

A theory of weak bisimulation for Core CML

WILLIAM FERREIRA*

*Computing Laboratory, University of Cambridge,
Cambridge, UK*

MATTHEW HENNESSY and ALAN JEFFREY†

*School of Cognitive and Computing Sciences, University of Sussex,
Brighton, UK*

Abstract

Concurrent ML (CML) is an extension of Standard ML of New Jersey with concurrent features similar to those of process algebra. In this paper, we build upon John Reppy's reduction semantics for CML by constructing a compositional operational semantics for a fragment of CML, based on higher-order process algebra. Using the operational semantics we generalise the notion of weak bisimulation equivalence to build a semantic theory of CML. We give some small examples of proofs about CML expressions, and show that our semantics corresponds to Reppy's up to weak first-order bisimulation.

Capsule Review

This paper is a very nice exercise in giving a solid theoretical understanding of a reasonably important programming language. CML provides a powerful form of concurrency and communication, integrated with ML's higher-order features; it is probably more comprehensible and practical than, say, the pi-calculus. This paper gives a theoretical treatment of CML, fairly familiar to readers versed in bisimulation. Indeed, the treatment of the familiar problem that weak bisimulations are not congruences with respect to choice is addressed in a novel way, using sensitive and insensitive relations, which may be applicable to other calculi as well.

1 Introduction

There have been various attempts to extend standard programming languages with concurrent or distributed features (Giacalone *et al.*, 1989; Holmström, 1983; Nikhil, 1990). Concurrent ML (CML) (Reppy, 1991a; Reppy, 1992; Panangaden and Reppy, 1996) is a practical and elegant example. The language Standard ML is extended with two new type constructors, one for generating communication channels, and the other for delayed computations, and a new function for spawning concurrent threads of computation. Thus the language has all the functional and higher-order features of ML, but in addition programs also have the ability to communicate with each other by transmitting values along communication channels.

* William Ferreira was funded by a CASE studentship from British Telecom.

† This work is carried out in the context of EC BRA 7166 CONCUR 2.

In Reppy (1992), a reduction style operational semantics is given for a subset of CML called λ_{cv} , which may be viewed as a concurrent version of the call-by-value λ -calculus of Plotkin (1975). Reppy's semantics gives reduction rules for whole programs, not for program fragments. It is not *compositional*, in that the semantics of a program is not defined in terms of the semantics of its subterms. Reppy's semantics is designed to prove properties about programs (for example type safety), and not about program fragments (for example equational reasoning).

In this paper we construct a compositional operational semantics in terms of a labelled transition system, for a core subset of CML which we call μ CML. This semantics not only describes the evaluation steps of programs, as in Reppy (1992), but also their communication potentials in terms of their ability to input and output values along communication channels. This semantics extends the semantics of higher-order processes (Thomsen, 1995) with types and first-class functions.

We then proceed to demonstrate the usefulness of this semantics by using it to define a version of *weak bisimulation* (Milner, 1989), suitable for μ CML. We prove that, modulo the usual problems associated with the choice operator of CCS, our chosen equivalence is preserved by all μ CML contexts and therefore may be used as the basis for reasoning about CML programs. In this paper, we do not investigate in detail the resulting theory but confine ourselves to pointing out some of its salient features; for example, standard identities one would expect of a call-by-value λ -calculus are given, and we also show that certain algebraic laws common to process algebras (Milner, 1989) hold.


We now explain in more detail the contents of the remainder of the paper.

In section 2 we describe μ CML, a monomorphically typed core subset of CML, which nonetheless includes base types for channel names, booleans and integers, and type constructors for pairs, functions and delayed computations which are known as *events*. μ CML also includes a selection of the constructs and constants for manipulating event types, such as *transmit* and *receive* for constructing basic events for sending and receiving values, *wrap* for combining delayed computations, *choose* for selecting between delayed computations, and a function *spawn* for launching new concurrent threads of computation within a program. The major omission is that μ CML has no facility for generating new channel names. However, we believe that this can be remedied by using techniques common to the π -calculus (Milner, 1991; Milner *et al.*, 1992; Sangiorgi, 1992).

In the remainder of this section we present the operational semantics of μ CML in terms of a labelled transition system. To describe all possible states which can arise during the computation of a well-typed μ CML program, we need to extend the language. This extension is twofold. The first consists in adding the constants of event type used by Reppy (1992) to define λ_{cv} , i.e. constants to denote certain delayed computations. This extended language, which we call μ CML^{cv}, essentially coincides with the λ_{cv} , the language used by Reppy (1992), except for the omissions cited above. However, to obtain a compositional semantics we make further extensions to μ CML^{cv}. We add a parallel operator \parallel , commonly used in process algebras, which allows us to use programs in place of the multisets of programs of Reppy (1992).

The final addition is more subtle; we include in μ CML^{cv} expressions which

correspond to the synced versions of Reppy’s constants for representing delayed computations. Thus the labelled transition system uses as states programs from a language which we call μCML^+ . This language is a superset of μCML^{cv} , which is our version of Reppy’s λ_{cv} , which in turn is a superset of μCML , our mini-version of CML. The following diagram indicates the relationships between these languages:



In section 3 we discuss semantic equivalences defined on the labelled transition of section 2. We demonstrate the inadequacies of the obvious adaptations of *strong* and *weak* bisimulation equivalence (Milner, 1989), and then consider adaptations of *higher-order* and *irreflexive* bisimulations from Thomsen (1995). Finally, we suggest a new variation called *hereditary* bisimulation equivalence, which overcomes some of the problems encountered with using higher-order and irreflexive bisimulations.

In section 4 we show that hereditary bisimulation is preserved by all μCML contexts. This is an application of the proof method originally suggested by Howe (1989), but the proof is further complicated by the fact that hereditary bisimulations are defined in terms of pairs of relations satisfying mutually dependent properties.

In section 5 we briefly discuss the resulting algebraic theory of μCML expressions. This paper is intended only to lay the foundations of this theory, and so here we simply indicate that our theory extends both that of call-by-value λ -calculus (Plotkin, 1975) and process algebras (Milner, 1989).

In section 6 we show that, up to weak bisimulation equivalence, our semantics coincides with the reduction semantics for λ_{cv} presented in Reppy (1992). This technical result applies only to the common sub-language, namely μCML^{cv} .

In section 7 we briefly consider other approaches to the semantics of CML and related languages, and we end with some suggestions for further work.

2 The language

In this section we introduce our language μCML , a subset of Concurrent ML (Reppy, 1991a; Reppy, 1992; Panangaden and Reppy, 1996). We describe the syntax, including a typing system, and an operational semantics in terms of a labelled transition system.

Unfortunately, there is not enough space in this paper to provide an introduction to programming in CML: see Panangaden and Reppy (1996) for a discussion of the design and philosophy of CML.

The type expressions for our language are given by:

$$A ::= \text{unit} \mid \text{bool} \mid \text{int} \mid A \text{ chan} \mid A * A \mid A \rightarrow A \mid A \text{ event}$$

Thus we have three base types, unit, bool and int; the latter two are simply examples of useful base types and one could easily include more. These types are closed under

four constructors: pairing, function space, and the less common *chan* and *event* type constructors. Our language may be viewed as a typed λ -calculus augmented with the type constructors A *chan* for *communication channels* sending and receiving data of type A , and A *event* for constructing *delayed computations* of type A .

Let $Chan_A$ be a type-indexed family of disjoint sets of channel names, ranged over by k , and let Var denote a set of variables ranged over by x, y and z . The expressions of μ CML are given by the following abstract syntax:

$$\begin{aligned} e, f, g \in Exp & ::= v \mid ce \mid \text{if } e \text{ then } e \text{ else } e \mid (e, e) \mid \text{let } x = e \text{ in } e \mid ee \\ v, w \in Val & ::= \text{fix}(x = \text{fn } y \Rightarrow e) \mid x \mid \text{true} \mid \text{false} \mid k \mid () \mid 0 \mid 1 \mid \dots \\ c \in Const & ::= \text{fst} \mid \text{snd} \mid \text{add} \mid \text{mul} \mid \text{leq} \mid \text{transmit} \mid \text{receive} \\ & \quad \mid \text{choose} \mid \text{spawn} \mid \text{sync} \mid \text{wrap} \mid \text{never} \mid \text{always} \end{aligned}$$

The main syntactic category is that of Exp which look very much like the set of expressions for an applied *call-by-value* version of the λ -calculus. There are the usual pairing, let-binding and branching constructors, and two forms of application: the application of one expression to another, ee , the application of a constant to an expression, ce .

There is also a syntactic category of *value* expressions Val , used in giving a semantics to call-by-value functions and communicate-by-value channels. They are restricted in form: either a variable, a recursively defined function, $\text{fix}(x = \text{fn } y \Rightarrow e)$, or a predefined literal value for the base types. We will use some syntax sugar, writing $\text{fn } y \Rightarrow e$ for $\text{fix}(x = \text{fn } y \Rightarrow e)$ when x does not occur in e , and $e;f$ for $\text{let } x = e \text{ in } f$ when x does not occur in f .

Finally, there is a small collection of constant functions, consisting of a representative sample of constants for manipulating objects of base type, $\text{add}, \text{mul}, \text{leq}$, which could easily be extended, the projection functions fst and snd , together with the set of constants for manipulating *delayed computations* taken directly from Reppy (1992):

- transmit and receive , for constructing delayed computations which can send and receive values,
- choose , for constructing alternatives between delayed computations,
- spawn , for spawning new computational threads,
- sync , for launching delayed computations,
- wrap , for combining delayed computations,
- never , for a delayed computation which always deadlocks, and
- always , for a delayed computation which immediately terminates with a value.

Note that there is no method for generating channel names other than using the predefined set of names $Chan_A$.

There are two constructs in the language which bind occurrences of variables, $\text{let } x = e_1 \text{ in } e_2$ where free occurrences of x in e_2 are bound and $\text{fix}(x = \text{fn } y \Rightarrow e)$ where free occurrences of both x and y in e are bound. We will not dwell on the precise definitions of free and bound variables, but simply use $fv(e)$ to denote the set of variables which have free occurrences in e . If $fv(e) = \emptyset$ then e is said to be

fst	: $A * B \rightarrow A$	transmit	: $A \text{ chan} * A \rightarrow \text{unit event}$
snd	: $A * B \rightarrow B$	receive	: $A \text{ chan} \rightarrow A \text{ event}$
add	: $\text{int} * \text{int} \rightarrow \text{int}$	choose	: $A \text{ event} * A \text{ event} \rightarrow A \text{ event}$
mul	: $\text{int} * \text{int} \rightarrow \text{int}$	spawn	: $(\text{unit} \rightarrow \text{unit}) \rightarrow \text{unit}$
leq	: $\text{int} * \text{int} \rightarrow \text{bool}$	wrap	: $A \text{ event} * (A \rightarrow B) \rightarrow B \text{ event}$
sync	: $A \text{ event} \rightarrow A$	never	: $\text{unit} \rightarrow A \text{ event}$
always	: $A \rightarrow A \text{ event}$		

Fig. 1a. Type rules for μCML constant functions.

$$\begin{array}{c}
 \frac{}{\Gamma, x : A \vdash x : A} \quad \frac{\Gamma \vdash y : B}{\Gamma, x : A \vdash y : B} [x \neq y] \\
 \\
 \frac{}{\Gamma \vdash \text{true} : \text{bool}} \quad \frac{}{\Gamma \vdash \text{false} : \text{bool}} \quad \frac{}{\Gamma \vdash k : A \text{ chan}} [k \in \text{Chan}_A] \\
 \\
 \frac{}{\Gamma \vdash () : \text{unit}} \quad \frac{}{\Gamma \vdash n : \text{int}} \quad \frac{\Gamma, x : A \rightarrow B, y : A \vdash e : B}{\Gamma \vdash \text{fix}(x = \text{fn } y \Rightarrow e) : A \rightarrow B} \\
 \\
 \frac{\Gamma \vdash e : A}{\Gamma \vdash ce : B} [c : A \rightarrow B] \quad \frac{\Gamma \vdash e : A \rightarrow B \quad \Gamma \vdash f : A}{\Gamma \vdash ef : B} \quad \frac{\Gamma \vdash e : A \quad \Gamma \vdash f : B}{\Gamma \vdash (e, f) : A * B} \\
 \\
 \frac{\Gamma \vdash e : \text{bool} \quad \Gamma \vdash f : A \quad \Gamma \vdash g : A}{\Gamma \vdash \text{if } e \text{ then } f \text{ else } g : A} \quad \frac{\Gamma \vdash e : A \quad \Gamma, x : A \vdash f : B}{\Gamma \vdash \text{let } x = e \text{ in } f : B}
 \end{array}$$

Fig. 1b. Type rules for μCML expressions.

a *closed* expression, which we sometimes refer to as a *program*. We also use the standard notation of $e[v/x]$ to denote the substitution of the value v for all free occurrences of x in e where bound names may be changed to avoid the capture of free variables in v . (Since we are modelling a call-by-value language, we have limited substitution to values $e[v/x]$ rather than the more general $e[f/x]$. To model alpha-conversion, we have therefore included variables as possible values.)

We now examine briefly the type system for this language. The types for the constant functions of the language are given in figure 1; this is in agreement with the typing rules given in Reppy (1992) for λ_{cv} . Note that many of the constants (such as $\text{choose} : A \text{ event} * A \text{ event} \rightarrow A \text{ event}$) have a family of types.

This assignment of types to constant functions is used to infer types for arbitrary expressions in the standard way, using a type inference system. A *typing judgement* $\Gamma \vdash e : A$ consists of a *type assignment* Γ , an expression e and a type A such that $fv(e) \subseteq \{x_1, \dots, x_n\}$. A *type assignment* is a sequence of the form $x_1 : t_1, \dots, x_n : t_n$, where each t_i is a type. Intuitively a type judgement should be read as ‘in the type assignment Γ the expression e has type A ’. The type inference system is given in figure 1 and is straightforward. There are two structural rules, literals are assigned their natural types, while the types of functional values are inferred using a minor modification of the standard rule for functional abstractions. The remaining constructs are also handled using standard inference rules (Gunter, 1992).

We now turn our attention to the operational semantics. In Reppy (1992) and

Berry *et al.* (1992) a reduction semantics is given to λ_{cv} , and since μCML^{cv} is a subset of λ_{cv} , this induces a reduction semantics for μCML^{cv} ; this is discussed in full in section 6. The judgements in this reduction semantics are of the form:

$$C \xrightarrow{\tau} C'$$

where C, C' are configurations which combine a closed expression with a run-time environment necessary for its evaluation, and τ is Milner's notation for a silent action. However this semantics is not compositional as the reductions of an expression can not be deduced directly from the reductions of its constituent components. Here we give a compositional operational semantics with four kinds of judgements:

- $e \xrightarrow{\tau} e'$, representing a one step evaluation or reduction,
- $e \xrightarrow{\sqrt{v}} e'$, representing the production of the value v , with a side effect e' ,
- $e \xrightarrow{k?x} e'$, representing the potential to input a value x along the channel k , and
- $e \xrightarrow{k!v} e'$, representing the output of the value v along the channel k .

These are formally defined in figure 2, but we first give an informal overview. To define these relations we introduce extra syntactic constructs. These are introduced as required in the overview but are summarized in figure 2.

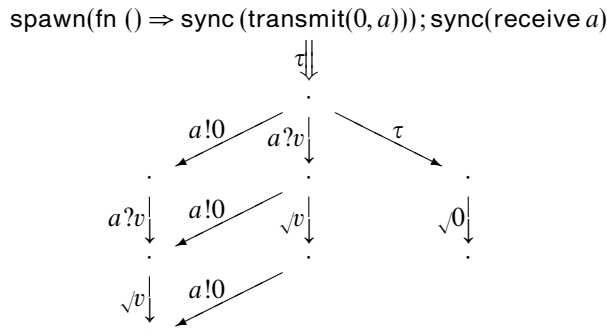
The rules for one step evaluation or reduction have much in common with those for a standard call-by-value λ -calculus. But in addition, a closed expression e of type A should evaluate to a value of type A , and it is this production of values which is the subject of the second kind of judgement. However μCML expressions can spawn subprocesses before returning a value, so we have to allow expressions to continue evaluation even after they have returned a result. For example in the expression:

$$\text{spawn}(\text{fn } () \Rightarrow \text{sync}(\text{transmit}(0, a))); \text{sync}(\text{receive } a)$$

one possible reduction is (where $\xRightarrow{\tau}$ indicates a sequence of τ -reductions):

$$\text{spawn}(\text{fn } () \Rightarrow \text{sync}(\text{transmit}(0, a))); \text{sync}(\text{receive } a) \xRightarrow{\tau} \xrightarrow{a?1} \xrightarrow{\sqrt{1}} \xrightarrow{a!0}$$

where the process returns the value 1 before outputting 0. For this reason we need a reduction $e \xrightarrow{\sqrt{v}} e'$ rather than the more usual termination $e \downarrow v$. The following diagram illustrates all of the possible transitions from this expression:



When giving an operational semantics to a language with side-effects, there are two standard approaches to retaining the information necessary to interpret them. The

first, used for example in Berry *et al.* (1992) and Reppy (1992), is to define a notion of *state* or *configuration*; these contain the program being evaluated together with auxiliary state information, and the judgements of the operational semantics apply to these configurations. The second, more common in work on process algebras (Bergstra and Klop, 1985; Milner, 1989), extends the syntax of the language being interpreted to encompass configurations. We choose the latter approach and one extra construct we add to the language is a parallel operator, $e \parallel f$. This has the same operational rules as in CCS, allowing reduction of both processes:

$$\frac{e \xrightarrow{\alpha} e'}{e \parallel f \xrightarrow{\alpha} e' \parallel f} \quad \frac{f \xrightarrow{\alpha} f'}{e \parallel f \xrightarrow{\alpha} e \parallel f'}$$

and communication between the processes:

$$\frac{e \xrightarrow{k!v} e' \quad f \xrightarrow{k?x} f'}{e \parallel f \xrightarrow{\tau} e' \parallel f'[v/x]} \quad \frac{e \xrightarrow{k?x} e' \quad f \xrightarrow{k!v} f'}{e \parallel f \xrightarrow{\tau} e'[v/x] \parallel f'}$$

The asymmetry is introduced by termination (a feature missing from CCS). A CML process has a *main thread of control*, and only the main thread can return a value. By convention, we write the main thread on the right, so the rule is:

$$\frac{f \xrightarrow{\lambda v} f'}{e \parallel f \xrightarrow{\lambda v} e \parallel f'}$$

There is no corresponding symmetric rule. For example:

$$() \parallel 1 \xrightarrow{\lambda 1} () \parallel \Lambda \quad () \parallel 1 \not\xrightarrow{\lambda ()} \Lambda \parallel 1$$

Since the only difference between concurrent processes is which term can return a value, concurrency is associative and symmetric on the left, so $e \parallel f \parallel g$ is bisimilar to $f \parallel e \parallel g$. In general, we can regard $n + 1$ processes in parallel:

$$e_1 \parallel \dots \parallel e_n \parallel f$$

as being a multiset of spawned threads e_1, \dots, e_n plus one main thread of control f , corresponding to the use of multi-sets in the reduction semantics of Berry *et al.* (1992) and Reppy (1992).

Concurrent processes are generated using the constant application $\text{spawn } e$. A first attempt to write the semantics for $\text{spawn } e$ would be the rule:

$$\text{spawn}(\text{fn } y \Rightarrow e) \xrightarrow{\tau} (\text{fn } y \Rightarrow e) () \parallel ()$$

One step in the evaluation of $\text{spawn}(\text{fn } y \Rightarrow e)$ leads to two expressions running in parallel, one being the spawned function application $(\text{fn } y \Rightarrow e) ()$ and the other the default value $()$ which results from every application of spawn . However, this rule for $\text{spawn } e$ is not general enough. First, it ignores the fact that the expression e may need to perform some computation before returning a function, which is captured by instantiating the static rule for constant application as:

$$\frac{e \xrightarrow{\alpha} e'}{\text{spawn } e \xrightarrow{\alpha} \text{spawn } e'}$$

Secondly, e may have spawned some concurrent processes before returning a func-

tion, and these should carry on evaluation, so we use the silent rule for constant application:

$$\frac{e \xrightarrow{\lambda v} e'}{\text{spawn } e \xrightarrow{\tau} e' \parallel v() \parallel ()}$$

The well-typedness of the operational semantics will ensure that v is a function of the appropriate type, $\text{unit} \rightarrow \text{unit}$.

With this method of representing newly created computation threads more of the rules corresponding to β -reduction in a call-by-value λ -calculus may now be given. To evaluate an application expression ef , first e is evaluated to a value of functional form and then the evaluation of f is initiated. This is represented by the rules:

$$\frac{e \xrightarrow{\alpha} e'}{ef \xrightarrow{\alpha} e'f} \quad \frac{e \xrightarrow{\lambda (fn\ y \Rightarrow g)} e'}{ef \xrightarrow{\tau} e' \parallel \text{let } y = f \text{ in } g}$$

(In fact, we use a slightly more complicated version of the latter rule as functions are allowed to be recursive.) Continuing with the evaluation of ef , we now evaluate f to a value which is then substituted into g for y . This is represented by the two rules:

$$\frac{f \xrightarrow{\tau} f'}{\text{let } x = f \text{ in } g \xrightarrow{\tau} \text{let } x = f' \text{ in } g} \quad \frac{f \xrightarrow{\lambda v} f'}{\text{let } x = f \text{ in } g \xrightarrow{\tau} f' \parallel g[v/x]}$$

The evaluation of the application expression cf is similar; f is evaluated to a value and then the constant c is applied to the resulting value. This is represented by the two rules

$$\frac{f \xrightarrow{\tau} f'}{cf \xrightarrow{\tau} cf'} \quad \frac{f \xrightarrow{\lambda v} f'}{cf \xrightarrow{\tau} f' \parallel \delta(c, v)}$$

Here, borrowing the notation of Reppy (1992), we use the function δ to represent the effect of applying the constant c to the value v . This effect depends upon the constant in question, and we have already seen one instance of this rule, for the constant `spawn`, which result from the fact that $\delta(\text{spawn}, v) = v() \parallel ()$. The definition of δ for all constants in the language is given in figure 2. For the constants associated with the base types this is self-explanatory; the others will be explained below as the constant in question is considered. Note that because of the introduction of \parallel into the language we can treat all constants uniformly, unlike Reppy (1992), where `spawn` and `sync` have to be considered in a special manner.

To implement the standard left-to-right evaluation of pairs of expressions, we introduce a new value $\langle v, w \rangle$ representing a pair which has been fully evaluated. Then, to evaluate (e, f) :

- first allow e to evaluate:

$$\frac{e \xrightarrow{\alpha} e'}{(e, f) \xrightarrow{\alpha} (e', f)}$$

- then when it terminates, start the evaluation of f :

$$\frac{e \xrightarrow{\lambda v} e'}{(e, f) \xrightarrow{\tau} e' \parallel \text{let } x = f \text{ in } \langle v, x \rangle}$$

These value pairs may then be used by being applied to functions of type $A * B$. For

$$\begin{array}{c}
 \frac{e \xrightarrow{\alpha} e'}{c e \xrightarrow{\alpha} c e'} \quad \frac{e \xrightarrow{\alpha} e'}{e f \xrightarrow{\alpha} e' f} \quad \frac{e \xrightarrow{\alpha} e'}{(e, f) \xrightarrow{\alpha} (e', f)} \\
 \frac{e \xrightarrow{\alpha} e'}{\text{if } e \text{ then } f \text{ else } g \xrightarrow{\alpha} \text{if } e' \text{ then } f \text{ else } g} \quad \frac{e \xrightarrow{\alpha} e'}{\text{let } x = e \text{ in } f \xrightarrow{\alpha} \text{let } x = e' \text{ in } f} \\
 \frac{e \xrightarrow{\alpha} e'}{e \parallel f \xrightarrow{\alpha} e' \parallel f} \quad \frac{f \xrightarrow{\alpha} f'}{e \parallel f \xrightarrow{\alpha} e \parallel f'} \quad \frac{f \xrightarrow{\sqrt{v}} f'}{e \parallel f \xrightarrow{\sqrt{v}} e \parallel f'}
 \end{array}$$

Fig. 2a. Operational semantics: static rules.

$$\frac{g e_1 \xrightarrow{\alpha} e}{g e_1 \oplus g e_2 \xrightarrow{\alpha} e} \quad \frac{g e_2 \xrightarrow{\alpha} e}{g e_1 \oplus g e_2 \xrightarrow{\alpha} e} \quad \frac{g e \xrightarrow{\alpha} e}{g e \Rightarrow v \xrightarrow{\alpha} v e}$$

Fig. 2b. Operational semantics: dynamic rules.

$$\begin{array}{c}
 \frac{e \xrightarrow{\sqrt{v}} e'}{c e \xrightarrow{\tau} e' \parallel \delta(c, v)} \quad \frac{e \xrightarrow{\sqrt{\text{true}}} e'}{\text{if } e \text{ then } f \text{ else } g \xrightarrow{\tau} e' \parallel f} \quad \frac{e \xrightarrow{\sqrt{\text{false}}} e'}{\text{if } e \text{ then } f \text{ else } g \xrightarrow{\tau} e' \parallel g} \\
 \frac{e \xrightarrow{\sqrt{v}} e'}{(e, f) \xrightarrow{\tau} e' \parallel \text{let } x = f \text{ in } \langle v, x \rangle} \quad \frac{e \xrightarrow{\sqrt{v}} e'}{e f \xrightarrow{\tau} e' \parallel \text{let } y = f \text{ in } g[v/x]} [v = \text{fix}(x = \text{fn } y \Rightarrow g)] \\
 \frac{e \xrightarrow{\sqrt{v}} e'}{\text{let } x = e \text{ in } f \xrightarrow{\tau} e' \parallel f[v/x]} \quad \frac{e \xrightarrow{k!v} e' \quad f \xrightarrow{k?x} f'}{e \parallel f \xrightarrow{\tau} e' \parallel f'[v/x]} \quad \frac{e \xrightarrow{k?x} e' \quad f \xrightarrow{k!v} f'}{e \parallel f \xrightarrow{\tau} e' \parallel f'[v/x]}
 \end{array}$$

Fig. 2c. Operational semantics: silent rules.

$$\overline{v \xrightarrow{\sqrt{v}} \Lambda} \quad \overline{k!v \xrightarrow{k!v} ()} \quad \overline{k? \xrightarrow{k?x} x} \quad \overline{A v \xrightarrow{\tau} v}$$

Fig. 2d. Operational semantics: axioms.

$$a ::= k!v \mid k?x \quad \alpha ::= a \mid \tau \quad l ::= \alpha \mid \sqrt{v}$$

Fig. 2e. Operational semantics: grammar of labels.

$$\begin{array}{ll}
 \delta(\text{fst}, \langle v, w \rangle) = v & \delta(\text{snd}, \langle v, w \rangle) = w \\
 \delta(\text{add}, \langle m, n \rangle) = m + n & \delta(\text{mul}, \langle m, n \rangle) = m \times n \\
 \delta(\text{leq}, \langle m, n \rangle) = m \leq n & \\
 \\
 \delta(\text{transmit}, \langle k, v \rangle) = [k!v] & \delta(\text{receive}, k) = [k?] \\
 \delta(\text{choose}, \langle [g e_1], [g e_2] \rangle) = [g e_1 \oplus g e_2] & \delta(\text{wrap}, \langle [g e], v \rangle) = [g e \Rightarrow v] \\
 \delta(\text{never}, ()) = [\Lambda] & \delta(\text{always}, v) = [A v] \\
 \delta(\text{spawn}, v) = v() \parallel () & \delta(\text{sync}, [g e]) = g e
 \end{array}$$

Fig. 2f. Operational semantics: reduction of constants.

$$\begin{aligned}
e, f, g \in \text{Exp} & ::= v \mid ce \mid \text{if } e \text{ then } e \text{ else } e \mid (e, e) \mid \text{let } x = e \text{ in } e \mid ee \\
v, w \in \text{Val} & ::= \text{fix}(x = \text{fn } y \Rightarrow e) \mid x \mid \text{true} \mid \text{false} \mid k \mid () \mid 0 \mid 1 \mid \dots \\
c \in \text{Const} & ::= \text{fst} \mid \text{snd} \mid \text{add} \mid \text{mul} \mid \text{leq} \mid \text{transmit} \mid \text{receive} \\
& \quad \mid \text{choose} \mid \text{spawn} \mid \text{sync} \mid \text{wrap} \mid \text{never} \mid \text{always}
\end{aligned}$$

Fig. 3a. Syntax of μCML .

$$\begin{aligned}
v, w \in \text{Val} & ::= \dots \mid \langle v, v \rangle \mid [ge] \\
ge \in \text{GExp} & ::= v!v \mid v? \mid ge \Rightarrow v \mid ge \oplus ge \mid \Lambda \mid \mathbf{A}v
\end{aligned}$$

Fig. 3b. Syntax of μCML^{cv} .

$$e, f, g \in \text{Exp} ::= \dots \mid ge \mid e \dashv\dashv e$$

Fig. 3c. Syntax of μCML^+ .

example the following inferences result from the definition of the function δ for the constants `fst` and `mul`:

$$\frac{e \xrightarrow{\langle v, w \rangle} e'}{\text{fst } e \xrightarrow{\tau} e' \dashv\dashv v} \quad \frac{e \xrightarrow{\langle m, n \rangle} e'}{\text{mul } e \xrightarrow{\tau} e' \dashv\dashv m \times n}$$

It remains to explain how *delayed computations*, i.e. programs of type A event, are handled. It is important to realize that expressions of type A event represent *potential* rather than actual computations and this potential can only be activated by an application of the constant `sync`, of type A event $\rightarrow A$. Thus, for example, the expression `receive k` is of type A event, and represents a delayed computation which has the potential to receive a value of type A along the channel k . The expression `sync(receive k)` can actually receive such a value v along channel k , or more accurately, can evaluate to such a value, provided some other computation thread can send the value along channel k .

The semantics of `sync` is handled by introducing a new constructor for values. For certain kinds of expressions ge of type A , which we call *guarded expressions*, let $[ge]$ be a value of type A event; this represents a *delayed computation* which when launched initiates a new computation thread which evaluates the expression ge . Then the expression `sync $[ge]$` reduces in one step to the expression ge . More generally the evaluation of the expression `sync e` proceeds as follows:

- First evaluate e until it can produce a value:

$$\frac{e \xrightarrow{\tau} e'}{\text{sync } e \xrightarrow{\tau} \text{sync } e'}$$

- then launch the resulting delayed computation:

$$\frac{e \xrightarrow{[ge]} e'}{\text{sync } e \xrightarrow{\tau} e' \dashv\dashv ge}$$

Note that here, as always, the production of a value may have as a side-effect the generation of a new computation thread e' and this is launched concurrently with the delayed computation ge . Also, both of these rules are instances of more general rules already considered. The first is obtained from the rule for the evaluation of applications of the form ce , and the second by defining $\delta(\text{sync}, [ge])$ to be ge .

The precise syntax for guarded expressions will emerge by considering what types of values of the form $[e]$ can result from the evaluation of expressions of type event from the basic language μCML . The constant receive is of type $A \text{ chan} \rightarrow A \text{ event}$ and therefore the evaluation of the expression $\text{receive } e$ proceeds by first evaluating e to a value of type $A \text{ chan}$ until it returns a value k , and then returning a delayed computation consisting of an event which can receive any value of type A on the channel k . To represent this event, we extend the syntax further by letting $k?$ be a guarded expression for any k and A , with the associated rule:

$$\frac{e \xrightarrow{k} e'}{\text{receive } e \xrightarrow{\tau} e' \parallel [k?]}$$

The construct transmit is handled in a similar manner, using guarded expressions of the form $k!v$:

$$\frac{e \xrightarrow{\langle k,v \rangle} e'}{\text{transmit } e \xrightarrow{\tau} e' \parallel [k!v]}$$

It is these two new expressions $k?$ and $k!v$ which perform communication between computation threads. Formally, $k!v$ is of type unit and we have the axiom:

$$\overline{k!v \xrightarrow{k!v} ()}$$

Intuitively, this may be read as $k!v$ evaluates in one step to the expression $()$ and this evaluation has as a side effect the transmission of the value v to the channel k . The semantics we consider for input is the *late* semantics, where the reduction rule binds a new variable x :

$$\overline{k? \xrightarrow{k?x} x}$$

Therefore, in general input moves are of the form $e \xrightarrow{k?x} f$ where $\vdash e : B$ and $x : A \vdash f : B$. Communication can now be modelled as in CCS by the simultaneous occurrence of input and output actions:

$$\frac{e \xrightarrow{k?x} e' \quad f \xrightarrow{k!v} f'}{e \parallel f \xrightarrow{\tau} e'[v/x] \parallel f'}$$

There remain four constructs for *delayed computations* to be explained. The first, never of type $\text{unit} \rightarrow A \text{ event}$, is handled by the introduction of the guarded expression Λ , representing a deadlocked evaluation, together with the inference rule:

$$\frac{e \xrightarrow{\langle \rangle} e'}{\text{never } e \xrightarrow{\tau} e' \parallel [\Lambda]}$$

obtained, once more, by defining $\delta(\text{never}, ())$ to be $[\Lambda]$.

The constant wrap is of type $A \text{ event} * (A \rightarrow B) \rightarrow B \text{ event}$. The evaluation of $\text{wrap } e$ proceeds in the standard way by evaluating e until it produces a value, which must be of the form $\langle [ge], v \rangle$, where ge is a guarded expression of type A and v has

type $A \rightarrow B$. Then the evaluation of $\text{wrap } e$ continues by the construction of the new *delayed computation* $[ge \Rightarrow v]$. Bearing in mind the fact that the production of values can generate new computation threads, this is formally represented by the inference rule:

$$\frac{e \xrightarrow{\langle [ge], v \rangle} e'}{\text{wrap } e \xrightarrow{\tau} e' \Vdash [ge \Rightarrow v]}$$

The guarded expression $ge \Rightarrow v$ is a *wrapper* which applies v to the result of evaluating ge :

$$\frac{ge \xrightarrow{\alpha} e}{ge \Rightarrow v \xrightarrow{\alpha} v e}$$

The *always* construct, of type $A \rightarrow A$ event, evaluates its argument to a value v , and then returns a trivial delayed computation; this computation, when activated, immediately evaluates to the value v . In order to represent these trivial computations we introduce a new constructor for guarded expressions, \mathbf{A} and the semantics of *always* is then captured by the rule:

$$\frac{e \xrightarrow{\langle v \rangle} e'}{\text{always } e \xrightarrow{\tau} e' \Vdash [\mathbf{A}v]}$$

Since $\mathbf{A}v$ immediately evaluates to the constant v we have:

$$\overline{\mathbf{A}v \xrightarrow{\tau} v}$$

The choice construct $\text{choose } e$ is a choice between *delayed computations* as choose has the type $A \text{ event} * A \text{ event} \rightarrow A \text{ event}$. To interpret it we introduce a new choice constructor $ge_1 \oplus ge_2$ where ge_1 and ge_2 are guarded expressions of the same type. Then $\text{choose } e$ proceeds by evaluating e until it can produce a value, which must be of the form $\langle [ge_1], [ge_2] \rangle$, and the evaluation continues by constructing the *delayed computation* $[ge_1 \oplus ge_2]$. This is represented by the rule:

$$\frac{e \xrightarrow{\langle [ge_1], [ge_2] \rangle} e'}{\text{choose } e \xrightarrow{\tau} e' \Vdash [ge_1 \oplus ge_2]}$$

The notation \oplus , introduced by Reppy (1992), is unfortunate, as it is used in Hennessy (1988) to represent the *internal choice* between processes, whereas here it represents *external choice*: we have the following auxiliary rules, which are the same as CCS summation:

$$\frac{ge_1 \xrightarrow{\alpha} e}{ge_1 \oplus ge_2 \xrightarrow{\alpha} e} \quad \frac{ge_2 \xrightarrow{\alpha} e}{ge_1 \oplus ge_2 \xrightarrow{\alpha} e}$$

This ends our informal description of the operational semantics of μCML . We now summarize, giving the precise definitions of the new syntax. For the purposes of comparison with the reduction semantics of λ_{cv} (Reppy, 1992), it is convenient to view the extension to μCML in two stages. The first is obtained by adding the new syntactic category of guarded expressions, and two new constructors for values:

$$\begin{aligned} v \in \text{Val} & ::= \dots \mid \langle v, v \rangle \mid [ge] \\ ge \in \text{GExp} & ::= v!v \mid v? \mid ge \Rightarrow v \mid ge \oplus ge \mid \Lambda \mid \mathbf{A}v \end{aligned}$$

The resulting language we call μCML^{cv} , as it corresponds very closely to Reppy's

$$\begin{array}{c}
 \frac{\Gamma \vdash v : A \quad \Gamma \vdash w : B}{\Gamma \vdash \langle v, w \rangle : A * B} \quad \frac{\Gamma \vdash ge : A}{\Gamma \vdash [ge] : A \text{ event}} \\
 \frac{\Gamma \vdash v : A \text{ chan} \quad \Gamma \vdash w : A}{\Gamma \vdash v!w : \text{unit}} \quad \frac{\Gamma \vdash v : A \text{ chan}}{\Gamma \vdash v? : A} \quad \frac{\Gamma \vdash ge : A \quad \Gamma \vdash v : A \rightarrow B}{\Gamma \vdash ge \Rightarrow v : B} \\
 \frac{\Gamma \vdash ge_1 : A \quad \Gamma \vdash ge_2 : A}{\Gamma \vdash ge_1 \oplus ge_2 : A} \quad \frac{}{\Gamma \vdash \Lambda : A} \quad \frac{\Gamma \vdash v : A}{\Gamma \vdash \mathbf{A}v : A} \\
 \frac{\Gamma \vdash e : A \quad \Gamma \vdash f : B}{\Gamma \vdash e \# f : B}
 \end{array}$$

Fig. 4. Type rules for extra μCML^+ constructs.

λ_{cv} . A precise comparison is given in section 6. The final language, μCML^+ , is obtained by extending μCML^{cv} with:

$$e \in \text{Exp} ::= \dots \mid ge \mid e \# e$$

and type judgements for all the extra constructs appear in figure 4.

The operational semantics is given as a set of transition relations over closed expressions from μCML^+ . These transition relations have as labels *Label*:

$$a ::= k!v \mid k?x \quad \alpha ::= a \mid \tau \quad l ::= \alpha \mid \surd v$$

which are typed with judgements $\vdash l : A$ in figure 5, and are defined to be the least relations satisfying the rules in figure 2. The rules are divided into three parts. The first gives the set of context rules, showing when moves may be propagated through certain contexts; the second give the reduction rules while the third contains the axioms.

It is worth pointing out that the context rules are asymmetric for the propagation of value production though the context $\#$; in $e \# f$ only the computation thread f can produce a value. This is in agreement with the reduction semantics of Reppy (1992) where in a given state represented by a multi-set of expressions only one distinguished expression is allowed to produce a value. Also in the rule for application, the evaluation of ef is somewhat more complicated than previously stated; values of functional type all involve the fix point operator and these fix points are automatically unfolded at the point of application.

We end this section with a Subject Reduction Theorem for our semantics:

Theorem 2.1

For every closed expression $\vdash e : A$ in μCML^+

- if $e \xrightarrow{\tau} e'$ then $\vdash e' : A$,
- if $e \xrightarrow{\surd v} e'$ then $\vdash e' : A$ and $\vdash v : A$,
- if $e \xrightarrow{k?x} e'$ and $k \in \text{Chan}_B$ then $x : B \vdash e' : A$, and
- if $e \xrightarrow{k!v} e'$ and $k \in \text{Chan}_B$ then $\vdash e' : A$ and $\vdash v : B$.

Proof

By rule induction on the inferences. \square

$$\frac{}{\Gamma \vdash \tau : A} \quad \frac{\Gamma \vdash v : A}{\Gamma \vdash \surd v : A} \quad \frac{}{\Gamma \vdash k?x : A} \quad \frac{\Gamma \vdash w : B}{\Gamma \vdash k!w : A} [k \in Chan_B]$$

Fig. 5. Type rules for labels.

3 Weak bisimulation equivalence

In this section we demonstrate the usefulness of our operational semantics by providing μCML^+ with an appropriate version of bisimulation equivalence. We discuss a range of possible bisimulation based equivalences, and eventually propose a new variation called *hereditary bisimulation equivalence*, which we feel is most suited to μCML^+ .

We first show how to adapt the notion of strong bisimulation equivalence to μCML^+ . Since our language is typed, it is more convenient to define the equivalence in terms of *type-indexed* families of relations. Moreover, since the operational semantics uses actions of the form $e \xrightarrow{k?x} f$ where f may be an open expression we need to consider relations over open expressions. Let an *open type-indexed* relation \mathcal{R} be a family of relations $\mathcal{R}_{\Gamma, A}$ such that if $e \mathcal{R}_{\Gamma, A} f$ then $\Gamma \vdash e : A$ and $\Gamma \vdash f : A$. We will often elide the subscripts from relations, for example writing $e \mathcal{R} f$ for $e \mathcal{R}_{\Gamma, A} f$ when context makes the type obvious. Let a *closed type-indexed* relation \mathcal{R} be an open type-indexed relation where Γ is everywhere the empty context, and can therefore be elided. For any closed type-indexed relation \mathcal{R} , let its *open extension* \mathcal{R}° be defined as:

$$e \mathcal{R}_{\vec{x}:\vec{A}, B}^\circ f \text{ iff } e[\vec{v}/\vec{x}] \mathcal{R}_B f[\vec{v}/\vec{x}] \text{ for all } \vdash \vec{v} : \vec{A}.$$

A closed type-indexed relation \mathcal{R} is *structure preserving* iff:

- if $v \mathcal{R}_A w$ and A is a base type then $v = w$,
- if $\langle v_1, v_2 \rangle \mathcal{R}_{A_1 * A_2} \langle w_1, w_2 \rangle$ then $v_i \mathcal{R}_{A_i} w_i$,
- if $[ge_1] \mathcal{R}_{A \text{ event}} [ge_2]$ then $ge_1 \mathcal{R}_A ge_2$, and
- if $v \mathcal{R}_{A \rightarrow B} v'$ then for all $\vdash w : A$ we have $v w \mathcal{R}_B v' w$.

With this notation we can now define strong bisimulations over μCML^+ expressions. A closed type-indexed relation \mathcal{R} is a *first-order strong simulation* iff it is structure-preserving and the following diagram can be completed:

$$\begin{array}{ccc} e_1 & \mathcal{R} & e_2 \\ \downarrow l & & \downarrow l \\ e'_1 & & e'_2 \end{array} \quad \text{as} \quad \begin{array}{ccc} e_1 & \mathcal{R} & e_2 \\ \downarrow l & & \downarrow l \\ e'_1 & \mathcal{R}^\circ & e'_2 \end{array}$$

Note the use of the open extension \mathcal{R}° . This means, for example, that if $e_1 \mathcal{R} e_2$ we require that the move $e_1 \xrightarrow{k?x} f_1$ be matched by a move $e_2 \xrightarrow{k?x} f_2$ where f_2 is such that for all values v of the appropriate type $f_1[v/x] \mathcal{R} f_2[v/x]$. Thus, in the terminology of Milner *et al.* (1992) our definition corresponds to the *late* version of bisimulation. (An alternative would be *early* bisimulation where input moves

are labelled with closed values rather than variables. This is computationally more appealing, but it is an open problem whether the techniques of the next section can be applied to early bisimulation.)

\mathcal{R} is a *first-order strong bisimulation* iff \mathcal{R} and \mathcal{R}^{-1} are first-order strong simulations. Let \sim^1 be the largest first-order strong bisimulation.

Proposition 3.1

\sim^1 is an equivalence.

Proof

Use diagram chases to show that if \mathcal{R} is a first-order strong simulation then so are the identity relation I and the relation composition $\mathcal{R}\mathcal{R}$. The result follows. \square

Unfortunately, \sim^1 is not a congruence for μCML^+ , since we have:

$$\text{add}(1, 2) \sim^1 \text{add}(2, 1)$$

however, sending the thunked expressions on channel k we get:

$$k!(\text{fn } x \Rightarrow \text{add}(1, 2)) \not\sim^1 k!(\text{fn } x \Rightarrow \text{add}(2, 1))$$

since the definition of strong bisimulation demands that the actions performed by expressions match up to syntactic identity. This counter-example can also be reproduced using only μCML contexts:

$$\text{sync}(\text{transmit}(k, \text{fn } x \Rightarrow \text{add}(1, 2))) \not\sim^1 \text{sync}(\text{transmit}(k, \text{fn } x \Rightarrow \text{add}(2, 1)))$$

since the left hand side can perform the move:

$$\text{sync}(\text{transmit}(k, \text{fn } x \Rightarrow \text{add}(1, 2))) \xrightarrow{\tau} \xrightarrow{k!(\text{fn } x \Rightarrow \text{add}(1, 2))} ()$$

but this can only be matched by the right-hand side up to strong bisimulation:

$$\text{sync}(\text{transmit}(k, \text{fn } x \Rightarrow \text{add}(2, 1))) \xrightarrow{\tau} \xrightarrow{k!(\text{fn } x \Rightarrow \text{add}(2, 1))} ()$$

In fact, it is easy to verify that the only first-order strong bisimulation which is a congruence for μCML is the identity relation.

To find a satisfactory treatment of bisimulation for μCML , we need to look to *higher-order bisimulation*, where the structure of the labels is accounted for. To this end, given a closed type-indexed relation \mathcal{R} , define its *extension to labels* \mathcal{R}^l as:

$$\frac{}{\tau \mathcal{R}_A^l \tau} \quad \frac{v \mathcal{R}_A w}{\sqrt{v} \mathcal{R}_A^l \sqrt{w}} \quad \frac{}{k?x \mathcal{R}_A^l k?x} \quad \frac{v \mathcal{R}_B w}{k!v \mathcal{R}_A^l k!w} [k \in \text{Chan}_B]$$

Then \mathcal{R} is a *higher-order strong simulation* iff it is structure-preserving and the following diagram can be completed:

$$\begin{array}{ccc} e_1 & \mathcal{R} & e_2 \\ \downarrow l_1 & & \\ e'_1 & & \end{array} \quad \text{as} \quad \begin{array}{ccc} e_1 & \mathcal{R} & e_2 \\ \downarrow l_1 & & \downarrow l_2 \\ e'_1 & \mathcal{R}^\circ & e'_2 \end{array} \quad \text{where } l_1 \mathcal{R}^l l_2$$

Let \sim^h be the largest higher-order strong bisimulation.

Proposition 3.2

\sim^h is a congruence.

Proof


Use a similar technique to the proof of Proposition 3.1 to show that \sim^h is an equivalence. To show that \sim^h is a congruence, define \mathcal{R} as:

$$\mathcal{R} = \{(C[e], C[f]) \mid e \sim^h f\}$$

and then show by induction on C that \mathcal{R} is a simulation. The result follows. \square

For many purposes, strong bisimulation is too fine an equivalence as it is sensitive to the number of reductions performed by expressions. This means it will not even validate elementary properties of β -reduction such as $(\text{fn } x \Rightarrow x)0 = 0$. We require the coarser *weak bisimulation* which allows τ -actions to be ignored.

This in turn requires some more notation. Let $\xRightarrow{\epsilon}$ be the reflexive transitive closure of $\xrightarrow{\tau}$, and let \xRightarrow{l} be $\xRightarrow{\epsilon} \xrightarrow{l}$ (i.e. any sequence of silent actions followed by an l action). Note that we are *not* allowing silent actions after the l action. Let \xRightarrow{l} be $\xRightarrow{\epsilon}$ if $l = \tau$ and \xRightarrow{l} otherwise. Then \mathcal{R} is a *first-order weak simulation* iff it is structure-preserving and the following diagram can be completed:



Let \approx^1 be the largest first-order weak bisimulation.


Proposition 3.3

\approx^1 is an equivalence.

Proof

Similar to the proof of Proposition 3.1. \square

Unfortunately, \approx^1 is not a congruence, for the same reason as \sim^1 , and so we can attempt the same modification. \mathcal{R} is a *higher-order weak simulation* iff it is structure-preserving and the following diagram can be completed:



Let \approx^h be the largest higher-order weak bisimulation.

Proposition 3.4

\approx^h is an equivalence.

Proof

Similar to the proof of Proposition 3.1. \square

However, \approx^h is still not a congruence, for the usual reason that weak bisimulation equivalence \approx is not a congruence for CCS summation. Recall from Milner (1989) that in CCS $\mathbf{0} \approx \tau.\mathbf{0}$ but $a.\mathbf{0} + \mathbf{0} \not\approx a.\mathbf{0} + \tau.\mathbf{0}$. We can duplicate this counter-example in μCML^+ since the CCS operator $+$ corresponds to the μCML^+ operator \oplus and $\mathbf{0}$ corresponds to Λ . However, \oplus may only be applied to *guarded expressions* and therefore we need a *guarded expression* which behaves like $\tau.\mathbf{0}$; the required expression is $\mathbf{A}[\Lambda] \Rightarrow \text{sync}$. Thus:

$$\Lambda \approx^h \mathbf{A}[\Lambda] \Rightarrow \text{sync}$$

since the right-hand side has only one reduction:

$$\begin{aligned} \mathbf{A}[\Lambda] \Rightarrow \text{sync} \\ \xrightarrow{\tau} \text{sync}[\Lambda] \\ \xrightarrow{\tau} \Lambda \end{aligned}$$

but:

$$\Lambda \oplus k!0 \not\approx^h (\mathbf{A}[\Lambda] \Rightarrow \text{sync}) \oplus k!0$$

because the only reduction of $\Lambda \oplus k!0$ is $\Lambda \oplus k!0 \xrightarrow{k!0} \Lambda \oplus \Lambda$ and:

$$\begin{aligned} (\mathbf{A}[\Lambda] \Rightarrow \text{sync}) \oplus k!0 \\ \xrightarrow{\tau} \text{sync}[\Lambda] \\ \xrightarrow{\tau} \Lambda \end{aligned}$$

This counter-example can also be replicated using the restricted syntax of μCML . We have:

$$\text{never}() \approx^h \text{wrap}(\text{always}(\text{never}()), \text{sync})$$

since the left-hand side has only one reduction:

$$\text{never}() \xrightarrow{\sqrt{[\Lambda]}} \Lambda$$

and the right-hand side can match this with:

$$\text{wrap}(\text{always}(\text{never}()), \text{sync}) \xrightarrow{\sqrt{[\mathbf{A}[\Lambda] \Rightarrow \text{sync}]}} \Lambda$$

and we have seen:

$$\Lambda \approx^h \mathbf{A}[\Lambda] \Rightarrow \text{sync}.$$

However:

$$\begin{aligned} \text{sync}(\text{choose}(\text{never}(), \text{transmit}(k, 0))) \\ \not\approx^h \text{sync}(\text{choose}(\text{wrap}(\text{always}(\text{never}()), \text{sync}), \text{transmit}(k, 0))) \end{aligned}$$

since the left hand side has only one reduction:

$$\begin{aligned} \text{sync}(\text{choose}(\text{never}(), \text{transmit}(k, 0))) \\ \xrightarrow{\tau} \Lambda \oplus k!0 \end{aligned}$$

whereas the right-hand side has the reduction:

$$\begin{aligned} & \text{sync}(\text{choose}(\text{wrap}(\text{always}(\text{never}()), \text{sync}), \text{transmit}(k, 0))) \\ & \xrightarrow{\tau} (\mathbf{A}[\Lambda] \Rightarrow \text{sync}) \oplus k!0 \end{aligned}$$

A first attempt to rectify this is to adapt Milner’s observational equivalence for μCML , and to define $=^h$ as the smallest symmetric relation such that the following diagram can be completed:

$$\begin{array}{ccc} e_1 & =^h & e_2 \\ \downarrow l_1 & & \downarrow l_1 \\ e'_1 & & e'_1 \end{array} \quad \text{as} \quad \begin{array}{ccc} e_1 & =^h & e_2 \\ \downarrow l_1 & & \downarrow l_2 \\ e'_1 & \approx^h & e'_2 \end{array} \quad \text{where } l_1 \approx^{h!} l_2$$

Proposition 3.5
 $=^h$ is an equivalence.

Proof
 Similar to the proof of Proposition 3.1. \square

This attempt fails, however, since it only looks at the first move of a process, and not at the first moves of any processes in its transitions. Thus, the above μCML counter-example for \approx^h being a congruence also applies to $=^h$, i.e.

$$\text{never}() =^h \text{wrap}(\text{always}(\text{never}()), \text{sync})$$

but:

$$\begin{aligned} & \text{sync}(\text{choose}(\text{never}(), \text{transmit}(k, 0))) \\ & \neq^h \text{sync}(\text{choose}(\text{wrap}(\text{always}(\text{never}()), \text{sync}), \text{transmit}(k, 0))) \end{aligned}$$

This failure was first noted in Thomsen (1995) for CHOCS.

Thomsen’s solution to this problem is to require that τ -moves can always be matched by at least one τ -move, which produces his definition of an *irreflexive simulation* as a structure-preserving relation where the following diagram can be completed:

$$\begin{array}{ccc} e_1 & \mathcal{R} & e_2 \\ \downarrow l_1 & & \downarrow l_1 \\ e'_1 & & e'_1 \end{array} \quad \text{as} \quad \begin{array}{ccc} e_1 & \mathcal{R} & e_2 \\ \downarrow l_1 & & \downarrow l_2 \\ e'_1 & \mathcal{R} & e'_2 \end{array} \quad \text{where } l_1 \mathcal{R}^! l_2$$

Let \approx^i be the largest irreflexive bisimulation.

Proposition 3.6
 \approx^i is a congruence.


Proof

The proof that \approx^i is an equivalence is similar to the proof of Proposition 3.1. The proof that it is a congruence is similar to the proof of Theorem 4.7 in the next section. \square


However this relation is rather too strong for many purposes, for example $\text{add}(1, 2) \not\approx^i \text{add}(1, \text{add}(1, 1))$ since the right hand side can perform more τ -moves than the left hand side. This is similar to the problem in CHOCS where $a.\tau.P \not\approx^i a.P$.

To find an appropriate definition of bisimulation for μCML , we observe that μCML only allows \oplus to be used on *guarded expressions*, and not on arbitrary expressions. We can thus ignore the initial τ -moves of all expressions *except* for guarded expressions. For this reason, we have to provide *two* equivalences: one on terms where we are not interested in initial τ -moves, and one on terms where we are.

A pair of closed type-indexed relations $\mathcal{R} = (\mathcal{R}^n, \mathcal{R}^s)$ form a *hereditary simulation* (we call \mathcal{R}^n an *insensitive simulation* and \mathcal{R}^s a *sensitive simulation*) iff \mathcal{R}^s is structure-preserving and we can complete the following diagrams:



and:



Let (\approx^n, \approx^s) be the largest hereditary bisimulation. Note that we require \mathcal{R}^s to be structure-preserving because it is used to compare the labels in transitions, which may contain abstractions or guarded events.

In the operational semantics of μCML expressions, guarded expressions can only appear in labels, and not as the residuals of transitions. This explains why in the definition of \approx^n labels are compared with respect to the sensitive relation \approx^s , whereas the insensitive relation is used for the residuals. For example, if $ge_1 \approx^n \not\approx^s ge_2$ then we have:

$$(\text{fn } x \Rightarrow ge_1) \approx^n (\text{fn } x \Rightarrow ge_2)$$

since once either side is applied to an argument, their first action will be a τ -step. On the other hand:

$$[ge_1] \not\approx^n [ge_2]$$

since $[\]$ is precisely the construct which allows us to embed ge_1 and ge_2 in a \oplus context.

Theorem 3.7


\approx^s is a congruence for μCML^+ , and \approx^n is a congruence for μCML .

Proof

The proof that \approx^s and \approx^n are equivalences is similar to the proof of Proposition 3.1. The proof that they form congruences is the subject of the next section. \square

Proposition 3.8

The equivalences on μCML^+ have the following strict inclusions:



Proof

For each inclusion, show that the first bisimulation satisfies the condition required to be the second form of bisimulation. To show that the inclusions are strict, we use the following examples:

$(\text{fn } x \Rightarrow \text{add}(1, 2))$	$\sim^h \not\sim^1$	$(\text{fn } x \Rightarrow \text{add}(2, 1))$
	$1 \approx^1 \not\sim^1$	$\text{let } x = 1 \text{ in } x$
$\text{choose}(\text{receive } k, \text{tau}(\text{receive } k))$	$\approx^i \not\sim^h$	$\text{tau}(\text{receive } k)$
$\text{add}(1, 2)$	$\approx^s \not\sim^i$	$\text{add}(1, \text{add}(1, 1))$
	$1 \approx^n \not\sim^s$	$\text{let } x = 1 \text{ in } x$
$\text{never}()$	$\approx^h \not\sim^n$	$\text{tau}(\text{never}())$
	$1 \approx^h \neq^h$	$\text{let } x = 1 \text{ in } x$

where:

$$\text{tau} = \text{fn } x \Rightarrow \text{wrap}(\text{always } x, \text{sync})$$

(Note that this settles an open question (Thomsen, 1995) as to whether \approx^i is the largest congruence contained in \approx^h .) \square

It is the operator \oplus which differentiates between the two equivalences \approx^n and \approx^h . However, to demonstrate the difference we need to be able to apply \oplus to guarded expressions which can spontaneously evolve, i.e. perform τ -moves. The only μCML^+ constructor for guarded expressions which allows this is **A**, and in turn occurrences of this can only be generated by the μCML constructor **always**. Therefore:

Proposition 3.9

For the subset of μCML^+ without **always** and **A**, \approx^n is the same as \approx^h , and \approx^s is the same as $=^h$.

Proof

From Proposition 3.8 $\approx^n \subseteq \approx^h$.

For the subset of μCML^+ without always and **A**, define \mathcal{R}^s as:

$$\{(v, w) \mid v \approx^h w\} \cup \{(ge_1, ge_2) \mid ge_1 \approx^h ge_2\} \cup \{(v_1 w, v_2 w) \mid v_1 \approx^h v_2\}$$

Then, since no event without **A** can perform a τ -move, and since the only initial moves of $v_i w$ are β -reductions, we can show that $(\approx^h, \mathcal{R}^s)$ forms an hereditary bisimulation, and so $\approx^h \subseteq \approx^n$. From this it is routine to show that $\approx^s = \approx^h$. \square

Unfortunately, we have not been able to show that \approx^n is the largest μCML congruence contained in weak higher-order bisimulation equivalence. However, we do have the following characterization:

Theorem 3.10

\approx^n is the largest higher-order weak bisimulation which respects μCML contexts.

Proof

By definition, \approx^n is a higher-order weak bisimulation, and we have shown that it respects μCML contexts. All that remains is to show that it is the largest such.

Let \mathcal{R} be a higher-order weak bisimulation which respects μCML contexts. Then define:

$$\begin{aligned} \mathcal{R}^n &= \mathcal{R} \cup \{(v_1 w, e_2) \mid v_1 \mathcal{R} v_2, v_2 w \xrightarrow{\tau} e_2\} \cup \{(e_1, v_2 w) \mid v_1 \mathcal{R} v_2, v_1 w \xrightarrow{\tau} e_1\} \\ \mathcal{R}^s &= \{(v, w) \mid v \mathcal{R} w\} \cup \{(ge_1, ge_2) \mid [ge_1] \mathcal{R} [ge_2]\} \cup \{(v_1 w, v_2 w) \mid v_1 \mathcal{R} v_2\} \end{aligned}$$

We will now show that $(\mathcal{R}^n, \mathcal{R}^s)$ forms an hereditary simulation, from which we can deduce $\mathcal{R} \subseteq \mathcal{R}^n \subseteq \approx^n$.

First, we note that \mathcal{R}^s is structure preserving, and that $\mathcal{R}^{sl} = \mathcal{R}^l$.

Then we show that we can complete the required diagrams for $(\mathcal{R}^n, \mathcal{R}^s)$ to be an hereditary simulation. The only tricky case is if:

$$\begin{array}{ccc} ge_1 & \mathcal{R}^s & ge_2 \\ \downarrow l_1 & & \\ e_1 & & \end{array}$$

in which case, by the definition of \mathcal{R}^s , $[ge_1] \mathcal{R} [ge_2]$, and since \mathcal{R} respects μCML contexts we have (for fresh k):

$$\begin{array}{ccc} \text{choose}([ge_1], \text{receive } k) & \mathcal{R} & \text{choose}([ge_2], \text{receive } k) \\ \Downarrow \sqrt{[ge_1 \oplus k?]} & & \Downarrow \sqrt{[ge_2 \oplus k?]} \\ \Lambda & \mathcal{R} & \Lambda \end{array}$$

and since \mathcal{R} is a higher-order weak bisimulation, we have:

$$\begin{array}{c} ge_1 \oplus k? \ \mathcal{R} \ ge_2 \oplus k? \\ \downarrow l_1 \\ e_1 \end{array}$$

which can be completed as:

$$\begin{array}{ccc} ge_1 \oplus k? \ \mathcal{R} \ ge_2 \oplus k? & & \\ \downarrow l_1 & \hat{l}_2 \Downarrow & \text{where } l_1 \ \mathcal{R}^l \ l_2 \\ e_1 \ \mathcal{R} \ e_2 & & \end{array}$$

but since $e_1 \not\stackrel{k?x}{\approx}$ and $l_1 \neq k?x$, we have $e_2 \not\stackrel{k?x}{\approx}$ and $l_2 \neq k?x$, and so:

$$\begin{array}{ccc} ge_1 \ \mathcal{R}^s \ ge_2 & & \\ \downarrow l_1 & l_2 \Downarrow & \text{where } l_1 \ \mathcal{R}^{sl} \ l_2 \\ e_1 \ \mathcal{R} \ e_2 & & \end{array}$$

The other cases are simpler, and so $(\mathcal{R}^n, \mathcal{R}^s)$ is an hereditary bisimulation. Thus $\mathcal{R} \subseteq \mathcal{R}^n \subseteq \approx^n$, and so \approx^n is the largest higher-order weak bisimulation which respects μCML contexts. \square

This Theorem should be contrasted with the case of CCS. In Milner (1989, section 7.2), it is shown that the largest congruence contained in weak bisimulation is not itself a weak bisimulation.

4 Bisimulation as a congruence

To serve as the basis of a useful semantic theory of μCML , bisimulation should be preserved by all of the constructs of the language. In this section we will show that \approx^s is a congruence for μCML^+ , and that \approx^n is a congruence for μCML .

Unfortunately, this proof is not straightforward, due to the higher-order nature of hereditary bisimulation. The problem is not unique to μCML , and it occurs in many higher-order languages, for example typed λ -calculi (Gordon, 1995), the untyped λ -calculus (Howe, 1989) and the Calculus of Higher-Order Communicating Systems (CHOCS) (Thomsen, 1995).

The difficulty is in finding the right form of induction to use, when all of the standard inductions (for example, on structure of terms, on number of τ -moves, on structure of proof) fail. For example, the proof of congruence for CHOCS Prop. 6.6 adapts Milner's technique Theorem 8, p. 155, but uses a non-well-founded induction. It seems that any inductive proof that weak bisimulation is a congruence for higher-order languages requires an induction on both syntax *and* proof structure. The usual

methods of performing nested induction fail in this case, and so another method of performing simultaneous induction is required. Fortunately, this is achieved by a technique developed for the lazy λ -calculus (Howe, 1989).

We shall apply Howe's technique to show that \approx^s is a congruence for μCML^+ , and that \approx^n is a congruence for μCML^+ without $[ge]$ and $ge_1 \oplus ge_2$. This particular application is made complicated by the fact that we have to deal a pair of relations, (\approx^n, \approx^s) which are defined in terms of each other. So although we follow the general proof method used by Howe (1989) and the notation of Gordon (1995), the various technical definitions about relations which follow will apply to pairs of relations of the form $\mathcal{R} = (\mathcal{R}^n, \mathcal{R}^s)$ with $\mathcal{R}^s \subseteq \mathcal{R}^n$. We will continue to apply the usual operations associated with relations, such as composition, under the assumption that such operations are applied pointwise.

Define a *context* to be given by the grammar:

$$\begin{aligned} C ::= & \cdot_i \mid e \mid c C \mid \text{if } C \text{ then } C \text{ else } C \mid (C, C) \mid \text{let } x = C \text{ in } C \\ & \mid C C \mid \text{fix}(x = \text{fn } y \Rightarrow C) \mid \langle C, C \rangle \\ & \mid [C] \mid C!C \mid C? \mid C \Rightarrow C \mid C \oplus C \mid \mathbf{A}C \mid C \# C \end{aligned}$$

Let $C[\vec{e}]$ be the term given by replacing each 'hole' \cdot_i by the term e_i (unlike substitution, we allow for capture of free variables). An equivalence \mathcal{R} is a *congruence* iff $e_i \mathcal{R} f_i$ implies $C[\vec{e}] \mathcal{R} C[\vec{f}]$.

Define an *uneventful context* to be one which does not use $[C]$ or $C \oplus C$, that is one given by the grammar:

$$\begin{aligned} C_n ::= & \cdot_i \mid e \mid c C_n \mid \text{if } C_n \text{ then } C_n \text{ else } C_n \mid (C_n, C_n) \mid \text{let } x = C_n \text{ in } C_n \\ & \mid C_n C_n \mid \text{fix}(x = \text{fn } y \Rightarrow C_n) \mid \langle C_n, C_n \rangle \\ & \mid C_n!C_n \mid C_n? \mid C_n \Rightarrow C_n \mid \mathbf{A}C_n \mid C_n \# C_n \end{aligned}$$

An equivalence \mathcal{R} is an *uneventful congruence* iff $e_i \mathcal{R} f_i$ implies $C_n[\vec{e}] \mathcal{R} C_n[\vec{f}]$. Note that any μCML context is an uneventful context, and so any uneventful congruence is a congruence for μCML . So we concentrate on showing that \approx^s is a congruence, and \approx^n is an uneventful congruence.

Define the *one-level deep* contexts with the grammar:

$$\begin{aligned} D ::= & x \mid l \mid c \cdot_1 \mid \text{if } \cdot_1 \text{ then } \cdot_2 \text{ else } \cdot_3 \mid (\cdot_1, \cdot_2) \mid \text{let } x = \cdot_1 \text{ in } \cdot_2 \\ & \mid \cdot_1 \cdot_2 \mid \text{fix}(x = \text{fn } y \Rightarrow \cdot_1) \mid \langle \cdot_1, \cdot_2 \rangle \\ & \mid [\cdot_1] \mid \cdot_1! \cdot_2 \mid \cdot_1? \mid \cdot_1 \Rightarrow \cdot_2 \mid \cdot_1 \oplus \cdot_2 \mid \mathbf{A}\cdot_1 \mid \cdot_1 \# \cdot_2 \end{aligned}$$

Let D_n range over uneventful one-level deep contexts.

For any pair of relations $\mathcal{R} = (\mathcal{R}^n, \mathcal{R}^s)$ with $\mathcal{R}^s \subseteq \mathcal{R}^n$, let its *compatible refinement*, $\widehat{\mathcal{R}}$ be defined:

$$\begin{aligned} \widehat{\mathcal{R}}^n &= \{(D_n[\vec{e}], D_n[\vec{f}]) \mid e_i \mathcal{R}^n f_i\} \cup \widehat{\mathcal{R}}^s \\ \widehat{\mathcal{R}}^s &= \{(D[\vec{e}], D[\vec{f}]) \mid e_i \mathcal{R}^s f_i\} \\ &\quad \cup \{(\text{fix}(x = \text{fn } y \Rightarrow e), \text{fix}(x = \text{fn } y \Rightarrow f)) \mid e \mathcal{R}^n f\} \end{aligned}$$

This definition is rather different from Howe’s and Gordon’s definition of $\widehat{\mathcal{R}} = \{(D[\vec{e}], D[\vec{f}]) \mid e_i \mathcal{R} f_i\}$. The differences are that:

- \approx^n is not a congruence, it is only an uneventful congruence, so we only close $\widehat{\mathcal{R}}^n$ under uneventful one-level deep contexts rather than arbitrary one-level deep contexts,
- we want to maintain the invariant that for all pairs of relations we consider, $\mathcal{R}^s \subseteq \mathcal{R}^n$, hence we include $\widehat{\mathcal{R}}^s$ in the definition of $\widehat{\mathcal{R}}^n$, and
- if two insensitive bisimilar expressions are thunked, the resulting expressions are sensitive bisimilar; for this reason the proof of Theorem 4.7 requires $\text{fix}(x = \text{fn } y \Rightarrow e) \widehat{\mathcal{R}}^s \text{fix}(x = \text{fn } y \Rightarrow f)$ when $e \mathcal{R}^n f$.

Proposition 4.1

If \mathcal{R} is an equivalence and $\widehat{\mathcal{R}} \subseteq \mathcal{R}$, then \mathcal{R}^s is a congruence and \mathcal{R}^n is an uneventful congruence.

Proof

A variant of the proof in Gordon (1995) and Howe (1989). Show by induction on C that if $e_i \mathcal{R}^s f_i$ then $C[\vec{e}] \mathcal{R}^s C[\vec{f}]$. Either $C = \cdot_i$, in which case the result is immediate, or $C = D[\vec{C}]$ and by induction $C_i[\vec{e}] \mathcal{R}^s C_i[\vec{f}]$, so by definition $C[\vec{e}] = D[\vec{C}[\vec{e}]] \widehat{\mathcal{R}}^s D[\vec{C}[\vec{f}]] = C[\vec{f}]$. It follows that \mathcal{R}^s is a congruence. The proof that \mathcal{R}^n is an uneventful congruence is similar. \square

For any \mathcal{R} , its *compatible closure*, \mathcal{R}^\bullet , is given by:

$$\frac{e \widehat{\mathcal{R}}^\bullet e' \mathcal{R}^\circ e''}{e \mathcal{R}^\bullet e''}$$

Note that $\mathcal{R}^{\bullet s} \subseteq \mathcal{R}^{\bullet n}$.

This definition of \mathcal{R}^\bullet is specifically designed to facilitate simultaneous inductive proof on syntax (since the definition involves one-level deep contexts) and on reductions (since the definition involves inductive use of \mathcal{R}°). This form of induction is precisely what is required to show the desired congruence results.

Its relevant properties are summed up in the following proposition.

Proposition 4.2

If \mathcal{R}° is a preorder then \mathcal{R}^\bullet is the smallest relation satisfying:

1. $\mathcal{R}^\bullet \mathcal{R}^\circ \subseteq \mathcal{R}^\bullet$,
2. $\widehat{\mathcal{R}}^\bullet \subseteq \mathcal{R}^\bullet$, and
3. $\mathcal{R}^\circ \subseteq \mathcal{R}^\bullet$.

Proof

A variant of the proof in Gordon (1995).

First we show that \mathcal{R}^\bullet is reflexive, by showing by structural induction on e that $e \mathcal{R}^{\bullet s} e$. Find $D[\vec{e}]$ such that $e = D[\vec{e}]$, so by induction $e_i \mathcal{R}^{\bullet s} e_i$, so by definition of $\widehat{\mathcal{R}}^\bullet$, $e = D[\vec{e}] \widehat{\mathcal{R}}^{\bullet s} D[\vec{e}] \mathcal{R}^{\circ s} D[\vec{e}] = e$.

Then we show the required properties:

1. $\mathcal{R}^\bullet \mathcal{R}^\circ \subseteq \widehat{\mathcal{R}}^\bullet \mathcal{R}^\circ \mathcal{R}^\circ \subseteq \widehat{\mathcal{R}}^\bullet \mathcal{R}^\circ \subseteq \mathcal{R}^\bullet$.

2. $\widehat{\mathcal{R}}^\bullet \subseteq \widehat{\mathcal{R}}^\bullet \mathcal{R}^\circ \subseteq \mathcal{R}^\bullet$.
3. $\mathcal{R}^\circ \subseteq \mathcal{R}^\bullet \mathcal{R}^\circ \subseteq \mathcal{R}^\bullet$.

To see that \mathcal{R}^\bullet is the smallest relation satisfying these properties we show that if \mathcal{S} satisfies these properties, then $\widehat{\mathcal{S}} \mathcal{R}^\circ \subseteq \mathcal{S} \mathcal{R}^\circ \subseteq \mathcal{S}$, and so $\mathcal{R}^\bullet \subseteq \mathcal{S}$. \square

Since $\widehat{\mathcal{R}}^\bullet \subseteq \mathcal{R}^\bullet$, we know from Proposition 4.1 that if \mathcal{R}^\bullet is an equivalence then \mathcal{R}^s is a congruence and \mathcal{R}^n is an uneventful congruence. However, we can show a stronger result than that, which is that \mathcal{R}^\bullet is closed under substitution of closed values:

Proposition 4.3

If \mathcal{R} is a preorder then for any $v \mathcal{R}^{s^*} w$:

1. if $e \mathcal{R}^{s^*} f$ then $e[v/x] \mathcal{R}^{s^*} f[w/x]$, and
2. if $e \mathcal{R}^{n^*} f$ then $e[v/x] \mathcal{R}^{n^*} f[w/x]$.

Proof

A variant of the proof in (Gordon, 1995; Howe, 1989). To prove the first part, we proceed by induction on e .

- If $e = x$ then $x \mathcal{R}^{s^*} f$, so $e[v/x] = v \mathcal{R}^{s^*} w \mathcal{R}^{s^*} f[w/x]$ so by Proposition 4.2 $e[v/x] \mathcal{R}^{s^*} f[w/x]$.
- If $e = \text{fix}(y = \text{fn } z \Rightarrow e_1)$ then we can find a g_1 such that $e_1 \mathcal{R}^{n^*} g_1$ and $\text{fix}(y = \text{fn } z \Rightarrow g_1) \mathcal{R}^{s^*} f$, so by induction $e_1[v/x] \mathcal{R}^{n^*} g_1[w/x]$, so $e[v/x] = \text{fix}(y = \text{fn } z \Rightarrow e_1[v/x]) \widehat{\mathcal{R}}^{s^*} \text{fix}(y = \text{fn } z \Rightarrow g_1[w/x]) \mathcal{R}^{s^*} f[w/x]$, so by definition of \mathcal{R}^\bullet , $e[v/x] \mathcal{R}^{s^*} f[w/x]$.
- Otherwise, we have $e = D[\tilde{e}]$ and $D[\tilde{e}][v/x] = D[\tilde{e}[v/x]]$, so we can find \tilde{g} such that $\tilde{e} \mathcal{R}^{s^*} \tilde{g}$ and $D[\tilde{g}] \mathcal{R}^{s^*} f$, so by induction $e_i[v/x] \mathcal{R}^{s^*} f_i[w/x]$, hence $e[v/x] = D[\tilde{e}][v/x] = D[\tilde{e}[v/x]] \widehat{\mathcal{R}}^{s^*} D[\tilde{f}[w/x]] = D[\tilde{f}[w/x]] \mathcal{R}^{s^*} f[w/x]$, so by definition of \mathcal{R}^\bullet , $e[v/x] \mathcal{R}^{s^*} f[w/x]$.

The proof of the second part is similar. \square

Our proof strategy is to show that \approx° and \approx^\bullet coincide. Since $\approx^\circ \subseteq \approx^\bullet$, this amounts to showing that $\approx^\bullet \subseteq \approx^\circ$, which we do by proving that \approx^\bullet , when restricted to programs, is an hereditary simulation.

Proposition 4.4

When restricted to closed expressions of μCML^+ , \approx^\bullet is an hereditary simulation.

Proof

We have to show that \approx^{s^*} is structure-preserving, and that the diagrams for an hereditary simulation can be completed.

Showing that \approx^{s^*} is structure preserving is a routine structural induction.

If:

$$\begin{array}{ccc}
 e & \approx^{n^*} & f \\
 \downarrow l_1 & & \\
 e' & &
 \end{array}$$

then we proceed by induction on e to show that we can complete the diagram as:

$$\begin{array}{ccc} e & \approx^{\bullet n} & f \\ \downarrow l_1 & & \Downarrow \hat{l}_2 \\ e' & \approx^{\bullet n} & f' \end{array}$$

where $l_1 \approx^{\bullet sl} l_2$, and similarly for $\approx^{\bullet s}$. We shall show three of the more interesting cases, the others are similar but more routine:

- if we have:

$$\begin{array}{ccc} e \equiv \text{let } x = e_1 \text{ in } e_2 \widehat{\approx}^n \text{let } x = g_1 \text{ in } g_2 \approx^n & f \\ \downarrow \tau & & \downarrow \tau \\ e' \equiv e'_1 \# e_2[v/x] & & \end{array}$$

where $e_i \approx^{\bullet n} g_i$ and $e_1 \xrightarrow{v} e'_1$, then by induction $g_1 \xrightarrow{w} g'_1$, $v \approx^{\bullet s} w$ and $e'_1 \approx^{\bullet n} g'_1$, so using Proposition 4.3, we have:

$$\begin{array}{ccc} e \equiv \text{let } x = e_1 \text{ in } e_2 \widehat{\approx}^n \text{let } x = g_1 \text{ in } g_2 \approx^n & f \\ \downarrow \tau & & \downarrow \tau \\ e' \equiv e'_1 \# e_2[v/x] \approx^{\bullet n} g'_1 \# g_2[w/x] \approx^n & f' \end{array}$$

- if we have:

$$\begin{array}{ccc} e \equiv e_1 e_2 & \widehat{\approx}^n & g_1 g_2 \approx^n & f \\ \downarrow \tau & & \downarrow \tau & \\ e' \equiv e'_1 \# \text{let } y = e_2 \text{ in } e_3[v/x] & & & \end{array}$$

where $e_i \approx^{\bullet n} g_i$, $e_1 \xrightarrow{v} e'_1$, and $v = \text{fix}(x = \text{fn } y \Rightarrow e_3)$ then by induction $g_1 \xrightarrow{w} g'_1$, $v \approx^{\bullet s} w$, up to α -conversion $w = \text{fix}(x = \text{fn } y \Rightarrow g_3)$, and $e'_1 \approx^{\bullet n} g'_1$. Then by the definition of \approx^{\bullet} , we can find an $v' = \text{fix}(x = \text{fn } y \Rightarrow h_3)$ such that $e_3 \approx^{\bullet n} h_3$ and $v' \approx^s w$, so by Proposition 4.3, $e_3[v/x] \approx^{\bullet n} h_3[v'/x] \approx^{n^\circ} v' y \approx^{n^\circ} w y \approx^{n^\circ} g_3[w/x]$, and so:

$$\begin{array}{ccc} e \equiv e_1 e_2 & \widehat{\approx}^n & g_1 g_2 \approx^n & f \\ \downarrow \tau & & \downarrow \tau & \\ e' \equiv e'_1 \# \text{let } y = e_2 \text{ in } e_3[v/x] \approx^{\bullet n} g'_1 \# \text{let } y = g_2 \text{ in } g_3[w/x] \approx^n & & & f' \end{array}$$

- if we have:

$$\begin{array}{ccc}
 e & \xlongequal{\quad} & \text{fix}(x = \text{fn } y \Rightarrow e_1) \widehat{\approx}^n \text{fix}(x = \text{fn } y \Rightarrow g_1) \approx^n & f \\
 \downarrow \sqrt{e} & & \downarrow \sqrt{e} & \\
 e' & \xlongequal{\quad} & \Lambda &
 \end{array}$$

where $e_1 \approx^{n^*} g_1$ then let $v = \text{fix}(x = \text{fn } y \Rightarrow g_1)$, so:

$$\begin{array}{ccccccc}
 e & \xlongequal{\quad} & \text{fix}(x = \text{fn } y \Rightarrow e_1) \widehat{\approx}^n \text{fix}(x = \text{fn } y \Rightarrow g_1) \approx^n & & f \\
 \downarrow \sqrt{e} & & \downarrow \sqrt{e} & & \downarrow \sqrt{v} & & \downarrow \sqrt{w} \\
 e' & \xlongequal{\quad} & \Lambda & \approx^{n^*} & \Lambda & \approx^n & f'
 \end{array}$$

and $e \approx^{n^*} v \approx^s w$.

Thus \approx^{\bullet} is an hereditary simulation. \square

We now have that \approx^{\bullet} is a simulation, and we would like to show that it is a bisimulation, for which it suffices to show that \approx^{\bullet} is symmetric. Unfortunately, this is not easy to prove directly, and so we use a result of Howe (1989) (pointed out to the authors by Andrew Pitts) which allows us to show that \approx^{\bullet^*} is symmetric.

Proposition 4.5

If \mathcal{R} is an equivalence then \mathcal{R}^{\bullet^*} is symmetric.

Proof

A variant of the proof in Howe (1989).

It suffices to show that if $e \mathcal{R}^{\bullet^*} f$ then $f \mathcal{R}^{\bullet^*} e$, and that if $e \mathcal{R}^n f$ then $f \mathcal{R}^{n^*} e$, which we show by induction on e . If $e \mathcal{R}^s f$, then either:

- $e = D[\hat{e}] \widehat{\mathcal{R}}^s D[\hat{f}] \mathcal{R}^{\circ} f$ and $e_i \mathcal{R}^s f_i$, so by induction $f_i \mathcal{R}^{s^*} e_i$, so $f \widehat{\mathcal{R}}^s D[\hat{f}] D \widehat{\mathcal{R}}^{s^*} [\hat{e}] = e$, or
- $e = \text{fix}(x = \text{fn } y \Rightarrow e') \widehat{\mathcal{R}}^s \text{fix}(x = \text{fn } y \Rightarrow f') \mathcal{R}^{\circ} f$ and $e' \mathcal{R}^n f'$, so by induction $f' \mathcal{R}^{n^*} e'$, so $f \widehat{\mathcal{R}}^s \text{fix}(x = \text{fn } y \Rightarrow f') \mathcal{R}^{s^*} \text{fix}(x = \text{fn } y \Rightarrow e') = e$.

The proof for \mathcal{R}^n is similar. \square

We can use this result to show that \approx^{\bullet^*} is a bisimulation.

Proposition 4.6

When restricted to closed expressions of μCML^+ , \approx^{\bullet^*} is an hereditary bisimulation.

Proof

By Proposition 4.4, \approx^{\bullet} is an hereditary simulation, and so \approx^{\bullet^*} is an hereditary simulation. By Proposition 4.5, \approx^{\bullet} is symmetric, and so \approx^{\bullet} is an hereditary bisimulation. \square

This gives us the result we set out to prove.

Theorem 4.7

\approx^s is a congruence, and \approx^n is an uneventful congruence.

Proof

From Proposition 4.6, \approx^\bullet is an hereditary bisimulation, so $\approx^\bullet \subseteq \approx^\circ$, and by Proposition 4.2 $\approx^\circ \subseteq \approx^\bullet$, so \approx^\bullet and \approx° are the same relation. Since $\widehat{\approx^\bullet} \subseteq \approx^\bullet$, we have the desired result by Proposition 4.1. \square

5 Properties of weak bisimulation

In this section, we show some results about program equivalence up to hereditary weak bisimulation. Some of these equivalences are easy to show, but some are trickier, and require properties about the transition systems generated by μCML^+ . Although much remains to be done on elaborating the algebraic theory of μCML programs we hope that the results in this section indicate that this equivalence can form the basis of a useful theory which generalizes those associated with process algebras and functional programming.

We have given an operational semantics to μCML by extending it with new constructs, most of which correspond to constructs found in standard process algebras. These include a choice operator \oplus , a parallel operator \parallel and suitable versions of input and output prefixing (Milner, 1989). The prefixes in μCML^{cv} have a slightly unusual syntax – their equivalents in CCS are given as:

<i>CCS prefix</i>	<i>μCML^{cv} equivalent</i>
$k?x.P$	$k? \Rightarrow \text{fn } x \Rightarrow P$
$k!v.P$	$k!v \Rightarrow \text{fn } x \Rightarrow P$
$\tau.P$	$\mathbf{A}() \Rightarrow \text{fn } x \Rightarrow P$

We now examine the extent to which \oplus and \parallel act like choice and parallel operators from a process algebras

We can find bisimulations for the following (and hence they are sensitive bisimilar):

$$\begin{aligned} \Lambda \parallel e &\sim^1 e \\ (e_1 \parallel e_2) \parallel e_3 &\sim^1 e_1 \parallel (e_2 \parallel e_3) \\ (e_1 \parallel e_2) \parallel e_3 &\sim^1 (e_2 \parallel e_1) \parallel e_3 \end{aligned}$$

Thus \parallel satisfies many of the standard laws associated with a parallel operator in a process algebra. However, it is not in general symmetric because of its interaction with the production of values:

$$v \parallel e \not\sim^1 e$$

For example:

$$1 \parallel \Lambda \sim^1 \Lambda \quad \Lambda \parallel 1 \not\sim^1 1$$

This means that we can view the parallel composition of processes as being of the form:

$$(\parallel_i e_i) \parallel f$$

where the order of the e_i is unimportant. Note that it is important which is the right-most expression in a parallel composition, since it is the main thread of computation, and so can return a value, which none of the other expressions can.

The choice operator of μCML^+ also satisfies the expected laws from process algebras, those of a commutative monoid, although it can only be applied to guarded expressions:

$$\begin{aligned} \Lambda \oplus ge &\sim^1 ge \\ (ge_1 \oplus ge_2) \oplus ge_3 &\sim^1 ge_1 \oplus (ge_2 \oplus ge_3) \\ ge_1 \oplus ge_2 &\sim^1 ge_2 \oplus ge_1 \end{aligned}$$

This means that we can view the sum of guarded expressions as being of the form:

$$\bigoplus_i ge_i$$

where the order of the ge_i is unimportant.

In fact, guarded expressions can be viewed in a manner quite similar to the *sum forms* used in the development of the algebraic theory of CCS (Milner, 1989). We can find bisimulations for the following (and hence they are sensitive bisimilar):

$$\begin{aligned} (ge_1 \oplus ge_2) \Rightarrow v &\sim^1 (ge_1 \Rightarrow v) \oplus (ge_2 \Rightarrow v) \\ ge \Rightarrow \text{fn } x \Rightarrow x &\approx^s ge \\ \mathbf{A}v &\approx^s \mathbf{A}() \Rightarrow \text{fn } x \Rightarrow v \end{aligned}$$

From this, we can show, by structural induction on syntax that all guarded expressions are of a given form:

$$ge \approx^s \bigoplus_i ge_i \Rightarrow v_i$$

where each ge_i is either $k_i!v_i$, $k_i?$ or $\mathbf{A}()$. From this and:

$$cv \approx^1 \delta(c, v)$$

we can show that all values $\vdash v : A$ event are of the form:

$$v \approx^n \text{choose}[\text{wrap}(e_1, v_1), \dots, \text{wrap}(e_n, v_n)]$$

where e_n is either $\text{transmit}(k_i, v_i)$, $\text{receive } k_i$, or $\text{always}()$.

We could continue in this manner emulating the algebraic theory of CCS, for example with expansion theorems for guarded expressions or values of event type. However we leave this for future work.

We now turn our attention to μCML viewed as a functional language. One would not expect β -reduction in its full generality in a language with side-effects such as μCML but we do obtain an appropriate call-by-value version:

$$(\text{fn } y \Rightarrow e)v \approx^1 e[v/y]$$

We also have expected laws such as:


$$\begin{aligned} \text{fst}(e, v) &\approx^1 e \\ \text{snd}(v, e) &\approx^1 e \end{aligned}$$

$$\begin{aligned}
 (\text{fix}(x = \text{fn } y \Rightarrow e))v &\approx^1 e[\text{fix}(x = \text{fn } y \Rightarrow e)/x][v/y] \\
 \text{let } x = v \text{ in } e &\approx^1 e[v/x] \\
 \text{let } y = (\text{let } x = e \text{ in } f) \text{ in } g &\approx^1 \text{let } x = e \text{ in } (\text{let } y = f \text{ in } g) \quad \text{where } x \notin \text{fv}(g)
 \end{aligned}$$

The last two equations are of particular interest, since they are exactly the left unit and associativity axioms of the monadic metalanguage (Moggi, 1991). The right unit equation:

$$\text{let } x = e \text{ in } x \approx^n e$$

is not so simple to show, and indeed, if e were an arbitrary labelled transition system then it would not be true, as can be seen by:




(This is the same example which makes *SKIP* not act as a right unit for sequential composition in CSP (Hoare, 1985) and **exit** not act as a right unit for \gg in LOTOS (ISO 8807, 1989).) Fortunately, we can show that our operational semantics for μ CML satisfies four properties which allow us to show the right unit equation.


A labelled transition system is *single-valued* iff:

$$\text{if } e \xrightarrow{v} e' \text{ then } e' \xrightarrow{w}$$

It is *value deterministic* iff:




It is *forward commutative* iff:



Note that since α may be an input move, e'' may be an open term, so we need to take the open extension $\xrightarrow{v^\circ}$ of the termination relation.

It is backward commutative iff:



Note in particular that LOTOS and CSP do not satisfy forward commutativity, which is why their sequential composition operators do not have a right unit. However, μ CML does satisfy these conditions.

Proposition 5.1

μ CML satisfies single-valuedness, value determinacy, forward commutativity and backward commutativity.

Proof

A routine induction on syntax. \square

The important property which such lts's satisfy is the following, where we assume the existence of the operator \dashv .

Proposition 5.2

In any single-valued, value deterministic, forward commutative, backward commutative lts, if $e \xrightarrow{\sqrt{v}} e'$ then $e \approx^1 e' \dashv v$.

Proof

Use the properties of the lts to establish that the following is a first-order weak bisimulation:

$$\{(e, e' \dashv v) \mid e \xrightarrow{\sqrt{v}} e'\} \cup \{(e', e' \dashv \Lambda) \mid e \xrightarrow{\sqrt{v}} e'\}$$

The result follows. \square

As a corollary to this proposition, it is routine to show that the following is a first-order weak bisimulation:

$$\{(e, \text{let } x = e \text{ in } x)\} \cup \approx^1$$

So we have the right unit equation we were looking for:

$$e \approx^1 \text{let } x = e \text{ in } x$$

These equations enable us to define a categorical model for μ CML where:

- objects are types,
- morphisms between A and B are typed expressions with one free variable $x : A \vdash e : B$, viewed up to weak bisimulation,
- the identity morphism is $x : A \vdash x : A$, and
- composition is $(x : A \vdash e : B); (y : B \vdash f : C) = (x : A \vdash \text{let } y = e \text{ in } f : C)$.

The equations for weak bisimulation discussed above show that morphism composition is associative and has the identity as both a left unit and right unit. Thus μ CML forms a category.

Again we leave the investigation of the properties of this category to future work, but we should point out that so far we have been unable to cast it as an instance of general categorical framework of Moggi (1991).

6 Comparing μCML^+ and λ_{cv}

In section 2 we presented the operational semantics of a subset of CML, as a labelled transition system, so that we might investigate its behavioural properties. In this section we shall make a formal connection between this semantics and the reduction semantics for λ_{cv} presented by Reppy (1992). We have not considered λ_{cv} in its entirety and so the comparison will be confined to the common subset, namely μCML^{cv} . We first reproduce, as faithfully as possible, the reduction semantics of Reppy as it applies to μCML . From this reduction semantics we then derive a labelled transition system for μCML expressions and our main theorem states that this labelled transition system (up to first-order weak bisimulation) is the same as ours. In fact, the more technical results we derive connecting the two semantics would support a much closer relationship, but expressing it would involve developing yet another bisimulation based equivalence.

Before presenting the operational semantics and our main theorem we clarify the differences between λ_{cv} and μCML^{cv} :

- We do not consider the λ_{cv} constructs `guard` and `wrapAbort`. We conjecture that the operational semantics of μCML would need to be considerably altered to cope with translating these constructs.
- We omit the λ_{cv} construct `chan x in e` since we cannot encode unique channel name generation in μCML , although it should not be difficult to add it using operational rules à la π -calculus. However, this would require using a bisimulation similar to Sangiorgi's (1992) context bisimulation for the higher-order π -calculus.
- We have added recursive function types to μCML^{cv} because in Reppy (1992) recursion is encoded using process creation and unique channel name generation.
- In λ_{cv} , constant functions such as `wrap` are values, where in μCML they have to be coded as `(fn x => wrap x)`. This restriction has no effect on the expressive power of μCML , and makes it simpler to reason about the operational semantics, since any value of type $A \rightarrow B$ must be of the form `fix(x = fn y => e)`.

We now present Reppy's reduction semantics for μCML^{cv} . In Reppy (1992) this is represented by a transition relation between multi-sets of μCML^{cv} , or more generally λ_{cv} expressions. Instead of multi-sets we use *configurations* of μCML^{cv} expressions given by the grammar:

$$C \in \text{Conf} ::= e \mid C \# C \mid \Lambda$$

Note that configurations are restricted forms of μCML^+ expressions. This will facilitate the comparison between the two semantics, since it can be carried out for configurations rather than μCML expressions.

The semantics of Reppy (1992) is expressed as a reduction relation \Longrightarrow between configurations and reductions have four independent sources. The first involves a sequential reduction within an individual μ CML expression, and this in turn is defined using another reduction relation \mapsto ; the second is the spawning of new *computation threads*, which results in an increase in the number of components of the configuration; the third is communication between two expressions and the last is required to handle the `always` construct. We need notation for each of these and we consider them in turn.

The operational rules for sequential reduction are defined *in context* in the style of Wright and Felleisen (1991), and the contexts that permit reduction are given by the following grammar:

$$E ::= [\cdot] \mid E e \mid v E \mid c E \mid (E, e) \mid (v, E) \mid \text{let } x = E \text{ in } e \mid \text{if } E \text{ then } e \text{ else } e$$

The relation \mapsto is defined to be the least relation satisfying the following rules:

$$\begin{array}{lll} E[cv] & \mapsto & E[\delta(cv)] \quad (c \notin \{\text{spawn, sync}\}) \quad \text{const} \\ E[(\text{fix}(x = \text{fn } y \Rightarrow e))v] & \mapsto & E[e[\text{fix}(x = \text{fn } y \Rightarrow e)/x][v/y]] \quad \text{beta} \\ E[\text{let } x = v \text{ in } e] & \mapsto & E[e[v/x]] \quad \text{let} \\ E[(v, w)] & \mapsto & E[\langle v, w \rangle] \quad \text{pair} \end{array}$$

Here each rule corresponds to a basic computation step in a sequential call-by-value language. We should point out that the last rule does not appear in Reppy (1992), it is implicit in Reppy’s statement “the syntactic class of the term (v_1, v_2) is either *Exp* or *Val*; this ambiguity is resolved in favour of *Val*”. We have made the grammar unambiguous, and have added an explicit reduction rule for resolving ambiguity.

Note that the definition of \mapsto is not compositional: the reductions of an expression are not defined in terms of the reductions of its sub-expressions. The following lemma will be useful in later proofs and shows that we can recover compositionality.

Lemma 6.1

If $e \mapsto e'$ then $E[e] \mapsto E[e']$.

Proof

By examination of the proof of the transition $e \mapsto e'$. \square

To capture reductions which involve communication it is necessary to define a notion of when two guarded expressions may give rise to a communication. For any k the relation:

$$ge \stackrel{k}{\bowtie} ge' \text{ with } (e, e')$$

read as ‘ ge matches ge' on k with result (e, e') ’ is defined to be the least relation satisfying the rules in figure 6. Intuitively this means that two concurrent threads e_1, e_2 of the form $e_1 = E_1[\text{sync}[ge]], e_2 = E_2[\text{sync}[ge']]$ may communicate in one step on the channel k with $E_1[e]$ and $E_2[e']$ being the result of this communication.

To handle reductions caused by `always` we need to formalise when guarded expressions such as `Av` can immediately return values. This is given by Reppy’s relation $ge \triangleright e$, is defined in Figure 6.

$$\begin{array}{c}
\frac{}{k!v \stackrel{k}{\bowtie} k? \mathbf{with} ((),v)} \quad \frac{ge \stackrel{k}{\bowtie} ge' \mathbf{with} (e,e')}{ge \stackrel{k}{\bowtie} ge' \Rightarrow v \mathbf{with} (e,v e')} \\
\frac{ge \stackrel{k}{\bowtie} ge' \mathbf{with} (e,e')}{ge \stackrel{k}{\bowtie} ge' \oplus ge'' \mathbf{with} (e,e')} \quad \frac{ge \stackrel{k}{\bowtie} ge'' \mathbf{with} (e,e'')}{ge \stackrel{k}{\bowtie} ge' \oplus ge'' \mathbf{with} (e,e'')} \\
\frac{ge \stackrel{k}{\bowtie} ge' \mathbf{with} (e,e')}{ge' \stackrel{k}{\bowtie} ge \mathbf{with} (e',e)}
\end{array}$$

Fig. 6a. The rules for matching events.

$$\frac{}{Av \triangleright v} \quad \frac{ge \triangleright e}{ge \Rightarrow v \triangleright v e} \quad \frac{ge \triangleright e}{ge \oplus ge' \triangleright e} \quad \frac{ge' \triangleright e'}{ge \oplus ge' \triangleright e'}$$

Fig. 6b. The rules for immediate evaluation of events.

We can now formally present the reduction relation \Longrightarrow between configurations. It is defined to be the least relation satisfying the rules:

$$\begin{array}{c}
\frac{e_i \mapsto e'_i}{(e_1 \Downarrow \cdots \Downarrow e_i \Downarrow \cdots \Downarrow e_n) \Longrightarrow (e_1 \Downarrow \cdots \Downarrow e_i \Downarrow \cdots \Downarrow e_n)} \quad \underline{\text{seq}} \\
\frac{}{(e_1 \Downarrow \cdots \Downarrow E[\text{spawn } v] \Downarrow \cdots \Downarrow e_n) \Longrightarrow (e_1 \Downarrow \cdots \Downarrow v() \Downarrow E[()] \Downarrow \cdots \Downarrow e_n)} \quad \underline{\text{spawn}} \\
\frac{ge \stackrel{k}{\bowtie} ge' \mathbf{with} (e,e')}{(e_1 \Downarrow \cdots \Downarrow E[\text{sync}[ge]] \Downarrow \cdots \Downarrow E'[\text{sync}[ge']] \Downarrow \cdots \Downarrow e_n) \Longrightarrow (e_1 \Downarrow \cdots \Downarrow E[e] \Downarrow \cdots \Downarrow E'[e'] \Downarrow \cdots \Downarrow e_n)} \quad \underline{\text{comm}} \\
\frac{ge \triangleright e}{(e_1 \Downarrow \cdots \Downarrow E[\text{sync}[ge]] \Downarrow \cdots \Downarrow e_n) \Longrightarrow (e_1 \Downarrow \cdots \Downarrow E[e] \Downarrow \cdots \Downarrow e_n)} \quad \underline{\text{eval}}
\end{array}$$

This completes our exposition of Reppy's semantics as it applies to μCML^{cv} , which for convenience we call the μCML^{cv} semantics. We refer to that in section 2 as the μCML^+ semantics and we now compare them. To do this, we extract a labelled transition system from the μCML^{cv} semantics by defining:

$$\begin{array}{l}
C \mapsto^{\tau} C' \quad \text{iff} \quad C \Longrightarrow C' \\
C \mapsto^v C' \quad \text{iff} \quad C = C'' \Downarrow v \text{ and } C' = C'' \Downarrow \Lambda \text{ (up to } \Downarrow \text{ associativity and } \Lambda \text{ left unit)} \\
C \mapsto^{kv} C' \quad \text{iff} \quad C \Downarrow k? \Longrightarrow C' \Downarrow v \\
C \mapsto^{k?x} C' \quad \text{iff} \quad C \Downarrow k!x \Longrightarrow C' \Downarrow ()
\end{array}$$

We then show that this labelled transition system is weakly bisimilar to the μCML^+ lts:

Theorem 6.2

The μCML^{cv} semantics of a configuration is weakly bisimilar to its μCML^+ semantics.

The remainder of this section is devoted to proving this result. Although the style of presentation of these two semantics are very different, the resulting relations are very similar and there are essentially only two sources for the differences. The first is that certain reductions in μCML^{cv} , when modelled in the μCML^+ semantics, require in addition some ‘housekeeping’ reductions. A typical example is the reduction:

$$(\text{fn } x \Rightarrow e)v \mapsto e[v/x].$$

In μCML^+ this requires two reductions:

$$(\text{fn } x \Rightarrow e)v \xrightarrow{\tau} \text{let } x = v \text{ in } e \xrightarrow{\tau} e[v/x]$$

This problem is handled by identifying the set of ‘housekeeping’ reductions, such as the second reduction above, within the μCML^+ semantics. These turn out to be very simple and we can work with ‘housekeeping normal forms’ in which no further housekeeping reductions can be made.

The second divergence between the semantics concerns the treatment of `spawn`; expressions in μCML^+ may spawn new processes which give rise to parallel processes occurring as sub-terms of the expression. For example, the reductions of `(spawn v, e)` in μCML^+ and μCML^{cv} are:

$$\begin{aligned} (\text{spawn } v, e) &\xrightarrow{\tau} (\Lambda \parallel v () \parallel (), e) \\ (\text{spawn } v, e) &\mapsto v () \parallel ((), e) \end{aligned}$$

This difference is handled by working with the μCML^{cv} semantics up to a syntactically defined equivalence; this equivalence is contained in strong bisimulation equivalence and it also preserves housekeeping reductions.

We now explain in some more detail these technical developments; most of the associated proofs are relegated to an Appendix. House-keeping reductions are ones derived using the rules:

$$\begin{array}{c} \frac{e \xrightarrow{[ge]} e'}{\text{sync } e \xrightarrow{\tau} e' \parallel ge} \quad \frac{e \xrightarrow{[v]} e'}{(e, f) \xrightarrow{\tau} e' \parallel \text{let } x = f \text{ in } \langle v, x \rangle} \\ \frac{e \xrightarrow{[v]} e'}{ef \xrightarrow{\tau} e' \parallel \text{let } y = f \text{ in } g[v/x]} [v = \text{fix}(x = \text{fn } y \Rightarrow g)] \end{array}$$

We shall write $e \xrightarrow{\tau_H} e'$ whenever $e \xrightarrow{\tau} e'$ is a housekeeping reduction.

It is routine to verify that the housekeeping moves are ‘semantically unimportant’, as is captured by the next proposition:

Proposition 6.3


If $e \xrightarrow{\tau_H} e'$ then $e \approx^1 e'$.

Proof

Construct a weak bisimulation for each case. \square

Moreover, we can show a confluence result for the μCML^+ semantics about house-keeping moves:

Proposition 6.4



Proof

First show by induction on ge that $ge \not\rightarrow^{\tau_H}$. Then prove by induction on e , using forward commutativity, that if $e \xrightarrow{\tau_H} e'$ and $e \xrightarrow{l} e''$ are distinct reductions then we can find e''' such that $e' \xrightarrow{l} e'''$ and $e'' \xrightarrow{\tau_H} e'''$. The result follows. \square

Call a term ‘tidy’ if it has no housekeeping reductions. Then we can show that every μCML^+ term has a unique tidy normal form.

Proposition 6.5

For any μCML^+ term e there is a unique tidy e' such that $e \xrightarrow{\tau_H}^* e'$.

Proof

Show by induction on e that there is some tidy e' such that $e \xrightarrow{\tau_H}^* e'$. From Proposition 6.4, this e' is unique. \square


We now turn our attention to the syntactic equivalence used to handle the different treatments of `spawn`. To define the equivalence \equiv it is convenient to introduce reduction contexts for μCML^+ , equivalent to those for μCML^{cv} :

$$E^+ ::= [\cdot] \mid E^+ e \mid c E^+ \mid (E^+, e) \mid \text{let } x = E^+ \text{ in } e \mid \text{if } E^+ \text{ then } e \text{ else } e \mid E^+ \# e \mid e \# E^+$$

In the Appendix, we show that these satisfy the natural properties one would expect of reduction contexts. Let \equiv be the smallest equivalence given by:

$$\overline{E^+[\Lambda \# e]} \equiv \overline{E^+[e]} \quad \overline{E_1^+[E_2^+[e \# f]]} \equiv \overline{E_1^+[e \# E_2^+[f]]}$$

The equivalence \equiv is a strong first-order bisimulation which respects housekeeping, that is a relation \mathcal{R} where we can complete the diagram:



and similarly for \mathcal{R}^{-1} .

Proposition 6.6

\equiv is a strong first-order bisimulation which respects housekeeping.


Proof

See the Appendix. \square


We can also show a very strong correspondence between reductions of μCML^{cv} configurations, and their tidy normal forms.

Proposition 6.7

If $C \xrightarrow{\tau_H^*} e$ and e is tidy, then the following diagrams can be completed:



and:



Proof

See the Appendix. \square


With these technical results we can now prove the main result showing the correspondence between the two semantics:

Theorem 6.8

The μCML^{cv} semantics of a configuration is weakly bisimilar to its μCML^+ semantics.


Proof

Intuitively we know, from Proposition 6.3, that μCML^+ expressions are semantically equivalent to their tidy forms, and Proposition 6.7 can be used to transform μCML^{cv} moves from an expression into μCML^+ moves of its tidy form up to \equiv , and vice-versa. Formally, we show that $\xrightarrow{\tau_H^*} \equiv \xleftarrow{\tau_H^*}$ is a weak bisimulation by completing the diagram:




by using Proposition 6.5 to find e_1 's tidy form e_2 , and then using Propositions 6.4,


6.6 and 6.7 to show:



and by completing the diagram:



by using Proposition 6.5 to find e_1 's tidy form e_3 and then using Propositions 6.4, 6.6 and 6.7 to show:



The result follows. \square

7 Conclusions

In this paper, we have defined a compositional operational semantics for a core subset of CML, called μ CML, and used it to develop at least the beginnings of an algebraic theory of CML programs based on an appropriate version of weak bisimulation equivalence. The operational semantics required an extension of the language to μ CML⁺, although it is worth pointing out that all of the added constructs can be defined in the core language μ CML up to weak bisimulation equivalence.

Much research remains to be done. The algebraic theory of μ CML, started in

section 5, needs to be developed to the extent that it can be used to characterize the semantic equivalence \approx^n . More generally, both the operational semantics and the semantic equivalence should be extended to incorporate more of the features of CML. Of particular interest is the generation of new channel names. We believe that our operational semantics can be adapted to handle new channel generation but the semantic equivalence would need to be changed to an appropriate adaptation of *context bisimulation equivalence* (Sangiorgi, 1992).

As pointed out in section 3 our semantic equivalence, \approx^n , is based on the *late* version of bisimulations (Milner *et al.*, 1992). This fits in quite well with the functional nature of CML, but nevertheless, it would be of interest to consider other variations. One can easily define an *early* version of \approx^n or versions where silent moves are allowed to occur after a matching $\xrightarrow{1}$ move. However, we have been unable to adapt Howe's method to show that these equivalences are preserved by μ CML contexts.

In section 3, we were forced to develop the theory of hereditary bisimulations because of the usual problems of τ actions resolving choice. In the sublanguage without *always* and **A**, we showed that weak bisimulation coincided with insensitive hereditary bisimulation, and so has a simpler and more elegant theory. This theory has been investigated by the first author (Ferreira, 1995). In this theory, it is possible to use CSP rather than CCS summation, and so weak bisimulation is respected by all contexts. As a side-effect of this, it is possible to remove the syntactic restriction that $[ge]$ can only be applied to guarded expressions. The third author has shown (Jeffrey, 1995) that the resulting semantics can be presented in terms of computational monads (Moggi, 1991).

There has already been a considerable amount of research into the foundations of CML and related languages. Much of this is concerned with developing more detailed type systems, where types contain information on the behaviour of expressions as they evolve (Nielson and Nielson, 1993; Nielson and Nielson, 1996). Here we confine our remarks to work directly concerned with the development of semantic theories. We have already given a detailed comparison with the operational semantics given in Reppy (1991b, 1992) and Panangaden and Reppy (1996). This semantics has been used in Berry *et al.* (1996) to study an implementation of ML reference types using process generation. If we extend our approach to include channel generation then we could hope to give an algebraic treatment of their results. In Bolignano and Debabi (1994) and Debabi (1994), there are a number of different semantics given to languages related to CML. A denotational semantics is given using the concept of 'dynamic types', but it has not yet been related to any operationally-based equivalence. An operational semantics is also given for a language called *FPI*. This contains many CML features, but the author notes that accommodating any *spawn* or *fork* operator would be difficult. In Havelund (1994) and Baeten and Vaandrager (1992), the *spawn* operator is studied within the context of process algebras. The former gives a two-level operational semantics for a simple 'pure' process algebra with *fork*, and uses this to develop a semantic equivalence based on strong bisimulation; an axiomatization is also given using an auxiliary operator called *forked*. The latter shows how the various algebraic theories of *ACP* can be adapted to support the addition of a *spawn* operator. This contains an lts based

operational semantics for $ACP + spawn$, and their treatment of $spawn$ has been used by Ferreira and Hennessy (1995) to give an operational semantics of a language which can be considered to be an untyped version of μCML . However, bisimulation based equivalences are not developed by Ferreira and Hennessy (1995); instead, a testing equivalence is defined (Hennessy, 1988) and a fully-abstract denotational semantics based on Acceptance Trees is given.

Other languages which contain much in common with CML include *CHOCS* (Thomsen, 1995), *FACILE* (Giacalone *et al.*, 1989), *PICT* (Pierce and Turner, 1995), *ACTORS* (Agha *et al.*, 1994) and *HO π* (Sangiorgi, 1992). Most of these are endowed with an operational semantics, some of which are similar in spirit to ours. However, we feel that our treatment of $spawn$ and delayed computations is novel, and hope that it can be used to good effect with other languages. Many of these languages also have associated with them bisimulation based semantic equivalences. Section 3 may be viewed as an extension of the research Thomsen (1995), and the new equivalence \approx^n can easily be adopted to languages such as *CHOCS* and *FACILE*. We have also already indicated that when we extend μCML to include channel generation, it will be necessary to adopt the *context bisimulation equivalence*, originally developed in Sangiorgi (1992). In short, although semantic theories are being developed independently for these languages, many of the techniques developed will find more general application.

Appendix

This section is devoted to the proof of Proposition 6.6 and Proposition 6.7, but first we need some auxiliary results. The following three propositions state elementary properties of the reduction contexts for μCML^+ , introduced in section 6 and we omit the proofs; they all use structural induction on contexts:

Proposition A.1

If $e \xrightarrow{\alpha} e'$ then $E^+[e] \xrightarrow{\alpha} E^+[e']$.

Proposition A.2

If $E_1^+[e] \xrightarrow{l} f$ then either:

- $f = E_2^+[e]$ and for all g , $E_1^+[g] \xrightarrow{l} E_2^+[g]$, or
- $f = E_2^+[e']$, $e \xrightarrow{l'} e'$, and for all $g \xrightarrow{l'} g'$, $E_1^+[g] \xrightarrow{l} E_2^+[g']$.

Proposition A.3

For any E there is an E^+ such that for all e , $E[e] \xrightarrow{\tau_H}^* E^+[e]$.

With these we can now prove Proposition 6.6:

Proposition A.4

\equiv is a strong first-order bisimulation which respects housekeeping.

Proof

First observe that an alternative definition of \equiv is as the smallest equivalence given by:

$$\begin{array}{c} \overline{\Lambda \dashv\vdash e \equiv e} \quad \overline{(e \dashv\vdash f) \dashv\vdash g \equiv e \dashv\vdash (f \dashv\vdash g)} \quad \overline{e \dashv\vdash (f \dashv\vdash g) \equiv f \dashv\vdash (e \dashv\vdash g)} \\ \overline{(e \dashv\vdash f)g \equiv e \dashv\vdash (fg)} \quad \overline{c(e \dashv\vdash f) \equiv e \dashv\vdash (cf)} \quad \overline{(e \dashv\vdash f, g) \equiv e \dashv\vdash (f, g)} \\ \overline{\text{let } x = e \dashv\vdash f \text{ in } g \equiv e \dashv\vdash \text{let } x = f \text{ in } g} \quad \overline{\text{if } e \dashv\vdash f \text{ then } g \text{ else } h \equiv e \dashv\vdash \text{if } f \text{ then } g \text{ else } h} \\ \frac{e \equiv f}{E[e] \equiv E[f]} \end{array}$$

Then show by induction on the proof of this alternative that \equiv satisfies the required properties to be a first-order strong bisimulation which preserves housekeeping. \square

The next result shows that the auxiliary predicates used in the reduction semantics of μCML^{cv} , \implies , have their exact counterparts in the μCML^+ semantics:

Proposition A.5

1. $ge \xrightarrow{k!v} e$ iff $ge \bowtie^k k? \mathbf{with} (e, v)$,
2. $ge \xrightarrow{k?x} e$ iff $ge \bowtie^k k!x \mathbf{with} (e, ())$,
3. $ge \xrightarrow{\tau} e$ iff $ge \triangleright e$, and
4. if $ge_1 \bowtie^k ge_2 \mathbf{with} (e_1, e_2)$ then $ge_i \xrightarrow{k!v} e_i$ and $ge_j \xrightarrow{k?v} e_j$.


Proof

A routine structural induction. \square


With these results we can now give the proof of Proposition 6.7, which for convenience we restate:

Proposition A.6

If $C \xrightarrow{\tau_H^*} e$ and e is tidy, then the following diagrams can be completed:



and:



Proof

The first diagram is completed by case analysis of $C \mapsto^l C'$. We shall prove some of the cases, as the others are similar.

- If $C \mapsto^\tau C'$ from the **const** rule, then $C = E_1^+[E_2[cv]]$ and $C' = E_1^+[E_2[cv]]$. Then by Propositions A.1 and A.3:

$$\begin{array}{ccccc}
 C & \equiv & E_1^+[E_2[cv]] & \xrightarrow{\tau_H^*} & E_1^+[E_2^+[cv]] & \equiv & e \\
 \downarrow \tau & & \downarrow \tau & & & & \downarrow \tau \\
 C' & \equiv & E_1^+[E_2[\delta(cv)]] & \xrightarrow{\tau_H^*} & E_1^+[E_2^+[\delta(cv)]] & \equiv & E_1^+[E_2^+[\Lambda \# \delta(cv)]]
 \end{array}$$

- If $C \mapsto^{\sqrt{v}} C'$ then $C = C'' \# v$ and $C' = C'' \# \Lambda$, so:

$$\begin{array}{ccccc}
 C & \equiv & C'' \# v & \xrightarrow{\tau_H^*} & e'' \# v & \equiv & e \\
 \downarrow \sqrt{v} & & \downarrow \sqrt{v} & & & & \downarrow \sqrt{v} \\
 C' & \equiv & C'' \# \Lambda & \xrightarrow{\tau_H^*} & e'' \# \Lambda & \equiv & e'' \# \Lambda
 \end{array}$$

- If $C \mapsto^{k!v} C'$ then (from the definition of $C \mapsto^{k!v} C'$ and the **comm** rule) $C = E_1^+[E_2[\text{sync}[ge]]]$, $C = E_1^+[E_2[e]]$, and $ge \bowtie^k k?$ **with** (e, v) , so by Proposition A.5, $ge \mapsto^{k!v} e$, and so by Propositions A.1 and A.3:

$$\begin{array}{ccccc}
 C & \equiv & E_1^+[E_2[\text{sync}[ge]]] & \xrightarrow{\tau_H^*} & E_1^+[E_2^+[ge]] & \equiv & e \\
 \downarrow k!v & & \downarrow k!v & & & & \downarrow k!v \\
 C' & \equiv & E_1^+[E_2[e]] & \xrightarrow{\tau_H^*} & E_1^+[E_2^+[e]] & \equiv & E_1^+[E_2^+[e]]
 \end{array}$$

The second diagram is completed by induction on C . We shall prove some of the cases, as the others are similar.

If $C = E[f]$, E is a one-level deep reduction context for both μCML^+ and μCML^{cv} , $e = E[g]$, $f \xrightarrow{\tau_H^*} g$, $e' = E[g']$ and $g \xrightarrow{\alpha} g'$ then by induction $f \mapsto^l C' \xrightarrow{\tau_H^*} f' \equiv g'$ and we can show by induction on E that $E[g] \mapsto^l \equiv E[C]$ so by Propositions 6.6:

$$\begin{array}{ccccc}
 C & \equiv & E[f] & \xrightarrow{\tau_H^*} & E[g] & \equiv & e \\
 \downarrow \alpha & & & & \downarrow \alpha & & \downarrow \alpha \\
 C'' & \equiv & E[f'] & & & & \\
 & \searrow \tau_H^{\circ} & & \searrow \tau_H^{\circ} & & & \\
 & & f'' & \equiv & E[g'] & \equiv & e'
 \end{array}$$

Otherwise:

- If $C = cf$ then $f \xrightarrow{\tau_H^*} g$, g is tidy and $cg \xrightarrow{\tau_H^*} e$, so either:

- $c = \text{sync}$, $e = g' \parallel ge$, $g \xrightarrow{\surd/[ge]} g'$, and $f \xrightarrow{\tau_H^*} g$, so by induction and the definition of $\xrightarrow{\surd/v}$, $f = g = [ge]$ and $g' = \Lambda$, so $e' = \Lambda \parallel g''$ and $ge \xrightarrow{\alpha} g''$, so by Proposition A.5, $\text{sync}[ge] \xrightarrow{\alpha} g''$, and so:

$$\begin{array}{ccccc}
 C & \equiv & \text{sync}[ge] & \xrightarrow{\tau_H^*} & \Lambda \parallel ge & \equiv & e \\
 \downarrow \alpha & & & & \downarrow \alpha & & \downarrow \alpha \\
 g'' & \xrightarrow{\tau_H^{*\circ}} & g'' & \equiv & \Lambda \parallel g'' & \equiv & e'
 \end{array}$$

- $c = \text{spawn}$, $e = \text{spawn } g$, $e' = g' \parallel v() \parallel ()$ and $g \xrightarrow{\surd/v} g'$, so by induction and the definition of $\xrightarrow{\surd/v}$, $f = g = v$ and $g' = \Lambda$, and so:

$$\begin{array}{ccccc}
 C & \equiv & \text{spawn } v & \xrightarrow{\tau_H^*} & \text{spawn } g & \equiv & e \\
 \downarrow \tau & & & & \downarrow \tau & & \downarrow \tau \\
 v() \parallel () & \xrightarrow{\tau_H^{*\circ}} & v() \parallel () & \equiv & \Lambda \parallel v() \parallel () & \equiv & e'
 \end{array}$$

- or $e' = g' \parallel \delta(c, v)$ and $g \xrightarrow{\surd/v} g'$, so by induction and the definition of $\xrightarrow{\surd/v}$, $f = g = v$ and $g' = \Lambda$, and so:

$$\begin{array}{ccccc}
 C & \equiv & cv & \xrightarrow{\tau_H^*} & \text{spawn } g & \equiv & e \\
 \downarrow \tau & & & & \downarrow \tau & & \downarrow \tau \\
 \delta(c, v) & \xrightarrow{\tau_H^{*\circ}} & \delta(c, v) & \equiv & \Lambda \parallel \delta(c, v) & \equiv & e'
 \end{array}$$

- If $C = f_1 f_2$ then $f_1 \xrightarrow{\tau_H^*} g_1 \xrightarrow{\surd/v} g'_1$ where $v = \text{fix}(x = \text{fn } y \Rightarrow g_3)$, $f_2 \xrightarrow{\tau_H^*} g_2$, $e = g'_1 \parallel \text{let } y = g_2 \text{ in } g_3[v/x]$, so by induction and the definition of $\xrightarrow{\surd/v}$, $f_1 = g_1 = v$ and $g'_1 = \Lambda$, and so either:

- $e' = g'_1 \parallel \text{let } y = g'_2 \text{ in } g_3[v/x]$ and $g_2 \xrightarrow{\alpha} g'_2$ so by induction (up to associativity of \parallel and Λ being a left unit), $f_2 \xrightarrow{\alpha} C' \parallel f'_2 \xrightarrow{\tau_H^{*\circ}} f_3 \parallel f''_2 \equiv g'_2$, and so:

$$\begin{array}{ccccc}
 C & \equiv & v f_2 & \xrightarrow{\tau_H^*} & \Lambda \parallel \text{let } y = g_2 \text{ in } g_3[v/x] & \equiv & e \\
 \downarrow \alpha & & & & \downarrow \alpha & & \downarrow \alpha \\
 C' \parallel v f_2 & \xrightarrow{\tau_H^{*\circ}} & f_3 \parallel \text{let } y = f''_2 \text{ in } g_3[v/x] & \equiv & \Lambda \parallel \text{let } y = g'_2 \text{ in } g_3[v/x] & \equiv & e'
 \end{array}$$

- or $e' = g'_1 \parallel g'_2 \parallel g_3[v/x][w/y]$ and $g_2 \xrightarrow{\surd/w} g'_2$, so by induction and the

definition of \mapsto^v , $f_2 = g_2 = w$ and $g'_2 = \Lambda$, and so:

$$\begin{array}{ccccc}
 C & \xlongequal{\quad} & v w & \xrightarrow{\tau_H^*} & \Lambda \# \text{let } y = w \text{ in } g_3[v/x] & \xlongequal{\quad} & e \\
 \downarrow \tau & & & & \downarrow \tau & & \downarrow \tau \\
 g_3[v/x][w/y] & \xrightarrow{\tau_H^*} & g_3[v/x][w/y] & \xlongequal{\quad} & \Lambda \# \Lambda & \# & g_3[v/x][w/y] & \xlongequal{\quad} & e'
 \end{array}$$

The result follows. \square

References

- Agha, G., Mason, I., Smith, S. and Talcott, C. (1994) A foundation for actor computation. *J. Functional Programming*.
- Baeten, J. C. M. and Vaandrager, F. W. (1992) An algebra for process creation. *Acta Informatica*, **29**(4), 303–334.
- Bergstra, J. A. and Klop, J. W. (1985) Algebra of communicating processes with abstraction. *Theoret. Comput. Sci.*, **37**, 77–121.
- Berry, D., Milner, R. and Turner, D. N. (1992) A semantics for ML concurrency primitives. *Proc. POPL 92*.
- Bolignano, D. and Debabi, M. (1994) A semantic theory for concurrent ML. *Proc. TACS '94*.
- Debabi, M. (1994) Integration de paradigmes de programmation paralle, fonctionnelle et imperative. *PhD thesis*, Universite D'Orsay, France.
- Ferreira, W. (1995) Semantic theories for concurrent ML. *DPhil thesis*, COGS, Sussex University.
- Ferreira, W. and Hennessy, M. (1995) Towards a semantic theory of CML. *Proc. MFCS 95: Lecture Notes in Computer Science 969*. Springer-Verlag.
- Giacalone, A., Mishra, P. and Prasad, S. (1989) Facile: A symmetric integration of concurrent and functional programming. *Proc. Tapsoft 89: Lecture Notes in Computer Science 352*, pp. 184–209. Springer-Verlag.
- Gordon, A. D. (1995) Bisimilarity as a theory of functional programming. *Proc. MFPS 95: Electronic Notes in Computer Science 1*. Springer-Verlag.
- Gunter, C. (1992) *Semantics of Programming Languages*. MIT Press.
- Havelund, K. (1994) The fork calculus: Towards a logic for concurrent ML. *PhD thesis*, École Normale Superieur, Paris.
- Hennessy, M. (1988) *Algebraic Theory of Processes*. MIT Press.
- Hoare, C. A. R. (1985) *Communicating Sequential Processes*. Prentice-Hall.
- Holmström, S. (1983) PFL: A functional language for parallel programming. *Proc. Declarative Programming Workshop*, pp. 114–139.
- Howe, D. (1989) Equality in lazy computation systems. *Proc. LICS 89*, pp. 198–203.
- Howe, D. (1992) *Proving congruence of simulation orderings in functional languages*. Unpublished manuscript.
- ISO 8807 (1989) *LOTOS – a formal description technique based on the temporal ordering of observational behaviour*.
- Jeffrey, A. S. A. (1995) A fully abstract semantics for a concurrent functional language with monadic types. *Proc. LICS 95*, pp. 255–264.
- Milner, R. (1989) *Communication and Concurrency*. Prentice-Hall.

- Milner, R. (1991) The polyadic π -calculus: a tutorial. *Proc. International Summer School on Logic and Algebra of Specification*.
- Milner, R., Parrow, J. and Walker, D. (1992) A calculus of mobile processes. *Inform. & Comput.*, **100**(1), 1–77.
- Moggi, E. (1991) Notions of computation and monad. *Inform. & Comput.*, **93**, 55–92.
- Nielson, F. and Nielson, H. R. (1993) *From CML to process algebras*. Report DAIMI FN-19, Department of Computer Science, Aarhus University.
- Nielson, F. and Nielson, H. R. (1996) Communication analysis for concurrent ML. *ML with Concurrency*. Springer-Verlag.
- Nikhil, S. (1990) *Id reference manual*. MIT Laboratory for Computer Science.
- Panangaden, P. and Reppy, J. (1996). The essence of concurrent ML. *ML with Concurrency*. Springer-Verlag.
- Pierce, B. C. and Turner, D. N. (1995) *Pict: A programming language based on the pi-calculus*. Technical report in preparation. (Available electronically from <http://www.cl.cam.ac.uk/users/bcp1000/ftp/index.html>)
- Plotkin, G. (1975) Call-by-name, call-by-value and the lambda-calculus. *Theoret. Comput. Sci.*, **1**, 125–159.
- Reppy, J. (1991a) A higher-order concurrent language. *Proc. SIGPLAN 91*, pp. 294–305.
- Reppy, J. (1991b) *An operational semantics of first-class synchronous operations*. Technical report TR 91-1232, Department of Computer Science, Cornell University.
- Reppy, J. (1992) Higher-order concurrency. *PhD thesis*, Cornell University.
- Sangiorgi, D. (1992) Expressing mobility in process algebras: First-order and higher-order paradigms. *PhD thesis*, LFCS, Edinburgh University.
- Thomsen, B. (1995) A theory of higher order communicating systems. *Inform. & Comput.*, **116**(1), 38–57.
- Wright, A. and Felleisen, M. (1991) *A syntactic approach to type soundness*. Technical report TR91-160, Department of Computer Science, Rice University.