

Distributed Control Flow with Classical Modal Logic *

Tom Murphy VII Karl Crary Robert Harper
Carnegie Mellon University (tom7, crary, rwh@cs.cmu.edu)

Abstract

In previous work we presented a foundational calculus for spatially distributed computing based on intuitionistic modal logic. Through the modalities \Box and \Diamond we were able to capture two key invariants: the mobility of portable code and the locality of fixed resources.

This work investigates issues in distributed control flow through a similar propositions-as-types interpretation of *classical* modal logic. The resulting programming language is enhanced with the notion of a network-wide continuation, through which we can give computational interpretation of classical theorems (such as $\Box A \equiv \neg \Diamond \neg A$). Such continuations are also useful primitives for building higher-level constructs of distributed computing. The resulting system is elegant, logically faithful, and computationally reasonable.

1 Introduction

This paper is an exploration of distributed control flow using a propositions-as-types interpretation of classical modal logic. We build on our previous intuitionistic calculus, Lambda 5 [8], which is a simple programming language (and associated logic) for distributed computing. Lambda 5 focuses particularly on the spatial distribution of programs, and allows the programmer to express the *place* in which computation occurs using modal typing judgments. Through the modal operators \Box and \Diamond we are then able to express invariants about mobility and locality of resources. Our new calculus, C5, extends Lambda 5 with network-wide first-class continuations, which arise naturally from the underlying clas-

sical logic. Network-wide continuations create a new relationship between the modalities \Box and \Diamond , which we see with several examples, and serve as building-blocks for other useful primitives. Before we introduce C5, we begin with a short reprise of Lambda 5.

1.1 Lambda 5

The Lambda 5 programming model is a network with many different *places*, or *nodes*. In order to be faithful to this model, we use a style of logic that has the ability to reason simultaneously from multiple perspectives, namely, modal logic. Compared to propositional logic, which is concerned with *truth*, modal logic deals with truth at different *worlds*. These worlds are related by an *accessibility relation*, which affects the strength of the modal connectives; different assumptions about accessibility give rise to different modal logics. For modelling a network where the worlds are nodes, we choose Intuitionistic S5 [15], whose relation is reflexive, symmetric, and transitive—every world is related to every other world. Therefore, except when comparing it to other systems, we essentially dispense with the accessibility relation altogether. This leads to a simple explanation of the judgments and connectives, which is as follows.

A true $A@w$ is the basic judgment, meaning that the proposition A is true at the world w (we abbreviate this to $A@w$). There are two new proposition forms for quantifying over worlds. $\Box A$ is the statement that A is true at every world. $\Diamond A$ means that A is true at *some* world. Because we think of these worlds as places in the network, operationally we interpret type $\Box A$ as representing mobile code or data of type A , and the type $\Diamond A$ as an address of code or data of type A .

Propositions must be situated at a world in order to be judged true, so it is important to distinguish between the proposition $\Box A$ and the judgment $\Box A@w$, the latter meaning that A is true in every world *from the perspective of* w . In S5, every world has the same perspective with regard to statements about *all* or *some* world(s). But operationally this will be significant, as

*The ConCert Project is supported by the National Science Foundation under grant ITR/SY+SI 0121633: “Language Technology for Trustless Software Dissemination”.

there is no true “global” code, only mobile code that currently exists at some world.

Though the logic distinguishes between $\Box A @ \omega$ and $\Box A @ \omega'$, both have precisely the same immediate consequences. The typical rule for eliminating \Box , for instance as given by Simpson [15] is

$$\frac{\Box A @ \omega}{A @ \omega'}$$

With this rule, it *never really matters* where $\Box A$ exists, since we can eliminate it instantly to any world. However, we really do care operationally where mobile code resides, and so we adjust the rules of Lambda 5 to reflect this bias. The logic features a novel *decomposition* into locally-acting introduction and elimination rules, i.e.

$$\frac{\Box A @ \omega}{A @ \omega}$$

as well as *motion* rules for moving between worlds, i.e.

$$\frac{\Box A @ \omega}{\Box A @ \omega'}$$

We argue [8] that this results in a more appropriate operational interpretation. Our classical system also features this decomposition, and like Lambda 5, we are able to retain a crisp connection to the underlying logic.

Although distributed computing problems are often thought of as being concurrent, both Lambda 5 and our new calculus are sequential. We consider concurrency an orthogonal issue, although we give remarks on how it can be accomplished in Section 6.

1.2 Classical Control Flow

The notion that control operators such as Scheme’s `call/cc` or Felleisen’s \mathcal{C} can be given logical meaning via classical logic is well known. Essentially, if we interpret the type $\neg A$ as a continuation expecting a value of type A , then the types of these operators are classical tautologies. Griffin first proposed this in 1990 [5] with later refinements by (for example) Murthy [9]. Parigot’s $\lambda\mu$ -calculus [10] takes this idea and develops it into a full-fledged natural deduction system for classical logic.¹ It began to become clear that this was no accident—classical logic *is* the logic of control flow.

Therefore, a natural next step is to look at *classical* S5 to see what kind of programming language it gives us, which is the topic of this paper. We find that the notion of a network-wide continuation arises naturally, giving a computational explanation to (intuitionistically ridiculous) classical theorems such as

¹Our calculus is quite similar to his (extended to the modal case!), although we prefer to present it with an emphasis on *truth* and *falsehood* judgments.

$\Box A \equiv \neg \Diamond \neg A$. We also believe that such primitives can be useful for building distributed computing mechanisms such as synchronous message passing.

The paper proceeds as follows. We first present classical S5 judgmentally, giving a natural deduction system and intuition for its operational behavior. We then verify that our proof system really is faithful by establishing a correspondence with a dual sequent calculus that admits cut. Next we give proof terms for some classical theorems, to elucidate the new connection between \Box and \Diamond made possible by network-wide continuations. In order to make these intuitions concrete, we then give an operational semantics based on an abstract network. We follow with some ideas about concurrency and how network-wide continuations can be used by distributed applications, and conclude with a discussion of related work.

2 Classical S5

Because we wish to take a propositions-as-types interpretation of modal logic, a judgmental proof theory for our logic is critical. In this section we give such a presentation of Classical S5.

Because modal logic is concerned with truth relativized to worlds, our judgments must reflect that. We have two main judgments in our proof theory. The first,

$$A \text{ true } @ \omega$$

simply states that the proposition A is true at the world ω . Dually, we have

$$A \text{ false } \star \omega$$

which says that the proposition A is *false* at the world ω . Although these two judgments are dual, the natural deduction system is biased towards the first; it is primarily concerned with deducing that propositions are true. We will only make assumptions about falsehood for the purpose of deriving a contradiction. As is standard, we reify the hypotheses about truth and falsehood into contexts (eliding `true` and `false`), and the central judgment of our proof theory becomes

$$\Gamma; \Delta \vdash A @ \omega$$

where we deduce that A is true at world ω under truth assumptions of the form $B @ \omega'$ appearing in Γ and falsehood assumptions of the form $C \star \omega''$ appearing in Δ . We also have hypotheses about the existence of worlds. It is cumbersome to write a context of world assumptions and conditions on world existence in every judgment. Instead we use pure hypothetical notation

$$\Gamma; \Delta \vdash A @ \omega$$

$$\begin{array}{c}
\frac{\Gamma, x:A@{\omega}; \Delta \vdash M : B@{\omega}}{\Gamma; \Delta \vdash \lambda x.M : A \supset B@{\omega}} \supset I \\
\frac{\Gamma; \Delta \vdash M : A@{\omega'}}{\Gamma; \Delta \vdash \mathbf{box} \omega'.M : \Box A@{\omega}} \Box I \\
\frac{\Gamma, x:A@{\omega'}; \Delta \vdash N : B@{\omega}}{\Gamma; \Delta \vdash M : \Diamond A@{\omega}} \Diamond I \\
\frac{\Gamma; \Delta, u:A\star\omega \vdash M : A@{\omega}}{\Gamma; \Delta \vdash \mathbf{letcc} u \mathbf{in} M : A@{\omega}} bc \\
\frac{\Gamma; \Delta \vdash M : A@{\omega} \quad \Gamma; \Delta \vdash N : B@{\omega}}{\Gamma; \Delta \vdash \langle M, N \rangle : A \wedge B@{\omega}} \wedge I \\
\frac{\Gamma; \Delta \vdash N : A@{\omega}}{\Gamma; \Delta \vdash M : A \supset B@{\omega}} \supset E \\
\frac{\Gamma; \Delta \vdash M : \Box A@{\omega}}{\Gamma; \Delta \vdash \mathbf{unbox} M : A@{\omega}} \Box E \\
\frac{\Gamma; \Delta \vdash M : A@{\omega}}{\Gamma; \Delta \vdash \mathbf{here} M : \Diamond A@{\omega}} \Diamond I \\
\frac{\Gamma; \Delta, u:A\star\omega \vdash M : A@{\omega}}{\Gamma; \Delta, u:A\star\omega \vdash \mathbf{throw} M \mathbf{to} u : C@{\omega'}} \# \\
\frac{\Gamma; \Delta \vdash M : A_1 \wedge A_2@{\omega}}{\Gamma; \Delta \vdash \pi_i M : A_i@{\omega}} \wedge E_i \\
\frac{\Gamma, x:A@{\omega}, \Gamma'; \Delta \vdash x : A@{\omega}}{\Gamma, x:A@{\omega}, \Gamma'; \Delta \vdash x : A@{\omega}} hyp \\
\frac{\Gamma; \Delta \vdash M : \Box A@{\omega'} \quad \omega'}{\Gamma; \Delta \vdash \mathbf{get}_{\Box}[\omega']M : \Box A@{\omega}} \Box M \\
\frac{\Gamma; \Delta \vdash M : \Diamond A@{\omega'} \quad \omega'}{\Gamma; \Delta \vdash \mathbf{get}_{\Diamond}[\omega']M : \Diamond A@{\omega}} \Diamond M \\
\frac{\Gamma; \Delta \vdash M : \perp@{\omega'} \quad \omega'}{\Gamma; \Delta \vdash \mathbf{rpc}[\omega']M : C@{\omega}} \perp E
\end{array}$$

Figure 1: Classical S5 natural deduction

to express a judgment hypothetical in the existence of world ω' (which may be the same as ω !). We also take the common shortcut of only permitting mention of worlds that exist. Therefore, all judgments are hypothetical in at least *some* world (the world at which the conclusion is formed), until we introduce world constants in Section 5.

Operationally, we will think of a falsehood assumption $A\star\omega$ as a continuation, living at world ω , that expects something of type A .

Our natural deduction system appears in Figure 1. These rules include proof terms, which we will explain shortly. Aside from the falsehood context, the rules for \Box , \Diamond and \supset are the same as in Lambda 5. The new connectives \perp (discussed below) and \wedge are treated as they would be in the intuitionistic case. The major additions are the structural rules *bc* (by contradiction) and *#* (contradict), which enable classical reasoning.

The *bc* rule is read as follows: In order to prove $A@{\omega}$, we can assume that A is false at ω . This corresponds directly to the classical axiom $(\neg A \supset A) \supset A$. Operationally, this grabs the current continuation and binds the falsehood variable to it. The *#* rule may be alarming at first glance, because it requires the assumption $A\star\omega$ to appear in the conclusion. This is because the *#* rule is actually the *hypothesis* rule for falsehood assumptions, and will have a corresponding substitution principle.² The rule simply states that if we have the

²A theory of *hypothetical hypotheticals* would be able to express this in a less awkward—but perhaps no less alarming—way. Abel [1] for instance gives such a third-order encoding of the $\lambda\mu$ -

assumption that A is false and are able to prove that A is true (at the same world), then we can deduce a contradiction and thus any proposition. The *#* rule is realized operationally as a **throw** of an expression to a matching continuation. Note that continuations are *global*—we can throw from any world to a remote continuation $A\star\omega$, provided that we are able to construct a proof of $A@{\omega}$.

The rules for \Box and \Diamond are important to review. \Box elimination is the easiest to understand: If we know that $\Box A$ is true at some world, then we know A is true at the same world. To prove $\Box A$, we must prove A at a hypothetical world about which nothing is known (rule $\Box I$). Operationally, we realize $\Box A$ as a piece of suspended code, with the hypothetical world ω' bound within it. Introduction of \Diamond is simple; if we know A then we know that A is true *somewhere* (namely here). Operationally this will record the value in a table and return an address that witnesses its existence. Elimination of \Diamond is as follows: if we know $\Diamond A$, then we know there is some world where A is true (but we don't know anything else about it). Call this world ω' and assume $A@{\omega'}$ in order to continue reasoning. Finally, we provide *motion* rules (as per our decomposition) $\Box M$ and $\Diamond M$. Both simply allow knowledge of $\Box A$ or $\Diamond A$ at one world to be transported to another. Operationally these move the values between worlds.

Bottom has no introduction form, but we allow the *remote elimination* of it (rule $\perp E$). This is similar to the motion rules for \Box and \Diamond , except that nothing is

ever returned, because there is no canonical value of type \perp . For this reason we call the proof term `rpc`, as it invokes a sort of remote procedure call on the target world.³

For each kind of hypothesis we have a substitution theorem.

Theorem 1 (Truth Substitution)

If $\Gamma; \Delta \vdash M : A @ \omega$
 and $\Gamma, x:A @ \omega; \Delta \vdash N : B @ \omega'$
 then $\Gamma; \Delta \vdash [M/x]N : B @ \omega'$.

Theorem 2 (Falsehood Substitution)

If $\forall C, \omega''. \Gamma, x:A @ \omega; \Delta \vdash M : C @ \omega''$
 and $\Gamma; \Delta, u:A \star \omega \vdash N : B @ \omega'$
 then $\Gamma; \Delta \vdash \llbracket x.M/u \rrbracket N : B @ \omega'$.

Here, truth substitution $[M/x]N$ is defined in the standard way. Note again that a proof of type $B @ \omega'$ can contain sub-expressions that are well-typed at *other* worlds! In the operational semantics we will always ship these expressions to their home world before evaluating them. Theorem 2, however, warrants special attention. This principle is dual to the $\#$ rule just as Theorem 1 is dual to *hyp*. The $\#$ rule contradicts an $A \star \omega$ with an $A @ \omega$, so to eliminate a falsehood assumption by substitution we are able to assume $A @ \omega$ and must produce another contradiction. Reading \vdash as logical consequence, we have that if A false gives B , and A true gives C (for all C), then B . This can easily be seen as a consequence of excluded middle. We write this substitution as $\llbracket x.M/u \rrbracket N$ where x is a binder (with scope through M) that stands for the value thrown to u . It is defined pointwise on N except for a use of the $\#$ rule on u :

$$\llbracket x.M/u \rrbracket \text{throw } N' \text{ to } u \doteq [N'/x]M$$

This principle is close to what Parigot calls *structural substitution* for the $\lambda\mu$ -calculus. Operationally, we see this as replacing the throw with some other handler for A . Since the new handler must have parametric type, typically it is a `throw` to some other continuation, perhaps after performing some computation on the proof of A .

Proof of Theorem 2 is by a straightforward induction on the derivation of B , appealing to Theorem 1 in the case above. \square

We wish to know that our proof theory (specially constructed to give rise to a good operational semantics) is not simply ad hoc; that it really embodies classical S5, and is globally sound. To do so we prove in the next section a correspondence to a straightforward sequent formulation of classical S5 with the subformula

³We could have equivalently had a `get \perp` and a local `abort`, but there appears to be no practical use to this decomposition.

property. We'll use the sequent calculus as intuition as we develop proof terms for some classically true propositions. Following that is a discussion of the operational semantics (Section 5), which does not depend on the sequent calculus.

3 Sequent Calculus

Our sequent calculus is motivated by simplicity and duality alone, because we will not give it a computational interpretation. One traditional way of doing classical theorem proving is to negate the target formula and prove a contradiction from it. We base our sequent calculus around this view: the sequent

$$\Gamma \# \Delta$$

means that the truth assumptions in Γ and the falsehood assumptions in Δ are mutually contradictory.⁴ The calculus is given in Figure 2. We treat contexts as unordered multisets, so the *action* can occur anywhere in either context.

These rules should be read bottom-up, as if during proof search. The *contra* rule allows us to form a contradiction whenever a proposition is both true and false at the same world. The $\square T$ rule says that if we know $\square A @ \omega$, then we know $A @ \omega'$ for any ω' that exists. On the other hand, if we know that $\square A$ is false, then we know A is false at some world ω' . However, we must treat this world as hypothetical and fresh since we don't know which one it is. The rules for \diamond are perfect mirror images of the rules for \square . For implication, we use the classical truth tables to provide rules of inference. If we know that $A \supset B$ is false, then we know A is true but B is false. If we know that $A \supset B$ is true, then we know that *either* A is false or B is true.

A key feature of the sequent calculus is the subformula property: every step (when read bottom-up) proceeds by decomposing exactly one connective. This means that proofs in the sequent calculus work by only examining the structure of the proposition at hand; this gives us a nice orthogonality condition for the connectives in our logic.

The translation from natural deduction to the sequent calculus requires a lemma (which should not be a rule of inference because it violates the subformula property). In an intuitionistic calculus this would be *cut*; for the symmetric classical calculus it turns out to be the familiar classical notion of *excluded middle*.

Theorem 3 (Excluded Middle)

If $\Gamma, A @ \omega \# \Delta$
 and $\Gamma \# A \star \omega, \Delta$
 then $\Gamma \# \Delta$.

⁴Our rules are also consistent with the more typical multiple-conclusion reading, "if all of Γ are true, then one of Δ is true."

$$\begin{array}{c}
\frac{\Gamma, A @ \omega \# A \star \omega, \Delta \text{ contra}}{\Gamma, A @ \omega \# \Delta} \perp T \\
\frac{\Gamma, A \supset B @ \omega, B @ \omega \# \Delta}{\Gamma, A \supset B @ \omega \# A \star \omega, \Delta} \supset T \\
\frac{\Gamma, \Box A @ \omega, A @ \omega' \# \Delta}{\Gamma, \Box A @ \omega \# \Delta} \Box T \\
\frac{\Gamma, \Diamond A @ \omega, A @ \omega' \# \Delta}{\Gamma, \Diamond A @ \omega \# \Delta} \Diamond T \\
\frac{\Gamma, A \wedge B @ \omega, A @ \omega, B @ \omega \# \Delta}{\Gamma, A \wedge B @ \omega \# \Delta} \wedge T
\end{array}
\quad
\begin{array}{c}
\frac{\Gamma, \perp @ \omega \# \Delta \perp T}{\Gamma \# A \supset B \star \omega, D} \supset F \\
\frac{\Gamma \# A \star \omega', \Box A \star \omega, \Delta}{\Gamma \# \Box A \star \omega, \Delta} \Box F \\
\frac{\Gamma \# A \star \omega', \Diamond A \star \omega, \Delta}{\Gamma \# \Diamond A \star \omega, \Delta} \Diamond F \\
\frac{\Gamma \# A \star \omega, A \wedge B \star \omega, \Delta}{\Gamma \# A \wedge B \star \omega, \Delta} \wedge F
\end{array}$$

Figure 2: Classical S5 sequent calculus

Proof of Theorem 3 is by lexicographic induction on the proposition A and then on the two derivations. \square

With excluded middle, we can prove the correspondence between natural deduction and the sequent calculus.

Theorem 4 (Equivalence)

(a) If $\Gamma; \Delta \vdash M : A @ \omega$
then $\Gamma \# A \star \omega, \Delta$.

(b) If $\Gamma \# \Delta$
then $\exists M. \forall C, \omega. \Gamma; \Delta \vdash M : C @ \omega$.

It is easy to see why 4(b) is the right statement. Since we think of $\Gamma \# \Delta$ as a proof of contradiction, this corresponds to a derivation that proves any proposition at any world in natural deduction. Theorem 4(a) is more subtle. We show that if A is true under assumptions Γ and Δ , then A being false at the same world is contradictory with those assumptions. Computationally, we can think of this as the “final continuation” to which the result computed in natural deduction is passed. Putting these two theorems together, we have that $\Gamma; \Delta \vdash M : A @ \omega$ gives $\Gamma \# A \star \omega, \Delta$, which then gives $\forall C, \omega'. \Gamma; \Delta, u : A \star \omega \vdash M' : C @ \omega'$. In particular, we choose $C = A$ and $\omega' = \omega$, and then by application of bc we have the original judgment (with perhaps a different proof term $\text{letcc } u \text{ in } M'$). Thus \vdash and $\#$ are really equivalent.

The proof of Theorem 4(a) is by straightforward induction on the derivation, using Theorem 3 where necessary. (The structural rules bc and $\#$ just become uses of contraction and weakening in the sequent calculus.) Proof of 4(b) is more tricky. Uses of T rules are easy; they correspond directly to the elimination rules⁵

⁵Except for implication, which is phrased differently in the

in natural deduction. But since our natural deduction is biased towards manipulating truth rather than falsehood, the F rules are more difficult and make nontrivial use of the falsehood substitution theorem. For instance, in the $\wedge F$ case we have by induction:

$$\begin{array}{l}
\Gamma; \Delta, up : A \wedge B \star \omega, ua : A \star \omega \vdash N_1 : C @ \omega' \quad (\forall C, \omega') \\
\Gamma; \Delta, up : A \wedge B \star \omega, ub : B \star \omega \vdash N_2 : C @ \omega' \quad (\forall C, \omega')
\end{array}$$

By two applications of Theorem 2, we get that the following proof term has any type at any world:

$$\llbracket x. \llbracket y. \text{throw } \langle x, y \rangle \text{ to } up/ub \rrbracket N_2 / ua \rrbracket N_1$$

First, we form a **throw** of the pair $\langle x, y \rangle$ to our pair continuation up . This has free truth hypotheses $x : A$ and $y : B$. Therefore, we can use it to substitute away the ub continuation in N_2 (any throw of M to ub becomes a throw of $\langle x, M \rangle$ to up). Finally, we can use this new term to substitute away ua in N_1 , giving us a term that depends only on the pair continuation up . This pattern of *prepending* work onto continuations through substitution is characteristic of this proof, and reflects our bias towards the truth judgment in natural deduction. As another example, in the case for the $\Diamond F$ rule we have by induction:

$$\Gamma; \Delta, u : A \star \omega', ud : \Diamond A \star \omega \vdash N : C @ \omega'' \quad (\forall C, \omega'')$$

Our proof term in natural deduction is then:

$$\llbracket x. \text{throw}(\text{get}_{\Diamond}[\omega'](\text{here } x)) \text{ to } ud/u \rrbracket N$$

Simply enough, if u is ever thrown to, then we instead take that term’s address (which lives at ω'), move it to ω , and throw it to our $\Diamond A$ continuation ud .

sequent calculus.

Finally, the case for $\Box F$ is interesting because it involves a `letcc`.⁶ By induction we have:

$$\forall \omega'. \Gamma; \Delta, u:A\star\omega', ub:\Box A\star\omega \vdash N : C@{\omega''} \quad (\forall C, \omega'')$$

Then the proof term witnessing the theorem here is:

`throw(box ω' . letcc u in N) to ub`

It is not possible to use falsehood substitution on u in this case. To do so we would need to turn a term of type $A@{\omega'}$ into a $\Box A@{\omega}$ to throw to ub . Although at a meta-level we know that we can choose any ω' , it won't be possible to internalize this in order to create a $\Box A$. Instead we must introduce a new box, and choose ω' to be the new hypothetical world that the $\Box I$ rule introduces. At that point we use `letcc` to create a real $A\star\omega'$ assumption to discharge u . The remaining cases are similar or straightforward. \square

All of the proofs in this section have been formalized in the Twelf system [12] and verified by its metatheorem checker [14].⁷ The most natural LF encoding of falsehood is 3rd-order [1]; we use a 2nd-order encoding (proving the falsehood substitution theorem by hand) because third-order proof checking is not yet in the distribution.

4 Examples

In this section we give proof terms showing the new connection between \Box and \Diamond made possible by network-wide continuations. Because the examples we'll look at involve negation ($\neg A$), we'll need to briefly explain how we treat it.

4.1 Negation

Although we have not given the rules for the negation connective, it is easily added to the system. Here we equivalently take the standard shortcut of treating $\neg A$ as an abbreviation for $A \supset \perp$. We computationally read $\neg A@{\omega}$ as a continuation expecting A , although this should be distinguished from primitive continuations u with type $A\star\omega$: the former is formed by lambda abstraction and eliminated by application, while the latter is formed with `letcc` and eliminated by a `throw` to it. The two are related in that we can reify a continuation assumption $u:A\star\omega$ as a negated formula $\neg A$ by lambda abstracting a throw to it: $\lambda a. \text{throw } a \text{ to } u$. Likewise, we can get a falsehood assumption from a term M of type $\neg A$, namely $M(\text{letcc } u \text{ in } \dots)$.

⁶In fact, this is the only place in the proof where a `letcc` is necessary. This suggests a normal form for natural deduction terms where `letcc` appears only once at the outermost scope and immediately inside each `box`.

⁷They can be found at <http://www.cs.cmu.edu/~concert/>.

Finally, note that we have derived sequent calculus rules $\neg T$ and $\neg F$. Each just flips the proposition under negation to the other side of the sequent, as expected. (The reader can verify that these are indeed the rules derived from $\supset T$ and $\supset F$ if the antecedent of implication is \perp .)

4.2 Classical Axioms

Our first example comes from the standard practice in classical modal logic of defining \Box in terms of \Diamond :

$$\Box A \equiv \neg \Diamond \neg A$$

From left to right the implication is intuitionistically valid, so we'll look at the proof of the implication right to left. We begin with the sequent calculus proof, to show why this is clearly true classically. We elide any residual assumptions that go unused.

$$\frac{\frac{\frac{\frac{\frac{\frac{}{\neg, A@{\omega'} \# A@{\omega'}, -}}{\neg F}}{\neg \# \neg A@{\omega'}, A@{\omega'}, -}}{\Diamond F}}{\neg \# \Diamond \neg A@{\omega}, A@{\omega'}, -}}{\neg \# \Diamond \neg A@{\omega}, \Box A@{\omega}} \Box F^{\omega'}}{\neg \Diamond \neg A@{\omega} \# \Box A@{\omega}} \neg T}{\# \neg \Diamond \neg A \supset \Box A@{\omega}} \supset F$$

Critically, we are using $\Box F$ to get the hypothetical world at which $\Box A$ is false. From there, we can learn $\not\Box A$ at the same world, which leads to a contradiction. In natural deduction, the proof tells an interesting story:

```

 $\lambda dc.$                                 ( $dc : \neg \Diamond \neg A@{\omega}$ )
  box  $\omega'$ .                            (need to return  $A$ )
    letcc  $u$  in rpc[ $\omega$ ]                (applying  $dc$  will yield  $\perp$ )
      dc(get $_{\Diamond}$ [ $\omega$ ](here( $\lambda a. \text{throw } a \text{ to } u$ )))

```

In each example, we'll assume that the whole term lives at the world ω . Operationally, the reading of $\neg \Diamond \neg A \supset \Box A$ is that given a continuation dc (expecting the address of an A continuation), we will return a boxed A that is well-formed anywhere. The proof term given accomplishes this by creating a box that, when opened, grabs the current continuation u , which has type $A\star\omega'$. With the continuation in hand, we travel back to ω (where dc lives), and apply dc to the address of a function that throws to u . In short, at the moment the box is opened we have a *lack of* an A , which we can grab with `letcc` and then take the address of with `here`. This is enough to send to the continuation that we're provided.

Dually we can define \Diamond in terms of \Box . Again, one direction is intuitionistically valid. The other,

$$\neg \Box \neg A \supset \Diamond A$$

is asked to conjure up an address of an arbitrary A given a continuation. It is implemented by the following proof term:

$$\begin{array}{ll} \lambda bc. & (bc: \neg \Box \neg A @ \omega) \\ \text{letcc } u \text{ in} & (u: \Diamond A \star \omega) \\ bc(\text{box } \omega'. \lambda a. & (a: A @ \omega')) \\ \text{throw}(\text{get}_{\Diamond}[\omega'](\text{here } a)) \text{ to } u) & \end{array}$$

Here, we immediately do a `letcc`, grabbing the $\Diamond A$ continuation at ω . We then form a `box` to pass to the continuation bc . It contains a function of type $A \supset \perp$, which takes the address of its argument and throws it to the saved continuation u . Thus the location of A that we return is any world that invokes the $\neg A$ that we've boxed up.

We've left disjunction out of our calculus. Theoretically it poses no problem, although operationally it requires us to perform some tricks to avoid a strange "remote case analysis." In Section 7 however, we see that we can encode it using the de Morgan translation into \neg and \wedge . Regardless of how we implement it, disjunction is a source of a wealth of interesting programs.

Without being as formal, let's take a look at a program implementing the classical axiom $\Box A \vee \Diamond \neg A$. We'll assume constructors `inl` and `inr` for injecting into the disjunction. Those familiar with the implementation of the axiom $A \vee \neg A$ might guess that this returns the address of an A continuation, as in that case. Actually, this doesn't work! We can't build a $\Box A$ by accumulating evidence for A at different worlds. Instead, we return again a box that does something when opened.

$$\begin{array}{ll} \text{letcc } uo \text{ in} & (uo: \Box A \vee \Diamond \neg A \star \omega) \\ \text{inl}(\text{box } \omega'. & \\ \text{letcc } u \text{ in} & \\ \text{throw}(\text{inr}(\text{get}_{\Diamond}[\omega'] \text{here}(\lambda a. \text{throw } a \text{ to } u))) & \\ \text{to } uo) & \end{array}$$

First we save the current continuation as uo , since we will need to "change our minds" about which clause we return! Initially, we return a box whose body also grabs the continuation (of type $A \star \omega'$) as ub . Suppose the box is opened at the world ω' . We then throw to the remote continuation uo a program that comes back to ω' , forms a term of $\neg A$, publishes it, and moves it to the first world.

To summarize, when asked for $\Box A \vee \Diamond \neg A$, the program

1. initially says $\Box A$.
2. if the box is opened, the program uses the lack of A to produce a $\Diamond \neg A$, time travels back to when it was asked about the disjunction, and returns this different answer.

3. if the $\neg A$ continuation is ever invoked, the program goes back and uses the A to fulfill the outstanding lack of A at the world where the box was opened.

In the style of sci-fi storytelling popular when describing such things, we conclude our examples with the following story (with apologies to Wadler [16]):

A magician who purports to be from the future is making bold claims. Asking for a volunteer, he offers the following prize to anyone who comes on stage:

"I'm going to hand you a box that has *you* inside it!"

"Either that, or I'll give you the address of a place with a magical time travelling portal."

Being questionably brave, you volunteer and walk onto the stage. The magician hands you your prize—a large cardboard box. Noting your skepticism, he adds, "You can open it anywhere, and you'll be inside."

You decide to take the box home. It's much too light to have anything in it, let alone yourself! You open the box and look inside, wondering what sort of gag he has planned. But suddenly you find that the box has disappeared, and you're standing on stage waiting for him to tell you what you've won, again.

"The address of the time-travelling portal is," he begins, rattling off your home address. You are startled that he could have known your address, but when you later arrive home, you see an open cardboard box waiting. Is this supposed to be the portal? Knowing it to be harmless, but insisting upon proving the magician to be a fraud, you step into it.

A hot flash of embarrassment passes over you as you realize that you are now standing in a cardboard box, in your house, as promised.

5 Type System and Operational Semantics

Our deductive proof theory begets a natural programming language whose syntax is the proof terms from Figure 1. In order to give this language an operational interpretation, we need to introduce a number of syntactic constructs, which are given in Figure 3.

As in Lambda 5, the behavior of a program is specified in terms of an abstract network that steps from state to state. The network is built out of a fixed number of worlds, whose names we write as bold \mathbf{w} . Because we can now mention specific worlds in addition to hypothetical worlds ω , we introduce world expressions, which are written with a Roman w . A network state \mathbb{N} has two parts. First is a world configuration \mathbb{W} which identifies

types	$A, B ::= A \supset B \mid \Box A \mid \Diamond A \mid A \wedge B \mid \perp$
networks	$\mathbb{N} ::= \mathbb{W}; R$
configs	$\mathbb{W} ::= \{\mathbf{w}_1 : \langle \mathcal{C}_1, \mathcal{B}_1 \rangle, \dots\}$
cursors	$R ::= \mathbf{w} : [k \prec v] \mid \mathbf{w} : [k \succ M]$
tables	$b ::= \bullet \mid b, \ell = v$
cont tables	$\mathcal{C} ::= \bullet \mid \mathcal{C}, \mathbf{k} = k$
config types	$\Sigma ::= \{\mathbf{w}_1 : \langle \mathfrak{E}_1, \mathfrak{B}_1 \rangle, \dots\}$
table types	$\mathfrak{B} ::= \bullet \mid \mathfrak{B}, \ell : A$
ctable types	$\mathfrak{E} ::= \bullet \mid \mathfrak{E}, \mathbf{k} : A$
world exps	$w ::= \omega \mid \mathbf{w}$
world vars	ω world names \mathbf{w}
labels	ℓ value vars x, y
cont labs	\mathbf{k} cont vars u
values	$v ::= \lambda x.M \mid \mathbf{box} \omega.M \mid \mathbf{w}.\ell \mid \langle v, v' \rangle$
conts	$k ::= \mathbf{return} Z \mid \mathbf{finish}$ $\mid \mathbf{abort} \mid k \triangleleft f$
cont exps	$Z ::= \mathbf{w}.\mathbf{k} \mid u$
frames	$f ::= \circ N \mid v \circ \mid \mathbf{here} \circ \mid \mathbf{unbox} \circ$ $\mid \mathbf{letd} \omega.x = \circ \mathbf{in} N \mid \pi_n \circ$ $\mid \langle \circ, N \rangle \mid \langle v, \circ \rangle$
exps	$M, N ::= v \mid MN \mid x \mid \ell \mid \mathbf{get}_{\Box}[\mathbf{w}]M$ $\mid \mathbf{here} M \mid \mathbf{get}_{\Diamond}[\mathbf{w}]M$ $\mid \mathbf{unbox} M \mid \mathbf{letd} \omega.x = M \mathbf{in} N$ $\mid \mathbf{rpc}[\mathbf{w}]M \mid \mathbf{letcc} u \mathbf{in} M$ $\mid \mathbf{throw} M \mathbf{to} Z \mid \langle M, N \rangle \mid \pi_n M$

Figure 3: Syntax of type system

two tables with each world \mathbf{w}_i present. The first table \mathcal{C}_i stores network-wide continuations by mapping continuation labels \mathbf{k} to literal continuations k . The second table b_i maps value labels ℓ to values in order to store values whose address we have published. These tables have types \mathfrak{E} and \mathfrak{B} respectively (which map labels \mathbf{k} and ℓ to types), and so we can likewise construct the type of an entire configuration, written Σ .

Aside from the current world configuration, a network state also contains a *cursor* denoting the current focus of computation. The cursor either takes the form $\mathbf{w} : [k \prec v]$ (returning the value v to the continuation k) or $\mathbf{w} : [k \succ M]$ (evaluating the expression M in continuation k). In either case it selects a world \mathbf{w} where the computation is taking place.

Continuations themselves are stacks of frames (expressions with a “hole,” written \circ) with a bottommost **return**, **finish** or **abort**. The **finish** continuation represents the end of computation, so a network state whose cursor is returning a value to **finish** is called *terminal*. The **abort** continuation will be unreachable, and **return** will send the received value to a remote continuation.

Most of the expressions and values are straightforward.

As in Lambda 5, the canonical value for \Box abstracts over the hypothetical world and leaves its argument unevaluated ($\mathbf{box} \omega'.M$). The canonical form for \Diamond is a pair of a world name and a label $\mathbf{w}.\ell$, which addresses a table entry at that world. Such an address is well-formed anywhere (assuming that \mathbf{w} 's table has a label ℓ containing a value of type A) and has type $\Diamond A @ \mathbf{w}'$. On the other hand we have another sort of label, written just ℓ , which is disembodied from its world. These labels arise from the **letd** construct, which deconstructs an address $\mathbf{w}.\ell$ into its components \mathbf{w} and ℓ (see the $\Diamond E$ rule from Figure 1). Disembodied labels only make sense at a single world—here ℓ would have type $A @ \mathbf{w}$.

Although the external language only allows a **throw** to a continuation variable, intermediate states of evaluation require that these be replaced with the continuation expression $\mathbf{w}.\mathbf{k}$, which pairs a continuation label with the world at which it lives. These continuation expressions are filled in by **letcc**.

Judgment	Reading
$\Sigma; \Gamma; \Delta \vdash M : A @ \mathbf{w}$	The expression M has type A at world \mathbf{w}
$\Sigma \vdash k : A \star \mathbf{w}$	The continuation k expects a value of type A at world \mathbf{w}
$\Sigma; \Delta \vdash Z : A \star \mathbf{w}$	The continuation expression Z is well-formed with type A at \mathbf{w}
$\Sigma \vdash b @ \mathbf{w}$	The value table b is well-formed at the world named \mathbf{w}
$\Sigma \vdash \mathcal{C} \star \mathbf{w}$	The continuation table \mathcal{C} is well-formed at the world named \mathbf{w}
$\Sigma \vdash R$	The cursor is well-formed
$\Sigma \vdash \mathbb{N}$	The network is well-formed

Figure 4: Index of Judgments. In each judgment Σ is a configuration typing, Γ is a context of truth hypotheses, and Δ is a context of falsehood hypotheses

The type system is given in Figure 5 (we omit for space the rules that are the same as in Figure 1 except for the configuration typing Σ). The index of judgments in Figure 4 may be a useful reference in understanding them.

The rules *addr* and *lab* are used to type run-time artifacts of address publishing. In either case, we look up the type in the appropriate table typing \mathfrak{B} . As mentioned, *throw* allows a continuation expression Z , which

$$\begin{array}{c}
\frac{\Sigma(\mathbf{w}) = \langle \mathfrak{E}, \beta \rangle \quad \beta(\ell) = A}{\Sigma; \Gamma; \Delta \vdash \mathbf{w}.\ell : \diamond A @ \mathbf{w}'} \text{ addr} \\
\frac{\Sigma; \Gamma; \Delta \vdash M : A @ \mathbf{w} \quad \Sigma; \Delta \vdash Z : A \star \mathbf{w}}{\Sigma; \Gamma; \Delta \vdash \text{throw } M \text{ to } Z : C @ \mathbf{w}'} \text{ throw} \\
\frac{\Sigma(\mathbf{w}) = \langle \mathfrak{E}, \beta \rangle \quad \mathfrak{E}(\mathbf{k}) = A}{\Sigma; \Delta \vdash \mathbf{w}.\mathbf{k} : A \star \mathbf{w}} \text{ addr}^* \\
\frac{}{\Sigma \vdash \text{abort} : \perp \star \mathbf{w}} \text{ kabort} \\
\frac{A = \square A' \text{ or } \diamond A' \quad \Sigma; \cdot \vdash Z : A \star \mathbf{w}'}{\Sigma \vdash \text{return } Z : A \star \mathbf{w}} \text{ kret} \\
\frac{\Sigma \vdash k : B \star \mathbf{w} \quad \Sigma; \cdot \vdash N : A @ \mathbf{w}}{\Sigma \vdash k \triangleleft \circ N : A \supset B \star \mathbf{w}} \text{ kapp}_1 \\
\frac{\Sigma \vdash k : C \star \mathbf{w} \quad \Sigma; x : A @ \omega; \cdot \vdash N : C @ \mathbf{w}}{\Sigma \vdash k \triangleleft \text{letd } \omega.x = \circ \text{ in } N : \diamond A \star \mathbf{w}} \text{ kletd} \\
\frac{\Sigma \vdash k : A \wedge B \star \mathbf{w} \quad \Sigma; \cdot \vdash N : B @ \mathbf{w}}{\Sigma \vdash k \triangleleft \langle \circ, N \rangle : A \star \mathbf{w}} \text{ kpair}_1 \\
\frac{\beta = (\ell_1 : A_1, \dots) \quad \Sigma; \cdot \vdash v_1 : A_1 @ \mathbf{w} \quad \dots}{\underbrace{\{\dots, \mathbf{w} : \langle \mathfrak{E}, \beta \rangle, \dots\}}_{\Sigma} \vdash \underbrace{\ell_1 = v_1, \dots}_{b} @ \mathbf{w}} \text{ b} \\
\frac{\mathbf{w} \in \text{dom}(\Sigma) \quad \Sigma; \cdot \vdash v : A @ \mathbf{w} \quad \Sigma \vdash k : A \star \mathbf{w}}{\Sigma \vdash \mathbf{w} : [k \prec v]} \text{ ret} \\
\frac{\Sigma \vdash R \quad \Sigma \vdash \mathfrak{e}_i @ \mathbf{w}_i \dots \quad \Sigma \vdash b @ \mathbf{w}_i \dots}{\Sigma \vdash \{\mathbf{w}_1 : \langle \mathfrak{e}_1, b_1 \rangle, \dots, \mathbf{w}_m : \langle \mathfrak{e}_m, b_m \rangle\}; R} \text{ net} \\
\frac{\Sigma(\mathbf{w}) = \langle \mathfrak{E}, \beta \rangle \quad \beta(\ell) = A}{\Sigma; \Gamma; \Delta \vdash \ell : A @ \mathbf{w}} \text{ lab} \\
\frac{\Sigma; \Gamma; \Delta, u : A \star \mathbf{w} \vdash M : A @ \mathbf{w}}{\Sigma; \Gamma; \Delta \vdash \text{letcc } u \text{ in } M : A @ \mathbf{w}} \text{ letcc} \\
\frac{}{\Sigma; \Delta, u : A \star \mathbf{w} \vdash u : A \star \mathbf{w}} \text{ hyp}^* \\
\frac{}{\Sigma \vdash \text{finish} : A \star \mathbf{w}} \text{ kfinish} \\
\frac{\Sigma \vdash k : \diamond A \star \mathbf{w}}{\Sigma \vdash k \triangleleft \text{here } \circ : A \star \mathbf{w}} \text{ khere} \\
\frac{\Sigma \vdash k : B \star \mathbf{w} \quad \Sigma; \cdot \vdash v : A \supset B @ \mathbf{w}}{\Sigma \vdash k \triangleleft v \circ : A \star \mathbf{w}} \text{ kapp}_2 \\
\frac{\Sigma \vdash k : A \star \mathbf{w}}{\Sigma \vdash k \triangleleft \text{unbox } \circ : \square A \star \mathbf{w}} \text{ kunbox} \\
\frac{\Sigma \vdash k : A \wedge B \star \mathbf{w} \quad \Sigma; \cdot \vdash v : A @ \mathbf{w}}{\Sigma \vdash k \triangleleft \langle v, \circ \rangle : B \star \mathbf{w}} \text{ kpair}_2 \\
\frac{\mathfrak{E} = (\mathbf{k}_1 : A_1, \dots) \quad \Sigma \vdash k_1 : A_1 \star \mathbf{w} \quad \dots}{\underbrace{\{\dots, \mathbf{w} : \langle \mathfrak{E}, \beta \rangle, \dots\}}_{\Sigma} \vdash \underbrace{\mathbf{k}_1 = k_1, \dots}_{\mathfrak{e}} \star \mathbf{w}} \mathfrak{e} \\
\frac{\mathbf{w} \in \text{dom}(\Sigma) \quad \Sigma; \cdot \vdash M : A @ \mathbf{w} \quad \Sigma \vdash k : A \star \mathbf{w}}{\Sigma \vdash \mathbf{w} : [k \succ v]} \text{ eval}
\end{array}$$

Figure 5: Type System

must take the form of a variable (typed with hyp^* , as in the logic) or address into a continuation table.

Typing of literal continuations k is fairly unsurprising. Note that the judgment $\Sigma \vdash k : A \star \mathbf{w}$ means that the continuation k *expects* a value of type A at \mathbf{w} . The **return** continuation arises only from a get_{\diamond} or get_{\square} , and so it allows only values of type $\diamond A$ or $\square A$. We reuse the network continuation mechanism here to refer to the outstanding get_{\diamond} or get_{\square} on the remote machine.

For an entire network to be well-formed (rule net), all of the tables must have the type indicated by the configuration type Σ , which means that they must have exactly the same labels, and the values or continuations must be well-typed at the specified types (rules b and \mathfrak{e}). Finally, the cursor must be well-formed: it must select a world that exists in the network, and there must exist a type A such that its continuation and value or expression both have type A and are closed.

Having set up the syntax and type system, we can

now give the operational semantics and type safety theorem. After the following section we remark on how the semantics can be made concurrent, and some thoughts on applications of distributed continuations.

5.1 Operational Semantics

The operational semantics of our language is given in Figure 6, as a binary relation \mapsto between network states. The semantics evaluates programs sequentially, though we give a concurrent semantics in Section 6.

As should be obvious, the semantics is continuation-based. At any step, the cursor is selecting a world and continuation, with a value to return to it or an expression to evaluate. The rules generally fall into a few categories, as exemplified by the (standard) rules for \supset : There are **push** rules, in which we begin evaluating a subexpression of some M , pushing the context into the continuation, **swap** rules, where we have finished evaluating one sub-expression and move onto the next, and

$\supset_e\text{-p}$	$\mathbb{W}; \mathbf{w} : [k \succ MN]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \triangleleft \circ N \succ M]$	
$\supset_e\text{-s}$	$\mathbb{W}; \mathbf{w} : [k \triangleleft \circ N \prec v]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \triangleleft v \circ \succ N]$	
$\supset_e\text{-r}$	$\mathbb{W}; \mathbf{w} : [k \triangleleft (\lambda x.M) \circ \prec v]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \succ [v/x]M]$	
$\wedge_i\text{-p}$	$\mathbb{W}; \mathbf{w} : [k \succ \langle M, N \rangle]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \triangleleft \langle \circ, N \rangle \succ M]$	
$\wedge_i\text{-s}$	$\mathbb{W}; \mathbf{w} : [k \triangleleft \langle \circ, N \rangle \prec v]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \triangleleft \langle v, \circ \rangle \succ N]$	
$\wedge_i\text{-r}$	$\mathbb{W}; \mathbf{w} : [k \triangleleft \langle v_1, \circ \rangle \prec v_2]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \prec \langle v_1, v_2 \rangle]$	
$\wedge_{e_n}\text{-p}$	$\mathbb{W}; \mathbf{w} : [k \succ \pi_n M]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \triangleleft \pi_n \circ \succ M]$	
$\wedge_{e_n}\text{-r}$	$\mathbb{W}; \mathbf{w} : [k \triangleleft \pi_n \prec \langle v_1, v_2 \rangle]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \prec v_n]$	
$\square_i\text{-v}$	$\mathbb{W}; \mathbf{w} : [k \succ \mathbf{box} \omega.M]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \prec \mathbf{box} \omega.M]$	
$\diamond_i\text{-v}$	$\mathbb{W}; \mathbf{w} : [k \succ \mathbf{w}'.\ell]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \prec \mathbf{w}'.\ell]$	
$\supset_i\text{-v}$	$\mathbb{W}; \mathbf{w} : [k \succ \lambda x.M]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \prec \lambda x.M]$	
$\diamond_i\text{-p}$	$\mathbb{W}; \mathbf{w} : [k \succ \mathbf{here} M]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \triangleleft \mathbf{here} \circ \succ M]$	
$\diamond_i\text{-r}$	$\{\mathbf{w} : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w} : [k \triangleleft \mathbf{here} \circ \prec v]$	\mapsto	$\{\mathbf{w} : \langle \mathcal{E}, (b, \ell = v) \rangle, \dots\}; \mathbf{w} : [k \prec \mathbf{w}.\ell]$	(ℓ fresh)
$\ell\text{-r}$	$\{\mathbf{w} : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w} : [k \succ \ell]$	\mapsto	$\{\mathbf{w} : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w} : [k \prec v]$	($b(\ell) = v$)
$\diamond_e\text{-p}$	$\mathbb{W}; \mathbf{w} : [k \succ \mathbf{letd} \omega.x = M \mathbf{in} N]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \triangleleft \mathbf{letd} \omega.x = \circ \mathbf{in} N \succ M]$	
$\diamond_e\text{-r}$	$\mathbb{W}; \mathbf{w} : [k \triangleleft \mathbf{letd} \omega.x = \circ \mathbf{in} N \prec \mathbf{w}'.\ell]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \succ [\ell/x][\mathbf{w}'/\omega]N]$	
$\square_e\text{-p}$	$\mathbb{W}; \mathbf{w} : [k \succ \mathbf{unbox} M]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \triangleleft \mathbf{unbox} \circ \succ M]$	
$\square_e\text{-r}$	$\mathbb{W}; \mathbf{w} : [k \triangleleft \mathbf{unbox} \circ \prec \mathbf{box} \omega.M]$	\mapsto	$\mathbb{W}; \mathbf{w} : [k \succ [\mathbf{w}/\omega]M]$	
\mathbf{letcc}	$\{\mathbf{w} : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w} : [k \succ \mathbf{letcc} u \mathbf{in} M]$	\mapsto	$\{\mathbf{w} : \langle \langle \mathcal{E}, \mathbf{k} = k \rangle, b \rangle, \dots\}; \mathbf{w} : [k \succ [\mathbf{w}.\mathbf{k}/u]M]$	(\mathbf{k} fresh)
\mathbf{throw}	$\{\mathbf{w}' : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w} : [k \succ \mathbf{throw} M \mathbf{to} \mathbf{w}'.\mathbf{k}]$	\mapsto	$\{\mathbf{w}' : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w}' : [k' \succ M]$	($\mathcal{E}(\mathbf{k}) = k'$)
\mathbf{rpc}	$\mathbb{W}; \mathbf{w} : [k \succ \mathbf{rpc}[\mathbf{w}']M]$	\mapsto	$\mathbb{W}; \mathbf{w}' : [\mathbf{abort} \succ M]$	($\mathbf{w} \in \text{dom}(\mathbb{W})$)
\square_m	$\{\mathbf{w} : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w} : [k \succ \mathbf{get}_{\diamond}[\mathbf{w}']M]$	\mapsto	$\{\mathbf{w} : \langle \langle \mathcal{E}, \mathbf{k} = k \rangle, b \rangle, \dots\}; \mathbf{w}' : [\mathbf{return} \mathbf{w}.\mathbf{k} \succ M]$	(\mathbf{k} fresh)
\diamond_m	$\{\mathbf{w} : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w} : [k \succ \mathbf{get}_{\square}[\mathbf{w}']M]$	\mapsto	$\{\mathbf{w} : \langle \langle \mathcal{E}, \mathbf{k} = k \rangle, b \rangle, \dots\}; \mathbf{w}' : [\mathbf{return} \mathbf{w}.\mathbf{k} \succ M]$	(\mathbf{k} fresh)
\mathbf{ret}	$\{\mathbf{w} : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w}' : [\mathbf{return} \mathbf{w}.\mathbf{k} \prec v]$	\mapsto	$\{\mathbf{w} : \langle \mathcal{E}, b \rangle, \dots\}; \mathbf{w} : [k \prec v]$	($\mathcal{E}(\mathbf{k}) = k$)

Figure 6: Operational Semantics

reduction rules, where we have a value and actually do something with it. Every well-typed machine state will be closed with respect to truth, falsehood, and world hypotheses, so we don't have rules for variables and can specialize some rules.

The first interesting rule is $\diamond_i\text{-r}$. It publishes the value v and returns its address by generating a new label, mapping that label to v within its value table, and returning the pair $\mathbf{w}.\ell$, where \mathbf{w} is the current world. Whenever we try to evaluate a label (rule $\ell\text{-r}$), we look it up in the current world's value table in order to fetch the value. A key consequence of type safety (Theorems 5, 6) is that labels are only evaluated in the correct world. To eliminate an address (rule $\diamond_e\text{-r}$) we substitute the constituent world and label through the body of the \mathbf{letd} . Note that this step is slightly non-standard, because we substitute the *expression* ℓ for a variable rather than

some value. But because the variable is in general at a different world, we are not in a position to get its value yet. We instead wait until the expression ℓ is sent to its home world (perhaps as part of some larger expression) to be looked up. The rules for \square are much simpler: $\mathbf{box} \omega.M$ is already a value (rule $\square_i\text{-v}$), and to \mathbf{unbox} we simply substitute the current world for the hypothetical one (rule $\square_e\text{-r}$).

When encountering a \mathbf{letcc} , we grab the current continuation k . Because the continuation may be referred to from elsewhere in the network, we publish it in a table and form a global address for it (of the form $\mathbf{w}.\mathbf{k}$), just as we did for \diamond addresses. This value is substituted for the falsehood variable u using standard substitution—*not* the special falsehood substitution we used in Section 2. The latter was a proof-theoretic notion used to eliminate uses of the hypothesis; here we

want the use of the hypothesis (**throw**) to have run-time significance. A point of comparison is the above paragraph, where we substituted the expression ℓ for a variable because we wanted to delay the operation until the time the variable is “looked up.”

Throwing to a continuation (rule *throw*) is handled straightforwardly. The continuation expression will be closed, and therefore of the form $\mathbf{w}'.\mathbf{k}$. We look up the label \mathbf{k} in \mathbf{w}' —or rather, *cause* \mathbf{w}' to look it up—and pass the expression M to it. Note that we do not evaluate the argument before throwing it to the remote continuation. In general we *can not* evaluate it, because it is only well-typed at the remote world, which may be different from the world we’re in.

Finally, we have the rules that move between worlds. The rule for **rpc** is easiest; since the target world expression must be closed it will be a world constant in the domain of \mathbb{W} . We simply move the cursor to that world (destroying the current continuation, which can never be reached), and begin evaluating the expression M under the unreachable continuation **abort**. The rules for **get $_{\diamond}$** and **get $_{\square}$** work similarly, but they need to save the current continuation since they will be returned to! These steps push a **return** frame, which reduces like **throw**. In contrast, however, the argument (of type $\square A$ or $\diamond A$) will be eagerly evaluated, because such values are portable. (After all, the whole point is to create the box at one world and then move it to another.)

In order for our language to make sense it must be type safe; any well-typed program must have a well-defined meaning as a sequence of steps in the abstract network. Type safety is stated as usual in terms of progress and preservation:

Theorem 5 (Progress)

If $\Sigma \vdash \mathbb{N}$
then either \mathbb{N} is terminal or $\exists \mathbb{N}'. \mathbb{N} \mapsto \mathbb{N}'$.

Theorem 6 (Preservation)

If $\Sigma \vdash \mathbb{N}$ and $\mathbb{N} \mapsto \mathbb{N}'$
then $\exists \Sigma'. \Sigma' \supseteq \Sigma$ and $\Sigma' \vdash \mathbb{N}'$.

Progress says that any well-formed network state can take another step, or is done. (Recall a *terminal* network is one where the cursor is returning a value to a **finish** continuation.) Preservation says that any well-typed network state that takes a step results in another well-typed state (perhaps in an extended configuration typing Σ'^8). By iterating alternate applications of these theorems we see that any well-typed program is able to step repeatedly and remain well-formed, or else eventually comes to rest in a terminal state.

⁸ $\Sigma' \supseteq \Sigma$ iff Σ' and Σ each describe the same set of worlds, and for each world, if $\mathfrak{C}'(\mathbf{k}) = A$ then $\mathfrak{C}(\mathbf{k}) = A$, and likewise for \mathfrak{B}' and \mathfrak{B} .

6 Concurrency and Communication

Many distributed computing problems benefit from concurrency, with one or more processes running on each node in the network. This section gives some brief thoughts on concurrency in our classical calculus.

First-class continuations are often used in the implementation of coroutines. With primitives for recursion and state we could also implement coroutines in C5, however, such an implementation is silly because it would require the implementation of a global scheduler, and would anyway defeat the purpose of concurrency on multiple nodes—only one coroutine would be running at any given time!

Fortunately, our operational semantics admits ad hoc concurrency easily. If we simply replace the cursor R in our network state “ $\mathbb{W}; R$ ” with a multiset of cursors \mathfrak{R} , then we can permit a step on any one of these cursors essentially according to the old rules. Formally,

$$\begin{aligned} \mathbb{W}; \mathfrak{R} &\mapsto^c \mathbb{W}'; \mathfrak{R}' \\ \text{iff } \mathfrak{R} &= R \uplus \mathfrak{R}_{\text{rest}} \\ \text{and } \mathbb{W}; R &\mapsto \mathbb{W}'; R' \\ \text{and } \mathfrak{R}' &= R' \uplus \mathfrak{R}_{\text{rest}} \end{aligned}$$

We can then add primitives as desired to spawn new cursors. A very simple one evaluates M and N in parallel and returns each one to the same continuation.

$$\frac{\Gamma; \Delta \vdash M : A @ \mathbf{w} \quad \Gamma; \Delta \vdash N : A @ \mathbf{w}}{\Gamma; \Delta \vdash M | N : A @ \mathbf{w}} \text{ par}$$

$$\mathbb{W}; \mathfrak{R} \uplus \mathbf{w} : [k \succ M | N] \mapsto^c \mathbb{W}; \mathfrak{R} \uplus \mathbf{w} : [k \succ M] \uplus \mathbf{w} : [k \succ N]$$

A suitable extension of type safety holds for \mapsto^c .

With concurrency in place we can implement CML-style channels [13] with the help of continuations (and a few other features for developing mutable recursive structures). The type of a channel carrying values of type A is:

$$A \text{ chan} \doteq \diamond(A \text{ queue} \wedge (\neg A) \text{ queue})$$

Here a channel is represented as the address of a pair of queues. The first holds the values that have been sent on the channel and not yet received, the second holds the continuations of outstanding **recvs**. To implement **recv** (assuming no values are waiting in the first queue), we grab the current continuation, enqueue it, and abort.

This is a standard technique; the point here is to emphasize the utility of continuations as primitives for implementing useful distributed computing features.

7 Disjunction

To add disjunction to C5, we need to use the following elimination form in order to preserve the correspon-

dence with classical S5:

$$\frac{\begin{array}{l} \Gamma; \Delta \vdash M : A \vee B @ \omega' \\ \Gamma, x:A @ \omega'; \Delta \vdash N_1 : C @ \omega \\ \Gamma, x:B @ \omega'; \Delta \vdash N_2 : C @ \omega \end{array}}{\text{case } M \text{ of} \\ \Gamma; \Delta \vdash \begin{array}{l} \text{inl } x \Rightarrow N_1 : C @ \omega \\ | \text{inr } x \Rightarrow N_2 \end{array}} \vee E$$

This rule is completely unsurprising except that the case object M is at a *different world*, ω' . In our logic we've tried hard to avoid this sort of *action-at-a-distance*, instead preferring to have our introduction and elimination rules compute locally. However, a motion rule for disjunction is out of the question, because it is unsound: it is not the case that if $\Gamma; \Delta \vdash A \vee B @ \omega$ then necessarily $\Gamma; \Delta \vdash A \vee B @ \omega'$. In our previous paper we speculated that this rule could be implemented nonetheless by sending back merely a *bit* telling the case-analyzing world which branch it should enter, but this requires some specious operational machinery. The same is true in the classical case, which is why we have avoided treating disjunction so far.

The problem with a local rule (where all four worlds are the same) comes when translating the sequent calculus rule

$$\frac{\begin{array}{l} \Gamma, A \vee B @ \omega, A @ \omega \# \Delta \\ \Gamma, A \vee B @ \omega, B @ \omega \# \Delta \end{array}}{\Gamma, A \vee B @ \omega \# \Delta} \vee T$$

into natural deduction for Theorem 4(b). We have by induction

$$\begin{array}{l} \Gamma, x:A \vee B @ \omega, a:A @ \omega; \Delta \vdash N_1 : C @ \omega' (\forall C, \omega') \\ \Gamma, x:A \vee B @ \omega, b:B @ \omega; \Delta \vdash N_2 : C @ \omega' (\forall C, \omega') \end{array}$$

and a derivation of $\Gamma, x:A \vee B @ \omega; \Delta \vdash x : A \vee B @ \omega$ by the *hyp* rule. However, we cannot apply a local \vee -elimination rule, because it would require its object x to be at the same world as its conclusion. Thus we are able to prove C at ω , but not at *all* worlds.

As it turns out, support for disjunction and remote disjunction elimination is already present in C5, thanks to one of de Morgan's laws. Suppose that we define $A \vee B$ as follows

$$A \vee B \doteq \neg(\neg A \wedge \neg B)$$

That is, $A \vee B$ becomes a continuation that takes *two* continuations: one if the disjunct is A , and one if the disjunct is B . This technique is well-known for CPS conversion, and first-class continuations let us do it without having to CPS-convert the entire program. Encoding the injections is easy:

$$\begin{array}{l} \text{inl } M \doteq \lambda ab.(\pi_1 ab)M \\ \text{inr } M \doteq \lambda ab.(\pi_2 ab)M \end{array}$$

Encoding local case analysis is standard:

$$\text{case } M \text{ of} \quad \text{letcc } u \\ \begin{array}{l} \text{inl } x \Rightarrow N_1 \\ | \text{inr } x \Rightarrow N_2 \end{array} \doteq \text{in } M \langle \lambda x. \text{throw } N_1 \text{ to } u, \\ \lambda x. \text{throw } N_2 \text{ to } u \rangle$$

Finally, because we have locally CPS-converted, we can do the case analysis remotely, and rely on **throw** to get us back:

$$\text{case } M \text{ of} \quad \text{letcc } u \\ \begin{array}{l} \text{inl } x \Rightarrow N_1 \\ | \text{inr } x \Rightarrow N_2 \end{array} \doteq \text{in rpc}[\omega']M \langle \lambda x. \text{throw } N_1 \text{ to } u, \\ \lambda x. \text{throw } N_2 \text{ to } u \rangle$$

This has exactly the same typing conditions as the remote rule above; x is bound to the remote type $A @ \omega'$, even though the expression N_1 is evaluated at ω .

Classical logic is ripe with possibilities for definition. It is interesting to consider their implications. Recall that in Section 4 we proved $\diamond A$ equivalent to $\neg \square \neg A$. This means that, as classicists typically do, we could then just consider $\diamond A$ as a derived form. This would amount to a roundabout way of using the continuation table to publish values rather than the value table. Clearly, we could also take the even stranger route of defining $\square A$ in terms of \diamond , which gives us a mobile code “server” that sends code to our continuation whenever we like.

8 Conclusions

8.1 Related Work

Parigot's $\lambda\mu$ -calculus has inspired many computational proof systems for classical logic, including Wadler's dual calculus [16]. The calculus is sequent-oriented and contains *cut* as a computational primitive, emphasizing the duality of computing with values and covalues (continuations). For programming in C5, we choose a natural deduction system which is deliberately *non-dual*. We bias the logic towards truth, which corresponds to computing mainly with values (as is typical) rather than covalues. Nevertheless, we expect that a dual version of classical S5 could be easily made to work, perhaps starting from the sequent calculus presented in Section 3.

Recently, others have used modal logic to describe distributed tasks or as the basis for programming languages, although we know of no modal systems that feature distributed continuations.

Borghuis and Feijs give an early computational interpretation of modal logic in their intuitionistic Modal Type System for Networks [2]. Their calculus and logic describe programs on a network with stationary services (e.g., printers, file converters) and mobile data (e.g., documents). They use \square , annotated with a location, to

represent services. For example, $\Box^o(A \supset B)$ means a function from A to B at the location o . However, the calculus has no way of internalizing mobility as a proposition, so mobile data is limited to base types. Services are similarly restricted to depth-one arrow types. By using \Box for mobile code and \Diamond for stationary resources, we believe our resulting calculus is both simpler and more general.

Moody [7] gives a system based on the constructive modal logic S4 due to Pfenning and Davies [11]. This language is based on judgments $A \text{ true}$ (here), $A \text{ poss}$ (somewhere), and $A \text{ valid}$ (everywhere) rather than the “explicit worlds” formulation of Lambda 5 and C5. The operational semantics of his system takes the form of a process calculus with nondeterminism, concurrency and synchronization; a significantly different approach from our sequential abstract machine. Interpreted in the Kripke model, S4’s accessibility relation satisfies only reflexivity and transitivity, not symmetry. Moody uses the limited accessibility to express process interdependence rather than—as we do—connections between actual network locations. Programs are therefore somewhat higher-level and express *potential mobility* instead of explicit code motion as in our mobility rules. In particular, due to the lack of symmetry it is not possible to return to world after a remote procedure call from it, except by returning a value.

Jia and Walker [6] give a judgmental account of an S5-like system based on intuitionistic hybrid logic. Hybrid logics internalize worlds inside propositions by including a *proposition* that a value of type A resides at world ω , “ $A \text{ at } \omega$.” This leads to a technically different logic and language although they give a similar interpretation to the modalities. Their rules for \Box and \Diamond are non-local, which means that they rely heavily on global actions. Like Moody, they give their network semantics as a process calculus with passive synchronization across processes as a primitive notion.

8.2 Future Work

Our language now has a full arsenal of connectives and control operators, each connected to logic. Much work remains before C5 can be a practical programming language rather than exploratory calculus. Some are routine—adding extra-logical primitives like recursion and references—and some are difficult—compilation of mobile code fragments, distributed garbage collection, failure recovery, and certification.

Although we believe that C5 accomodates concurrency easily, it would be nice to have a logically-inspired account of it. Some other directions remain open to try. Proof search in linear logic sequent calculus [4] is known to admit an interpretation as concurrent com-

putation [3]. Perhaps linear S5 in sequent style would be able to elegantly express both spatial properties and concurrency in logic?

We have presented a proof theory and corresponding programming language, C5, based on the classical modal logic S5. By exploiting the modalities we are able to give a logical account of mobility and locality, and thus an expressive programming language for distributed computing. From the logic’s classical nature we derive the mechanism of distributed continuations, which creates a new connection between the \Box and \Diamond connectives, and forms a basis for the implementation of distributed computing primitives.

Acknowledgements

Thanks to Jason Reed and Frank Pfenning for their suggestions and expertise with LF encodings.

References

- [1] Andreas Abel. A third-order representation of the $\lambda\mu$ -Calculus. In S.J. Ambler, R.L. Crole, and A. Momigliano, editors, *Electronic Notes in Theoretical Computer Science*, volume 58. Elsevier, 2001.
- [2] Tijn Borghuis and Loe M. G. Feijs. A constructive logic for services and information flow in computer networks. *The Computer Journal*, 43(4):274–289, 2000.
- [3] Jean-Yves Girard. Towards a geometry of interaction. *Contemporary Mathematics*, 92:69–108.
- [4] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–102, January 1987.
- [5] Timothy G. Griffin. The formulae-as-types notion of control. In *Conf. Record 17th Annual ACM Symp. on Principles of Programming Languages, POPL’90, San Francisco, CA, USA, 17–19 Jan 1990*, pages 47–57. ACM Press, New York, 1990.
- [6] Limin Jia and David Walker. Modal proofs as distributed programs. *13th European Symposium on Programming*, pages 219–223, March 2004.
- [7] Jonathan Moody. Modal logic as a basis for distributed computation. Technical Report CMU-CS-03-194, Carnegie Mellon University, Oct 2003.
- [8] Tom Murphy, VII, Karl Cray, Robert Harper, and Frank Pfenning. A symmetric modal lambda calculus for distributed computing. In *Proceedings of the 19th IEEE Symposium on Logic in Computer Science (LICS 2004)*. IEEE Press, July 2004.
- [9] Chetan Murthy. Classical proofs as programs: How, what and why. Technical Report TR91-1215, Cornell University, 1991.
- [10] Michel Parigot. $\lambda\mu$ -Calculus: An algorithmic interpretation of classical natural deduction. In Andrei Voronkov, editor, *Logic Programming and Automated*

Reasoning, International Conference LPAR'92, St. Petersburg, Russia, July 15-20, 1992, Proceedings, volume 624 of *Lecture Notes in Computer Science*. Springer, 1992.

- [11] Frank Pfenning and Rowan Davies. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science*, 11:511–540, 2001. Notes to an invited talk at the *Workshop on Intuitionistic Modal Logics and Applications (IMLA'99)*, Trento, Italy, July 1999.
- [12] Frank Pfenning and Carsten Schürmann. System description: Twelf – a meta-logical framework for deductive systems. In Harald Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction*, pages 202–206, Trento, Italy, July 1999. Springer-Verlag. LNAI 1632.
- [13] John H. Reppy. *Concurrent Programming in ML*. Cambridge University Press, Cambridge, England, 1999.
- [14] Carsten Schürmann and Frank Pfenning. A coverage checking algorithm for LF. In D. Basin and B. Wolff, editors, *Proceedings of the 16th International Conference on Theorem Proving in Higher Order Logics (TPHOLs 2003)*, pages 120–135, Rome, Italy, September 2003. Springer-Verlag LNCS 2758.
- [15] Alex Simpson. *The Proof Theory and Semantics of Intuitionistic Modal Logic*. PhD thesis, University of Edinburgh, 1994.
- [16] Philip Wadler. Call-by-value is dual to call-by-name. In *Proceedings of the 8th International Conference on Functional Programming (ICFP)*. ACM Press, August 2003.