P and Q whenever in some computation of P the actions on certain intended channels are different from those of any computation of Q .

Definition 16. (slight modification of [16, Definition 3.1]) A process

$$P = P_1 | \dots | P_n$$

is an electoral system if every computation of P can be extended (if necessary) to a computation which declares a leader i by outputting \bar{o}_i , and where no computation declares two different leaders.

The intuition behind the following theorem is that priorities give us something of the power of one-many (broadcast) communication, in that a single process can simultaneously interrupt several other processes. By contrast, π -calculus communication is always one-one.

Theorem 5. There is no uniform encoding of CPG into the π -calculus preserving a reasonable semantics.

Proof. (Sketch) We follow the proof of Ene and Muntian's result that the broadcast π -calculus cannot be encoded in the π -calculus [6]. The network of CPG processes $P_1 | \dots | P_n$ is an electoral system, where $P_i \stackrel{\text{df}}{=} u : a.(\bar{u}|o_i)|\bar{a}$. If P_i manages to communicate on a then P_i declares itself the leader. No other process can now do this, since P_i is preventing all the other processes from performing a by offering \bar{u} .

The rest of the proof is as in [6]. Suppose that we have an encoding $\llbracket \cdot \rrbracket$ of CPG into the π -calculus. Let $\sigma(o_i) = o_{m+i}$, with σ the identity otherwise. Consider $P_1 | \dots | P_{m+n}$. This is an electoral system and so the encoding $\llbracket P_1 | \dots | P_{m+n} \rrbracket$ must be also. But

$$\begin{aligned} \llbracket P_1 | \dots | P_{m+n} \rrbracket &= \llbracket (P_1 | \dots | P_m) | \sigma(P_1 | \dots | P_n) \rrbracket \\ &= \llbracket P_1 | \dots | P_m \rrbracket \parallel \llbracket \sigma(P_1 | \dots | P_n) \rrbracket \end{aligned}$$

So we have two electoral systems of m and n processes respectively, which can be run independently in the π -calculus to produce two winners. Contradiction. \square

Since CCS can be encoded in π -calculus, it follows that CPG has greater expressive power than CCS. It also follows that we can add expressive power to the π -calculus by adding priority guards.

Theorem 6. *There is no uniform encoding of the π -calculus into CPG preserving a reasonable semantics.*

Proof. Much as in [16], where it is shown for CCS rather than CPG. \square

The results of this section apply equally to Camilleri-Winskel and Cleaveland-Hennessy-style priority.

10 Conclusions

We have introduced priority guards into CCS to form the language CPG. We have defined both strong and weak bisimulation equivalences and seen that they are conservative over the CCS equivalences, and that they are congruences. We have given complete equational laws for finite CPG in both the strong and weak cases. Conservation over CCS has the consequence that in verifying CPG systems we can often use standard CCS reasoning, as long as we take some care with actions in the set of prioritised actions Pri .

CPG overcomes the asymmetry between inputs and outputs present both in Camilleri and Winskel's calculus and in the corresponding calculus of Cleaveland, Lüttgen and Natarajan.

Finally, we have seen that priority guards add expressiveness to both CCS and the π -calculus.

We wish to thank Philippa Gardner, Rajagopal Nagarajan, Catuscia Palamidessi, Andrew Phillips, Irek Ulidowski, Maria Grazia Vigliotti, Nobuko Yoshida and the anonymous referees for helpful discussions and suggestions.

References

1. J.C.M. Baeten, J. Bergstra, and J.-W. Klop. Syntax and defining equations for an interrupt mechanism in process algebra. *Fundamenta Informaticae*, 9:127–168, 1986.
2. J. Bergstra and J.-W. Klop. Process algebra for synchronous communication. *Information and Computation*, 60:109–137, 1984.
3. J. Camilleri and G. Winskel. CCS with priority choice. *Information and Computation*, 116(1):26–37, 1995.
4. R. Cleaveland and M.C.B. Hennessy. Priorities in process algebra. *Information and Computation*, 87(1/2):58–77, 1990.
5. R. Cleaveland, G. Lüttgen, and V. Natarajan. Priority in process algebra. In J.A. Bergstra, A. Ponse, and S.A. Smolka, editors, *Handbook of Process Algebra*. Elsevier, 2001.
6. C. Ene and T. Muntian. Expressiveness of point-to-point versus broadcast communications. In *FCT '99*, volume 1684 of *Lecture Notes in Computer Science*, pages 258–268. Springer-Verlag, 1999.
7. C.J. Fidge. A formal definition of priority in CSP. *ACM Transactions on Programming Languages and Systems*, 15(4):681–705, 1993.
8. H. Hansson and F. Orava. A process calculus with incomparable priorities. In *Proceedings of the North American Process Algebra Workshop*, pages 43–64, Stony Brook, New York, 1992. Springer-Verlag Workshops in Computer Science.

9. C.A.R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
10. A. Jeffrey. A typed, prioritized process algebra. Technical Report 13/93, Dept. of Computer Science, University of Sussex, 1993.
11. C.-T. Jensen. *Prioritized and Independent Actions in Distributed Computer Systems*. PhD thesis, Aarhus University, 1994.
12. R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
13. R. Milner. *Communicating and Mobile Systems: the π -calculus*. Cambridge University Press, 1999.
14. R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes. *Information and Computation*, 100:1–77, 1992.
15. 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, 14th Conference*, volume 880 of *Lecture Notes in Computer Science*, pages 217–230. Springer-Verlag, 1994.
16. C. Palamidessi. Comparing the expressive power of the synchronous and the asynchronous π -calculus. In *Proceedings of the 25th Annual Symposium on Principles of Programming Languages, POPL '97*, pages 256–265. ACM, 1997.