

A Modal Language for the Safety of Mobile Values

Sungwoo Park

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

Abstract. In the context of distributed computations, local resources give rise to an issue not found in stand-alone computations: the safety of mobile code. One approach to the safety of mobile code is to build a modal type system with the modality \Box that corresponds to necessity of modal logic. We argue that the modality \Box is not expressive enough for safe communications in distributed computations, in particular for the safety of mobile values. We present a modal language which focuses on the safety of mobile values rather than the safety of mobile code. The safety of mobile values is achieved with a new modality \Box which expresses that given code evaluates to a mobile value. We demonstrate the use of the modality \Box with a communication construct for remote procedure calls.

1 Introduction

A distributed computation is a cooperative process taking place in a network of nodes. Each node is capable of performing a stand-alone computation and also communicating with other nodes to distribute and collect code and data. Thus a distributed computation has the potential to make productive use of all the nodes in the network simultaneously.

Usually a distributed computation assumes a heterogeneous group of nodes with different *local resources*. A local resource can be either a permanent/physical object available at a particular node (*e.g.*, printer, database) or an ephemeral/semantic object created during a stand-alone computation (*e.g.*, heap cell, abstract data type). Local resources are accessed via their references (*e.g.*, handle for a database file, pointer to a heap cell).

Local resources, however, give rise to an issue not found in stand-alone computations: the safety of *mobile code*, or in our terminology, the safety of *mobile terms* where a term represents a piece of code. In essence, a node cannot access remote resources in the same way that it accesses its own local resources, but it may receive mobile terms in which references to remote resources are exposed. Therefore the safety of mobile terms is achieved by supporting direct access to remote resources (*e.g.*, remote file access, remote memory access), as in Obliq [1], by transmitting copies of local resources along with mobile terms, as in Facile [2], by preventing references to remote resources from being dereferenced, as in Mobile UNITY [3], or by allowing all of these methods, as in λ dist [4]. Our paper focuses on the third case where we reject mobile terms containing references to remote resources.

One approach to the safety of mobile terms is to build a modal type system with the modality \Box [5,6,7,8] which is based on a spatial interpretation of necessity of modal logic such as S4 and S5. The basic idea is that a value of modal type $\Box A$ contains a

mobile term that can be evaluated at any node. By requiring that a mobile term be from a value of type $\Box A$, we ensure its safety without recourse to runtime checks.

A type system augmented with the modality \Box is not, however, expressive enough for safe communications of *values*, *i.e.*, the safety of *mobile values*. In other words, we cannot rely solely on modal types $\Box A$ to verify that a value communicated from one node to another is mobile (*e.g.*, when a remote procedure call returns, or when a value is written to a channel). The reason is that in general, a value of type $\Box A$ contains *not a mobile value but a mobile term*. The evaluation of such a mobile term (with the intention of obtaining a mobile value) may result in a value that is not necessarily mobile because of references to local resources created during the evaluation.

As an example, consider a term of type $\text{int} \rightarrow \text{int}$ in an ML-like language:

```
let
  val new_reference = ref 0
  val f = fn x => x + !new_reference
in
  f
end
```

The above term can be evaluated at any node and thus may be used in building a mobile term of type $\Box(\text{int} \rightarrow \text{int})$. The resultant value f , however, is not mobile because it accesses a local resource `new_reference`. In contrast, the following term, also of type $\text{int} \rightarrow \text{int}$, cannot be used in building a mobile term of type $\Box(\text{int} \rightarrow \text{int})$, but the resultant value is mobile because it does not access any local resource:

```
let
  val v = !some_existing_reference
  val f = fn x => x + v
in
  f
end
```

Hence the modality \Box is irrelevant to the safety of mobile values, which should now be verified by programmers themselves.

This paper investigates a new modality \Box which expresses that a given term evaluates to a mobile value. The basic idea is that a term contained in a value of modal type $\Box A$ evaluates to a value that is valid at any node. For example, the first term above cannot be used in building a term of type $\Box(\text{int} \rightarrow \text{int})$, but the second term above may be used in building such a term. To obtain a value to be communicated to other nodes, we evaluate a term contained in a value of type $\Box A$. In this way, we achieve the safety of mobile values.

While the mobility of a term is independent of the mobility of the value to which it evaluates, the modality \Box is weaker than the modality \Box in that we can emulate \Box with \Box . For example, we may define $\Box A$ as $\Box(\text{unit} \rightarrow A)$, in which case we check the mobility of a term M of type A by checking the mobility of a value $\text{fn } _ => M$ of type $\text{unit} \rightarrow A$. Thus \Box is inherently more expressive than \Box , and the use of \Box practically eliminates the need for \Box . The converse is not the case, however: we cannot

emulate \boxplus with \boxdot because the modality \boxplus requires at least a distinction between terms and values, which is not provided by the type system for the modality \boxdot .

Since the modality \boxdot is inadequate for ensuring the safety of mobile values, safe communications are restricted to mobile terms if the underlying type system uses only \boxdot . Such a restriction leads to an unusual implementation of common communication constructs in distributed computations. For example, in λ_{rpc} by Jia and Walker [7] and *Lambda 5* by Murphy *et al.* [8], a remote procedure call returns a mobile term (instead of a mobile value) which the caller node needs to further evaluate in order to obtain the final result of the remote procedure call. By focusing on mobile values rather than mobile terms, the modality \boxplus avoids such anomalies and gives a faithful implementation of common communication constructs.

In Sections 2 and 3, we develop a call-by-value language λ_{\boxplus} with mutable references and the modality \boxplus . We choose mutable references as a representative example of local resources; other kinds of local resources can be treated in an analogous way. We formulate its type system in the natural deduction style by giving introduction and elimination rules for each connective and modality. The modality \boxplus requires us to introduce a typing judgment differentiating values from terms. The type system takes into account *primitive types* (such as boolean values and integers) for which mobility is an inherent property.

In Section 4, we develop λ_{\boxplus} into λ_{\boxplus}^N which has a *network operational semantics* and is thus capable of modeling distributed computations. We demonstrate the use of modal types with a communication construct for remote procedure calls. The safety of mobile terms and mobile values is shown by type safety of λ_{\boxplus}^N , *i.e.*, its progress and type preservation properties.

Section 5 compares λ_{\boxplus}^N with other modal languages for distributed computations. Section 6 concludes with future work. Due to space limitations, we refer the reader to our technical report [9] for details of all proofs.

2 Call-by-Value Language λ with Mutable References

This section reviews the type system of λ , a typical call-by-value language with mutable references, in the context of distributed computations. Figure 1 shows the definition of λ .

The syntax of λ uses metavariables A, B, C for types and M, N for terms. $()$ is an expression of type unit, which we include as an example of a primitive type. $\lambda x : A. M$ and $M M'$ are a λ -abstraction and a λ -application, respectively. $\text{ref } M$ allocates a fresh reference, $!M$ dereferences an existing reference, and $M := M'$ assigns a new value to a reference; a location l , of type $\text{ref } A$, is a value for a reference.

A variable x with binding $x : A$ is assumed to hold a value because λ uses the call-by-value strategy. We use a typing judgment $\Gamma \mid \Psi \vdash M : A$ to mean that term M has type A under typing context Γ and store typing Ψ ; a store typing judgment $\Psi \vdash \psi \text{ okay}$ means that store ψ conforms to store typing Ψ .

The operational semantics of λ uses a reduction judgment $M \mid \psi \longrightarrow M' \mid \psi'$ to mean that term M with store ψ reduces to term M' with store ψ' where $\psi = \psi'$ is allowed. A β -reduction judgment $M \longrightarrow_{\beta} M'$ uses a capture-avoiding substitution $[V/x]M$ defined in a standard way. We write $\phi[[M]]$ for a term obtained by filling the hole \square in an

type	$A ::= \text{unit} \mid A \rightarrow A \mid \text{ref } A$
term	$M ::= () \mid x \mid \lambda x:A. M \mid M M \mid \text{ref } M \mid !M \mid M := M \mid l$
value	$V ::= () \mid \lambda x:A. M \mid l$
typing context	$\Gamma ::= \cdot \mid \Gamma, x : A$
store typing	$\Psi ::= \cdot \mid \Psi, l \mapsto A$
store	$\psi ::= \cdot \mid \psi, l \mapsto V$

$\frac{}{\Gamma \mid \Psi \vdash () : \text{unit}} \text{Unit}$		
$\frac{x : A \in \Gamma}{\Gamma \mid \Psi \vdash x : A} \text{Var}$	$\frac{\Gamma, x : A \mid \Psi \vdash M : B}{\Gamma \mid \Psi \vdash \lambda x:A. M : A \rightarrow B} \rightarrow_l$	$\frac{\Gamma \mid \Psi \vdash M : A \rightarrow B \quad \Gamma \mid \Psi \vdash N : A}{\Gamma \mid \Psi \vdash M N : B} \rightarrow_E$
$\frac{\Gamma \mid \Psi \vdash M : A}{\Gamma \mid \Psi \vdash \text{ref } M : \text{ref } A} \text{Ref}$	$\frac{\Gamma \mid \Psi \vdash M : \text{ref } A}{\Gamma \mid \Psi \vdash !M : A} \text{Deref}$	$\frac{\Gamma \mid \Psi \vdash M : \text{ref } A \quad \Gamma \mid \Psi \vdash N : A}{\Gamma \mid \Psi \vdash M := N : \text{unit}} \text{Assign}$
$\frac{\Psi(l) = A}{\Gamma \mid \Psi \vdash l : \text{ref } A} \text{Loc}$	$\frac{\text{dom}(\Psi) = \text{dom}(\psi) \quad \cdot \mid \Psi \vdash \psi(l) : \Psi(l) \text{ for every } l \in \text{dom}(\psi)}{\Psi \vdash \psi \text{ okay}} \text{Store}$	
$(\lambda x:A. M) V \longrightarrow_\beta [V/x]M$		
evaluation context $\phi ::= [] \mid \phi M \mid V \phi \mid \text{ref } \phi \mid !\phi \mid \phi := M \mid V := \phi$		
$\frac{M \longrightarrow_\beta M'}{\phi[M] \mid \psi \longrightarrow \phi[M'] \mid \psi} \text{Red}_\beta$		
$\frac{l \notin \text{dom}(\psi)}{\phi[\text{ref } V] \mid \psi \longrightarrow \phi[l] \mid \psi, l \mapsto V} \text{Ref}$		
$\frac{\psi(l) = V}{\phi[!l] \mid \psi \longrightarrow \phi[V] \mid \psi} \text{Deref}$		
$\frac{}{\phi[l := V] \mid \psi \longrightarrow \phi[()] \mid [l \mapsto V]\psi} \text{Assign}$		

Fig. 1. Definition of the language λ

evaluation context ϕ with M . $[l \mapsto V]\psi$ replaces $l \mapsto V'$ in ψ by $l \mapsto V$; we write $\psi(l)$ and $\Psi(l)$ for the value and the type to which l is mapped under ψ and Ψ , respectively.

In the context of distributed computations, $x : A$ in a typing context Γ means that variable x holds a value of type A that is valid at a hypothetical node where typechecking takes place, which we call the *current node* throughout the paper. Then a typing judgment $\Gamma \mid \Psi \vdash M : A$ means that if both typing context Γ and store typing Ψ are satisfied, the evaluation of term M at the current node returns a value V of type A . It does not, however, tell us if M is a mobile term that can be evaluated at other nodes. More importantly, it does not tell us if V is a mobile value that is valid at other nodes. Therefore the above type system is not expressive enough for the safety of mobile terms and mobile values in distributed computations.

λ_\square extends λ with the modality \square which is concerned with *where we can use the result of evaluating a given term*. We call λ_\square a *modal language* because of its use of the modality \square in the type system. Although its type system addresses the safety of mobile values, λ_\square is still a language for stand-alone computations in which no communications between nodes actually take place. The modality \square does not originate from modal logic, but the type system of λ_\square reuses typing judgments for necessity of modal logic by Pfenning and Davies [10].

3 Modal Language λ_\square

The idea behind the modality \square is two-fold. First, if a term M is well-typed under an empty typing context and an empty store typing, *i.e.*, $\cdot \mid \cdot \vdash M : A$, we can evaluate it at

any node. Intuitively M is valid at any node, or *globally valid*, because its evaluation depends on no existing local resources. As a special case, if a value V satisfies $\cdot \mid \vdash V : A$, it is globally valid because it does not contain references (to local resources). Second the typing judgment $\Delta \mid \Psi \vdash M : A$ of λ is unable to express the property that the value to which term M evaluates is globally valid. Therefore we need an additional typing judgment for the type system of λ_{\square} so as to express such properties of terms.

In order to indicate that a variable stores a globally valid value, we introduce a *global typing context* Δ . Γ is now called a *local typing context*.

$$\begin{aligned} \text{global typing context } \Delta & ::= \cdot \mid \Delta, x \sim A \\ \text{local typing context } \Gamma & ::= \cdot \mid \Gamma, x : A \end{aligned}$$

A binding $x \sim A$ in Δ means that variable x holds a globally valid value of type A ; hence a global typing context does not affect the mobility of a term being typechecked.

We use a typing judgment $\Delta; \Gamma \mid \Psi \vdash M : A$ to mean that under global typing context Δ , local typing context Γ , and store typing Ψ , term M evaluates to a value of type A valid at the current node; it may be viewed as a typing judgment for λ where a typing context is split into Δ and Γ . We introduce a new form of typing judgment $\Delta; \Gamma \mid \Psi \vdash M \sim A$ to mean that M evaluates to a globally valid value of type A (which is also valid at the current node). By the definition of these typing judgments, the following typing rules hold independently of the syntax of λ_{\square} :

$$\frac{x \sim A \in \Delta}{\Delta; \Gamma \mid \Psi \vdash x \sim A} \text{GVar} \quad \frac{x \sim A \in \Delta}{\Delta; \Gamma \mid \Psi \vdash x : A} \text{GVar}' \quad \frac{\Delta; \cdot \mid \cdot \vdash V : A}{\Delta; \Gamma \mid \Psi \vdash V \sim A} \text{GVal}$$

The rule **GVar'** says that a globally valid variable x in $x \sim A$ is valid at the current node. The rule **GVal** conforms to the definition of the new typing judgment: the premise check if V is globally valid, in which case the conclusion holds because V is already a value.

The type system of λ_{\square} classifies types into three kinds: *primitive types* P , *potentially global types* G , and *local types* L . A primitive type is one for which mobility is an inherent property. For example, $()$, of type unit , is atomic and cannot contain references to local resources. Therefore values of type unit are always globally valid, which implies that unit is a primitive type. Formally we define primitive types as follows:

Definition 1. P is a primitive type if and only if $\Delta; \Gamma \mid \Psi \vdash V : P$ implies $\Delta; \cdot \mid \cdot \vdash V : P$.

By the definition of primitive types, $\Delta; \Gamma \mid \Psi \vdash M : P$ semantically implies $\Delta; \Gamma \mid \Psi \vdash M \sim P$. In order to relieve the programmer of the burden of explicitly expressing mobility for primitive types, λ_{\square} provides a separate typing rule for primitive types:

$$\frac{\Delta; \Gamma \mid \Psi \vdash M : P}{\Delta; \Gamma \mid \Psi \vdash M \sim P} \text{Prim}\sim$$

In contrast to primitive types, a local type L has no globally valid values associated with it. For example, locations, of type $\text{ref } A$, can never be globally valid, which implies that $\text{ref } A$ is a local type. Thus $\Delta; \Gamma \mid \Psi \vdash M \sim L$ never holds. (See Proposition 2.) A value of a potentially global type may or may not be globally valid depending on whether it contains references to local resources. For example, a function type $A \rightarrow B$ is a potentially global type.

λ_{\square} introduces two new terms $\text{box } M$ and $\text{letbox } x = M \text{ in } N$:

$$\begin{aligned} \text{term } M &::= \dots \mid \text{box } M \mid \text{letbox } x = M \text{ in } M \\ \text{value } V &::= \dots \mid \text{box } M \end{aligned}$$

$\text{box } M$ has a modal type $\square A$, and expects M to evaluate to a globally valid value. $\text{letbox } x = M \text{ in } N$ expects M to evaluate to $\text{box } M'$; then it evaluates M' before substituting the resultant value for x in N . The β -reduction rule for the modality \square uses a capture-avoiding substitution $[V/x]M$ extended in a standard way.

$$\begin{aligned} \text{letbox } x = \text{box } V \text{ in } M &\longrightarrow_{\beta} [V/x]M \\ \text{evaluation context } \phi &::= \dots \mid \text{letbox } x = \phi \text{ in } M \mid \text{letbox } x = \text{box } \phi \text{ in } M \end{aligned}$$

$\text{box } M$ corresponds to the introduction rule for the modality \square . Note that in $\text{letbox } x = M \text{ in } N$, the type of M does not determine the form of the typing judgment for the whole term. That is, regardless of the type of M , there are two possibilities for where the result of evaluating N is valid: at the current node and at any node. Therefore \square has one introduction rule and two elimination rules:

$$\begin{aligned} \frac{\Delta; \Gamma \mid \Psi \vdash M \sim A}{\Delta; \Gamma \mid \Psi \vdash \text{box } M : \square A} \square I \quad & \frac{\Delta; \Gamma \mid \Psi \vdash M : \square A \quad \Delta, x \sim A; \Gamma \mid \Psi \vdash N : C}{\Delta; \Gamma \mid \Psi \vdash \text{letbox } x = M \text{ in } N : C} \square E \\ & \frac{\Delta; \Gamma \mid \Psi \vdash M : \square A \quad \Delta, x \sim A; \Gamma \mid \Psi \vdash N \sim C}{\Delta; \Gamma \mid \Psi \vdash \text{letbox } x = M \text{ in } N \sim C} \square E' \end{aligned}$$

Figure 2 summarizes the definition of λ_{\square} . The operational semantics of λ_{\square} uses the same reduction judgment $M \mid \psi \longrightarrow M' \mid \psi'$ as in λ . Proposition 3 confirms that $\Delta; \Gamma \mid \Psi \vdash M \sim A$ is stronger than $\Delta; \Gamma \mid \Psi \vdash M : A$.

Proposition 2. *If no variables are bound to local types in Δ , then $\Delta; \Gamma \mid \Psi \vdash M \sim L$ is not derivable. That is, if $\Delta; \Gamma \mid \Psi \vdash M \sim A$, then A is not a local type.*

Proposition 3. *The rule $\frac{\Delta; \Gamma \mid \Psi \vdash M \sim A}{\Delta; \Gamma \mid \Psi \vdash M : A}$ Global is admissible.*

3.1 Example

We illustrate the use of the modality \square by rewriting in λ_{\square} the two examples in Introduction. We assume a primitive type int for integers and an infix operator $+$ for adding two integers. We encode a modal type $\square A$ as $\square(\text{unit} \rightarrow A)$, and check the mobility of a term M of type A by checking the mobility of a value $\lambda_{_} : \text{unit}. M$ of type $\text{unit} \rightarrow A$ where $_$ denotes a fresh variable.

The first example is written in λ_{\square} as follows:

$$M_1 = (\lambda r : \text{ref int}. \lambda x : \text{int}. x + !r) (\text{ref } 0)$$

We cannot use M_1 to build a term of type $\square(\text{int} \rightarrow \text{int})$ because there is no typing derivation of $\Delta; \Gamma \mid \Psi \vdash M_1 \sim \text{int} \rightarrow \text{int}$:

$$\frac{\frac{\text{(no typing rule applicable)}}{\Delta; \Gamma \mid \Psi \vdash (\lambda r : \text{ref int}. \lambda x : \text{int}. x + !r) (\text{ref } 0) \sim \text{int} \rightarrow \text{int}}}{\Delta; \Gamma \mid \Psi \vdash \text{box } (\lambda r : \text{ref int}. \lambda x : \text{int}. x + !r) (\text{ref } 0) : \square(\text{int} \rightarrow \text{int})} \square I$$

type	$A ::= P \mid G \mid L$
primitive type	$P ::= \text{unit}$
potentially global type	$G ::= A \rightarrow A \mid \Box A$
local type	$L ::= \text{ref } A$
term	$M ::= \dots \mid \text{box } M \mid \text{letbox } x = M \text{ in } M$
value	$V ::= \dots \mid \text{box } M$
global typing context	$\Delta ::= \cdot \mid \Delta, x \sim A$
local typing context	$\Gamma ::= \cdot \mid \Gamma, x : A$

$$\frac{x \sim A \in \Delta}{\Delta; \Gamma \mid \Psi \vdash x \sim A} \text{GVar} \quad \frac{x \sim A \in \Delta}{\Delta; \Gamma \mid \Psi \vdash x : A} \text{GVar}'$$

$$\frac{\Delta; \Gamma \mid \Psi \vdash M \sim A}{\Delta; \Gamma \mid \Psi \vdash \text{box } M : \Box A} \Box \quad \frac{\Delta; \Gamma \mid \Psi \vdash M : \Box A \quad \Delta, x \sim A; \Gamma \mid \Psi \vdash N : C}{\Delta; \Gamma \mid \Psi \vdash \text{letbox } x = M \text{ in } N : C} \Box E$$

$$\frac{\Delta; \Gamma \mid \Psi \vdash M : \Box A \quad \Delta, x \sim A; \Gamma \mid \Psi \vdash N \sim C}{\Delta; \Gamma \mid \Psi \vdash \text{letbox } x = M \text{ in } N \sim C} \Box E'$$

$$\frac{\Delta; \Gamma \mid \Psi \vdash M : P}{\Delta; \Gamma \mid \Psi \vdash M \sim P} \text{Prim}\sim \quad \frac{\Delta; \cdot \mid \cdot \vdash V : A}{\Delta; \Gamma \mid \Psi \vdash V \sim A} \text{GVal}$$

$$\text{letbox } x = \text{box } V \text{ in } M \longrightarrow_{\beta} [V/x]M$$

evaluation context $\phi ::= \dots \mid \text{letbox } x = \phi \text{ in } M \mid \text{letbox } x = \text{box } \phi \text{ in } M$

Fig. 2. Definition of the modal language λ_{\Box}

M_1 itself, however, is mobile because $\lambda_{-}.\text{unit}$. M_1 is mobile:

$$\frac{\vdots}{\frac{\Delta; \cdot \mid \cdot \vdash \lambda_{-}.\text{unit}. (\lambda r : \text{ref int}. \lambda x : \text{int}. x + !r) (\text{ref } 0) : \text{unit} \rightarrow (\text{int} \rightarrow \text{int})}{\Delta; \Gamma \mid \Psi \vdash \lambda_{-}.\text{unit}. (\lambda r : \text{ref int}. \lambda x : \text{int}. x + !r) (\text{ref } 0) \sim \text{unit} \rightarrow (\text{int} \rightarrow \text{int})} \text{GVal}}$$

The second example is written in λ_{\Box} as follows where variable r is bound to an existing reference of type ref int :

$$M_2 = \text{letbox } v = \text{box } !r \text{ in } \lambda x : \text{int}. x + v$$

We can use M_2 to build a term of type $\Box(\text{int} \rightarrow \text{int})$:

$$\frac{\frac{\frac{\Delta; \Gamma, r : \text{ref int} \mid \Psi \vdash r : \text{ref int}}{\Delta; \Gamma, r : \text{ref int} \mid \Psi \vdash !r : \text{int}} \text{Deref} \quad \frac{\vdots}{\Delta, v \sim \text{int}; x : \text{int} \mid \cdot \vdash x + v : \text{int}} \text{Var}}{\Delta, v \sim \text{int}; \cdot \mid \cdot \vdash \lambda x : \text{int}. x + v : \text{int} \rightarrow \text{int}} \rightarrow \text{I}}{\Delta, v \sim \text{int}; \Gamma, r : \text{ref int} \mid \Psi \vdash \lambda x : \text{int}. x + v \sim \text{int} \rightarrow \text{int}} \text{GVal}}{\frac{\Delta; \Gamma, r : \text{ref int} \mid \Psi \vdash \text{letbox } v = \text{box } !r \text{ in } \lambda x : \text{int}. x + v \sim \text{int} \rightarrow \text{int}}{\Delta; \Gamma, r : \text{ref int} \mid \Psi \vdash \text{box letbox } v = \text{box } !r \text{ in } \lambda x : \text{int}. x + v : \Box(\text{int} \rightarrow \text{int})} \Box \text{I}} \Box E$$

M_2 itself, however, is not mobile because $\lambda_{-}.\text{unit}$. M_2 is not mobile:

$$\frac{\text{(impossible to typecheck because of } r)}{\frac{\Delta; \cdot \mid \cdot \vdash \lambda_{-}.\text{unit}. \text{letbox } v = \text{box } !r \text{ in } \lambda x : \text{int}. x + v : \text{unit} \rightarrow (\text{int} \rightarrow \text{int})}{\Delta; \Gamma, r : \text{ref int} \mid \Psi \vdash \lambda_{-}.\text{unit}. \text{letbox } v = \text{box } !r \text{ in } \lambda x : \text{int}. x + v \sim \text{unit} \rightarrow (\text{int} \rightarrow \text{int})} \text{GVal}}$$

A more straightforward but less satisfactory translation of the second example uses a λ -application instead of a **letbox** construct:

$$M'_2 = (\lambda v : \text{int}. \lambda x : \text{int}. x + v) (!r)$$

M'_2 is operationally equivalent to M_2 , but cannot be used in building a term of type $\Box(\text{int} \rightarrow \text{int})$:

$$\frac{\frac{\text{(no typing rule applicable)}}{\Delta; \Gamma, r : \text{ref int} \mid \Psi \vdash (\lambda v : \text{int}. \lambda x : \text{int}. x + v) (!r) \sim \text{int} \rightarrow \text{int}}}{\Delta; \Gamma, r : \text{ref int} \mid \Psi \vdash \text{box} (\lambda v : \text{int}. \lambda x : \text{int}. x + v) (!r) : \Box(\text{int} \rightarrow \text{int})} \Box$$

The reason why **box** M'_2 fails to have type $\Box(\text{int} \rightarrow \text{int})$ is that the type of $\lambda v : \text{int}. \lambda x : \text{int}. x + v$, namely $\text{int} \rightarrow (\text{int} \rightarrow \text{int})$, fails to express that a mobile value of type $\text{int} \rightarrow \text{int}$ is returned. In fact, the inner λ -abstraction $\lambda x : \text{int}. x + v$ cannot be a mobile value anyway, since a binding $v \sim \text{int}$ is not added to a global typing context.

We could introduce a new type $A \Box \rightarrow B$ for those λ -abstractions taking a value of type A and returning a mobile value of type B . The typing rules for the new connective $\Box \rightarrow$ are given as follows:

$$\frac{\Delta; \Gamma, x : A \mid \Psi \vdash M \sim B}{\Delta; \Gamma \mid \Psi \vdash \lambda x : A. M : A \Box \rightarrow B} \rightarrow_{\Box} \quad \frac{\Delta; \Gamma \mid \Psi \vdash M : A \Box \rightarrow B \quad \Delta; \Gamma \mid \Psi \vdash N : A}{\Delta; \Gamma \mid \Psi \vdash M N \sim B} \rightarrow_{E_{\Box}}$$

Although the new connective $\Box \rightarrow$ allows more flexibility in programming, we decide not to include it in the definition of λ_{\Box} because a simple encoding of $A \Box \rightarrow B$ as $A \rightarrow \Box B$ suffices. For example, we can eliminate the rule \rightarrow_{\Box} by translating $\lambda x : A. M : A \Box \rightarrow B$ into $\lambda x : A. \text{box } M : A \rightarrow \Box B$ and the rule $\rightarrow_{E_{\Box}}$ by translating $M N$ into **letbox** $v = M N$ in v .

3.2 Type Safety of λ_{\Box}

The proof of type safety of λ_{\Box} is routine except for the formulation of the substitution theorem (Theorem 4). In the second clause of the substitution theorem, $\Delta; \cdot \mid \cdot \vdash V : A$ proves that V is a globally valid value of type A , which we substitute for variable x in term M . Corollary 5 follows from the definition of primitive types as given in Definition 1.

Theorem 4 (Substitution)

- If $\Delta; \Gamma \mid \Psi \vdash V : A$ and $\Delta; \Gamma, x : A \mid \Psi \vdash M : C$, then $\Delta; \Gamma \mid \Psi \vdash [V/x]M : C$.
 If $\Delta; \Gamma \mid \Psi \vdash V : A$ and $\Delta; \Gamma, x : A \mid \Psi \vdash M \sim C$, then $\Delta; \Gamma \mid \Psi \vdash [V/x]M \sim C$.
 If $\Delta; \cdot \mid \cdot \vdash V : A$ and $\Delta, x \sim A; \Gamma \mid \Psi \vdash M : C$, then $\Delta; \Gamma \mid \Psi \vdash [V/x]M : C$.
 If $\Delta; \cdot \mid \cdot \vdash V : A$ and $\Delta, x \sim A; \Gamma \mid \Psi \vdash M \sim C$, then $\Delta; \Gamma \mid \Psi \vdash [V/x]M \sim C$.

Corollary 5

- If $\Delta; \Gamma \mid \Psi \vdash V : P$ and $\Delta, x \sim P; \Gamma \mid \Psi \vdash M : C$, then $\Delta; \Gamma \mid \Psi \vdash [V/x]M : C$.
 If $\Delta; \Gamma \mid \Psi \vdash V : P$ and $\Delta, x \sim P; \Gamma \mid \Psi \vdash M \sim C$, then $\Delta; \Gamma \mid \Psi \vdash [V/x]M \sim C$.

Theorem 6 (Progress). Suppose that term M satisfies $\cdot; \cdot \mid \Psi \vdash M : A$ or $\cdot; \cdot \mid \Psi \vdash M \sim A$ for some store typing Ψ and type A . Then either:

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

Theorem 7 (Type preservation)

$$\text{Suppose } \left\{ \begin{array}{l} \cdot; \cdot \mid \Psi \vdash M : A \\ \Psi \vdash \psi \text{ okay} \\ M \mid \psi \longrightarrow M' \mid \psi' \end{array} \right. .$$

$$\text{Then there exists a store typing } \Psi' \text{ such that } \left\{ \begin{array}{l} \cdot; \cdot \mid \Psi' \vdash M' : A \\ \Psi \subset \Psi' \\ \Psi' \vdash \psi' \text{ okay} \end{array} \right. .$$

$$\text{Suppose } \left\{ \begin{array}{l} \cdot; \cdot \mid \Psi \vdash M \sim A \\ \Psi \vdash \psi \text{ okay} \\ M \mid \psi \longrightarrow M' \mid \psi' \end{array} \right. .$$

$$\text{Then there exists a store typing } \Psi' \text{ such that } \left\{ \begin{array}{l} \cdot; \cdot \mid \Psi' \vdash M' \sim A \\ \Psi \subset \Psi' \\ \Psi' \vdash \psi' \text{ okay} \end{array} \right. .$$

3.3 Logic for λ_{\square}

The type system for modal types $\square A$ is unusual in that it differentiates values (*i.e.*, terms in weak head normal form) from ordinary terms, as shown in the rule GVal . This differentiation implies that the logic corresponding to the modality \square via the Curry-Howard isomorphism requires a judgment that inspects not only hypotheses in a proof but also the proof structure itself (*e.g.*, inferences rules used in the proof). Thus the modality \square sets itself apart from other modalities and is not found in any other logic.

In the pure fragment of λ_{\square} without primitive types and local types, the modality \square shows similarities with modal possibility \diamond and lax modality \circ in [10]. Specifically a proof-theoretic analysis of \square gives rise to a new form of substitution $\langle M/x \rangle N$ which is defined inductively on the structure of the term being substituted (*i.e.*, M) instead of the term being substituted into (*i.e.*, N). Let us interpret a β -reduction rule as the reduction of a typing derivation in which an introduction rule is followed by a corresponding elimination rule. For example, the β -reduction rule for the connective \rightarrow may be seen as the reduction of the following typing derivation in the pure λ -calculus (where we omit store typings):

$$\frac{\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x : A. M : A \rightarrow B} \rightarrow I}{\Gamma \vdash (\lambda x : A. M) N : B} \rightarrow E \quad \longrightarrow_{\beta} \quad \Gamma \vdash [N/x]M : B$$

Likewise we obtain a β -reduction rule for \square from the reduction of a typing derivation in which the introduction rule $\square I$ is followed by the elimination rule $\square E$ or $\square E'$ (where we omit store typings):

$$\frac{\frac{\Delta; \Gamma \vdash M \sim A}{\Delta; \Gamma \vdash \text{box } M : \square A} \square I \quad \Delta, x \sim A; \Gamma \vdash N : C}{\Delta; \Gamma \vdash \text{letbox } x = \text{box } M \text{ in } N : C} \square E \quad \longrightarrow_{\beta} \quad \Delta; \Gamma \vdash \langle M/x \rangle N : C$$

To see why $\langle M/x \rangle N$ is defined inductively on the structure of M , observe that the reduction of $\text{letbox } x = \text{box } M \text{ in } N$ requires an analysis of M instead of N . The reason is that only a value can be substituted for x , but M may not be a value; therefore we

have to analyze M to decide how to transform the whole term so that x is eventually replaced by a value. Conceptually N should be replicated at those places within M where the evaluation of M is finished, so that M and N are evaluated exactly once and in that order. If M is already a value V , we reduce the whole term to $[V/x]N$. Thus we are led to define $\langle M/x \rangle N$ as follows:

$$\begin{aligned} \langle V/x \rangle N &= [V/x]N \\ \langle \text{letbox } x' = M' \text{ in } M''/x \rangle N &= \text{letbox } x' = M' \text{ in } \langle M''/x \rangle N \end{aligned}$$

Note that we *cannot* define $\langle M_1 M_2/x \rangle N$ because without primitive types and the rule $\text{Prim}\sim$, there is no typing derivation of $\Delta; \Gamma \vdash M_1 M_2 \sim A$ and thus $\text{box } M_1 M_2$ cannot be well-typed.

In the presence of primitive types, the β -reduction

$$\text{letbox } x = \text{box } M \text{ in } N \longrightarrow_{\beta} \langle M/x \rangle N$$

is no longer valid because $\text{letbox } x = \text{box } M \text{ in } N$ may typecheck while $\langle M/x \rangle N$ is undefined. For example, $M = M_1 M_2$ of type unit satisfies $\Delta; \Gamma \vdash M \sim \text{unit}$ by the rule $\text{Prim}\sim$, but $\langle M_1 M_2/x \rangle N$ is undefined. Intuitively the rule $\text{Prim}\sim$ disguises an un-analyzable term of a primitive type as an analyzable term. Thus, in order to reduce $\text{letbox } x = \text{box } M \text{ in } N$, the operational semantics of λ_{\square} is forced to reduce M into a value V first, instead of analyzing M to transform the whole term. Then an ordinary substitution $[V/x]N$ suffices for the reduction of $\text{letbox } x = \text{box } M \text{ in } N$.

We close this section with a brief discussion of the properties of the modality \square .

- $\square A \rightarrow A$ $\lambda x: \square A. \text{letbox } y = x \text{ in } y$
A mobile value is a special case of an ordinary term.
- $\square A \rightarrow \square \square A$ $\lambda x: \square A. \text{letbox } y = x \text{ in box box } y$
A mobile value itself is mobile.
- $\square(A \rightarrow B) \rightarrow \square A \rightarrow \square B$
A mobile λ -abstraction does not necessarily return a mobile value.

4 λ_{\square}^N with a Network Operational Semantics

While the type system of λ_{\square} is appropriate for understanding the role of the modality \square , it is not expressive enough for distributed computations which may generate terms whose type is determined by *remote nodes*. For example, a future construct $[11]$ initiates a stand-alone computation at a remote node and returns a pointer to the remote node; then the type of the pointer is determined by the term being evaluated at the remote node.

This section extends the type system of λ_{\square} so that we can typecheck such terms, and also develop a network operational semantics to model distributed computations. We refer to the resultant language as λ_{\square}^N . We incorporate a communication construct for remote procedure calls into λ_{\square} so as to allow communications between nodes to actually take place. Type safety of λ_{\square}^N ensures the safety of mobile terms and mobile values.

4.1 Extended Type System and Network Operational Semantics

We represent the state of a network with a *configuration* π which records term M and store ψ associated with each node γ . A *configuration typing* Π records the type of the term being evaluated at each node. We assume that no node appears more than once in π , and consider Π as an unordered set. As a new term, γ serves as a reference to a node.

term	$M ::= \dots \gamma$	(node reference)
configuration	$\pi ::= \cdot \pi, \{M \psi @ \gamma\}$	
configuration typing	$\Pi ::= \cdot \Pi, \gamma \sim A$	$(A \neq L)$

- $\{M | \psi @ \gamma\}$ in π means that node γ is currently evaluating term M with store ψ .
- $\gamma \sim A$ in Π means that the term at node γ evaluates to a value of type A . For the sake of simplicity, we require that every term in a network evaluates to a globally valid value.

The extended type system is formulated with a *configuration typing judgment* $\Pi \vdash \pi \text{ okay}$ which means that configuration π has configuration typing Π . In order to be able to typecheck a node reference γ (which is a term), we include a configuration typing Π in each typing judgment: $\Delta; \Gamma | \Psi | \Pi \vdash M : A$, $\Delta; \Gamma | \Psi | \Pi \vdash M \sim A$, and $\Psi | \Pi \vdash \psi \text{ okay}$. The rules for the extended type system are derived from (and given the same name as) the previous rules by including a configuration typing in every typing judgment. We need two additional typing rules **Node** and **Node'** for node references; the rule **Conf** may be regarded as the definition of the configuration typing judgment.

$$\begin{array}{c}
\frac{\gamma \sim A \in \Pi}{\Delta; \Gamma | \Psi | \Pi \vdash \gamma \sim A} \text{ Node} \quad \frac{\gamma \sim A \in \Pi}{\Delta; \Gamma | \Psi | \Pi \vdash \gamma : A} \text{ Node}' \\
\\
\frac{\begin{array}{c} \gamma \sim A \in \Pi \\ \text{dom}(\Pi) = \text{dom}(\pi) \quad \Psi | \Pi \vdash \psi \text{ okay} \quad \text{for every } \{M | \psi @ \gamma\} \in \pi \\ \cdot; \cdot | \Psi | \Pi \vdash M \sim A \end{array}}{\Pi \vdash \pi \text{ okay}} \text{ Conf}
\end{array}$$

The network operational semantics is formulated with a *configuration reduction judgment* $\pi \Longrightarrow \pi'$, which means that configuration π reduces (or evolves) to configuration π' . We provide two rules for the configuration reduction judgment:

$$\begin{array}{c}
\frac{M | \psi \longrightarrow M' | \psi'}{\pi, \{M | \psi @ \gamma\} \Longrightarrow \pi, \{M' | \psi' @ \gamma\}} \text{ Red} \\
\\
\frac{}{\pi, \{\phi \llbracket \gamma' \rrbracket | \psi @ \gamma\}, \{V | \psi' @ \gamma'\} \Longrightarrow \pi, \{\phi \llbracket V \rrbracket | \psi @ \gamma\}, \{V | \psi' @ \gamma'\}} \text{ Sync}
\end{array}$$

The rule *Red* says that stand-alone computations at individual nodes are part of a distributed computation. In the rule *Sync*, a node reference γ' suspends the stand-alone computation at node γ until it is replaced by a mobile value V through a synchronization operation with node γ' . That is, node reference γ' is *not* a value, but reduces to a mobile value (which is globally valid) only after node γ' has finished evaluating a term. Note that a configuration reduction is non-deterministic because the rule *Red* can choose an arbitrary node γ from a given configuration.

4.2 Communication Construct for Remote Procedure Calls

The network operational semantics becomes interesting only with communication constructs; without communication constructs, all nodes perform stand-alone computations independently of each other, and type safety holds trivially.

In designing communication constructs, we could begin with an existing modal logic, such as S4 and S5, and adhere to a spatial interpretation of the modalities in it. Then there arise a few logically motivated primitive operations, with which we can implement various communication constructs. This approach is appealing because of the strong logical foundation underlying communication constructs as well as the pleasant correspondence between modal logic and distributed computations.

A strict adherence to a spatial interpretation of modal logic, however, sacrifices flexibility in programming. For example, all previous work on modal languages for distributed computations [6,7,8] builds on the idea of using modal types $\Box A$ for mobile *terms* that may be evaluated at any node in the network, which is a typical spatial interpretation of the modality \Box . As modal types $\Box A$ do not ensure the safety of mobile values and safe communications are thus restricted to mobile terms, it is difficult to give a faithful implementation of those communication constructs (such as remote procedure calls, future constructs, and communication channels) that expect or return mobile values.

Since it is not concerned with “how” to transmit mobile values between nodes, the modality \Box itself does not specify a principle for the design of communication constructs in λ_{\Box}^N . Instead we have to design each communication construct individually by exploiting node references γ in conjunction with the rule *Sync*. We do not believe that our approach is *ad hoc*, since not every communication construct can be given a logical interpretation anyway (*e.g.*, communication channels). In fact, even λ_{rpc} [7] and *Lambda 5* [8], both of which are based on a spatial interpretation of modal logic S5, introduce primitive operations which are logically “motivated,” but do not actually have their counterparts in S5.

As an illustration, we develop a communication construct for remote procedure calls. We use the natural deduction style by giving introduction and elimination rules for it. (Thus our communication construct is also logically motivated to a certain extent.) We can use the same idea to develop similar communication constructs, such as future constructs, that transmit both mobile terms and mobile values.

A remote procedure call transmits a mobile term M to a remote node γ to initiate a stand-alone computation, and then waits for the result of evaluating M . In order to ensure the safety of the remote procedure call, M needs to satisfy the following two conditions:

- M itself is globally valid so that the evaluation of M at node γ is safe.
- M evaluates to a globally valid value so that the result of the remote procedure call is valid.

We can test if M satisfies the two conditions by typechecking `box box M`:

- Typechecking the outer `box` construct tests if `box M` evaluates to a globally valid value, in which case M is also globally valid because `box M` is already a value.
- Typechecking the inner `box` construct tests if M evaluates to a globally valid value.

Thus, under typing contexts Δ and Γ and store typing Ψ , we have to prove $\Delta; \cdot \vdash M \sim A$:

$$\frac{\frac{\Delta; \cdot \vdash M \sim A}{\Delta; \cdot \vdash \text{box } M : \Box A} \Box I}{\Delta; \Gamma \mid \Psi \vdash \text{box } M \sim \Box A} \text{GVal}}{\Delta; \Gamma \mid \Psi \vdash \text{box box } M : \Box \Box A} \Box I$$

We do not, however, use a term $\text{box box } M$ of type $\Box \Box A$ for a remote procedure call because there is no way to tell whether such a term is intended for a remote procedure call or just for creating a globally valid value. Instead we introduce a new modality \Box^2 specifically designed for remote procedure calls. As far as the type system is concerned, we may think of $\Box^2 A$ as an abbreviation of $\Box \Box A$. We use $\text{box}^2 M$ and $\text{letbox}^2 x = M \text{ in } N$ for the introduction and elimination rules for \Box^2 , respectively:

$$\begin{array}{ll} \text{potentially global type } G & ::= \dots \mid \Box^2 A \\ \text{term } M & ::= \dots \mid \text{box}^2 M \mid \text{letbox}^2 x = M \text{ in } M \\ \text{value } V & ::= \dots \mid \text{box}^2 M \end{array}$$

$$\frac{\Delta; \cdot \vdash M \sim A}{\Delta; \Gamma \mid \Psi \vdash \text{box}^2 M : \Box^2 A} \Box^2 I \quad \frac{\Delta; \Gamma \mid \Psi \vdash M : \Box^2 A \quad \Delta, x \sim A; \Gamma \mid \Psi \vdash N : C}{\Delta; \Gamma \mid \Psi \vdash \text{letbox}^2 x = M \text{ in } N : C} \Box^2 E$$

$$\frac{\Delta; \Gamma \mid \Psi \vdash M : \Box^2 A \quad \Delta, x \sim A; \Gamma \mid \Psi \vdash N \sim C}{\Delta; \Gamma \mid \Psi \vdash \text{letbox}^2 x = M \text{ in } N \sim C} \Box^2 E'$$

In the rules $\Box^2 E$ and $\Box^2 E'$, it helps to think of $\text{letbox}^2 x = M \text{ in } N$ as $\text{letbox } x' = M \text{ in letbox } x = x' \text{ in } N$ where x' is a fresh variable.

As for the operational semantics, $\Box^2 A$ diverges from $\Box \Box A$. $\text{letbox}^2 x = M \text{ in } N$ at node γ expects M to evaluate to $\text{box}^2 M'$; then it makes a remote procedure call by starting an evaluation of M' at a fresh node γ' and replacing M' by a node reference γ' . When the remote procedure call returns a (globally valid) mobile value V , we replace node reference γ' by V , and then reduce $\text{letbox}^2 x = \text{box}^2 V \text{ in } N$ to $[V/x]N$.

$$\text{letbox}^2 x = \text{box}^2 V \text{ in } M \longrightarrow_{\beta} [V/x]M$$

$$\text{evaluation context } \phi ::= \dots \mid \text{letbox}^2 x = \phi \text{ in } M \mid \text{letbox}^2 x = \text{box}^2 [] \text{ in } M$$

$$\frac{M \neq \gamma' \quad \text{fresh node reference } \gamma'}{\pi, \{ \phi \parallel \text{letbox}^2 x = \text{box}^2 M \text{ in } N \parallel \psi @ \gamma \} \implies \pi, \{ \phi \parallel \text{letbox}^2 x = \text{box}^2 \gamma' \text{ in } N \parallel \psi @ \gamma \}, \{ M \mid \cdot @ \gamma' \}} \text{RPC}$$

In the extended definition of evaluation contexts, an important restriction is that the hole in $\text{letbox}^2 x = \text{box}^2 [] \text{ in } M$ can be filled only with a node reference. For example, $(\text{letbox}^2 x = \text{box}^2 [] \text{ in } M)[\gamma]$ is allowed, but $(\text{letbox}^2 x = \text{box}^2 [] \text{ in } M)[N]$ is not allowed. Without this restriction, $\text{letbox}^2 x = \text{box}^2 N \text{ in } M$ reduces to $\text{letbox}^2 x = \text{box}^2 N' \text{ in } M$ without making a remote procedure call, if N reduces to N' . Note also that $\text{letbox}^2 x = \text{box}^2 \gamma' \text{ in } M$ does not reduce to $[\gamma'/x]M$ because node reference γ' is not a value.

Independently of the modality \Box^2 , the modality \Box is still useful for creating arguments to remote procedure calls. That is, we use \Box to compose mobile terms for remote

procedure calls, and \boxtimes^2 to transmit them to remote nodes. For example, the following term makes a remote procedure call to add two integers n_1 and n_2 , both of which are bound to variables x_1 and x_2 via **letbox** constructs:

$$\text{letbox } x_1 = \text{box } n_1 \text{ in letbox } x_2 = \text{box } n_2 \text{ in letbox}^2 v = \text{box}^2 x_1 + x_2 \text{ in } v$$

4.3 Type Safety of λ_{\boxtimes}^N

Type safety of λ_{\boxtimes}^N consists of *configuration progress* (Theorem 9) and *configuration typing preservation* (Theorem 11). Proofs of Theorems 9 and 11 use type safety for stand-alone computations in λ_{\boxtimes}^N (Theorems 8 and 10).

Theorem 8 (Progress). *Suppose that term M satisfies $\cdot; \cdot \mid \Psi \mid \Pi \vdash M : A$ or $\cdot; \cdot \mid \Psi \mid \Pi \vdash M \sim A$ for some store typing Ψ , configuration typing Π , and type A . Then one of the following holds:*

- (1) M is a value,
- (2) $M = \phi[\![\gamma]\!]$,
- (3) $M = \phi[\![\text{letbox}^2 x = \text{box}^2 M' \text{ in } N]\!]$,
- (4) for any store ψ such that $\Psi \mid \Pi \vdash \psi$ okay, there exist some term M' and store ψ' such that $M \mid \psi \longrightarrow M' \mid \psi'$.

Theorem 9 (Configuration progress). *Suppose $\Pi \vdash \pi$ okay. Then either:*

- (1) π consists only of
 - $\{V \mid \psi @ \gamma\}$,
 - $\{\phi[\![\gamma']]\! \mid \psi @ \gamma\}$,
 - $\{\phi[\![\text{letbox}^2 x = \text{box}^2 \gamma' \text{ in } N]\! \mid \psi @ \gamma\}$,
- (2) there exists π' such that $\pi \Longrightarrow \pi'$.

Theorem 10 (Type preservation)

Suppose $\begin{cases} \cdot; \cdot \mid \Psi \mid \Pi \vdash M : A \\ \Psi \mid \Pi \vdash \psi \text{ okay} \\ M \mid \psi \longrightarrow M' \mid \psi' \end{cases}$.

Then there exists a store typing Ψ' such that $\begin{cases} \cdot; \cdot \mid \Psi' \mid \Pi \vdash M' : A \\ \Psi \subset \Psi' \\ \Psi' \mid \Pi \vdash \psi' \text{ okay} \end{cases}$.

Suppose $\begin{cases} \cdot; \cdot \mid \Psi \mid \Pi \vdash M \sim A \\ \Psi \mid \Pi \vdash \psi \text{ okay} \\ M \mid \psi \longrightarrow M' \mid \psi' \end{cases}$.

Then there exists a store typing Ψ' such that $\begin{cases} \cdot; \cdot \mid \Psi' \mid \Pi \vdash M' \sim A \\ \Psi \subset \Psi' \\ \Psi' \mid \Pi \vdash \psi' \text{ okay} \end{cases}$.

Theorem 11 (Configuration typing preservation)

Suppose $\begin{cases} \Pi \vdash \pi \text{ okay} \\ \pi \Longrightarrow \pi' \end{cases}$.

Then there exists a configuration typing Π' such that $\begin{cases} \Pi \subset \Pi' \\ \Pi' \vdash \pi' \text{ okay} \end{cases}$.

Type safety of λ_{\square}^N implies that mobile terms and mobile values are both safe to use: well-typed terms never go wrong even in the presence of mobile terms and mobile values.

5 Related Work

Borghuis and Feijs [5] present a typed λ -calculus *MTSN* (Modal Type System for Networks) which assumes stationary services (*i.e.*, stationary code) and mobile data. An indexed modal type $\square^{\omega}(A \rightarrow B)$ represents services transforming data of type A into data of type B at node ω . *MTSN* is a task description language rather than a programming language, since services are all “black boxes” whose inner workings are unknown. For example, terms of type $tex \rightarrow dvi$ all describe procedures to convert *tex* files to *dvi* files. Thus reduction on terms is tantamount to simplifying procedures to achieve a certain task.

Jia and Walker [7] present a modal language λ_{rpc} which is based on hybrid logic [12] as every typing judgment explicitly specifies the current node where typechecking takes place. The modalities \square and \diamond are used for mobile terms that can be evaluated at any node and at a certain node, respectively.

Murphy *et al.* [8] present a modal language *Lambda 5* which addresses both code mobility and resource locality. It is based on modal logic S5 where all judgments are relativized to nodes, as in Simpson [13]. A value of type $\square A$ contains a mobile term that can be evaluated at any node, and a value of type $\diamond A$ contains a *label*, a reference to a local resource. A label may appear at remote nodes, but the type system guarantees that it is dereferenced only at the node where it is valid.

λ_{rpc} and *Lambda 5* are fundamentally different from λ_{\square}^N in their use of modal types $\square A$ for remote procedure calls. In both languages, a remote procedure call, by the pull construct in λ_{rpc} and by the *fetch* construct in *Lambda 5*, is given a specific node where the evaluation is to occur, and therefore *does not expect a term contained in a value of type* $\square A$. Instead it expects just a term of type $\square A$, which itself may not be mobile but eventually produces a mobile term valid at any node including the caller node. The resultant mobile term is delivered to (*i.e.*, pulled or fetched by) the caller node, which needs to further evaluate it to obtain a value. As such, both languages do not address the issue of value mobility. In contrast, a remote procedure call in λ_{\square}^N transmits a term *contained* in a value of type $\square^2 A$ and relies on the modality \square^2 to directly return a mobile value.

Moody [6] presents a system which is based on modal logic S4. The modality \square is used for mobile terms that can be evaluated at any node, and the modality \diamond is used for terms located at some node. As in λ_{rpc} and *Lambda 5*, remote procedure calls use modal types $\square A$ to transmit mobile terms to unknown remote nodes. Moody’s system uses the elimination rules for the modalities \square and \diamond to send mobile terms to remote nodes, and does not provide a separate construct for remote procedure calls.

Liblit and Aiken [14] give a type-theoretic analysis of pointers in distributed computations. Their type systems distinguish between global pointers (for global address space) and local pointers (for local address space), and deal with safety and performance issues with global pointer dereferencing. While the use of type qualifiers gives a powerful type inference system for minimizing the number of global pointers, their language

focuses on distributed data rather than mobile code, and it is not obvious whether their type systems can be extended to include mobile code.

6 Conclusion and Future Work

We present a modal language λ_{\square}^N for distributed computations which ensure the safety of both mobile terms and mobile values with a single modality \square . The modality \square is more expressive than the necessity modality \Box from modal logic, and enables us to achieve a more faithful implementation of common communication constructs. \square is, however, useful in λ_{\square}^N only because the unit of communication includes values. That is, if the unit of communication was just terms and did not include values, \Box would be enough and \square would be unnecessary.

A drawback of λ_{\square}^N is that references to local resources cannot be transmitted to remote nodes. As an example, consider a location l of type $\text{ref } A$ at node γ . Node γ wishes to share l among all its child nodes, *e.g.*, those nodes created by remote procedure calls. No child node, however, even knows the existence of l because references to local resources cannot escape their host nodes.

To overcome this drawback, we are currently investigating another modality \diamond which is similar to the modality \diamond of λ_{rpc} [7], *but focuses on values rather than terms*. The idea is that term M in $\text{dia } M$ of type $\diamond A$ evaluates to a value valid at a certain node that is unknown to the type system but known to the runtime system. The modality \diamond makes it possible for mobile terms to contain references to remote resources, thereby allowing more flexibility in programming for distributed computations.

References

1. Cardelli, L.: A language with distributed scope. In: Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, ACM Press (1995) 286–297
2. Knabe, F.C.: Language Support for Mobile Agents. PhD thesis, Department of Computer Science, Carnegie Mellon University (1995)
3. Mascolo, C., Picco, G.P., Roman, G.C.: A fine-grained model for code mobility. In: Proceedings of the 7th European Software Engineering Conference held jointly with the 7th ACM SIGSOFT International Symposium on Foundations of Software Engineering, Springer-Verlag (1999) 39–56
4. Sekiguchi, T., Yonezawa, A.: A calculus with code mobility. In: FMOODS '97: Proceeding of the IFIP TC6 WG6.1 International Workshop on Formal Methods for Open Object-based Distributed Systems, Chapman & Hall, Ltd. (1997) 21–36
5. Borghuis, T., Feijs, L.: A constructive logic for services and information flow in computer networks. *The Computer Journal* **43**(4) (2000) 275–289
6. Moody, J.: Modal logic as a basis for distributed computation. Technical Report CMU-CS-03-194, Carnegie Mellon University (2003)
7. Jia, L., Walker, D.: Modal proofs as distributed programs (extended abstract). In Schmidt, D., ed.: Proceedings of the European Symposium on Programming, LNCS 2986, Springer (2004) 219–233

8. Murphy, VII, T., Crary, K., Harper, R., Pfenning, F.: A symmetric modal lambda calculus for distributed computing. In: Proceedings of the 19th IEEE Symposium on Logic in Computer Science (LICS 2004), IEEE Press (2004)
9. Park, S.: A modal language for the safety of mobile values. Technical Report POSTECH-CSE-06-001, Department of Computer Science and Engineering, Pohang University of Science and Technology (2006)
10. Pfenning, F., Davies, R.: A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science* **11**(4) (2001) 511–540
11. Halstead, Jr., R.H.: Multilisp: a language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems* **7**(4) (1985) 501–538
12. Braüner, T.: Natural deduction for hybrid logic. *Journal of Logic and Computation* **14**(3) (2004) 329–353
13. Simpson, A.K.: The Proof Theory and Semantics of Intuitionistic Modal Logic. PhD thesis, Department of Philosophy, University of Edinburgh (1994)
14. Liblit, B., Aiken, A.: Type systems for distributed data structures. In: Proceedings of the 27th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, ACM Press (2000) 199–213