

Modal Proofs As Distributed Programs^{*}

(Extended Abstract)

Limin Jia David Walker

Princeton University
{ljia, dpw}@cs.princeton.edu
fax: 609-258-1771

Abstract. We develop a new foundation for distributed programming languages by defining an intuitionistic, modal logic and then interpreting the modal proofs as distributed programs. More specifically, the proof terms for the various modalities have computational interpretations as *remote procedure calls*, commands to *broadcast* computations to all nodes in the network, commands to use *portable* code, and finally, commands to invoke computational *agents* that can find their own way to safe places in the network where they can execute. We prove some simple meta-theoretic results about our logic as well as a safety theorem that demonstrates that the deductive rules act as a sound type system for a distributed programming language.

1 Introduction

One of the characteristics of distributed systems that makes developing robust software for them far more difficult than developing software for single stand-alone machines is *heterogeneity*. Different places in the system may have vastly different properties and resources. For instance, different machines may be attached to different hardware devices, have different software installed and run different services. Moreover, differing security concerns may cause different hosts to provide different interfaces to distributed programs, even when the underlying computational resources are similar.

In order to model such heterogeneous environments, programming language researchers usually turn to formalisms based on one sort of process algebra or another. Prime examples include the distributed join calculus [6] and the ambient calculus [3]. These calculi often come with rich theories of process equivalence and are useful tools for reasoning about distributed systems. However, a significant disadvantage of starting with process algebras as a foundation for distributed computing is that they immediately discard the wealth of logical principles that underlie conventional sequential programming paradigms and that form the sequential core of any distributed computation.

^{*} This research was supported in part by NSF Career Award CCR-0238328 and DARPA award F30602-99-1-0519.

In this paper, we develop a foundation for safe distributed programming, which rather than rejecting the logical foundations of sequential (functional) programming, extends them with new principles tuned to programming in heterogeneous distributed environments. More specifically, we develop an intuitionistic, modal logic and provide an operational interpretation of the logical proofs as distributed programs. Our logic has the property that at any given place, all of the intuitionistic propositional tautologies are provable. Consequently, our correspondence between proofs and programs implies we can safely execute any (closed) functional program at any place in our distributed programming environment.

We extend these simple intuitionistic proofs with modal connectives and provide computational interpretations of the connectives as operations for remote code execution:

- Objects with type $\tau @ p$ are return values with type τ . They are the results of *remote procedure calls* from the place p .
- Objects with type $\Box \tau$ are computations that run safely everywhere, and may be *broadcast* to all places in the network.
- Objects with type $\boxplus \tau$ are logically equivalent to those objects with type $\Box \tau$, but are treated operationally simply as *portable* code that can run everywhere, but is not actually broadcast.
- Objects with type $\Diamond \tau$ are computational *agents* that have internalized the place p where they may execute safely to produce a value with type τ .

The central technical contributions of our work may be summarized as follows.

- We develop an intuitionistic, modal logic from first principles following the logical design techniques espoused by Martin L of [12] and Frank Pfenning [14, 15]. Our logic is a relative of the hybrid logics, which are discussed in more detail at the end of the paper (see Section 4).
- Each connective in the logic is defined orthogonally to all others, is shown to be locally sound and complete, and supports the relevant substitution principles.
- This paper concentrates on the natural deduction formulation of the logic due to its correspondence with functional programs. However, we have also developed a sequent calculus that has cut elimination and shown that the sequent calculus is sound and complete with respect to the natural deduction system.
- We give an operational interpretation of the proofs in our logic as distributed functional programs. We prove that the logical deduction rules are sound when viewed as a type system for the programming language.

Due to space considerations, we have omitted many details from this extended abstract. Proofs of our theorems, language extensions (including an additional modality interpreted as place-relative remote procedure call, references and recursive functions), and further programming examples may be found in our extended technical report [9].

2 A Logic of Places

2.1 Preliminaries

The central purpose of our logic is to facilitate reasoning about heterogeneous distributed systems where different nodes may have different properties and may contain different resources. Hence, the primary judgment in the logic not only considers *whether* a formula is true, but also *where* it is true. More precisely, each primitive judgment has the form

$$\vdash^P F \text{ at } p$$

where F is a formula from the logic and p is a particular place in the system. These places are drawn from the set P . When P is unimportant or easy to guess from the context (i.e. most of the time) we omit them from the judgment and simply write $\vdash F \text{ at } p$. For any syntactic object X , $\text{FP}(X)$ denote the set of free places that appear in X . We consider judgments $\vdash^P F \text{ at } p$ to have no meaning if $\text{FP}(F) \cup \{p\} \not\subseteq P$.

To gain some intuition about this preliminary set up, consider a network of five computers (A, B, C, D, E). Each of these computers (A, B, C, etc.) may have different physical devices attached to it. For instance, E may have a printer attached to it and B may have a scanner attached. If *sc* (“there is a scanner here”) and *pt* (“there is a printer here”) are propositions in the logic, we might assert judgments such as $\vdash \text{pt at } E$ and $\vdash \text{sc at } B$ to describe this situation.

Hypothetical Judgments In order to engage in more interesting reasoning about distributed resources, we must define hypothetical judgments, facilities for reasoning from hypotheses and the appropriate notion of substitution. To begin with, hypothetical judgments have the form $\Delta \vdash^P F \text{ at } p$ where Δ is a list of (variable-labeled) assumptions:

$$\text{contexts } \Delta ::= \cdot \mid \Delta, x : F \text{ at } p$$

We do not distinguish between contexts that differ only in the order of assumptions. We use hypotheses according to the following inference rule (where L stands for use of a Local hypothesis).

$$\frac{}{\Delta, x : F \text{ at } p \vdash F \text{ at } p} L$$

Intuitionistic Connectives With the definition of hypothetical judgments in hand, we may proceed to give the meaning of the usual intuitionistic connectives for truth (\top), implication ($F_1 \rightarrow F_2$) and conjunction ($F_1 \wedge F_2$) in terms of their introduction and elimination rules.

$$\frac{}{\Delta \vdash \top \text{ at } p} \top I$$

$$\frac{\Delta, x : F_1 \text{ at } p \vdash F_2 \text{ at } p}{\Delta \vdash F_1 \rightarrow F_2 \text{ at } p} \rightarrow I \quad \frac{\Delta \vdash F_1 \rightarrow F_2 \text{ at } p \quad \Delta \vdash F_1 \text{ at } p}{\Delta \vdash F_2 \text{ at } p} \rightarrow E$$

$$\frac{\frac{\Delta \vdash F_1 \text{ at } p \quad \Delta \vdash F_2 \text{ at } p}{\Delta \vdash F_1 \wedge F_2 \text{ at } p} \wedge I}{\frac{\Delta \vdash F_1 \wedge F_2 \text{ at } p}{\Delta \vdash F_1 \text{ at } p} \wedge E1 \quad \frac{\Delta \vdash F_1 \wedge F_2 \text{ at } p}{\Delta \vdash F_2 \text{ at } p} \wedge E2}$$

None of the rules above are at all surprising: Each rule helps explain how one of the usual intuitionistic connectives operates at a particular place. Hence, if we limit the set of places to a single place “_” the logic will reduce to ordinary intuitionistic logic. So far, there is no interesting way to use assumptions at multiple different places, but we will see how to do that in a moment.

As a simple example, consider reasoning about the action of the printer at place E in the network we introduced earlier. Let **pdf** be a proposition indicating the presence of a PDF file waiting to be printed and **po** be a proposition indicating the presence of a printout. The following derivation demonstrates how we might deduce the presence of a printout at E. The context Δ referenced below is

$$f_E : \text{pdf at } E, \text{ptr}_E : \text{pt at } E, \text{print} : \text{pdf} \wedge \text{pt} \rightarrow \text{po at } E$$

This context represents the presence of a PDF file and printer at E as well as some software (a function) installed at E that can initiate the printing process.

$$\frac{\frac{\frac{\Delta \vdash \text{pdf} \wedge \text{pt} \rightarrow \text{po at } E}{\Delta \vdash \text{pdf} \wedge \text{pt at } E} L \quad \frac{\frac{\Delta \vdash \text{pdf at } E}{\Delta \vdash \text{pdf at } E} L \quad \frac{\Delta \vdash \text{pt at } E}{\Delta \vdash \text{pt at } E} L}{\Delta \vdash \text{pdf} \wedge \text{pt at } E} \wedge I}{\Delta \vdash \text{po at } E} \rightarrow E$$

2.2 Inter-place Reasoning

To reason about relationships between objects located at different places we introduce a modal connective, which describes objects in terms of their location in the system.

We derive our modal connective by internalizing the judgmental notion that a formula is true at a particular place, but not necessarily elsewhere. We write this new modal formula as $F @ p$. The introduction and elimination rules follow.

$$\frac{\Delta \vdash F \text{ at } p}{\Delta \vdash F @ p \text{ at } p'} @ I \quad \frac{\Delta \vdash F @ p \text{ at } p'}{\Delta \vdash F \text{ at } p} @ E$$

This connective allows us to reason about objects, software or devices “from a distance.” For instance, in our printer example, it is possible to refer to the printer located at E while reasoning at D; to do so we might assert $\Delta \vdash \text{pt} @ E \text{ at } D$. Moreover, we can relate objects at one place to objects at another. For instance, in order to share E’s printer, D needs to have software that can convert local PDF files at D to files that may be used and print properly at E (perhaps this software internalizes some local fonts, inaccessible to E). An assumption of the form $\text{DtoE} : \text{pdf} \rightarrow \text{pdf} @ E \text{ at } D$ would allow us to reason about

such software.¹ If Δ' is the assumption DtoE above together with an assumption $f_D : \text{pdf at } D$, the following derivation allows us to conclude that we can get the PDF file to E. We can easily compose this proof with the earlier one to demonstrate that PDF files at D can not only be sent to E, but actually printed there.

$$\frac{\frac{\overline{\Delta' \vdash \text{pdf} \rightarrow \text{pdf @ } E \text{ at } D} \quad L \quad \overline{\Delta' \vdash \text{pdf at } D} \quad L}{\Delta' \vdash \text{pdf @ } E \text{ at } D} \quad \text{@ } E}{\Delta' \vdash \text{pdf at } E} \rightarrow E$$

2.3 Global Reasoning

While our focus is on reasoning about networks with heterogeneous resources, we cannot avoid the fact that certain propositions are true *everywhere*. For instance, the basic laws of arithmetic do not change from one machine to the next, and consequently, we should not restrict the application of these laws to any particular place. Just as importantly, we might want to reason about distributed applications deployed over a network of machines, all of which support a common operating system interface. The functionality provided by the operating system is available everywhere, just like the basic laws of arithmetic, and use of the functionality need not be constrained to one particular place or another.

Global Judgments To support global reasoning, we generalize the judgment considered so far to include a second context that contains assumptions that are valid everywhere. Our extended judgments have the form $\Gamma; \Delta \vdash^P F \text{ at } p$ where Γ is a global context and Δ is the local context we considered previously.

$$\begin{aligned} \text{Global Contexts } \Gamma &::= \cdot \mid \Gamma, x : F \\ \text{Local Contexts } \Delta &::= \cdot \mid \Delta, x : F \text{ at } p \end{aligned}$$

Our extended logic contains two sorts of hypothesis rules, one for using each sort of assumption. Rule L is identical to the hypothesis rule used in previous sections (modulo the unused global context Γ). Rule G specifies how to use global hypotheses; they may be placed in any location and used there.

$$\frac{\overline{\Gamma; \Delta, x : F \text{ at } p \vdash F \text{ at } p} \quad L \quad \overline{\Gamma, x : F; \Delta \vdash F \text{ at } p} \quad G$$

All rules from the previous sections are included in the new system unchanged aside from the fact that Γ is passed unused from conclusion to premises.

Internalizing Global Truth The modal connective $\Box F$ internalizes the notion that the formula F is true everywhere. If a formula may be proven true at a new place p , which by definition can contain no local assumptions, then that formula must be true everywhere:²

¹ We assume that “@” binds tighter than implication or conjunction. When fully parenthesized, the assumption above has the following form: $(\text{pdf} \rightarrow (\text{pdf @ } E)) \text{ at } D$

² By convention, judgments are meaningless when they contain places not contained in the sets P that annotate the judgment. Consequently, this rule requires that $p \notin \text{FP}(\Gamma) \cup \text{FP}(\Delta) \cup \text{FP}(F)$.

$$\frac{\Gamma; \Delta \vdash^{P+p} F \text{ at } p}{\Gamma; \Delta \vdash^P \Box F \text{ at } p'} \Box I$$

Here, we use $P + p$ to denote the disjoint union $P \cup \{p\}$. If $p \in P$, we consider $P + p$, and any judgment containing such notation, to be undefined. If we can prove $\Box F$, we can assume that F is globally true in the proof of any other judgment F' at p' :

$$\frac{\Gamma; \Delta \vdash^P \Box F \text{ at } p \quad \Gamma, x : F; \Delta \vdash^P F' \text{ at } p'}{\Gamma; \Delta \vdash^P F' \text{ at } p'} \Box E$$

Returning to our printing example, suppose node E decides to allow *all* machines to send it PDF files. In order to avoid hiccups in the printing process, E intends to distribute software to all machines that allow them to convert local PDF files to files that will print properly on E 's printer. We might represent this situation with a hypothesis $\text{ToE} : \Box(\text{pdf} \rightarrow \text{pdf} @ E) \text{ at } E$. Now, given a PDF file at any other node q in the network (coded as the assumption $f_q : \text{pdf} \text{ at } q$), we can demonstrate that it is possible to send the PDF file to E using the following proof, where Δ'' contains assumptions ToE and f_q . The global context Γ contains the single assumption $\text{ToE}' : \text{pdf} \rightarrow \text{pdf} @ E$.

$$\frac{\mathcal{D} \quad \frac{\frac{\Gamma; \Delta'' \vdash \text{pdf} \rightarrow \text{pdf} @ E \text{ at } q \quad G \quad \frac{\Gamma; \Delta'' \vdash \text{pdf} \text{ at } q \quad L}{\Gamma; \Delta'' \vdash \text{pdf} @ E \text{ at } q} \rightarrow E}{\Gamma; \Delta'' \vdash \text{pdf} @ E \text{ at } q} @ E}{\Gamma; \Delta'' \vdash \text{pdf} \text{ at } E} \Box E}{\vdash; \Delta'' \vdash \text{pdf} \text{ at } E} \Box E$$

where the derivation $\mathcal{D} = \frac{\vdash; \Delta'' \vdash \Box(\text{pdf} \rightarrow \text{pdf} @ E) \text{ at } E \quad L}{\vdash; \Delta'' \vdash \Box(\text{pdf} \rightarrow \text{pdf} @ E) \text{ at } E} L$

The Truth is Out There The dual notion of a globally true proposition F is a proposition that is true *somewhere*, although we may not necessarily know where. We already have all the judgmental apparatus to handle this new idea; we need only internalize it in a connective ($\Diamond F$). The introduction rule states that if the formula holds at any particular place p in the network, then it holds somewhere. The elimination rule explains how we use a formula F that holds somewhere: We introduce a new place p and assume F holds there.

$$\frac{\Gamma; \Delta \vdash^P F \text{ at } p}{\Gamma; \Delta \vdash^P \Diamond F \text{ at } p'} \Diamond I \quad \frac{\Gamma; \Delta \vdash^P \Diamond F \text{ at } p' \quad \Gamma; \Delta, x : F \text{ at } p \vdash^{P+p} F' \text{ at } p''}{\Gamma; \Delta \vdash^P F' \text{ at } p''} \Diamond E$$

Modal Axioms Possibility and necessity satisfy the following standard modal axioms (taken from Huth and Ryan [8, p. 284]).

- K: $\vdash; \cdot \vdash \Box(F_1 \rightarrow F_2) \rightarrow (\Box F_1 \rightarrow \Box F_2) \text{ at } p$
- B: $\vdash; \cdot \vdash F \rightarrow \Box \Diamond F \text{ at } p$
- D: $\vdash; \cdot \vdash \Box F \rightarrow \Diamond F \text{ at } p$
- T: $\vdash; \cdot \vdash \Box F \rightarrow F \text{ at } p$
- 4: $\vdash; \cdot \vdash \Box F \rightarrow \Box \Box F \text{ at } p$
- 5: $\vdash; \cdot \vdash \Diamond F \rightarrow \Box \Diamond F \text{ at } p$

2.4 Properties

We have proved local soundness and completeness properties (i.e., subject reduction for beta-reduction and eta-expansion of proofs) for each of the connectives in our logic, as well as standard substitution properties. In addition, we have developed a sequent calculus and proved the following two theorem:

Theorem 1. *The natural deduction system is sound and complete with regard to the sequent calculus (SEQ).*

Theorem 2. *The cut rule is admissible in SEQ.*

3 λ_{rpc} : A Distributed Programming Language

The previous section developed an intuitionistic, modal logic capable of concisely expressing facts about the placement of various objects in a network. Here, we present the proof terms of the logic and show how they may be given an operational interpretation as a distributed programming language, which we call λ_{rpc} . The logical formulas serve as types that prevent distributed programs from “going wrong” by attempting to access resources that are unavailable at the place where the program is currently operating.

$$\begin{array}{l}
 \text{Types } \tau ::= \\
 \mathbf{b} \mid \top \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \wedge \tau_2 \mid \tau @ p \mid \square \tau \mid \boxplus \tau \mid \diamond \tau \\
 \\
 \text{Proof Terms/Programs } e ::= \\
 c \mid x \mid \mathbf{sync}(x) \mid \mathbf{run}(x [p]) \mid () \quad \text{const/var}/\top \\
 \mid \lambda x:\tau.e \mid e_1 e_2 \quad \text{functions}(\rightarrow) \\
 \mid \langle e_1, e_2 \rangle \mid \pi_i e \quad \text{pairs}(\wedge) \\
 \mid \mathbf{ret}(e, p) \mid \mathbf{rpc}(e, p) \quad \text{rpc}(@) \\
 \mid \mathbf{close}(\lambda p.e) \mid \mathbf{bc} e_1 \text{ at } p \text{ as } x \text{ in } e_2 \quad \text{broadcast}(\square) \\
 \mid \mathbf{port}(\lambda p.e) \mid \mathbf{pull} e_1 \text{ at } p \text{ as } x \text{ in } e_2 \quad \text{portable}(\boxplus) \\
 \mid \mathbf{agent}[e, p] \mid \mathbf{go} e_1 \text{ at } p \text{ return } x, q \text{ in } e_2 \text{ agent}(\diamond)
 \end{array}$$

Fig. 1. λ_{rpc} Syntax

3.1 Syntax and Typing

Figure 1 presents the syntax of programs and their types, and Figure 2 presents the typing rules for the language, which are the natural deduction-style proof rules for the logic.

Types and Typing Judgments. The types correspond to the formulas of the logic; we use the meta variable τ rather than F to indicate a shift in interpretation. We also included a set of base types (**b**).

Since we have discovered two different operational interpretations of $\Box F$, and we would like to explain both of them in this paper, we have extended the language of formulas (types) to include an extra modality $\boxplus\tau$ to handle the second interpretation. To support the two universal modalities, we also separate the logical context Γ into two parts, Γ_{\Box} and Γ_{\boxplus} , during type checking. Hence the overall type checking judgment has the following form.

$$\Gamma_{\Box}; \Gamma_{\boxplus}; \Delta \vdash^P e : \tau \text{ at } p$$

By deleting either $\Box\tau$ or $\boxplus\tau$ (and the associated context), we recover *exactly the same logic* as discussed in the previous section.³

Programs. The programs include an unspecified set of constants (c), and the standard introduction and elimination forms for unit, functions and pairs.

Variables from each different context are used in different ways. As a mnemonic for the different sorts of uses, we have added some syntactic sugar to the standard proof terms. Uses of local variables from Δ are just like ordinary uses of variables in your favorite (call-by-value) functional language so they are left undecorated. Variables in Γ_{\Box} refer to computations that have been broadcast at some earlier point. In order to use such a variable, the program must *synchronize* with the concurrently executing computation. Hence, we write **sync**(x) for such uses. Variables in Γ_{\boxplus} refer to portable closures. The use of a variable in this context corresponds to *running* the closure with the current place p as an argument. Hence, we write **run**($x [p]$) for such uses.

Our first modality $\tau @ p$ has an operational interpretation as a remote procedure call. The introduction form (**ret**(e, p)) constructs a “return value” for a remote procedure call. This “return value” can actually be an arbitrary expression e , which will be returned to and run at the place p . The elimination form (**rpc**(e, p')) is the remote procedure call itself. It sends the expression e to the remote site p' where e will be evaluated. If the expression is well-typed, it will eventually evaluate to **ret**(e', p): a return value that can be run safely at the caller’s place, which, in this case is place p .

The introduction form for $\Box F$ is **close**($\lambda p. e$). It creates a closure that may be *broadcast* by the elimination form (**bc** e_1 **at** p_1 **as** x **in** e_2) to every node in the network. More specifically, the elimination form executes e_1 at p_1 , expecting e_1 to evaluate to **close**($\lambda p. e$). When it does, the broadcast expression chooses a new universal reference for the closure, which is bound to x , and sends $\lambda p. e$ to every place in the network where it is applied to the current place and the resulting expression is associated with its universal reference. Finally expression e_2 is executed with the universal reference bound to x . Remote procedure calls or

³ We could have defined two languages, one language for each interpretation of $\Box F$, where the typing rules of each language correspond exactly to the logical inference rules. However, we obtain a more powerful result by showing that these two interpretations inter-operate.

$$\begin{array}{c}
\frac{}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta, x : \tau \text{ at } p \vdash^P x : \tau \text{ at } p} L \quad \frac{}{\Gamma_{\square}, x : \tau; \Gamma_{\boxplus}; \Delta \vdash^P \mathbf{sync}(x) : \tau \text{ at } p} G_{\square} \\
\frac{}{\Gamma_{\square}; \Gamma_{\boxplus}, x : \tau; \Delta \vdash^P \mathbf{run}(x[p]) : \tau \text{ at } p} G_{\boxplus} \\
\frac{}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P () : \top \text{ at } p} Unit \quad \frac{}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P c : \mathbf{b} \text{ at } p} Const \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta, x : \tau_1 \text{ at } p \vdash^P e : \tau_2 \text{ at } p}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \lambda x : \tau_1. e : \tau_1 \rightarrow \tau_2 \text{ at } p} \rightarrow I \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e_1 : \tau_1 \rightarrow \tau_2 \text{ at } p \quad \Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e_2 : \tau_1 \text{ at } p}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e_1 e_2 : \tau_2 \text{ at } p} \rightarrow E \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e_1 : \tau_1 \text{ at } p \quad \Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e_2 : \tau_2 \text{ at } p}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \langle e_1, e_2 \rangle : \tau_1 \times \tau_2 \text{ at } p} \wedge I \quad \frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e : \tau_1 \times \tau_2 \text{ at } p}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \pi_i e : \tau_i \text{ at } p} \wedge E \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e : \tau \text{ at } p}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \mathbf{ret}(e, p) : \tau @ p \text{ at } p'} @ I \quad \frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e : \tau @ p \text{ at } p'}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \mathbf{rpc}(e, p') : \tau \text{ at } p} @ E \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^{P+p} e : \tau \text{ at } p}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \mathbf{close}(\lambda p. e) : \square \tau \text{ at } p'} \square I \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e_1 : \square \tau \text{ at } p \quad \Gamma_{\square}, x : \tau; \Gamma_{\boxplus}; \Delta \vdash^P e_2 : \tau' \text{ at } p'}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \mathbf{bc} e_1 \text{ at } p \text{ as } x \text{ in } e_2 : \tau' \text{ at } p'} \square E \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^{P+p} e : \tau \text{ at } p}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \mathbf{port}(\lambda p. e) : \boxplus \tau \text{ at } p'} \boxplus I \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e_1 : \boxplus \tau \text{ at } p \quad \Gamma_{\square}; \Gamma_{\boxplus}, x : \tau; \Delta \vdash^P e_2 : \tau' \text{ at } p'}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \mathbf{pull} e_1 \text{ at } p \text{ as } x \text{ in } e_2 : \tau' \text{ at } p'} \boxplus E \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e : \tau \text{ at } p}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \mathbf{agent}[e, p] : \diamond \tau \text{ at } p'} \diamond I \\
\frac{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P e_1 : \diamond \tau \text{ at } p' \quad \Gamma_{\square}; \Gamma_{\boxplus}; \Delta, x : \tau \text{ at } p \vdash^{P+p} e_2 : \tau' \text{ at } p''}{\Gamma_{\square}; \Gamma_{\boxplus}; \Delta \vdash^P \mathbf{go} e_1 \text{ at } p' \mathbf{return} x, p \text{ in } e_2 : \tau' \text{ at } p''} \diamond E
\end{array}$$

Fig. 2. λ_{rpc} Typing

broadcasts generated during evaluation of e_2 may refer to the universal reference bound to x , which is safe, since x has been broadcast everywhere.

Objects of type $\sqsupset\tau$ are portable closures; they may be run anywhere. The elimination form (**pull** e_1 **at** p_1 **as** x **in** e_2) takes advantage of this portability by first computing e_1 at p_1 , which should result in a value with the form **port**($\lambda p. e$). Next, it pulls the closure $\lambda p. e$ from p_1 and substitutes it for x in e_2 . The typing rules will allow x to appear anywhere, including in closures in e_2 that will eventually be broadcast or remotely executed. Once again, this is safe since e is portable and runs equally well everywhere.

Our last connective $\diamond\tau$ is considered the type of a computational agent that is smart enough to know where it can go to produce a value with type τ . We introduce such an agent by packaging an expression with a place where the expression may successfully be run to completion. The elimination form (**go** e_1 **at** p_1 **return** x, p **in** e_2) first evaluates e_1 at p_1 , producing an agent (**agent**[e, p_2]). Next, it commands the agent to go to the hidden place p_2 and execute its encapsulated computation there. When the agent has completed its task, it synchronizes with the current computation and e_2 continues with p bound to p_2 and x bound to a value that is safe to use at p_2 .

Simple examples. To gain a little more intuition about how to write programs in this language, consider computational interpretations of some of the proofs from the previous section. The context Δ referenced below contains the following assumptions.

$$\begin{aligned} f_D &: \text{pdf at } D \quad \text{print} : \text{pdf} \wedge \text{pt} \rightarrow \text{po at } E \\ f_E &: \text{pdf at } E \quad \text{DtoE} : \text{pdf} \rightarrow \text{pdf @ } E \text{ at } D \\ \text{ptr}_E &: \text{pt at } E \quad \text{ToE} : \square(\text{pdf} \rightarrow \text{pdf @ } E) \text{ at } E \end{aligned}$$

Printing a PDF file (involving local computation only):

$$;\Delta \vdash \text{print}(f_E, \text{ptr}_E) : \text{po at } E$$

Fetching a PDF file (involving a remote procedure call in which the computation $\text{DtoE } f_D$ is executed at D):

$$;\Delta \vdash \text{rpc}(\text{DtoE } f_D, D) : \text{pdf at } E$$

Fetching then printing:

$$;\Delta \vdash (\lambda x:\text{pdf}.\text{print}(x, \text{ptr}_E))(\text{rpc}(\text{DtoE } f_D, D)) : \text{po at } E$$

Broadcasting E 's PDF conversion function to all nodes then fetching a PDF file from node q (recall that in general, uses of these global variables involves synchronizing with the broadcast expression; below the broadcast expression is a value, but we synchronize anyway):

$$;\Delta, f_q : \text{pdf at } q \vdash \text{bc ToE at } E \text{ as ToE}' \text{ in } \text{rpc}(\text{sync}(\text{ToE}') f_q, q) : \text{pdf at } E$$

Another way to manage PDF files is to make them portable. For instance, if C and D contain portable PDF files, then E can pull these files from their resident locations and print them on its local printer. Remember that portable values are

polymorphic closures that are “run” when used. In this case, the closure simply returns the appropriate PDF file.

$$\begin{aligned} & \vdash \Delta, f_C : \text{pdf at } C, f_D : \text{pdf at } D \vdash \\ & \quad \mathbf{pull } f_C \text{ at } C \text{ as } f'_C \text{ in} \\ & \quad \mathbf{pull } f_D \text{ at } D \text{ as } f'_D \text{ in} \\ & \quad \mathbf{let } _ = \mathbf{print}(\mathbf{run}(f'_C [E]), \mathbf{ptr}_E) \text{ in} \\ & \quad \mathbf{let } _ = \mathbf{print}(\mathbf{run}(f'_D [E]), \mathbf{ptr}_E) \text{ in} \\ & \quad \dots \qquad \qquad \qquad : \tau \text{ at } E \end{aligned}$$

3.2 Operational Semantics and Safety

When one defines an operational semantics for a language, it is essential to choose the correct level of abstraction. When the level of abstraction is too low, details get in the way of understanding the larger issues; when the level of abstraction is too high, there may be confusion as to how the semantics relates to reality.

In our case, we could give an operational semantics based directly on proof reduction, but this semantics would be at too high a level of abstraction to observe important details concerning the distribution of values over a network. In particular, we would not be able to distinguish between the two very different interpretations of \square . Consequently, we give an operational semantics at a lower level of abstraction than proof reduction by including an explicit, concrete network in the semantics as shown in Figure 3.2.⁴ Nevertheless, the basis for the semantics is the interaction of introduction and elimination rules as the proof theory suggests. The various new syntactic objects we use to specify our operational model are listed below.

<i>Networks</i>	$\mathcal{N} ::= (P, \mathcal{L})$
<i>Process Envs.</i>	$\mathcal{L} ::= \cdot \mid \mathcal{L}, \ell \rightarrow e \text{ at } p$
<i>Values</i>	$v ::= c \mid \lambda x:\tau.e \mid \langle v_1, v_2 \rangle \mid \mathbf{ret}(e, p)$ $\quad \mid \mathbf{close}(\lambda p.e) \mid \mathbf{port}(\lambda p.e) \mid \mathbf{agent}[e, p]$
<i>RT Terms</i>	$e ::= \dots \mid \mathbf{sync}(\ell) \mid \mathbf{run}(\lambda p.e [p_1]) \mid \mathbf{sync}(\mathbf{rpc}(\ell, p))$ $\quad \mid \mathbf{sync}(\mathbf{bc } \ell \text{ at } p \text{ as } x \text{ in } e_2) \mid \mathbf{sync}(\mathbf{pull } \ell \text{ at } p \text{ as } x \text{ in } e_2)$ $\quad \mid \mathbf{sync}_1(\mathbf{go } \ell \text{ at } p \text{ return } x, q \text{ in } e) \mid \mathbf{sync}_2(\mathbf{go } \ell \text{ at } p \text{ return } x, q \text{ in } e)$
<i>Contexts</i>	$C ::= [] \mid C e_2 \mid v_1 C \mid \langle C, e_2 \rangle \mid \langle v_1, C \rangle \mid \pi_i C$

Networks \mathcal{N} are pairs consisting of a set of places P , a distributed process environment \mathcal{L} . We have seen places before. The process environment \mathcal{L} is a

⁴ Our choice here is also informed by the history of operational interpretations of linear logic. Several researchers [1, 11] gave interpretations of linear logic using a store-less semantics derived directly from logical proof reduction. These semantics led to *significant* confusion about the memory management properties of the languages: How many pointers to a linear object could there be, and can linear objects be deallocated after they are used? It was impossible to say because the operational models did not contain pointers! Only when Chirimar et al. [4] and Turner and Wadler [18] later gave interpretations that deviated from the proof-theoretic interpretation by including an explicit store was the story made clear. Our interpretation of modal logic deviates from the proof theory in a similar way to these latter works as we include an explicit network.

$\mathcal{L} \mapsto \mathcal{L}'$	
sync OS	$\mathcal{L}, \ell' \rightarrow C[\mathbf{sync}(\ell)] \text{ at } p, \ell \rightarrow v \text{ at } p$ $\mapsto \mathcal{L}, \ell' \rightarrow C[v] \text{ at } p, \ell \rightarrow v \text{ at } p$
run OS	$\mathcal{L}, \ell \rightarrow C[\mathbf{run}(\lambda p.e[p_1])] \text{ at } p_2$ $\mapsto \mathcal{L}, \ell \rightarrow C[e[p_1/p]] \text{ at } p_2$
\rightarrow OS	$\mathcal{L}, \ell \rightarrow C[(\lambda x:\tau.e)v] \text{ at } p$ $\mapsto \mathcal{L}, \ell \rightarrow C[e[v/x]] \text{ at } p$
\wedge OS	$\mathcal{L}, \ell \rightarrow C[\pi_i\langle v_1, v_2 \rangle] \text{ at } p$ $\mapsto \mathcal{L}, \ell \rightarrow C[v_i] \text{ at } p$
$\textcircled{\@}$ OS1	$\mathcal{L}, \ell \rightarrow C[\mathbf{rpc}(e, p_1)] \text{ at } p_0$ $\mapsto \mathcal{L}, \ell \rightarrow C[\mathbf{sync}(\mathbf{rpc}(\ell_1, p_1))] \text{ at } p_0, \ell_1 \rightarrow e \text{ at } p_1$
$\textcircled{\@}$ OS2	$\mathcal{L}, \ell \rightarrow C[\mathbf{sync}(\mathbf{rpc}(\ell_1, p_1))] \text{ at } p_0, \ell_1 \rightarrow \mathbf{ret}(e, p_0) \text{ at } p_1$ $\mapsto \mathcal{L}, \ell \rightarrow C[e] \text{ at } p_0, \ell_1 \rightarrow \mathbf{ret}(e, p_0) \text{ at } p_1$
\square OS1	$\mathcal{L}, \ell \rightarrow C[\mathbf{bc} e_1 \text{ at } p_1 \text{ as } x \text{ in } e_2] \text{ at } p_0$ $\mapsto \mathcal{L}, \ell \rightarrow C[\mathbf{sync}(\mathbf{bc} \ell_1 \text{ at } p_1 \text{ as } x \text{ in } e_2)] \text{ at } p_0, \ell_1 \rightarrow e_1 \text{ at } p_1$
\square OS2	$\mathcal{L}, \ell \rightarrow C[\mathbf{sync}(\mathbf{bc} \ell_1 \text{ at } p_1 \text{ as } x \text{ in } e_2)] \text{ at } p_0, \ell_1 \rightarrow \mathbf{close}(\lambda p.e) \text{ at } p_1$ $\mapsto \mathcal{L}, \ell \rightarrow C[e_2[\ell_2/x]] \text{ at } p_0, \ell_1 \rightarrow \mathbf{close}(\lambda p.e) \text{ at } p_1,$ $\{\ell_2 \rightarrow e[q/p] \text{ at } q\} (\forall q \in P)$
\boxplus OS1	$\mathcal{L}, \ell \rightarrow C[\mathbf{pull} e_1 \text{ at } p_1 \text{ as } x \text{ in } e_2] \text{ at } p_0$ $\mapsto \mathcal{L}, \ell \rightarrow C[\mathbf{sync}(\mathbf{pull} \ell_1 \text{ at } p_1 \text{ as } x \text{ in } e_2)] \text{ at } p_0, \ell_1 \rightarrow e_1 \text{ at } p_1$
\boxplus OS2	$\mathcal{L}, \ell \rightarrow C[\mathbf{sync}(\mathbf{pull} \ell_1 \text{ at } p_1 \text{ as } x \text{ in } e_2)] \text{ at } p_0, \ell_1 \rightarrow \mathbf{port}(\lambda p.e) \text{ at } p_1$ $\mapsto \mathcal{L}, \ell \rightarrow C[e_2[\lambda p.e/x]] \text{ at } p_0, \ell_1 \rightarrow \mathbf{port}(\lambda p.e) \text{ at } p_1$
\diamond OS1	$\mathcal{L}, \ell \rightarrow C[\mathbf{go} e_1 \text{ at } p_1 \text{ return } x, q \text{ in } e_2] \text{ at } p_0$ $\mapsto \mathcal{L}, \ell \rightarrow C[\mathbf{sync}_1(\mathbf{go} \ell_1 \text{ at } p_1 \text{ return } x, q \text{ in } e_2)] \text{ at } p_0, \ell_1 \rightarrow e_1 \text{ at } p_1$
\diamond OS2	$\mathcal{L}, \ell \rightarrow C[\mathbf{sync}_1(\mathbf{go} \ell_1 \text{ at } p_1 \text{ return } x, q \text{ in } e_2)] \text{ at } p_0, \ell_1 \rightarrow \mathbf{agent}[e, p_2] \text{ at } p_1$ $\mapsto \mathcal{L}, \ell \rightarrow C[\mathbf{sync}_2(\mathbf{go} \ell_2 \text{ at } p_2 \text{ return } x, q \text{ in } e_2)] \text{ at } p_0,$ $\ell_1 \rightarrow \mathbf{agent}[e, p_2] \text{ at } p_1, \ell_2 \rightarrow e \text{ at } p_2$
\diamond OS3	$\mathcal{L}, \ell \rightarrow C[\mathbf{sync}_2(\mathbf{go} \ell_1 \text{ at } p_1 \text{ return } x, q \text{ in } e_2)] \text{ at } p_0, \ell_1 \rightarrow v \text{ at } p_1$ $\mapsto \mathcal{L}, \ell \rightarrow C[e_2[p_1/q][v/x]] \text{ at } p_0, \ell_1 \rightarrow v \text{ at } p_1$

Fig. 3. λ_{rpc} Operational Semantics

finite partial map from places p in P to process ids to expressions. We write these partial maps as lists of elements with the form $\ell \rightarrow e$ at p . We assume that no pair of place and location (p and ℓ) appears in two different components of the map. We do not distinguish between maps that differ only in the ordering of their elements. $\mathcal{L}(p)(\ell)$ denotes e when $\mathcal{L} = \mathcal{L}', \ell \rightarrow e$ at p . We use the notation $\mathcal{L} \setminus \ell$ to denote the mapping \mathcal{L} with all elements of the form $\ell \rightarrow e$ at p removed.

We also introduce new expressions (the RT terms above) that only occur at run time to give an operation semantics to our program. Run-time terms are used to represent expressions, which are suspended part-way through evaluation and are waiting to synchronize with remotely executing expressions.

Lastly, we define evaluation contexts C to specify the order of evaluation.

In order to show that the network is well-typed at every step in evaluation, we add typing rules to give types to the RT terms and we also give well-formedness

conditions for the network as a whole. The typing judgment for a network has the form $\vdash \mathcal{L} : \Gamma_{\square}; \cdot; \Delta$. See the technical report [9] for further details. *Operational Rules and Type Safety*. The state of a network $\mathcal{N} = (P, \mathcal{L})$ evolves according to the operational rules listed in Figure 3.2. These rules specify a relation with the form $\mathcal{L} \mapsto \mathcal{L}'$.

The type system is sound with respect to our operational semantics for distributed program evaluation. The proofs of Preservation and Progress theorems, stated below, follow the usual strategy.

Theorem 3 (Preservation). *If $\vdash \mathcal{L} : \Gamma_{\square}; \cdot; \Delta$ and $\mathcal{L} \mapsto \mathcal{L}'$ then there exists Γ'_{\square} and Δ' such that $\vdash \mathcal{L}' : \Gamma'_{\square}; \cdot; \Delta'$.*

Theorem 4 (Progress). *If $\vdash \mathcal{L} : \Gamma_{\square}; \cdot; \Delta$ then either*

- $\mathcal{L} \mapsto \mathcal{L}'$, or
- for all places p in P , and for all ℓ in $\text{Dom}(\mathcal{L}(p))$, $\mathcal{L}(p)(\ell)$ is a value.

4 Discussion

Extensions and Variations This paper presents a solid foundation on which to build a distributed functional programming language. However, in terms of language design, it is only the beginning. A few interesting extensions and variations are discussed briefly below.

- Values everywhere. Our interpretation of \square involves either broadcasting a closure or substituting a closure into local code. In each case, there is some computational overhead to manage the closure: Either we synchronize or we run the closure when it gets used. To avoid this overhead, we could place a value restriction on the expression in the introduction form for \square . One possibility that is definitely not an option is evaluating eagerly under \square before broadcasting or substitution: In the presence of references (and almost certainly other effects), this evaluation strategy is unsound.
- Dynamic network evolution. The current work assumes that the set of network places and the network topology is fixed. While this is a reasonable assumption for some distributed programming environments, others allow the topology to evolve. An interesting challenge for future work is to extend our logic and language with features that express evolution. We believe that the new name connectives developed in the context of nominal logics [16, 13] may be of help here.
- Synchronous and asynchronous variations. Just as ordinary sequential programming languages may be defined with different evaluation strategies (call-by-value, call-by-name, call-by-need), it appears possible to develop different operational interpretations of the modal connectives in which execution is more or less synchronized. For instance, when defining the operation of the \square -connective, we could wait until all broadcast expressions have completed evaluation before proceeding with the evaluation of the second expression e_2 . Likewise, remote procedure calls are synchronized: Evaluation does not

proceed until they have received the return value, even though the following computation does not necessarily require the value immediately. In the future, we plan to explore these nuances in greater detail.

Related Work Hybrid logics are an old breed of logic that date back to Arthur Prior’s work in the 1960s [17]. As does our logic, they mix modal necessity and possibility with formulas such as $F @ p$ that are built from pure names. More recently, researchers have developed a rich semantic theory for these logics and studied both tableau proofs and sequents; many resources on these topics and others are available off the hybrid logics web page.⁵ However, work on hybrid logics is usually carried out in a classical setting and we have not found an intuitionistic, natural deduction style proof theory like ours that can serve as a foundation for distributed functional programming languages. The reason for this is likely that Pfenning and Davies [15] have only recently shown how to assign a robust natural deduction-style proof theory to pure modal logic.

Cardelli and Gordon’s ambient logic [2] highlights the idea that modalities for possibility and necessity need not only be interpreted *temporally*, but can also be interpreted *spatially*, and this abstract idea was a central influence in our work. However, at a more technical level, the ambient logic is entirely different from the logic we develop here: The ambient logic has a widely different set of connectives, is classical as opposed to intuitionistic, and is defined exclusively by a sequent calculus rather than by natural deduction. Moreover, it does not serve as a type system for ambient programs; rather, it is a tool for reasoning about them.

Much effort has been invested in interpreting proof search in sequent calculus of linear logic as concurrent computation. Recently along this line of research, Kobayashi et al. have explained how modal linear logic can be viewed as a distributed concurrent programming language [10]. In this work, they introduce a formula similar to our $F @ p$ but they define it axiomatically rather than through introduction and elimination rules (or left and right rules), as we have done. Consequently, their formulation cannot serve as the basis for a functional programming language. They did not consider modal necessity and possibility.

Another major influence on our work is Pfenning and Davies’ judgmental reconstruction of modal logic [15], which is developed in accordance with Martin Löf’s design patterns for type theory [12]. Pfenning and Davies go on to interpret modal necessity temporally (as opposed to spatially) in their work on staged computation [5]. One obvious technical difference between our logic and theirs is that our logic is founded on local judgments that include the specific place where a proposition is true whereas their logic is not.

Concurrently with this research, Harper, Moody and Pfenning [7] have begun to investigate the role that modal logics may play in developing a foundation for distributed computing. More specifically, they interpret objects with type $\Box \tau$ as jobs that may be injected into the Grid and run anywhere. This interpretation

⁵ See <http://www.hylo.net>.

appears to resemble our portable code more closely than our broadcast code. In addition, in their Grid computing domain, every node is assumed to contain identical resources, so they are investigating type systems derived from pure modal logics rather than hybrid logics like the one we have presented here. Consequently, they have no analogue of our remote procedure calls.

References

1. S. Abramsky. Computational interpretations of linear logic. *Theoretical Computer Science*, 111:3–57, 1993.
2. L. Cardelli and A. Gordon. Anytime, anywhere: Modal logics for mobile ambients. In *Twenty-Seventh ACM Symposium on Principles of Programming Languages*, pages 365–377. ACM Press, Jan. 2000.
3. L. Cardelli and A. D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240(1):177–213, June 2000.
4. J. Chirimar, C. A. Gunter, and J. G. Riecke. Proving memory management invariants for a language based on linear logic. In *ACM Conference on Lisp and Functional Programming*, pages 139–150, Apr. 1992.
5. R. Davies and F. Pfenning. A modal analysis of staged computation. *Journal of the ACM*, 48(3):555–604, 2001.
6. C. Fournet, G. Gonthier, J.-J. Lévy, L. Maranget, and D. Rémy. A calculus of mobile agents. In *7th International Conference on Concurrency Theory (CONCUR'96)*, volume 1119 of *Lecture Notes in Computer Science*, pages 406–421, Pisa, Italy, Aug. 1996. Springer.
7. R. Harper. Trustless grid computing in concert, June 2003. Talk. <http://www-2.cs.cmu.edu/concert/talks/Harper2003Trustless/trustless.ppt>.
8. M. R. A. Huth and M. D. Ryan. *Logic in Computer Science: Modelling and Reasoning about Systems*. Cambridge University Press, Cambridge, England, 2000.
9. L. Jia and D. Walker. Modal proofs as distributed programs. Technical Report TR-671-03, Princeton University, 2003. Forthcoming.
10. N. Kobayashi, T. Shimizu, and A. Yonezawa. Distributed concurrent linear logic programming. *Theoretical Computer Science*, 227(1–2):185–220, Sept. 1999.
11. P. Lincoln and J. Mitchell. Operational aspects of linear lambda calculus. In *Proceedings 7th Annual IEEE Symp. on Logic in Computer Science, LICS'92, Santa Cruz, CA, USA, 22–25 June 1992*, pages 235–246. IEEE Computer Society Press, Los Alamitos, CA, 1992.
12. M. Löf. On the meanings of the logical constants and the justifications of the logical laws. Technical Report 2, University of Siena, 1985.
13. D. Miller and A. Tiu. A proof theory for generic judgments. In *IEEE Symposium on Logic in Computer Science*, pages 118–127, Ottawa, Canada, June 2003.
14. F. Pfenning. Logical frameworks. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, chapter 16, pages 977–1061. Elsevier Science and MIT Press, 2001.
15. F. Pfenning and R. Davies. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science*, 11:511–540, 2001.
16. A. M. Pitts. Nominal logic, a first order theory of names and binding. *Information and computation*, 2003. To appear.
17. A. Prior. *Past, present and future*. Oxford University press, 1967.
18. D. N. Turner and P. Wadler. Operational interpretations of linear logic. *Theoretical Computer Science*, 227:231–248, 1999. Special issue on linear logic.