

# Type-safe Higher-order Channels in ML-like Languages \*

Sungwoo Park

Pohang University of Science and Technology  
Republic of Korea  
gla@postech.ac.kr

## Abstract

As a means of transmitting not only data but also code encapsulated within functions, higher-order channels provide an advanced form of task parallelism in parallel computations. In the presence of mutable references, however, they pose a safety problem because references may be transmitted to remote threads where they are no longer valid.

This paper presents an ML-like parallel language with *type-safe* higher-order channels. By type safety, we mean that no value written to a channel contains references, or equivalently, that no reference escapes via a channel from the thread where it is created. The type system uses a typing judgment that is capable of deciding whether the value to which a term evaluates contains references or not. The use of such a typing judgment also makes it easy to achieve another desirable feature of channels, *channel locality*, that associates every channel with a unique thread for serving all values addressed to it.

Our type system permits mutable references in sequential computations and also ensures that mutable references never interfere with parallel computations. Thus it provides both flexibility in sequential programming and ease of implementing parallel computations.

**Categories and Subject Descriptors** D.3.1 [Formal Definitions and Theory]: Semantics, Syntax

**General Terms** Languages

**Keywords** Higher-order channels, Channel locality, Parallel languages, Distributed languages

## 1. Introduction

The advent of multicore processors and the impending demise of free lunch (Sutter 2005) have changed the conventional wisdom on computer hardware (Asanovic et al. 2006). There is little room for increasing clock frequency while increasing hardware parallelism is now the only viable way of improving processor performance. Such a radical change in the trend of computer hardware has revitalized research on language support for parallel programming, as

\*This work was supported by the Korea Science and Engineering Foundation (KOSEF) grant funded by the Korea government (MOST) (No. R01-2007-000-11087-0) and POSTECH BSRI research fund-2007.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICFP'07, October 1–3, 2007, Freiburg, Germany.  
Copyright © 2007 ACM 978-1-59593-815-2/07/0010...\$5.00

evidenced by recent parallel languages such as X10 (Charles et al. 2005), Fortress (Allan et al. 2007), Chapel (Cray Inc. 2005), Data Parallel Haskell (Chakravarty et al. 2007), and Manticore (Fluet et al. 2007).

Depending on the granularity of parallel computations, parallel programming models are divided into *data parallelism* or *task parallelism*. Data parallelism applies an independent operation to each element of a collection of homogeneous data and often results in massive parallel computations. A simple case of data parallelism is *flat data parallelism* in which the independent operation is sequential. A more sophisticated model called *nested data parallelism* (Blelloch 1996) allows the independent operation itself to be a parallel computation. Task parallelism executes cooperative threads in parallel which communicate via shared memory or channels in order to perform synchronization. The use of shared memory simplifies programming tasks, especially if a high-level abstraction such as software transactional memory (Shavit and Touitou 1995) is provided, but it entails the memory consistency problem. The use of channels empirically requires more effort than programming with shared memory (Hochstein et al. 2005), but it eliminates the memory consistency problem.

This paper is primarily concerned with an extension of an ML-like language, *i.e.*, a call-by-value functional language with mutable references, that offers task parallelism with *higher-order channels* (but without shared memory). Higher-order channels transmit not only data, such as integers and channel names, but also pieces of code encapsulated within functions. The capability to transmit code between threads opens a new range of communication constructs such as futures, remote evaluation, code on demand, and hot code replacement.

Because of the similarity between task parallelism in parallel computations and process parallelism in distributed computations, our target language can also be thought of as a distributed language in which processes share no global memory and communicate only via higher-order channels. Thus, although our work primarily aims at designing higher-order channels for parallel languages, it can be equally applied to distributed languages.

### 1.1 Type-safe higher-order channels

The main problem with higher-order channels is that in the presence of mutable references, code containing references may travel between threads, but in the absence of shared memory, a reference created by a thread cannot be dereferenced at another thread. Previous parallel or distributed languages with a similar programming model avoid this problem by dispensing with mutable references altogether, as in the parallel language Manticore (Fluet et al. 2007), or by designing the runtime system so as to create copies of heap cells whenever their references are transmitted, as in the distributed languages Facile (Knabe 1995) and JoCAML (Fournet et al. 2003).

This paper takes a different approach by developing *type-safe* higher-order channels. The type system ensures that no reference

escapes from the thread where it is created so that all values written to channels remain valid even after being transmitted to remote threads. To be specific, suppose that a thread is executing a channel write  $a!M$  where  $a$  is a channel for a certain type  $A$  and  $M$  is a term of the same type  $A$ . Evaluating  $M$  yields a value  $V$  which is then transmitted to another thread executing a channel read  $a?$ . Type safety of channels means that  $V$  contains no references and that  $V$  is valid at both threads regardless of its type  $A$ .

Conventional type systems, however, are inadequate to guarantee type safety of channels. The crux of the problem is that a typical typing judgment  $M : A$  does not indicate whether the result of evaluating  $M$  contains references or not. For example, a term

$$(\lambda v:\text{int}. \lambda x:\text{int}. x + v) 0$$

of type  $\text{int} \rightarrow \text{int}$  evaluates to a value containing no reference, but another term

$$(\lambda r:\text{ref int}. \lambda x:\text{int}. x + !r) (\text{ref } 0)$$

of the same type evaluates to a value containing a reference. In order to guarantee type safety of channels, therefore, we need at least two kinds of typing judgments: an ordinary typing judgment asserting that a given term evaluates to a value that may contain references, or a *local value*, and another stronger typing judgment asserting that a given term evaluates to a value containing no references, or a *global value*.

Our work is based on the type system in our previous work (Park 2006) which uses two typing judgments to distinguish between local values and global values in the context of distributed computations. We simplify the type system by combining the two typing judgments into a single typing judgment  $M : A_{@L}$  where  $L$  is a *locality* indicating whether the value to which  $M$  evaluates is local ( $L = \text{L}$ ) or global ( $L = \text{G}$ ). The connection between two localities  $\text{G}$  and  $\text{L}$  is established by a new construct  $\text{box } M$  and a modal type  $\Box A$  such that  $M : A_{@G}$  implies  $\text{box } M : \Box A_{@L}$  (i.e., if  $M$  evaluates to a global value of type  $A$ , then  $\text{box } M$  has type  $\Box A$ ). Now a channel write  $a!M$  typechecks only if  $M : A_{@G}$  holds. As a result, a channel read  $a?$  always returns a global value and thus  $a? : A_{@G}$  automatically holds.

## 1.2 Channel locality

The type system developed for higher-order channels makes it easy to achieve an important feature of task parallelism, *channel locality*, which states that every channel is associated with a unique thread for reading off all values sent to it. Channel locality obviates the need for a sophisticated mechanism in the runtime system for dynamically determining the destination of each value written to a channel.

Most of the previous approaches to enforcing channel locality (Fournet et al. 1996; Amadio 1997; Yoshida and Hennessy 1999; Schmitt and Stefani 2003) have been developed for calculi for concurrent processes based on the pi-calculus. We find that these approaches are difficult to apply to our setting which uses as a base language the simply-typed lambda calculus with mutable references instead of the pi-calculus or its variant.

The main idea for achieving channel locality in our work is to split channels into two kinds, *read channels* and *write channels*, and treat read channels as local values but write channels as global values.  $\text{new}_{\langle A \rangle}$ , a construct for creating channels for type  $A$ , evaluates to a pair of read channel  $a^r$  and write channel  $a^w$  such that  $a^r$  accepts only those values written to  $a^w$ . Since  $a^r$  is a local value, channel reads from  $a^r$  are allowed only at the thread where  $a^r$  is created. Channel writes to  $a^w$  are, however, allowed at any thread because  $a^w$  is a global value. Thus a pair of  $a^r$  and  $a^w$  can open unidirectional communications from any thread to the thread to which  $a^r$  belongs.

It is easy to classify read channels as local values and write channels as global values. First we assign different types, namely *read channel types* and *write channel types*, to read channels and write channels, respectively. Then we treat read channel types like reference types so that read channels cannot be part of a global value, but define write channel types as primitive types whose values are all inherently global (like integers). For this reason, our decision to distinguish between read channels and write channels has a different motivation from previous work in which read channels and write channels are also distinguished, but channel locality is not enforced (Odersky 1995) or enforced only syntactically (Zhang and Potter 2002).

## 1.3 Contributions

We develop an ML-like parallel language  $\lambda_{\Box}^{\text{PC}}$  which features type-safe higher-order channels with channel locality. Its type system inherits from the type system of (Park 2006) the ability to distinguish between local values and global values. Its operational semantics models parallel computations where multiple threads communicate via higher-order channels. Channel locality is a direct consequence of assigning different types to read channels and write channels. To the best of our knowledge, type-safe higher-order channels in the presence of mutable references have not been investigated.

Although  $\lambda_{\Box}^{\text{PC}}$  deals primarily with type safety for task parallelism, providing type safety for data parallelism is also straightforward by virtue of its ability to distinguish between local values and global values. As an example, consider a parallel map construct  $\text{mapP}$  which applies a function  $f$  to each element of an array in parallel. By requiring that  $f$  be a global value so that child threads share no references, we can prevent  $\text{mapP}$  from running into the memory consistency problem. In this regard,  $\lambda_{\Box}^{\text{PC}}$  serves as a unified framework for achieving type safety for both data parallelism and task parallelism when mutable references are allowed.

Removing mutable references simplifies the implementation of a parallel language because lack of mutable references implies automatic data separation in parallel computations. For example, Manticore (Fluet et al. 2007) excludes mutable references in its base language (a subset of Standard ML) in order to simplify its implementation. Parallel dialects of Haskell, such as pH (Nikhil and Arvind 2001) and Data Parallel Haskell (Chakravarty et al. 2007), also benefit from lack of mutable references. In comparison,  $\lambda_{\Box}^{\text{PC}}$  permits mutable references in sequential computations, but its type system ensures that mutable references never interfere with parallel computations. Thus the type system of  $\lambda_{\Box}^{\text{PC}}$  wins us both flexibility in sequential programming and ease of implementing parallel computations.

## 1.4 Organization of the paper

Section 2 presents the base language  $\lambda_{\Box}$  for our work. Although the type system differentiates global values from local values,  $\lambda_{\Box}$  is just a sequential language to which global values are of no use. Hence we develop a parallel operational semantics for modeling parallel computations. Section 3 presents the resultant language  $\lambda_{\Box}^{\text{P}}$  which comes with a new construct  $\text{spawn } M$  for creating new threads. Section 4 extends  $\lambda_{\Box}^{\text{P}}$  with new constructs for higher-order channels to obtain  $\lambda_{\Box}^{\text{PC}}$ . Section 5 illustrates how to implement various communication constructs in  $\lambda_{\Box}^{\text{PC}}$  such as futures, bidirectional communications, shared references, remote evaluation, code on demand, and hot code replacement. Section 6 discusses two extensions of  $\lambda_{\Box}^{\text{PC}}$  (including type safety for data parallelism). Section 7 discusses related work and Section 8 concludes.

type	$A, B, C$	$::=$	$\text{unit} \mid A \rightarrow A \mid A \times A \mid A+A \mid \text{ref } A \mid \Box A$
primitive type	$P$		
term	$M, N$	$::=$	$x \mid \lambda x:A. M \mid M M \mid$ $(M, M) \mid \text{fst } M \mid \text{snd } M \mid$ $\text{inl } M \mid \text{inr } M \mid \text{case } M \text{ of inl } x \Rightarrow M \mid \text{inr } x \Rightarrow M \mid$ $() \mid \text{fix } x:A. M \mid$ $\text{ref } M \mid !M \mid M := M \mid l \mid$ $\text{box } M \mid \text{letbox } x = M \text{ in } M$
value	$V$	$::=$	$() \mid \lambda x:A. M \mid (V, V) \mid \text{inl } V \mid \text{inr } V \mid l \mid \text{box } M$
locality	$L$	$::=$	$\text{G} \mid \text{L}$
typing context	$\Gamma$	$::=$	$\cdot \mid \Gamma, x : A_{\text{@}L}$
store	$\psi$	$::=$	$\cdot \mid \psi, l \mapsto V$
store typing	$\Psi$	$::=$	$\cdot \mid \Psi, l \mapsto A$

Figure 1. Abstract syntax for  $\lambda_{\Box}$

## 2. Base language $\lambda_{\Box}$

This section presents the base language  $\lambda_{\Box}$  which is a reformulation of the call-by-value language with modal types  $\Box A$  in our previous work (Park 2006). In the previous work, we use two separate typing judgments to distinguish between local values (which may contain references) and global values (which contains no references): an ordinary typing judgment  $M : A$  to mean that  $M$  evaluates to a local value, and a stronger typing judgment  $M \sim A$  to mean that  $M$  evaluates to global value. We combine the two typing judgments into a single typing judgment  $M : A_{\text{@}L}$  where a locality  $L$ , either L or G, indicates whether  $M$  evaluates to a local value or a global value. The simplification of typing judgments is essential to maintaining the complexity of the type system at a manageable level. For example, the introduction of sum types requires eight new typing rules in the type system of (Park 2006) whereas the base language  $\lambda_{\Box}$  here uses only three typing rules.

### 2.1 Definition of $\lambda_{\Box}$

Figure 1 shows the abstract syntax for  $\lambda_{\Box}$  which is based on the simply-typed  $\lambda$ -calculus with product types  $A \times A$ , sum types  $A+A$ , reference types  $\text{ref } A$ , and the fixed point construct  $\text{fix } x:A. M$ . (We will define the set of primitive types later.)  $\text{ref } M$  allocates a fresh reference,  $!M$  dereferences an existing reference, and  $M := N$  assigns a new value to a reference. A location  $l$ , of type  $\text{ref } A$ , is a value for a reference. A new construct  $\text{box } M$  has a modal type  $\Box A$ , and another new construct  $\text{letbox } x = M \text{ in } N$  expects  $M$  to be of type  $\Box A$ .

We use a typing judgment  $\Gamma \mid \Psi \vdash M : A_{\text{@}L}$  to mean that under typing context  $\Gamma$  and store typing  $\Psi$ , term  $M$  evaluates to a value of type  $A$  with locality  $L$ . The resultant value is local if  $L = \text{L}$  and global if  $L = \text{G}$ . (That is, L means “here only” and G “everywhere.”) A binding  $x : A_{\text{@}L}$  in a typing context  $\Gamma$  means that  $x$  holds a local value of type  $A$  if  $L = \text{L}$ , or a global value of type  $A$  if  $L = \text{G}$ . A store  $\psi$  maps locations to values, and a store typing  $\Psi$  maps locations to types of these values. We assume a relation  $\text{G} < \text{L}$  to reflect the fact that a global value may be used as a local value, but not vice versa.

Figure 2 shows the type system of  $\lambda_{\Box}$ . The rule Var uses  $L \leq L'$  to mean either  $L = L'$  or  $L < L'$ . The rules  $\rightarrow \text{I}$  through Loc are all derived from typing rules in the simply-typed  $\lambda$ -calculus by annotating each typing judgment with a locality L. Each of the rules  $\times \text{I}$ ,  $\times \text{E}_\text{L}$ , and  $\times \text{E}_\text{R}$  for product types uses the same unspecified locality  $L$  in its premise and conclusion. Note that these rules make sense only under the call-by-value semantics where  $(M, N)$  is a value only when both  $M$  and  $N$  are also values. If  $(M, N)$  is considered as a value for any two terms  $M$  and  $N$ , for example, all these rules become incorrect when  $L = \text{G}$ . Similarly the rules

$+ \text{I}_\text{L}$ ,  $+ \text{I}_\text{R}$ , and  $+ \text{E}$  for sum types make sense only under the call-by-value semantics.

The modality  $\Box$  has an introduction rule  $\Box \text{I}$  which states that  $M$  in  $\text{box } M$  always evaluates to a global value. The conclusion of the rule  $\Box \text{I}$  uses a locality L because  $\text{box } M$  itself may not be a global value if  $M$  contains references. For example, if  $l$  is a location of type  $\text{ref int}$ ,  $\text{box } !l$  is not a global value although it has a modal type  $\Box \text{int}$ . The elimination rule  $\Box \text{E}$  uses an unspecified locality  $L$  in the conclusion because  $M$  in  $\text{letbox } x = M \text{ in } N$  does not specify whether  $N$  evaluates to a local value or a global value. That is, regardless of the type of  $M$ , the value to which  $N$  evaluates can still be either local or global.

The rule GVal is the main rule for connecting the two localities L and G. Its premise uses the following definition of  $\Gamma_\text{G}$  which extracts bindings for variables holding global values from  $\Gamma$ :

$$\Gamma_\text{G} = \{x : A_{\text{@}G} \mid x : A_{\text{@}G} \in \Gamma\}$$

Then the premise  $\Gamma_\text{G} \mid \cdot \vdash V : A_{\text{@}L}$  states that  $V$  is a value that uses no local values (because of  $\Gamma_\text{G}$ ) and no references (because of an empty store typing); hence  $V$  is a mobile value. Since  $V$  is already a value and thus requires no further evaluation, we may say that  $V$  evaluates to a global value, which is expressed in the conclusion  $\Gamma \mid \Psi \vdash V : A_{\text{@}G}$ .

The rule Prim uses the notion of primitive type to provide another way to connect the two localities L and G. A type is primitive if all its values are inherently mobile. For example,  $\text{unit}$  is a primitive type because its only value,  $()$ , typechecks under any typing context and store typing, as shown in the rule Unit. (Other examples of primitive types would be  $\text{int}$  for integers and  $\text{bool}$  for boolean values.) Formally we define primitive types as follows:

**Definition 2.1.**  $P$  is a primitive type if  $\Gamma \mid \Psi \vdash V : P_{\text{@}L}$  implies  $\Gamma_\text{G} \mid \cdot \vdash V : P_{\text{@}L}$ .

Then terms of primitive types always evaluate to mobile values, and  $\Gamma \mid \Psi \vdash M : P_{\text{@}L}$  automatically implies  $\Gamma \mid \Psi \vdash M : P_{\text{@}G}$  as shown in the rule Prim.

Under the type system in Figure 2, the rule Unit justifies the use of  $\text{unit}$  as a primitive type. In addition,  $P_1 \times P_2$  and  $P_1 + P_2$  are primitive types if both  $P_1$  and  $P_2$  are primitive types, since values of product type  $P_1 \times P_2$  have the form  $(V_1, V_2)$ , and values of sum type  $P_1 + P_2$  have the form  $\text{inl } V$  or  $\text{inr } V$ . Thus we use the following set of primitive types for  $\lambda_{\Box}$ :

$$\text{primitive type } P ::= \text{unit} \mid P \times P \mid P + P$$

Proposition 2.2 justifies the relation  $\text{G} < \text{L}$ .

**Proposition 2.2.**

The rule  $\frac{\Gamma \mid \Psi \vdash M : A_{\text{@}G}}{\Gamma \mid \Psi \vdash M : A_{\text{@}L}}$  Global is admissible.

$$\begin{array}{c}
\frac{x : A_{\text{@L}} \in \Gamma \quad L \leq L'}{\Gamma \mid \Psi \vdash x : A_{\text{@L}'}} \text{Var} \quad \frac{\Gamma, x : A_{\text{@L}} \mid \Psi \vdash M : B_{\text{@L}}}{\Gamma \mid \Psi \vdash \lambda x : A. M : A \rightarrow B_{\text{@L}}} \rightarrow \mid \quad \frac{\Gamma \mid \Psi \vdash M : A \rightarrow B_{\text{@L}} \quad \Gamma \mid \Psi \vdash N : A_{\text{@L}}}{\Gamma \mid \Psi \vdash M N : B_{\text{@L}}} \rightarrow \text{E} \\
\frac{}{\Gamma \mid \Psi \vdash () : \text{unit}_{\text{@L}}} \text{Unit} \quad \frac{\Gamma, x : A_{\text{@L}} \mid \Psi \vdash M : A_{\text{@L}}}{\Gamma \mid \Psi \vdash \text{fix } x : A. M : A_{\text{@L}}} \text{Fix} \quad \frac{\Gamma \mid \Psi \vdash M : A_{\text{@L}}}{\Gamma \mid \Psi \vdash \text{ref } M : \text{ref } A_{\text{@L}}} \text{Ref} \\
\frac{\Gamma \mid \Psi \vdash M : \text{ref } A_{\text{@L}}}{\Gamma \mid \Psi \vdash !M : A_{\text{@L}}} \text{Drf} \quad \frac{\Gamma \mid \Psi \vdash M : \text{ref } A_{\text{@L}} \quad \Gamma \mid \Psi \vdash N : A_{\text{@L}}}{\Gamma \mid \Psi \vdash M := N : \text{unit}_{\text{@L}}} \text{Asn} \quad \frac{\Psi(l) = A}{\Gamma \mid \Psi \vdash l : \text{ref } A_{\text{@L}}} \text{Loc} \\
\frac{\Gamma \mid \Psi \vdash M : A_{\text{@L}} \quad \Gamma \mid \Psi \vdash M : B_{\text{@L}}}{\Gamma \mid \Psi \vdash (M, N) : A \times B_{\text{@L}}} \times \mid \quad \frac{\Gamma \mid \Psi \vdash M : A \times B_{\text{@L}}}{\Gamma \mid \Psi \vdash \text{fst } M : A_{\text{@L}}} \times \text{E}_L \quad \frac{\Gamma \mid \Psi \vdash M : A \times B_{\text{@L}}}{\Gamma \mid \Psi \vdash \text{snd } M : B_{\text{@L}}} \times \text{E}_R \\
\frac{\Gamma \mid \Psi \vdash M : A_{\text{@L}}}{\Gamma \mid \Psi \vdash \text{inl } M : A + B_{\text{@L}}} + \text{L} \quad \frac{\Gamma \mid \Psi \vdash M : B_{\text{@L}}}{\Gamma \mid \Psi \vdash \text{inr } M : A + B_{\text{@L}}} + \text{R} \quad \frac{\Gamma \mid \Psi \vdash M : A + B_{\text{@L}'}}{\Gamma \mid \Psi \vdash \text{case } M \text{ of inl } x \Rightarrow N \mid \text{inr } x' \Rightarrow N' : C_{\text{@L}}} + \text{E} \\
\frac{\Gamma \mid \Psi \vdash M : A_{\text{@G}}}{\Gamma \mid \Psi \vdash \text{box } M : \square A_{\text{@L}}} \square \mid \quad \frac{\Gamma \mid \Psi \vdash M : \square A_{\text{@L}} \quad \Gamma, x : A_{\text{@G}} \mid \Psi \vdash N : C_{\text{@L}}}{\Gamma \mid \Psi \vdash \text{letbox } x = M \text{ in } N : C_{\text{@L}}} \square \text{E} \\
\frac{\Gamma_{\text{G}} \mid \cdot \vdash V : A_{\text{@L}}}{\Gamma \mid \Psi \vdash V : A_{\text{@G}}} \text{GVal} \quad \frac{\Gamma \mid \Psi \vdash M : P_{\text{@L}}}{\Gamma \mid \Psi \vdash M : P_{\text{@G}}} \text{Prim}
\end{array}$$

Figure 2. Type system of  $\lambda_{\square}$

$$\begin{array}{c}
\text{evaluation context } \phi ::= \square \mid \phi M \mid V \phi \mid (\phi, M) \mid (V, \phi) \mid \text{fst } \phi \mid \text{snd } \phi \mid \\
\text{inl } \phi \mid \text{inr } \phi \mid \text{case } \phi \text{ of inl } x \Rightarrow M \mid \text{inr } x \Rightarrow M \text{ ref } \phi \mid !\phi \mid \phi := M \mid V := \phi \mid \\
\text{letbox } x = \phi \text{ in } M \mid \text{letbox } x = \text{box } \phi \text{ in } M \\
(\lambda x : A. M) V \rightarrow_{\beta} [V/x]M \quad \text{fst } (V, V') \rightarrow_{\beta} V \\
\text{fix } x : A. M \rightarrow_{\beta} [\text{fix } x : A. M/x]M \quad \text{snd } (V, V') \rightarrow_{\beta} V' \\
\text{letbox } x = \text{box } V \text{ in } M \rightarrow_{\beta} [V/x]M \quad \text{case inl } V \text{ of inl } x \Rightarrow M \mid \text{inr } x' \Rightarrow M' \rightarrow_{\beta} [V/x]M \\
\text{case inr } V \text{ of inl } x \Rightarrow M \mid \text{inr } x' \Rightarrow M' \rightarrow_{\beta} [V/x']M' \\
\frac{M \rightarrow_{\beta} M'}{\phi[M] \mid \psi \rightarrow \phi[M'] \mid \psi} R_{\beta} \\
\frac{\text{fresh } l \notin \text{dom}(\psi)}{\phi[\text{ref } V] \mid \psi \rightarrow \phi[l] \mid \psi, l \mapsto V} \text{Ref} \quad \frac{\psi(l) = V}{\phi[l] \mid \psi \rightarrow \phi[V] \mid \psi} \text{Drf} \quad \frac{}{\phi[l := V] \mid \psi \rightarrow \phi[l] \mid \psi, [l \mapsto V]\psi} \text{Asn}
\end{array}$$

Figure 3. Operational semantics for  $\lambda_{\square}$

Figure 3 shows the operational semantics for  $\lambda_{\square}$ . It uses a reduction judgment  $M \mid \psi \rightarrow M' \mid \psi'$  which means that term  $M$  with store  $\psi$  reduces to term  $M'$  with store  $\psi'$ . A  $\beta$ -reduction  $M \rightarrow_{\beta} M'$  uses a capture-avoiding substitution  $[N/x]M$  defined in a standard way.  $\phi[M]$  fills the hole  $\square$  in evaluation context  $\phi$  with term  $M$ .  $[l \mapsto V]\psi$  replaces  $l \mapsto V'$  in store  $\psi$  by  $l \mapsto V$ .  $\psi(l)$  denotes the value to which  $l$  is mapped under  $\psi$ ;  $\text{dom}(\psi)$  denotes the set of locations mapped under  $\psi$ .

Since  $\text{letbox } x = \text{box } \phi \text{ in } N$  is an evaluation context,  $\text{letbox } x = \text{box } M \text{ in } N$  further evaluates  $M$  so as to substitute the resultant value for  $x$  in  $N$ , even though  $\text{box } M$  is already a value. Thus  $\text{letbox } x = M' \text{ in } N$  first evaluates  $M'$  to identify how to obtain a value to be substituted for  $x$ , and then obtains such a value by evaluating  $M$ , if  $M'$  evaluates to  $\text{box } M$ . Since  $M$  evaluates to a global value, all occurrences of  $x$  in  $N$  become bound to a global value, as required by the rules  $\square \text{E}$ .

## 2.2 Type safety of $\lambda_{\square}$

The proof of type safety of  $\lambda_{\square}$  needs a store typing judgment  $\Psi \vdash \psi \text{ okay}$  which means that store  $\psi$  conforms to store typing  $\Psi$ .  $\Psi(l)$  denotes the type to which  $l$  is mapped under  $\Psi$ ;  $\text{dom}(\Psi)$

denotes the set of locations mapped under  $\Psi$ .

$$\frac{\text{for every } l \in \text{dom}(\psi) \quad \text{dom}(\Psi) = \text{dom}(\psi) \quad \cdot \mid \Psi \vdash \psi(l) : \Psi(l)_{\text{@L}}}{\Psi \vdash \psi \text{ okay}} \text{Store}$$

The proof of progress (Theorem 2.3) uses a canonical forms lemma, as usual. The proof of type preservation (Theorem 2.6) uses a substitution theorem (Theorem 2.4). It also uses Lemma 2.5 whose proof uses the definition of primitive types (Definition 2.1).

**Theorem 2.3 (Progress).** *Suppose  $\cdot \mid \Psi \vdash M : A_{\text{@L}}$ . Then either:*

- (1)  $M$  is a value, or
- (2) for any store  $\psi$  such that  $\Psi \vdash \psi \text{ okay}$ , there exist some term  $M'$  and store  $\psi'$  such that  $M \mid \psi \rightarrow M' \mid \psi'$ .

**Theorem 2.4 (Substitution).**

*If  $\Gamma \mid \Psi \vdash N : A_{\text{@L}}$ , then*

$$\Gamma, x : A_{\text{@L}} \mid \Psi \vdash M : C_{\text{@L}} \text{ implies } \Gamma \mid \Psi \vdash [N/x]M : C_{\text{@L}}.$$

*If  $\Gamma_{\text{G}} \mid \cdot \vdash V : A_{\text{@L}}$ , then*

$$\Gamma, x : A_{\text{@G}} \mid \Psi \vdash M : C_{\text{@L}} \text{ implies } \Gamma \mid \Psi \vdash [V/x]M : C_{\text{@L}}.$$

term	$M$	::=	$\dots \mid \text{spawn } M$
evaluation context	$\phi$	::=	$\dots \mid \text{spawn } \phi \mid \text{spawn box } \phi$
thread	$\gamma$		
configuration	$\pi$	::=	$\cdot \mid \pi, \{M \mid \psi @ \gamma\}$
configuration typing	$\Pi$	::=	$\cdot \mid \Pi, \gamma : A$

$\frac{\Gamma \mid \Psi \vdash M : \Box(\text{unit} \rightarrow A)_{\text{@L}}}{\Gamma \mid \Psi \vdash \text{spawn } M : \text{unit}_{\text{@L}}}$ Spawn	
$dom(\Pi) = dom(\pi)$	for each $\{M \mid \psi @ \gamma\} \in \pi$ , there exist $A_\gamma, \Psi_\gamma$ such that $\gamma : A_\gamma \in \Pi$ $\Psi_\gamma \vdash \psi \text{ okay}$ $\cdot \mid \Psi_\gamma \vdash M : A_{\gamma_{\text{@L}}}$
$\Pi \vdash \pi \text{ okay}$ Conf	
$\frac{M \mid \psi \rightarrow M' \mid \psi'}{\pi, \{M \mid \psi @ \gamma\} \Rightarrow \pi, \{M' \mid \psi' @ \gamma\}}$ Red	
$\frac{\text{fresh } \gamma'}{\pi, \{\phi[\text{spawn box } V] \mid \psi @ \gamma\} \Rightarrow \pi, \{\phi[()] \mid \psi @ \gamma\}, \{V () \mid \cdot @ \gamma'\}}$ Spawn	

**Figure 4.** Definition of  $\lambda_{\Box}^P$

**Lemma 2.5.** *If  $\cdot \mid \Psi \vdash V : A_{\text{@G}}$ , then  $\cdot \mid \cdot \vdash V : A_{\text{@L}}$ .*

*Proof.* By induction on the structure of the proof of  $\cdot \mid \Psi \vdash V : A_{\text{@G}}$ . We need to consider the rules Prim, GVal,  $\times L$ ,  $+L$ , and  $+R$ . In the case of the rule Prim,  $\cdot \mid \Psi \vdash V : A_{\text{@L}}$  holds where  $A$  is a primitive type, and  $\cdot \mid \cdot \vdash V : A_{\text{@L}}$  follows by Definition 2.1.  $\square$

**Theorem 2.6** (Type preservation).

*Suppose  $\cdot \mid \Psi \vdash M : A_{\text{@L}}$ ,  $\Psi \vdash \psi \text{ okay}$ , and  $M \mid \psi \rightarrow M' \mid \psi'$ . Then there exists a store typing  $\Psi'$  such that  $\cdot \mid \Psi' \vdash M' : A_{\text{@L}}$ ,  $\Psi \subset \Psi'$ , and  $\Psi' \vdash \psi' \text{ okay}$ .*

### 3. $\lambda_{\Box}^P$ with a parallel operational semantics

Although its type system uses localities and modal types to differentiate global values from local values,  $\lambda_{\Box}$  is just a base language for sequential computations to which global values are of no use. That is, its operational semantics focuses only on the sequential computation at a hypothetical thread and there is no way to exploit global values for communications between threads. This section extends  $\lambda_{\Box}$  with a parallel operational semantics which models multiple threads running concurrently (but not communicating with each other yet). The resultant language is called  $\lambda_{\Box}^P$ .

Figure 4 shows the definition for  $\lambda_{\Box}^P$ . At the syntax level,  $\lambda_{\Box}^P$  augments  $\lambda_{\Box}$  with a new construct  $\text{spawn } M$  whose typing rule Spawn requires  $M$  to be of type  $\Box(\text{unit} \rightarrow A)$ .  $\text{spawn } M$  starts a fresh thread with an empty store by evaluating a term obtained as follows. First it evaluates  $M$  to obtain  $\text{box } M'$ . Then it evaluates the inner term  $M'$  (as indicated by the evaluation context  $\text{spawn box } \phi$ ) to obtain  $\lambda_{\cdot} : \text{unit}$ .  $M''$  (or a similar form of thunk). Since  $\text{box } M'$  has type  $\Box(\text{unit} \rightarrow A)$ ,  $M'$  evaluates to a global value, which means that  $\lambda_{\cdot} : \text{unit}$ .  $M''$  contains no references. Therefore it is safe to evaluate  $(\lambda_{\cdot} : \text{unit}. M'') ()$  at a fresh thread with an empty store, which effectively starts to evaluate  $M''$  at the fresh thread. In a realistic language, we could use syntactic sugar  $\text{spawnT } \{M\}$  defined as  $\text{spawn box } \lambda_{\cdot} : \text{unit}. M$  in order to start a fresh thread by evaluating  $M$  without such intervening evaluations:

$$\frac{\Gamma_G \mid \cdot \vdash M : A_{\text{@L}}}{\Gamma \mid \Psi \vdash \text{spawnT } \{M\} : \text{unit}_{\text{@L}}} \text{SpawnTerm}$$

Note that the premise  $\Gamma_G \mid \cdot \vdash M : A_{\text{@L}}$  implies  $\Gamma \mid \Psi \vdash \text{box } \lambda_{\cdot} : \text{unit}. M : \Box(\text{unit} \rightarrow A)_{\text{@L}}$ .

A *configuration*  $\pi$ , as an unordered set, represents the state of a parallel computation by associating each thread  $\gamma$  with a term  $M$  being evaluated at  $\gamma$  and a store  $\psi$  allocated at  $\gamma$ , written as  $\{M \mid \psi @ \gamma\}$ . A *configuration typing*  $\Pi$  records the type of the term being evaluated at each thread. The type system of  $\lambda_{\Box}^P$  uses the same typing judgment as  $\lambda_{\Box}$ , except that it also uses a *configuration typing judgment*  $\Pi \vdash \pi \text{ okay}$  to mean that configuration  $\pi$  has configuration typing  $\Pi$ . The rule Conf may be regarded as the definition of the configuration typing judgment, where  $dom(\Pi)$  and  $dom(\pi)$  denote the set of threads in  $\Pi$  and  $\pi$ , respectively. Note that for each thread  $\gamma$  such that  $\{M \mid \psi @ \gamma\} \in \pi$ , it infers a store typing  $\Psi_\gamma$  by typechecking all locations present in  $M$  and  $\psi$ .

The parallel operational semantics uses a *configuration transition judgment*  $\pi \Rightarrow \pi'$  to mean that configuration  $\pi$  reduces (by the rule Red) or evolves (by the rule Spawn) to configuration  $\pi'$ . The rule Red says that a parallel computation consists primarily of sequential computations performed at individual threads. The rule Spawn starts a new thread  $\gamma'$  by evaluating  $V ()$ . Since  $\text{box } V$  guarantees that  $V$  is a global value, it is safe to evaluate  $V ()$  with an empty store at  $\gamma'$ . Note that a configuration transition is non-deterministic because a configuration may choose an arbitrary thread to which either rule is applied.

Type safety of  $\lambda_{\Box}^P$  consists of *configuration progress* (Theorem 3.1) and *configuration typing preservation* (Theorem 3.2). It is a corollary of type safety of  $\lambda_{\Box}^{\text{PC}}$  to be presented in the next section.

**Theorem 3.1** (Configuration progress). *Suppose  $\Pi \vdash \pi \text{ okay}$ . Then either:*

- (1)  $\pi$  consists only of  $\{V \mid \psi @ \gamma\}$ , or
- (2) there exists  $\pi'$  such that  $\pi \Rightarrow \pi'$ .

**Theorem 3.2** (Configuration typing preservation).

*Suppose  $\Pi \vdash \pi \text{ okay}$  and  $\pi \Rightarrow \pi'$ . Then there exists a configuration typing  $\Pi'$  such that  $\Pi \subset \Pi'$  and  $\Pi' \vdash \pi' \text{ okay}$ .*

### 4. $\lambda_{\Box}^{\text{PC}}$ with higher-order channels

This section extends  $\lambda_{\Box}^P$  with higher-order channels to obtain  $\lambda_{\Box}^{\text{PC}}$ . In order to achieve channel locality,  $\lambda_{\Box}^{\text{PC}}$  splits channels into two kinds: read channels and write channels. A read channel  $a^r$  is assigned a read channel type  $\langle A \rangle^r$ , and is treated as a local value so that it cannot escape the thread at which it is created. A write channel  $a^w$  is assigned a write channel type  $\langle A \rangle^w$ , and is treated as a global value so that it can be transmitted to another thread. Since values written to write channels travel between threads, they cannot be local values. That is, only global values are allowed to be written to write channels, and consequently all values read from read channels are also global.

Figure 5 shows the definition for  $\lambda_{\Box}^{\text{PC}}$ . A read channel  $a^r$  of type  $\langle A \rangle^r$  receives a global value of type  $A$  from its corresponding write channel  $a^w$  via a channel read  $a^r ?$ . A write channel  $a^w$  of type  $\langle A \rangle^w$  transmits a global value  $V$  of type  $A$  to its corresponding read channel  $a^r$  via a channel write  $a^w ! V$ . A new construct  $\text{new}_{\langle A \rangle}$  dynamically creates a pair of read channel  $a^r$  and write channel  $a^w$  to open unidirectional communications from  $a^w$  to  $a^r$ .  $a^r \mapsto A$  in a *read channel typing*  $\Phi^r$  means that  $a^r$  has type  $\langle A \rangle^r$ ; similarly  $a^w \mapsto A$  in a *write channel typing*  $\Phi^w$  means that  $a^w$  has type  $\langle A \rangle^w$ .

The type system of  $\lambda_{\Box}^{\text{PC}}$  uses a read channel typing  $\Phi^r$  and a write channel typing  $\Phi^w$  in each typing judgment:

$$\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash M : A_{\text{@L}}$$

type	$A ::= \dots \mid \langle A \rangle^r \mid \langle A \rangle^w$	
term	$M ::= \dots \mid \mathbf{new}_{\langle A \rangle} \mid a^r \mid a^w \mid M? \mid M!M$	
value	$V ::= \dots \mid a^r \mid a^w$	
evaluation context	$\phi ::= \dots \mid \phi? \mid \phi!M \mid V! \phi$	
read channel typing	$\Phi^r ::= \cdot \mid \Phi^r, a^r \mapsto A$	
write channel typing	$\Phi^w ::= \cdot \mid \Phi^w, a^w \mapsto A$	

  

$\frac{\Gamma_G \mid \Phi^w; \cdot; \cdot \vdash V : A_{\text{GL}}}{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash V : A_{\text{GL}}}$	GVal	$\frac{}{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash \mathbf{new}_{\langle A \rangle} : \langle A \rangle^r \times \langle A \rangle^w_{\text{GL}}}$	New
$\frac{a^r \mapsto A \in \Phi^r}{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash a^r : \langle A \rangle^r_{\text{GL}}}$	ChanR	$\frac{a^w \mapsto A \in \Phi^w}{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash a^w : \langle A \rangle^w_{\text{GL}}}$	ChanW
$\frac{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash M : \langle A \rangle^r_{\text{GL}}}{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash M? : A_{\text{GL}}}$	R?	$(L = G \text{ or } L = L)$	W!
$\frac{\text{for every } l \in \text{dom}(\psi)}{\text{dom}(\Psi) = \text{dom}(\psi) \quad \cdot \mid \Phi^w; \Phi^r; \Psi \vdash \psi(l) : \Psi(l)_{\text{GL}}}$		Store	
$\frac{\text{for each } \{M \mid \psi @ \gamma\} \in \pi, \text{ there exist } A_\gamma, \Phi_\gamma^r, \Psi_\gamma \text{ such that } \gamma : A_\gamma \in \Pi \quad \Phi^w; \Phi_\gamma^r; \Psi_\gamma \vdash \psi \text{ okay}}{\text{dom}(\Pi) = \text{dom}(\pi) \quad \cdot \mid \Phi^w; \Phi_\gamma^r; \Psi_\gamma \vdash M : A_{\gamma_{\text{GL}}}$		Conf	
$\frac{\text{fresh}(a^r, a^w)}{\pi, \{\phi[\mathbf{new}_{\langle A \rangle}]\} \mid \psi @ \gamma \Rightarrow \pi, \{\phi[(a^r, a^w)]\} \mid \psi @ \gamma}$		New	Sync
$\pi, \{\phi[a^r?]\} \mid \psi @ \gamma, \{\phi'[a^w!V]\} \mid \psi' @ \gamma' \Rightarrow \pi, \{\phi[V]\} \mid \psi @ \gamma, \{\phi'[\text{()}\} \mid \psi' @ \gamma'$			

**Figure 5.** Definition of  $\lambda_{\square}^{\text{PC}}$

All the previous typing rules are extended in a straightforward way by adding  $\Phi^r$  and  $\Phi^w$  to each typing judgment. The only exception is the rule GVal whose premise ( $\Gamma_G \mid \Phi^w; \cdot; \cdot \vdash V : A_{\text{GL}}$ ) uses an empty read channel typing because read channels are local values. Accordingly we redefine primitive types as follows:

**Definition 4.1.**

$P$  is a primitive type if and only if  $\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash V : P_{\text{GL}}$  implies  $\Gamma_G \mid \Phi^w; \cdot; \cdot \vdash V : P_{\text{GL}}$ .

Under the new definition of primitive types, a write channel type  $\langle A \rangle^w$  becomes a primitive type (see the rule ChanW):

$$\text{primitive type } P ::= \dots \mid \langle A \rangle^w$$

The rule R? uses an unspecified locality  $L$  in the conclusion so that Proposition 2.2 continues to hold. The rule W! says that channel writes accept only global values. The rule Store uses an extended store typing judgment  $\Phi^w; \Phi^r; \Psi \vdash \psi \text{ okay}$  to check that store  $\psi$  conforms to store typing  $\Psi$  under  $\Phi^r$  and  $\Phi^w$ .

The rule Conf is of particular importance in  $\lambda_{\square}^{\text{PC}}$  because it verifies channel locality in a given configuration. Note that an extended configuration typing judgment  $\Phi^w; \Pi \vdash \pi \text{ okay}$  shares a write channel typing  $\Phi^w$  for all threads (because write channels are global values), but does not assume a specific read channel typing. Instead, for each thread  $\gamma$  such that  $\{M \mid \psi @ \gamma\} \in \pi$ , it infers a *unique* read channel typing  $\Phi_\gamma^r$ , in addition to a store typing  $\Psi_\gamma$ , by typechecking all read channels present in  $M$  and  $\psi$ . The uniqueness of  $\Phi_\gamma^r$  for each thread  $\gamma$ , as stated in the third premise, implies that no two threads share common read channels, and thus amounts to channel locality in  $\pi$ .

The parallel operational semantics uses the same configuration transition judgment  $\pi \Rightarrow \pi'$  as in  $\lambda_{\square}^{\text{P}}$ . The premise of the rule

*New* means that read channel  $a^r$  and write channel  $a^w$  are unique across the entire set of threads. The rule *Sync* says that a channel read  $a^r?$  and a channel write  $a^w!V$  occur synchronously. (An asynchronous version of  $\lambda_{\square}^{\text{PC}}$  would require a different parallel operational semantics, but the same type system would continue to work.)

As with  $\lambda_{\square}^{\text{P}}$ , type safety of  $\lambda_{\square}^{\text{PC}}$  consists of configuration progress (Theorem 4.3) and configuration typing preservation (Theorem 4.5). Proofs of Theorems 4.3 and 4.5 use type safety for sequential computations in  $\lambda_{\square}^{\text{PC}}$  (Propositions 4.2 and 4.4). Note that in Theorem 4.5, a configuration transition  $\pi \Rightarrow \pi'$  preserves channel locality because of the use of extended configuration typing judgments.

**Proposition 4.2 (Progress).**

Suppose  $\cdot \mid \Phi^w; \Phi^r; \Psi \vdash M : A_{\text{GL}}$ . Then either:

- (1)  $M$  is a value,
- (2)  $M = \phi[\mathbf{new}_{\langle A \rangle}]$ ,
- (3)  $M = \phi[a^r?]$ ,
- (4)  $M = \phi[a^w!V]$ ,
- (5)  $M = \phi[\mathbf{spawn\ box\ } V]$ , or
- (6) for any store  $\psi$  such that  $\Phi^w; \Phi^r; \Psi \vdash \psi \text{ okay}$ , there exist some term  $M'$  and store  $\psi'$  such that  $M \mid \psi \rightarrow M' \mid \psi'$ .

**Theorem 4.3 (Configuration progress).**

Suppose  $\Phi^w; \Pi \vdash \pi \text{ okay}$ . Then either:

- (1)  $\pi$  consists only of  $\{V \mid \psi @ \gamma\}$ ,  $\{\phi[a^r?]\} \mid \psi @ \gamma$ ,  $\{\phi[a^w!V]\} \mid \psi @ \gamma$ , or
- (2) there exists  $\pi'$  such that  $\pi \Rightarrow \pi'$ .

**Proposition 4.4** (Type preservation).

Suppose  $\cdot \mid \Phi^w; \Phi^r; \Psi \vdash M : A_{\text{@L}}$ ,  $\Phi^w; \Phi^r; \Psi \vdash \psi$  okay, and  $M \mid \psi \rightarrow M' \mid \psi'$ . Then there exists a store typing  $\Psi'$  such that  $\cdot \mid \Phi^w; \Phi^r; \Psi' \vdash M' : A_{\text{@L}}$ ,  $\Psi \subset \Psi'$ , and  $\Phi^w; \Phi^r; \Psi' \vdash \psi'$  okay.

**Theorem 4.5** (Configuration typing preservation).

Suppose  $\Phi^w; \Pi \vdash \pi$  okay and  $\pi \Rightarrow \pi'$ . Then there exist a write channel typing  $\Phi'^w$  and a configuration typing  $\Pi'$  such that  $\Phi^w \subset \Phi'^w$ ,  $\Pi \subset \Pi'$ , and  $\Phi'^w; \Pi' \vdash \pi'$  okay.

## 5. Examples

This section presents examples of implementing various communication constructs in  $\lambda_{\square}^{\text{PC}}$ : futures, bidirectional communications, shared references, remote evaluation, code on demand, and hot code replacement. Throughout the examples, we use the following syntactic sugar:

$$\begin{aligned} \text{let } x = M \text{ in } N &\equiv (\lambda x : \_ . N) M \\ \text{let } (x, y) = M \text{ in } N &\equiv \text{let } z = M \text{ in} \\ &\quad \text{let } x = \text{fst } z \text{ in let } y = \text{snd } z \text{ in } N \\ \text{let rec } x = M \text{ in } N &\equiv \text{let } x = \text{fix } x : \_ . M \text{ in } N \end{aligned}$$

We also use syntactic sugar  $\text{spawnT } \{M\}$  defined in Section 3.

### 5.1 Futures

A future (Halstead, Jr. 1985) is a communication construct which can be thought of as a pointer to a thread. After the thread finishes its computation, the pointer may be dereferenced to obtain the result. Below we simulate futures using read channels.

In order to create a future, we first need to evaluate a term  $M$  at a fresh thread. Since  $M$  is to be evaluated at a remote thread, it must contain no references. We implement the future with a read channel  $a^r$  such that the result of evaluating  $M$  is written to the corresponding write channel  $a^w$ . Since a channel write entails a communication between threads,  $M$  must evaluate to a global value (in accordance with the rule W!). Hence future, our construct for futures, expects a value of type  $\square\square A$ , say  $\text{box box } M$ , and returns a read channel of type  $\langle A \rangle^r$  to which the result of evaluating  $M$  is sent:

$$\begin{aligned} \text{future} &: \square\square A \rightarrow \langle A \rangle^r_{\text{@G}} \\ \text{future} &= \\ &\lambda x : \square\square A. \\ &\quad \text{letbox } x' = x \text{ in} \\ &\quad \text{let } (y_r, y_w) = \text{new}_{\langle A \rangle} \text{ in} \\ &\quad \text{letbox } y_w = \text{box } y_w \text{ in} \\ &\quad \text{let } \_ = \text{spawnT } \{ \text{letbox } z = x' \text{ in } y_w ! z \} \text{ in} \\ &\quad y_r \end{aligned}$$

### 5.2 Bidirectional communications

Communications from a child thread back to its parent thread are easy to implement because write channels are global values — we start the child thread by evaluating a term containing write channels whose corresponding read channels belong to the parent thread. For communications in the other direction, however, the child thread needs to somehow return a write channel to the parent thread, which can be done only via another write channel originating from the parent thread. Thus, in order to receive a write channel of type  $\langle A \rangle^w$  from a child thread, the parent thread creates a pair of read channel  $a^r$  of type  $\langle \langle A \rangle^w \rangle^r$  and write channel  $a^w$  of type  $\langle \langle A \rangle^w \rangle^w$ . Then the child thread is passed  $a^w$ , through which a write channel of type  $\langle A \rangle^w$  is sent to the parent thread.

We design  $\text{spawn}^{\text{BI}}$ , our construct for bidirectional communications (or for unidirectional communications from a parent thread to a child thread), in such a way that given a value of type  $\square(\langle A \rangle^r \rightarrow C)$ , say  $\text{box } \lambda x_r : \langle A \rangle^r . M$ , it creates a child thread

evaluating  $M$  with  $x_r$  bound to a read channel  $a^r$  of type  $\langle A \rangle^r$  and returns a corresponding write channel  $a^w$  to the parent thread:

$$\begin{aligned} \text{spawn}^{\text{BI}} &: \square(\langle A \rangle^r \rightarrow C) \rightarrow \langle A \rangle^w_{\text{@G}} \\ \text{spawn}^{\text{BI}} &= \\ &\lambda z : \square(\langle A \rangle^r \rightarrow C). \\ &\quad \text{let } (x_r, x_w) = \text{new}_{\langle \langle A \rangle^w \rangle} \text{ in} \\ &\quad \text{letbox } x'_w = \text{box } x_w \text{ in} \\ &\quad \text{letbox } z' = z \text{ in} \\ &\quad \text{let } \_ = \\ &\quad \quad \text{spawnT } \{ \text{let } (y_r, y_w) = \text{new}_{\langle A \rangle} \text{ in} \\ &\quad \quad \quad \text{let } \_ = x'_w ! y_w \text{ in} \\ &\quad \quad \quad z' y_r \} \\ &\quad \text{in} \\ &\quad x_r ? \end{aligned}$$

Note that by the rule *Sync*, the channel write  $x'_w ! y_w$  blocks the child thread until it synchronizes with the corresponding channel read  $x_r ?$ . Thus the evaluation of  $z' y_r$ , the core of the child thread, may not start immediately after  $\text{new}_{\langle A \rangle}$  returns a pair of read and write channels. To start the evaluation of  $z' y_r$  immediately, the child thread could delegate the channel write to another new thread. In general, a channel write  $M ! N$  can be replaced by the following term which spawns a new thread for performing the channel write and returns immediately (unless an evaluation of  $M$  or  $N$  gets stuck with a channel read or a channel write):

$$\begin{aligned} &\text{letbox } x_M = \text{box } M \text{ in} \\ &\text{letbox } x_N = \text{box } N \text{ in} \\ &\text{spawnT } \{ x_M ! x_N \} \end{aligned}$$

The same idea can be applied to all instances of channel writes in the examples below.

### 5.3 Shared references

We implement a shared reference for type  $A$  as a thread with two components: a read channel of type  $\langle \langle A \rangle^w + A \rangle^r$ , as an interface, and a global value of type  $A$ , as its current content. Any thread may request the current content of the reference by sending a value  $\text{inl } a^w$  (of type  $\langle A \rangle^w + A$ ) to the read channel, in which case the current content is written back to  $a^w$ . To update the content of the reference, it sends a value  $\text{inr } V$  (also of type  $\langle A \rangle^w + A$ ) to the read channel, in which case the current content is updated with a new value  $V$ .

As an illustration, we use  $\text{spawn}^{\text{BI}}$  to create a shared reference for type  $\text{int}$  initialized with a value of 0:

$$\begin{aligned} &\text{let } \text{cell}_w = \\ &\quad \text{spawn}^{\text{BI}} \\ &\quad \quad \text{box } \lambda \text{cell}_r : \langle \langle \text{int} \rangle^w + \text{int} \rangle^r. \\ &\quad \quad \text{let rec } f = \lambda n : \text{int}. \\ &\quad \quad \quad \text{case } \text{cell}_r ? \text{ of } \text{inl } ch_w \Rightarrow \text{let } \_ = ch_w ! n \text{ in } f n \\ &\quad \quad \quad \quad \mid \text{inr } v \Rightarrow f v \\ &\quad \quad \text{in} \\ &\quad f 0 \end{aligned}$$

Then we use a write channel in  $\text{cell}_w$  to implement *get*, of type  $\text{unit} \rightarrow \text{int}$ , for requesting the current content of the reference and *set*, of type  $\text{int} \rightarrow \text{unit}$ , for updating the content of the reference; note that each call to *get* creates a fresh pair of read channel and write channel:

$$\begin{aligned} &\text{letbox } \text{cell}'_w = \text{box } \text{cell}_w \text{ in} \\ &\text{letbox } \text{get} = \text{box } \lambda \_ : \text{unit}. \text{let } (x_r, x_w) = \text{new}_{\langle \text{int} \rangle} \text{ in} \\ &\quad \text{let } \_ = \text{cell}'_w ! (\text{inl } x_w) \text{ in} \\ &\quad x_r ? \\ &\quad \text{in} \\ &\text{letbox } \text{set} = \text{box } \lambda n : \text{int}. \text{cell}'_w ! (\text{inr } n) \text{ in} \\ &\quad (\text{get}, \text{set}) \end{aligned}$$



call to the following construct  $\text{forward}'$  with arguments  $a_1^r$  and  $a_2^w$ :

```

forward' : ⟨A⟩r → (⟨A⟩w → unit)@G
forward' = λx : ⟨A⟩r. λy : ⟨A⟩w.
  letbox x' = box x in      (* does not typecheck *)
  letbox y' = box y in
  spawnT { let rec f = λ_. unit. let _ = y'!(x'?) in f () in
           f () }

```

Unfortunately  $\text{forward}'$  fails to typecheck because a read channel of type  $\langle A \rangle^r$  cannot be a global value.

Instead of spawning an ordinary thread, therefore, we spawn a *local thread* that starts by evaluating  $\text{forward}$   $a_1^r$   $a_2^w$ . The underlying assumption is that a thread consists of one or more local threads running concurrently which share *both the store and read channels*. (There arises the memory consistency problem, but it is a separate issue.) Hence both references and read channels can be part of a term for creating a new local thread. In other words, as far as creating local threads is concerned, both references and read channels can be regarded as global values. We do not, however, allow communications of read channels even between local threads, since the syntax for a channel write  $a^w!V$  does not indicate whether the corresponding read channel  $a^r$  resides at the same thread (in which case  $V$  may be another read channel) or at a remote thread (in which case  $V$  may not be another read channel). Thus channel writes still expect global values which do not include read channels. Fournet et al. (1996) make a similar assumption in their study of the distributed join-calculus: every channel has a unique reflexive solution, corresponding to a thread in  $\lambda_{\square}^{\text{PC}}$ , which may run multiple processes, corresponding to local threads in  $\lambda_{\square}^{\text{PC}}$ ; all these processes can interact with any message addressed to the channel.

We introduce a new construct  $\text{lthread } \{M\}$  for creating a local thread that starts by evaluating  $M$ :

$$\text{term } M ::= \dots \mid \text{lthread } \{M\}$$

Since local threads running at the same thread share references and read channels,  $\text{lthread } \{M\}$  typechecks whenever  $M$  typechecks:

$$\frac{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash M : A_{@L}}{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash \text{lthread } \{M\} : \text{unit}_{@L}} \text{LThread}$$

Now that a thread may run multiple local threads, the state of a thread  $\gamma$  is represented by  $\{M_1, \dots, M_n \mid \psi @ \gamma\}$  where each term  $M_i$ ,  $1 \leq i \leq n$ , is being evaluated by a local thread. All the previous rules for configuration transitions are extended in a straightforward way by rewriting  $\{M \mid \psi @ \gamma\}$  as  $\{M, M_1, \dots, M_n \mid \psi @ \gamma\}$ . An exception is the rule *Spawn* which creates a fresh thread consisting only of a single local thread, as shown in Figure 6. A new rule *Sync'* accounts for a channel read and a channel write occurring synchronously within the same thread  $\gamma$ . Another new rule *LThread* evaluates  $\text{lthread } \{M\}$  to creates a fresh local thread.

Now we rewrite  $\text{forward}$  by exploiting  $\text{lthread}$ :

```

forward = λx : ⟨A⟩r. λy : ⟨A⟩w.
  lthread { let rec f = λ_. unit. let _ = y!(x?) in f () in
           f () }

```

Another application of  $\text{lthread}$  is to simulate channels not subject to channel locality, *i.e.*, channels that can be read from and written to at any thread. The idea is the same as in implementing shared references in Section 5.3 except that each channel read from  $\text{cell}_r$  starts a new local thread. Here is an example of creating such a channel for type  $\text{int}$ . We call *read* for channel reads and *write* for channel writes, both of which are global values and thus can be

shared by all threads.

```

let cell_w =
  spawnBI
  box λcell_r : ⟨⟨int⟩w + int⟩r.
    let (c_r, c_w) = new(int) in
    let rec f = λ_. unit.
      case cell_r? of inl ch_w ⇒ lthread {ch_w!(c_r?)}
                    | inr v ⇒ lthread {c_w!v}
    in
    f ()
letbox cell'_w = box cell_w in
letbox read = box λ_. unit. let (x_r, x_w) = new(int) in
  let _ = cell'_w!(inl x_w) in
  x_r?
in
letbox write = box λx : int. cell'_w!(inr x) in
(read, write)

```

## 6.2 Type-safe data parallelism

The characteristic feature of the type system of  $\lambda_{\square}^{\text{PC}}$  is to be able to decide whether a given term evaluates to a global value or a local value. We can exploit this feature to provide type safety for data parallelism in  $\lambda_{\square}^{\text{PC}}$ , as illustrated below.

Consider a parallel map construct  $\text{mapP}$  which applies a function  $f$  to each element of an array  $e$  in parallel:

$$\text{mapP } f e$$

The parent thread evaluating  $\text{mapP } f e$  spawns multiple child threads each of which applies  $f$  to an element of  $e$ . In order to avoid the memory consistency problem, we require that  $f$  do not inherit references from the parent thread (because  $f$  may attempt to update such references). On the other hand, in order to provide more flexibility in programming, we allow  $f$  to allocate new references, which cannot be transmitted between child threads and thus are safe to use.

Checking if  $f$  satisfies the above two conditions is simple: we just test whether  $f$  is a global value or not. Note that although child threads do not share writable data, they can still share read-only data because variables bound to global values can serve as shared read-only data. That is, if a binding  $x : A_{@G}$  is available at the parent thread, all child threads may read variable  $x$  to obtain a global value.

A global value  $f$ , however, does not protect against write channels being shared by child threads. An easy solution is to introduce another locality  $G^*$  with a relation  $G^* < G$  and the following typing rule

$$\frac{\Gamma_{G^*} \mid \cdot; \cdot \vdash V : A_{@L}}{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash V : A_{@G^*}} \text{GVal}'$$

where

$$\Gamma_{G^*} = \{x : A_{@G^*} \mid x : A_{@G^*} \in \Gamma\}.$$

New constructs and types for  $G^*$  can be designed in an analogous way to those for  $G$ . Then a typing judgment  $f : A_{@G^*}$  ensures that  $f$  contains neither references nor write channels.

## 7. Related work

### 7.1 Mutable references in parallel or distributed languages

As there is no definitive standard for shared memory model in parallel or distributed languages, different policies for mutable references or similar constructs have been proposed. X10 (Charles et al. 2005) permits remote mutable variables belonging to remote threads, but accessing the contents of a remote mutable variable results in a runtime exception. Fortress (Allan et al. 2007) permits

$$\begin{array}{c}
\frac{\text{fresh } \gamma'}{\pi, \{\phi[\text{spawn box } V], M_1, \dots, M_n \mid \psi @ \gamma\} \Rightarrow \pi, \{\phi[()], M_1, \dots, M_n \mid \psi @ \gamma\}, \{V () \mid \cdot @ \gamma'\}} \text{Spawn} \\
\frac{\pi, \{\phi[a^r?], \phi'[a^w!V], M_1, \dots, M_n \mid \psi @ \gamma\} \Rightarrow \pi, \{\phi[V], \phi'[()], M_1, \dots, M_n \mid \psi @ \gamma\}}{\pi, \{\phi[\text{lthread } \{M\}], M_1, \dots, M_n \mid \psi @ \gamma\} \Rightarrow \pi, \{\phi[()], M, M_1, \dots, M_n \mid \psi @ \gamma\}} \text{LThread} \\
\text{Sync}'
\end{array}$$

**Figure 6.** Configuration transition rules for local threads

shared objects accessible to every thread, but at the cost of implicitly maintaining the sharedness (either *local* or *shared*) of every object. Titanium (Hilfinger et al. 2005), which extends Java, takes a different approach by allowing shared memory, but also distinguishing between local references and global references at the type level. UPC (El-Ghazawi et al. 2003), which extends C, takes a similar approach by using two kinds of pointers: private pointers and global pointers. Facile (Knabe 1995) and JoCAML (Fournet et al. 2003) dispense with remote references by sending copies of heap cells whenever their references are transmitted to remote threads. Erlang (Armstrong 1997) and Manticore (Fluet et al. 2007) do not use mutable references at all.

$\lambda_{\square}^{\text{PC}}$  is similar to Facile and JoCAML in that it permits mutable references but assumes no shared memory (and also in that it is a dialect of ML). The main difference is that Facile and JoCAML rely on the runtime system to avoid remote references whereas  $\lambda_{\square}^{\text{PC}}$  relies on the type system to forestall remote references.

## 7.2 Channel locality

The issue of channel locality has been studied mainly for the pi-calculus and its relatives. Amadio (1997) develops a type system for a fragment of the pi-calculus in which typing judgments use channels linearly to ensure channel locality. The distributed join-calculus (Fournet et al. 1996) assumes a syntactic restriction to enforce channel locality. Specifically it considers only syntactically well-formed DRCHAMs (distributed reflexive chemical machines) in which every channel is defined in at most one RCHAM (reflexive chemical machine), where a DRCHAM corresponds to a configuration and a RCHAM to a thread in  $\lambda_{\square}^{\text{PC}}$ . Schmitt and Stefani (2003) present a higher-order version of the distributed join calculus which uses a polymorphic type system to guarantee that the destination for each message is uniquely determined. It achieves channel locality in a slightly different sense than in  $\lambda_{\square}^{\text{PC}}$  because, for example, a message may be intercepted before reaching its destination. The dynamic join-calculus (Schmitt 2002) achieves channel locality in a similar vein, in which the destination for a message is determined according to its current position.

Yoshida and Hennessy (1999) present a calculus similar to  $\lambda_{\square}^{\text{PC}}$  in that it uses a *type system* to enforce channel locality. It combines the call-by-value lambda calculus and a higher-order extension of the pi-calculus, and allows processes themselves to be transmitted via channels. The type system introduces *sendable types* whose values can be transmitted between processes without destroying channel locality, and exploits a subtyping relation to enforce channel locality. The type system, however, is not a general solution applicable to our setting for two reasons. First there is a semantic restriction on function types which severely limits the generality of the calculus: for a function type  $\tau \rightarrow \sigma$ , if  $\sigma$  is sendable,  $\tau$  must also be sendable, in which case  $\tau \rightarrow \sigma$  is also automatically regarded as sendable. In our setting, such a restriction means, for example, that no term of type  $\text{ref int} \rightarrow \text{int}$  or  $\langle \text{int} \rangle^r \rightarrow \text{int}$  is even allowed and that every function of type  $\text{int} \rightarrow \text{int}$  must be global, neither of which is the case in  $\lambda_{\square}^{\text{PC}}$ . Second the type system ignores the call-by-value semantics of the underlying lambda calculus. Specifically

it includes a typing rule (TERM<sub>l</sub>) assigning a sendable type to a term whenever it typechecks under a typing context using sendable types only. In our setting, the typing rule would correspond to the following rule which fails to comply with the call-by-value semantics:

$$\frac{\Gamma_G \mid \Phi^w; \cdot; \cdot \vdash M : A_{\text{@L}}}{\Gamma \mid \Phi^w; \Phi^r; \Psi \vdash M : A_{\text{@G}}} \text{ (wrong)}$$

In comparison,  $\lambda_{\square}^{\text{PC}}$  imposes no syntactic or semantic restriction on function types and properly accounts for the call-by-value semantics in the presence of mutable references.

## 7.3 Splitting channels

The idea of using two kinds of channels, namely read channels and write channels (or input channels and output channels), is not new. For example, the type system of (Pierce and Sangiorgi 1993) can effectively distinguish between read channels and write channels because every channel is assigned a tag indicating its usage (read, write, or both). The polarized pi-calculus of (Odersky 1995) syntactically distinguishes between read channels and write channels. In the polarized pi-calculus, using read channels for channel writes or write channels for channel reads results in no reaction with other processes. The polar pi-calculus of (Zhang and Potter 2002) also syntactically distinguishes between read channels and write channels with an additional requirement that only write channels can be transmitted via channels.

$\lambda_{\square}^{\text{PC}}$  is, however, different from previous work in that it assigns different types to different kinds of channels and *exploits these types*, instead of the syntax, to enforce channel locality. In fact, such notions as local values, global values, and primitive types, all established in the base language  $\lambda_{\square}$ , have naturally led to the main idea for achieving channel locality in  $\lambda_{\square}^{\text{PC}}$ . Thus our decision to distinguish between read channels and write channels has a different motivation from previous work.

## 7.4 Modal types for parallel and distributed computations

There are a few distributed languages (Murphy et al. 2004; Jia and Walker 2004) whose type systems are based on the spatial interpretation of modal logic. Murphy et al. (2004) use the necessity modality  $\square$  in modal types for global terms which can be evaluated at any node in the network, and the possibility modality  $\diamond$  in modal types for references to local resources belonging to certain nodes. Jia and Walker (2004) use  $\square$  for the same purpose, but  $\diamond$  in modal types for terms that can be evaluated at certain nodes.

Our earlier work (Park 2006) is the first departure from the typical spatial interpretation of modal logic in that it uses a new modality  $\boxplus$  to focus on global values rather than global terms. (See (Park 2006) for a discussion on logical properties of the modality  $\boxplus$ . For example,  $\boxplus A \rightarrow A$  and  $\boxplus A \rightarrow \boxplus \boxplus A$  are both inhabited, but  $\boxplus(A \rightarrow B) \rightarrow \boxplus A \rightarrow \boxplus B$  is uninhabited if we ignore non-terminating terms.) Our present work adds communication constructs for higher-order channels, thereby providing improved flexibility in programming for parallel and distributed computations.

## 8. Conclusion and future work

We present an ML-like parallel language  $\lambda_{\square}^{\text{PC}}$  which features type-safe higher-order channels with channel locality. Sequential programming in  $\lambda_{\square}^{\text{PC}}$  may exploit mutable references as usual, since the type system ensures that mutable references never interfere with parallel computations. Thus implementing the parallel operational semantics of  $\lambda_{\square}^{\text{PC}}$  is no more complicated than implementing a similar parallel language without mutable references.

Our long-term goal is to build a programming system that supports three levels of parallelism within a unified framework. At the highest level, it implements distributed computations taking place in a network of nodes. Each node performs a stand-alone computation and also communicates with other nodes via higher-order channels. (Hence there is no distributed shared memory.) The next level implements task parallelism with higher-order channels as in  $\lambda_{\square}^{\text{PC}}$ . Thus a stand-alone computation at a node actually consists of multiple threads running in parallel. At the lowest level, multiple local threads with shared memory run within each thread. Task parallelism with shared memory as well as data parallelism can be implemented at this level. Then the type system of  $\lambda_{\square}^{\text{PC}}$  can be extended to provide type safety at all the three levels.

## Acknowledgments

The author is grateful to Hyeonseung Im and anonymous reviewers for their helpful comments.

## References

- Eric Allan, David Chase, Joe Hallet, Victor Luchangco, Jan-Willem Maessen, Sukyoung Ryu, Guy L. Steele Jr., and Sam Tobin-Hochstadt. The Fortress language specification version 1.0beta. Technical report, Sun Microsystems Inc., March 2007.
- Roberto Amadio. An asynchronous model of locality, failure, and process mobility. In *Proceedings of the 2nd International Conference on Coordination Languages and Models, LNCS 1282*, pages 374–391. Springer-Verlag, 1997.
- Joe Armstrong. The development of Erlang. In *Proceedings of the Second ACM SIGPLAN International Conference on Functional Programming*, pages 196–203. ACM Press, 1997.
- Krste Asanovic, Ras Bodik, Bryan Christopher Catanzaro, Joseph Gebis, Parry Husbands, Kurt Keutzer, David Patterson, William Plishker, John Shalf, Samuel Williams, and Katherine Yelick. The landscape of parallel computing research: A view from Berkeley. Technical Report UCB/Eecs-2006-183, Eecs Department, University of California, Berkeley, December 2006.
- Guy Blelloch. Programming parallel algorithms. *Communications of the ACM*, 39(3):85–97, 1996.
- Manuel Chakravarty, Roman Leshchinskiy, Simon Peyton Jones, Gabriele Keller, and Simon Marlow. Data Parallel Haskell: a status report. In *Proceedings of the ACM SIGPLAN Workshop on Declarative Aspects of Multicore Programming*, pages 11–18, 2007.
- Philippe Charles, Christian Grothoff, Vijay Saraswat, Christopher Donawa, Allan Kielstra, Kemal Ebcioglu, Christoph von Praun, and Vivek Sarkar. X10: an object-oriented approach to non-uniform cluster computing. In *Proceedings of the 20th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 519–538, 2005.
- Cray Inc. The Chapel language specification version 0.4. Technical report, February 2005.
- Tarek El-Ghazawi, William Carlson, and Jesse Draper. UPC language specifications, v1.1.1, October 2003.
- Matthew Fluet, Mike Rainey, John Reppy, Adam Shaw, and Yingqi Xiao. Manticore: A heterogeneous parallel language. In *Proceedings of the ACM SIGPLAN Workshop on Declarative Aspects of Multicore Programming*, pages 25–32, 2007.
- Cédric Fournet, Georges Gonthier, Jean-Jacques Lévy, Luc Maranget, and Didier Rémy. A calculus of mobile agents. In *CONCUR: 7th International Conference on Concurrency Theory, LNCS 1119*, pages 406–421. Springer, 1996.
- Cédric Fournet, Fabrice Le Fessant, Luc Maranget, and Alan Schmitt. JoCaml: a language for concurrent distributed and mobile programming. In Johan Jeuring and Simon Peyton Jones, editors, *Advanced Functional Programming, 4th International School, 2002*, volume 2638 of *Lecture Notes in Computer Science*, pages 129–158. Springer-Verlag, 2003.
- Robert Halstead, Jr. MULTILISP: a language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems*, 7(4):501–538, 1985.
- Paul Hilfinger, Dan Bonachea, Kaushik Datta, David Gay, Susan Graham, Benjamin Liblit, Geoffrey Pike, Jimmy Su, and Katherine Yelick. Titanium language reference manual, version 2.19. Technical Report UCB/Eecs-2005-15, Eecs Department, University of California, Berkeley, November 2005.
- Lorin Hochstein, Jeff Carver, Forrest Shull, Sima Asgari, Victor Basili, Jeffrey Hollingsworth, and Marvin Zelkowitz. Parallel programmer productivity: A case study of novice parallel programmers. In *Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*, page 35. IEEE Computer Society, 2005.
- Limin Jia and David Walker. Modal proofs as distributed programs (extended abstract). In *Proceedings of the European Symposium on Programming, LNCS 2986*, pages 219–233. Springer, 2004.
- Frederick Knabe. *Language Support for Mobile Agents*. PhD thesis, Department of Computer Science, Carnegie Mellon University, 1995.
- Tom Murphy, VII, Karl Crary, Robert Harper, and Frank Pfenning. A symmetric modal lambda calculus for distributed computing. In *Proceedings of the 19th IEEE Symposium on Logic in Computer Science*. IEEE Press, 2004.
- R. Nikhil and Arvind. *Implicit parallel programming in pH*. Morgan Kaufmann, 2001.
- Martin Odersky. Polarized name passing. In *Proceedings of the 15th Conference on Foundations of Software Technology and Theoretical Computer Science, LNCS 1026*, pages 324–337. Springer-Verlag, 1995.
- Sungwoo Park. A modal language for the safety of mobile values. In Naoki Kobayashi, editor, *Proceedings of the 4th Asian Symposium on Programming Languages and Systems, LNCS 4279*, pages 217–233. Springer, 2006.
- Benjamin Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. In *Proceedings of the 8th Annual IEEE Symposium on Logic in Computer Science*, pages 376–385, 1993.
- Alan Schmitt. Safe dynamic binding in the join calculus. In *TCS '02: Proceedings of the IFIP 17th World Computer Congress - TCI Stream / 2nd IFIP International Conference on Theoretical Computer Science*, pages 563–575. Kluwer, B.V., 2002.

- Alan Schmitt and Jean-Bernard Stefani. The M-calculus: a higher-order distributed process calculus. In *Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 50–61. ACM Press, 2003.
- Nir Shavit and Dan Touitou. Software transactional memory. In *Proceedings of the fourteenth ACM symposium on Principles of Distributed Computing*, pages 204–213. ACM Press, 1995.
- James Stamos and David Gifford. Remote evaluation. *ACM Transactions on Programming Languages and Systems*, 12(4): 537–564, 1990.
- Herb Sutter. The free lunch is over: a fundamental turn toward concurrency in software. *Dr. Dobbs's Journal*, 30(3), March 2005.
- Nobuko Yoshida and Matthew Hennessey. Subtyping and locality in distributed higher order processes (extended abstract). In *CONCUR: 10th International Conference on Concurrency Theory, LNCS 1664*, pages 557–572. Springer-Verlag, 1999.
- Xiaogang Zhang and John Potter. Responsive bisimulation. In *Proceedings of the 2nd IFIP International Conference on Theoretical Computer Science*, pages 601–612. Kluwer, B.V., 2002.