

Modal Types for Mobile Code (Thesis Proposal) *

Tom Murphy VII

December 27, 2005

Abstract

Modal logic is a family of logics with the ability to simultaneously reason about truth from multiple perspectives. Our previous work showed how the modal logic Intuitionistic S5 could form the basis for a simple lambda calculus for spatially distributed programs. I propose here a thesis project to demonstrate the efficacy and elegance of modal type systems for controlling spatially distributed resources in a programming language. The project has strong components of both theory and practice: the design of ML5, a new programming language for distributed computing, and its implementation. I present technical details for the language and implementation where they have been completed, and a plan for the work to be done where they have not.

1 Introduction

This research seeks to demonstrate that modal type systems provide an elegant and practical way to control spatially distributed resources in mobile computer programs. I propose to do this by designing and implementing a programming language with a modal type system. I will then demonstrate the language and implementation's effectiveness by building an example distributed application in the language.

The work is part of a larger research program, called ConCert, which is building a large-scale peer-to-peer "Grid Computing" platform based on certified code technology. For this proposal, I will begin by describing the current design of ConCert, which will serve to motivate the thesis work and argue for my qualifications to undertake the proposed research. I will then demonstrate how modal types are an appropriate tool for reasoning about spatial distribution, by reviewing a foundational calculus we developed [29, 28] based on modal logic. Given the problem and an appropriate tool, I then present the design for a new programming language for distributed computing called ML5.

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

Before going into these items in detail, a high-level summary of each point should serve to provide the necessary context to understand what follows.

1.1 Overview

ConCert. The ConCert project's goal is to build a large-scale peer-to-peer distributed network for trustless Grid computing [11]. Grid computing [17], which is named by analogy with electric power grids, is the practice of coupling diverse computational resources into a large shared computer. In this setting, code moves around the network, which presents a security issue: How does a code consumer know that a piece of code is safe to run? There are several functioning Grid computing systems already deployed [16, 3], and each relies on the establishment of trust relationships between the code producer and code consumer to address this security issue. In the ConCert system, we allow for the *trustless* dissemination of code by using certified code technology [30]. To achieve this, each piece of mobile code comes with a certificate demonstrating the relevant properties of the code. These are intrinsic properties of the code (*e.g.*, its type safety) rather than the extrinsic properties (*e.g.*, who wrote it) sometimes certified by cryptographic signatures.

The implementation of the proposed language, ML5, should also work by producing safety certificates. I will do this by building a type-directed certifying compiler [25, 33]. This means that the implementation must be carefully designed (see Section 6), but doing so has many research and engineering benefits.

Grid/ML. Grid/ML [27] is a high-level language that can be used to produce certified code for use on the ConCert network. It is a functional language based on Standard ML [23] with primitives for Grid computing. In Grid/ML, the programmer writes his whole Grid application as a single program. This program is made of two parts: the client, which runs only on the user's computer and interacts with him, and the mobile code, which runs on arbitrary hosts in the network.

Grid/ML is practical for a certain class of problems, and has been used for a few Grid programming experiments [?]. However, it has two major shortcomings. First, the mobile code is agnostic to where it is running, and so applications must treat the hosts in the network uniformly; it is not possible to make use of special resources only available at certain sites. Since distributed computing is often motivated by the desire to make use of resources other than mere processing power, this rules out an important set of applications.

Second, even with the assumption of uniformity, Grid/ML still has a distinction between client code and mobile code. Certain operations, such as I/O, can only be performed on the client and result in run-time failure if used in mobile code.

My solution to both problems is to enrich the ML type system with a concept of *place*. By doing so, the language will be able to support location-aware

programming, so that network locations need not be treated uniformly. As a consequence, the client will be a place like any other, and we will be able to distinguish its capabilities from the capabilities of other hosts. Because we use a type system, we can make these distinctions statically, excluding errors before the program is ever run.

Modal logic and modal types. Type systems for functional languages like ML have a close connection to logic, called the Curry-Howard isomorphism. Under this view, the propositions of intuitionistic logic are interpreted as the types of a programming language. Proofs of propositions then become programs inhabiting those types. Different logics, viewed through the lens of the Curry-Howard isomorphism, give rise to a variety of elegant and useful type systems for programming languages.

In order to develop a type system with a notion of place, we use a logic with the ability to reason spatially, namely modal logic [22]. The propositional logic upon which ML is based is concerned with the *truth* of propositions from a single universal viewpoint. Modal logic introduces the concept of truth from multiple different perspectives, which are called “worlds.” The logic is then able to reason simultaneously about truth in these worlds. Under the Curry-Howard view, these worlds become hosts in the network, and so our type system is correspondingly endowed with a notion of place.

Distributed programming with modal types. Because a proof in modal logic contains reasoning from multiple different worlds, programs in our modal programming language span multiple hosts. That is, each program consists of a series of nested expressions to be evaluated at various hosts in the network. These fragments make reference to the other sites in the network, and the resources particular to those sites. The main purpose of the type system is to track these references to remote resources so that they are only used in a safe way: Localized resources can only be used in the correct place. When code or data do not depend on their location, we are also able to indicate this in the type system. This allows them to be safely used anywhere. The main result of type safety of the language is that resources are only used in the correct place. This brings us to the thesis statement:

Thesis Statement. Modal type systems provide an elegant and practical way to control spatially distributed resources in mobile computer programs.

Although modal types naturally address spatial distribution, this is only one facet of distributed computing. In particular, a logical approach to concurrency and failure is beyond the scope of this thesis. However, the language will need to have basic support for these features in order to create a realistic application.

Implementation. The implementation of ML5 will consist of a type-directed compiler for the language and a distributed runtime system. The compiler is responsible for translating the high-level language into certified machine code. The runtime's primary task is to form the network of hosts on which the distributed programs run. To do this, it must be able to marshal code and data into raw bytes so that they can be transmitted on the network. In order to ensure safety, the marshalled representation must be accompanied by a certificate demonstrating its well-formedness. It is the runtime's responsibility to verify these certificates upon receipt, relative to the host's specific set of local resources. As the marshalled bytes are reconstituted into code and data, references to those local resources become linked to them.

The runtime has a few other responsibilities: it performs distributed garbage collection to reclaim resources, and provides rudimentary support for failure detection and concurrency (Section 7).

Application. To demonstrate the usefulness of the language and implementation, I will use it to build an example distributed application. Before choosing an appropriate application, I will need to have a better idea of how the language supports features like concurrency. Thus, this proposal does not commit to a specific application.

This concludes the high-level summary. The remainder proceeds as follows: I begin by explaining the proposed work in depth using the same outline as above (Sections 2–8). For the language design, I give both informal code examples and formal semantics for a simplified subset of the language. For the implementation of the compiler, I give formal translations for the core language constructs; these translations need only be extended to support a more full-featured language. For the compiler's back-end and runtime, I outline the research problems that still need to be solved. I conclude with a comparison to other languages for distributed computing (Section 9), and a timeline for the completion of the thesis work (Section 10).

2 ConCert

The current ConCert system is designed to harness one specific resource: idle CPU power from a large network of volunteered computers. A program designed to be run on ConCert is broken up into pieces of independent mobile code (called "cords") that may run in parallel. To the programmer, ConCert acts as a single, highly parallel computer with simple primitives for fork-join parallelism (see below).

A computer becomes part of the ConCert network by running a piece of software called the Conductor. The Conductor is responsible for maintaining contact with its peers, and for allocating cords to idle computers in order to be executed. This allocation is done using a "work-stealing" model: Each host maintains a queue of cords that are ready to be executed, and when a host is

idle, it “steals” work out of the queues of other hosts. This model is efficient [5], but because the location in which a cord will eventually run depends on the idle status of the machines in the network, it is not possible for the programmer to know in advance where his code will run. Moreover, to support fault tolerance (by restarting failed cords), a cord must be able to be run multiple times in different places and always produce the same result.

As a consequence, we provide cords with a uniform view of the network—a cord cannot tell what host it is running on, nor can it access any of that host’s local resources like permanent storage and I/O. This is acceptable because the only resource ConCert seeks to harness is idle CPU power. In contrast, the proposed work will allow programs to make use of any sort of distributed resource.

Adding language support for such local resources is tricky. To illustrate some of the issues, we look at a simple Grid/ML program.

```
let
  val f = openfile "numbers.txt"
  val f2 = openfile "factors.txt"

  (* append the list l to the output file *)
  fun writeresult l =
    write(f2, nums-to-string l ^ "\n")

  val inputs = readfile f

  (* prime factorization of n for n > 0 *)
  fun factor n =
    let
      fun trial l = n :: nil
        | trial m = if n mod m = 0
                    then factor m @ factor (n div m)
                    else trial (m - 1)
    in
      trial (floor (sqrt (real n)))
    end

  val cords =
    map (fn n => submit (fn () => factor n)) inputs

  val results = waitall cords
in
  app writeresult results
end
```

This Grid/ML program¹ computes the prime factorizations of the numbers

¹The example takes the liberty of assuming the existence of some library code, for instance

in the file "numbers.txt" and writes those to the file "factors.txt" (both files reside on the client). It does this by creating a cord for each factoring task and submitting them to the local work queue with the primitive `submit`. The Grid/ML primitives used have these types:

```
submit : (unit →  $\alpha$ ) →  $\alpha$  cord
waitall :  $\alpha$  cord list →  $\alpha$  list
```

A value of type α cord is a running computation that returns a result of type α . Grid/ML allows α to be instantiated with any type. In this case each job is an `int list cord`. After submitting the cords, the program then waits for them all to complete, and writes the results to the output file.

The file I/O that the program performs is an effect. Such effects are allowed only in the part of the Grid/ML program that runs on the client. This is because I/O interacts with the external world, and so it depends on the place in which the code is executed. Therefore, an effect inside cord code would violate the requirement that it be agnostic about the host on which it runs. In Grid/ML, such errors are detected only at run-time. For example, if we modify the program so that the body of `factor` tries to write to the file `f2` or read from the file `f`, then the program will abort at that operation. Although both files are in scope, the file descriptors do not make sense when executing at remote sites—even if we wanted to allow I/O in cord code, we would not be able to access *those* files once execution has left the client.

This does not necessarily mean that open file descriptors cannot ever leave the client. Suppose the program is modified to be higher-order:

```
let
  (* ... *)

  fun factor n =
    let
      (* ... *)
      val factors = trial (floor (sqrt (real n)))
    in
      (fn () => writeresult factors)
    end

  (* ... *)
in
  app (fn g => g ()) results
end
```

Now, instead of returning the list of factors directly, each cord returns a function that writes the result to the file on the client. The client consumes these results by calling the functions. (Writing the program this way is gratuitous

`nums-to-string` and `readfile`.

here, but higher-order mobile code is often genuinely useful!) This program has the same behavior as the first version. The reference to the local file safely makes a round trip from the client to the cord code and back in the environment of each function. Thus it is not the references themselves that we need to tame, but the way those references to local resources are *used*. To do this, the type system of ML5 will associate a place with each value, and then only allow a value to be consumed in that same place. However, values will be able to flow freely between places otherwise.

Associating a place with each value does ensure safety, but it can exclude too many programs; many values are actually portable to every location. We do not want such values to be unnecessarily restricted to the place where they are created. For example, in the program above the `factor` function refers to the functions `sqrt` and `@`, which are defined as part of the standard library on the client. This is safe because these functions do not use any local resources. An important feature of the ML5 type system is that it also accounts for values like these that can be used anywhere. Using these features the ML5 type system will allow us to statically typecheck both the original program and the higher-order version.

We have now motivated a type system enhanced with a notion of place. To summarize, the type system for ML5 will ensure that localized resources are only used in the correct place by associating with each sub-expression the place in which it will be evaluated, and associating with each bound variable the place where it may be used. We will also introduce support for bindings that are usable in any place, since this is very common. The foundation of this type system comes from modal logic, which is introduced in the next section.

3 Modal Logic

“Modal logic” refers to a family of logics with the ability to reason about truth from multiple simultaneous perspectives. For this proposal, we will start with the specific modal logic called Intuitionistic S5 and then extend it to suit our purposes. IS5 is most easily understood as an extension of Intuitionistic logic, whose basic judgment is

$$\Gamma \vdash A \text{ true}$$

where A is a proposition and Γ is a collection of hypotheses of the form B true for various propositions B . The judgment asserts that the proposition A is true, from a single universal viewpoint, assuming that the hypotheses in Γ are also true.

IS5 changes this judgment by indexing the notion of truth by a collection of “possible worlds.” These possible worlds are the different perspectives from which we judge the truth of propositions. All possible worlds must indeed be possible (they must not admit logical contradictions), but otherwise can disagree on what facts are true. As a real world example, in some possible worlds

the proposition “it is raining in Pittsburgh” is true and in others it is not. However, in no possible world is “it is raining in Pittsburgh and it is not raining in Pittsburgh” true. Modal logic gives us the ability to state that a proposition is true in some specific possible world. The basic judgment in IS5 is:

$$\Gamma \vdash A \text{ true@w}$$

where A is a proposition as before and w is a world. This judgment asserts that, assuming the hypotheses in Γ , the proposition A is true in the world w . World expressions w include both constants \mathbf{w} standing for specific worlds, and world variables, written ω . Truth in IS5 is always with respect to a world. The context Γ therefore consists of assumptions of the form $B \text{ true@w}'$.

Most of the rules from intuitionistic logic are imported into IS5, with the judgment simply changed from true to true@w. For example, we reason about conjunction as before:

$$\frac{\Gamma \vdash A \text{ true@w} \quad \Gamma \vdash B \text{ true@w}}{\Gamma \vdash A \wedge B \text{ true@w}} \wedge I \quad \frac{\Gamma \vdash A \wedge B \text{ true@w}}{\Gamma \vdash A \text{ true@w}} \wedge E_1$$

To conclude $A \wedge B$ at the world w , we must conclude both A and B at that same world. Likewise, if we know $A \wedge B$ at w , then we know A at that same world.

There are also two new connectives, which both use the indexed truth judgment in interesting ways. The first is \Box (read “box”), where $\Box A$ means that A is true in all possible worlds. The rules for \Box are:

$$\frac{\Gamma, \omega \text{ world} \vdash A \text{ true@}\omega}{\Gamma \vdash \Box A \text{ true@w}} \Box I \quad \frac{\Gamma \vdash \Box A \text{ true@w}}{\Gamma \vdash A \text{ true@w}'} \Box E$$

Because the meaning of $\Box A$ is that A is true in every world, we may conclude from $\Box A$ at some world w the proposition A at any world w' (rule $\Box E$). To prove that a proposition is true in every world (rule $\Box I$), we must prove that the proposition A is true at a new hypothetical world ω about which nothing is known. This world variable is bound in Γ .

The other new connective is \Diamond (read “diamond”). The proposition $\Diamond A$ means that A is true in *some* possible world. The \Diamond connective is defined as follows:

$$\frac{\Gamma \vdash A \text{ true@w}'}{\Gamma \vdash \Diamond A \text{ true@w}} \Diamond I \quad \frac{\Gamma \vdash \Diamond A \text{ true@w} \quad \Gamma, \omega \text{ world}, A \text{ true@}\omega \vdash C \text{ true@w}}{\Gamma \vdash C \text{ true@w}} \Diamond E$$

To prove $\Diamond A$ at w , we require a proof of A in any world w' (rule $\Diamond I$). To consume a proof of $\Diamond A$ (rule $\Diamond E$), we get to assume the existence of some world (about which nothing else is known) where A is true, in order to continue proving some other proposition C . The elimination of $\Diamond A$ happens at a different world (w') than the remainder of the reasoning (w). In our computational interpretation (Section 4), these worlds will be thought of as hosts in the network, and so this rule incurs a sort of “action at a distance” by involving two hosts in the operation. There, we will use a different but equivalent formulation of the rules that isolates this network communication to a single rule.

The at modality. The \Box and \Diamond connectives are the only new modalities of Intuitionistic S5. Our first extension to the logic adds a third modality, *at*. The proposition A *at* w means that A is true at the world w ; it is a direct internalization of the judgment A *true@w* as a proposition [19, 20]. Such internalizations are very common. For instance, implication ($A \supset B$) in propositional logic is the internalization of the hypothetical judgment (A *true* \vdash B *true*). The rules defining the connective are:

$$\frac{\Gamma \vdash A \text{ true@}w'}{\Gamma \vdash A \text{ at } w' \text{ true@}w} \text{ at I} \quad \frac{\Gamma \vdash A \text{ at } w' \text{ true@}w}{\Gamma \vdash A \text{ true@}w'} \text{ at E}$$

There are several reasons to extend IS5 with the *at* modality. One is that it is a natural and elegant consequence of formalizing IS5 with a world-indexed truth judgment. Because it directly internalizes the machinery of the proof system, it will also give us the ability to make more precise statements. In contrast, \Box and \Diamond are “lossy;” consider the following theorem of IS5:

$$\Box(A \supset B) \supset \Diamond A \supset \Diamond B$$

The reading of this proposition is as follows: Given that A implies B in every world, and given that A is true in some world, B is true in some world. This proposition is true by virtue of the fact that B is true at the same world that A is true (by using the implication). However, the proposition does not state that B is true at the same world; in fact, we are unable to state such a thing with the \Box and \Diamond connectives at all. The *at* modality allows us to make such statements. Once we introduce propositions that quantify over worlds, we can say

$$\Box(A \supset B) \supset \forall \omega. (A \text{ at } \omega) \supset (B \text{ at } \omega)$$

which is stronger than the above. In fact, the \Box and \Diamond modalities are definable in terms of *at* and quantification, where $\Box A$ is $\forall \omega. A \text{ at } \omega$ and $\Diamond A$ is $\exists \omega. A \text{ at } \omega$. Finally, the precision of the *at* modality will be necessary in the internals of the compiler for the purpose of closure conversion (Section 6.4).

Having completed the background discussion of IS5, we are now prepared to use it as the foundation of the ML5 language, which is presented in the next section. ML5 will modify the logic in several ways; for comparison, all of the standard rules are given in Appendix A.

4 The ML5 Language

I have designed and formalized in Twelf a simple calculus based on the modal logic from the previous section. (The calculus appears as the internal language of the implementation in Section 6.2). This calculus forms the basis of ML5, whose features are toured in this section. However, what is presented here is only that simple calculus dressed up to make it feel like a programming language; it is not a complete description of the language. The specification

of ML5 proper is part of the proposed work. Specifically, I do not cover the features that are treated the same as in SML (*e.g.* functions, datatypes, tuples). I also treat issues of concrete syntax loosely.

The language ML5 is designed around a Curry-Howard interpretation of Intuitionistic S5. To do this, we associate proof terms with the rules of the logic. These proof terms become the syntax of the programming language.

Its main typing judgment is

$$\Gamma \vdash M : A@w$$

which states that the expression M has type A at the world w , in the context Γ . As in the logic, each assumption in Γ is adorned with the world at which it is well-typed. Computationally, these worlds are the hosts in the network.

Because every subexpression of the program is typed at some world, an ML5 program consists of a tree of nested expressions to be evaluated at different hosts. Every variable in the context Γ is bound to a value that the program has access to, and these variables are typed at various different worlds. Therefore the program is able to manipulate both values that make sense to it (ones that are typed at the same world) and values that do not (ones that are typed at different worlds).

The worlds in the typing judgment are the static representations of hosts in the network. At runtime, we will need tokens with which to refer to these worlds. A value of type $w \text{ addr}$ is such a token, which acts as the address of a host. A host can compute its own address by using the `localhost()` primitive, which results in an address value of the form \bar{w} :

$$\frac{}{\Gamma \vdash \text{localhost}() : w \text{ addr} @w} \text{localhost} \quad \frac{}{\Gamma \vdash \bar{w} : w \text{ addr} @w'} \text{address}$$

The way an address is used is by switching the location of evaluation to the world described by the address. This primitive, called `get`, is ML5's only mechanism for moving control and data between hosts.

$$\frac{\begin{array}{l} \Gamma \vdash w' \text{ world} \\ \Gamma \vdash M : w' \text{ addr} @w \\ \Gamma \vdash N : A@w' \\ A \text{ mobile} \end{array}}{\Gamma \vdash \text{get}[w'; M]N : A@w} \text{get}$$

The `get` primitive transfers control from the current world w to the specified world w' . The judgment $\Gamma \vdash w' \text{ world}$ checks that the world expression w' is well-formed (Figure 2). The expression N , which is well-typed at w' , is evaluated there. The resulting value is then returned to w . For an arbitrary value, this motion is unsound; the value may be, or make use of, a resource that is local to w' . The final typing condition for `get` therefore restricts the primitive to types that are *mobile*. A mobile type is one whose values are always portable between worlds (Figure 1). Primitive types such as integers and strings are mobile. Addresses of worlds are mobile, as are data structures built from other

$$\begin{array}{c}
\frac{}{b \text{ mobile}} \qquad \frac{A \text{ mobile} \quad B \text{ mobile}}{A \times B \text{ mobile}} \\
\frac{}{\Box A \text{ mobile}} \qquad \frac{}{\Diamond A \text{ mobile}} \\
\frac{\alpha \text{ mobile}}{\mu\alpha.A \text{ mobile}} \qquad \frac{A \text{ mobile} \quad B \text{ mobile}}{A + B \text{ mobile}} \\
\frac{}{w \text{ addr mobile}}
\end{array}$$

Figure 1: The definition of the mobile judgment. The metavariable b ranges over ML base types like `int` and `string`. The rule for recursive types $\mu\alpha.A$ introduces a mobility assumption for the type variable, so mobility is actually a hypothetical judgment. However, no other type variables are ever considered mobile

$$\frac{}{\Gamma, \omega \text{ world}, \Gamma' \vdash \omega \text{ world}} \quad \frac{}{\Gamma \vdash w \text{ world}}$$

Figure 2: Definition of the judgment $\Gamma \vdash w \text{ world}$

mobile types. Abstract types, mutable references, and functions are not mobile, among other things. Some functions and abstract types may still be moved, however, by embedding them in the \Box modality, which is mobile.

Local action. The presence of the `get` primitive allows us to restrict all other operations to dynamically involve only one world. In doing so, we will retain all of the expressive power of the logic.

In the logic, the proposition $\Diamond A$ meant that A is true at some world. In the ML5 language, a value of type $\Diamond A$ is an address for some value of type A , which lives at an arbitrary unknown world. The rules for \Diamond in ML5 are as follows:

$$\frac{\Gamma \vdash M : A@w}{\Gamma \vdash \text{here } M : \Diamond A@w} \quad \Diamond I$$

$$\frac{\Gamma \vdash M : \Diamond A@w \quad \Gamma, \omega \text{ world}, a:\omega \text{ addr}@w, x:A@w \vdash N : C@w}{\Gamma \vdash \text{letd } \langle \omega, a, x \rangle = M \text{ in } N : C@w} \quad \Diamond E$$

For both of these primitives, all of the computation happens at the world w . We may only compute the address of the result of evaluating a local expression ($\Diamond I$). To use something of type $\Diamond A$, we must have it at the current world ($\Diamond E$). In the body of the `letd`, we have in scope a hypothetical world ω where the value is, an address a for that world, and the value itself x . In comparison, the rules from Section 3 allowed us to evaluate and take the address of remote

expressions, as well as eliminate remote addresses. We can derive those rules by composing the ML5 rules with `get`.

In the logic, proposition $\Box A$ meant that A is true in every world. In ML5, a value of type $\Box A$ is a suspended expression of type A that can be evaluated anywhere. The rules for the \Box constructor are also purely local:

$$\frac{\Gamma, \omega \text{ world} \vdash M : A@w}{\Gamma \vdash \text{box } \omega.M : \Box A@w} \Box I \quad \frac{\Gamma \vdash M : \Box A@w}{\Gamma \vdash \text{unbox } M : A@w} \Box E$$

Although the body of the `box` is typed at the hypothetical world ω ($\Box I$), this construct does not access any remote worlds at runtime. The code inside the `box` is suspended, so `box $\omega.M$` is already a value. Unboxing is straightforward ($\Box E$); the suspended expression is evaluated at this time. As before, the remote version of `unbox` can be derived by composing this rule with `get`.

We now have enough to build some small examples, which we will use to motivate the addition of a new judgment for global reasoning. The following proposition is one of the characteristic axioms of Intuitionistic S5:

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

As a type it describes a function that when given the address of some suspended portable code of type A we computes a value of type A . In ML5, a function `f` of this type is:

```
fun f (x :  $\Diamond \Box \alpha$ ) :  $\alpha$  =
  letd <w, a, y> = x
  in
    unbox(get [w;a] y)
  end
```

The function `f` works by getting the portable code from the world indicated by the address, and then running it locally.

A different characteristic axiom is problematic:

$$\Box A \supset \Box \Box A$$

In ML5, we might like to write this as follows (supposing we are at the world w):

```
(* ill-typed *)
fun g (x :  $\Box \alpha$ ) :  $\Box \Box \alpha$  =
  let val a : w addr = localhost()
  in box w'. get[w; a] x
  end
```

This code does not type-check. The idea is that we save the address of the local machine as `a`. We then create a second box, whose contents is an expression that gets the original box from the world w . The reason that this does not

type-check is that the variable `a` is bound at the world `w`. The `get`, which happens at the world `w'` but tries to use `a`, is ill-typed. Moreover, we cannot `get` the address itself from `w`, since to do so would require an address for `w` at `w'`, which is the very thing we lack! The solution will be to add a new judgment to classify values that are well-typed at any world. Addresses will be such values, which will make the above program well-typed with a small modification. The program also has a second problem: the boxed term that we create has to contact the home world `w` every time it is unboxed. Using the new judgment we will also be able to create a more efficient implementation of this function.

4.1 Validity

To enable global reasoning, we add a new judgment to the language:

$$\Gamma \vdash S \sim A@w$$

In this judgment, called the validity judgment, S is a *valid expression* with type A at the world w . S will evaluate (at w) to a value that has type A in any world. Assumptions of the form $u \sim A$ now appear in Γ as well. Since these assumptions stand for values that are well-typed in any world, they are not annotated with any world at all. A valid hypothesis can be used at any world:

$$\frac{}{\Gamma, u \sim A, \Gamma' \vdash u : A@w} \text{vhyp}$$

Valid expressions take on one of the two following forms:

$$\begin{array}{l} \text{valid expressions } S ::= \text{value } v \\ \quad \quad \quad \quad \quad \quad | \text{mobile } M \end{array}$$

The typing conditions are as follows:

$$\frac{\Gamma, \omega \text{ world} \vdash v : A@w}{\Gamma \vdash \text{value } v \sim A@w} \text{value} \quad \frac{\begin{array}{c} A \text{ mobile} \\ \Gamma \vdash M : A@w \end{array}}{\Gamma \vdash \text{mobile } M \sim A@w} \text{mobile}$$

For `value`, the argument is already a value. The condition is simply that the value be well-typed in any world, which is checked by typing it at a hypothetical world about which nothing is known. For `mobile`, the expression must be well-typed, and the type of the expression must be `mobile`. This ensures that after the expression is evaluated, the result will be well-typed at any world.

The Validity Modality. The valid judgment is internalized to define the circle modality $\circ A$. A value of type $\circ A$ is a suspended expression that, when evaluated, produces a value that is well typed anywhere [32]. It is defined as follows:

$$\frac{\Gamma \vdash S \sim A@w}{\Gamma \vdash \text{cir } S : \circ A@w} \circ I \quad \frac{\begin{array}{c} \Gamma \vdash M : \circ A@w \\ \Gamma, u \sim A \vdash N : C@w \end{array}}{\Gamma \vdash \text{let } v \ u = M \text{ in } N : C@w} \circ E$$

Because the valid expression is suspended, `cir S` is a value. At runtime, the `letv` construct proceeds by evaluating `M` to the form `cir value v` or `cir mobile M`, then evaluating `M` in the latter form, and then binding the resulting value to `u` for the purpose of evaluating `N`.

The \bigcirc modality allows us to type the above attempt at implementing $\square A \supset \square\square A$:

```
fun g (x :  $\square\alpha$ ) :  $\square\square\alpha$  =
  letv u ~ w addr = cir (mobile (localhost()))
  in
    box w'. get[w; u] x
  end
```

In this version, we are able to use the address for `w` at the world `w'` because it is a global hypothesis. A shortcoming remains: the box produced in this version must contact the world `w` every time it is unboxed. An improved version uses the fact that $\square\alpha$ is `mobile` to create a valid hypothesis `v` from the argument `x`, which can then be used directly at `w'`:

```
fun g (x :  $\square\alpha$ ) :  $\square\square\alpha$  =
  letv v ~  $\square\alpha$  = cir (mobile x)
  in
    box w'. v
  end
```

The type $\bigcirc A$ is not `mobile`, because the suspended expression of type `A` is typed with respect to the world where the circle is introduced. A portable suspended expression that produces a portable value would have type $\square\bigcirc A$.

A common idiom is to introduce the circle modality and then immediately eliminate it in order to bind a valid variable. ML5 therefore provides two derived forms:

$$\begin{aligned} \text{putv } u = v \text{ in } N &\doteq \text{letv } u = \text{cir}(\text{value } v) \text{ in } N \\ \text{putm } u = M \text{ in } N &\doteq \text{letv } u = \text{cir}(\text{mobile } M) \text{ in } N \end{aligned}$$

Polymorphism. Like Standard ML, ML5 supports prenex parametric polymorphism over types. It also supports polymorphism over worlds in an analogous way. We create a new syntactic class of polytypes, and two new forms of hypothesis that can appear in Γ :

$$\begin{aligned} \text{polytypes } P &::= \forall\alpha_1, \dots, \alpha_n, \omega_1, \dots, \omega_m. A \\ \text{contexts } \Gamma &::= \dots \mid x:P@w \mid x\sim P \end{aligned}$$

Only assumptions can have polymorphic type. Every use of a polymorphic assumption must completely instantiate the polymorphic type:

$$\frac{\Gamma, x:(\forall\alpha_1, \dots, \alpha_n, \omega_1, \dots, \omega_m. A)@w, \Gamma' \vdash}{x\langle A_1, \dots, A_n, w_1, \dots, w_m \rangle : ([^{A_1/\alpha_1, \dots, A_n/\alpha_n, w_1/\omega_1, \dots, w_m/\omega_m}]A)@w} \text{polyhyp}$$

Polymorphic variables are introduced by additional polymorphic versions of the `val`, `fun`, and `letv` bindings. For instance, the additional `val` rule is as follows:

$$\frac{\Gamma, \alpha_1 \text{ type}, \dots, \alpha_n \text{ type}, \omega_1 \text{ world}, \dots, \omega_m \text{ world} \vdash v : A@w \quad \Gamma, x:(\forall \alpha_1, \dots, \alpha_n, \omega_1, \dots, \omega_m. A)@w \vdash N : C@w}{\Gamma \vdash \text{let } \langle \alpha_1, \dots, \alpha_n; \omega_1, \dots, \omega_m \rangle \text{ val } x = v \text{ in } N : C@w} \text{ val}$$

ML5 has a value restriction, like SML: A binding can only be polymorphic if the right hand side is a value. This allows the type inference process to insert quantifiers without changing the dynamic meaning of the program. This proposal does not include a description of type inference, but I expect it to be a straightforward extension of the standard algorithm.

The at Modality. A value of type A at w is an encapsulated value, where that value is well-typed at another world.

$$\frac{\Gamma \vdash v : A@w'}{\Gamma \vdash \text{hold}_{w'} v : A \text{ at } w'@w} \text{ at I} \quad \frac{\Gamma \vdash M : A \text{ at } w'@w \quad \Gamma, x:A@w' \vdash N : C@w}{\Gamma \vdash \text{leta } x = M \text{ in } N : C@w} \text{ at E}$$

The `leta` binding consumes a value of type A at w' by putting an assumption $x:A@w'$ in the context. The value `holdw'v` represents the encapsulation of v . Its body must be a value because we are not in general at the correct world to evaluate it; consider the following ill-formed code:

```
(* ill-typed *)
let val x = hold[w'] (ref 0)
in leta y = x
  in
    M
  end
end
```

If this code is to be allowed, then within the expression M , y should be bound to an allocated reference at w' . There has been no opportunity to create such a reference, however. The static world w' might not even have a dynamic address available at this point, which would be necessary to make contact with it.

Dynamically, `hold` is just a type coercion and `leta` behaves just like `let`.

Sums. Sum types, which come from the disjunctive connective in logic, are not entirely straightforward. Like the elimination form for \diamond , the standard logical rule for disjunction allows the remote elimination of $A \vee B$:

$$\frac{\Gamma \vdash A \vee B \text{ true}@w' \quad \Gamma, A \text{ true}@w' \vdash C \text{ true}@w \quad \Gamma, B \text{ true}@w' \vdash C \text{ true}@w}{\Gamma \vdash C \text{ true}@w} \vee E$$

In this rule $A \vee B$ may be at a different world (w') than our current one (w). So far, all rules in ML5 other than `get` have acted solely on local data. We would like to maintain this, so that (for one thing) existing single-world ML code still runs without any overhead from the distributed features. The local rules for sum types in ML5 are:

$$\begin{array}{c}
\Gamma \vdash M : A + B@w' \\
\Gamma, x:A@w' \vdash N_1 : C@w \\
\Gamma, x:B@w' \vdash N_2 : C@w \\
\hline
\text{case } M \text{ of} \\
\Gamma \vdash \quad \text{inl } x \Rightarrow N_1 : C@w \\
\quad \quad | \text{ inr } x \Rightarrow N_2 \\
\hline
\frac{\Gamma \vdash M : A@w}{\Gamma \vdash \text{inl } M : A + B@w} +I_1 \quad \frac{\Gamma \vdash M : B@w}{\Gamma \vdash \text{inr } M : A + B@w} +I_2
\end{array}$$

For \square and \diamond , the `get` construct makes the nonlocal logical rules derivable. This does not work as easily for disjunction, because we cannot `get` a value of arbitrary sum type ($A \vee B$ is only mobile if both A and B are mobile.) However, the nonlocal case analysis is still derivable in ML5. Here is one way to do it; the rule

$$\begin{array}{c}
\Gamma \vdash O : w' \text{ addr } @w \\
\Gamma \vdash M : A + B@w' \\
\Gamma, x:A@w' \vdash N_1 : C@w \\
\Gamma, x:B@w' \vdash N_2 : C@w \\
\hline
\text{case}[w'; O] M \text{ of} \\
\Gamma \vdash \quad \text{inl } x \Rightarrow N_1 : C@w \\
\quad \quad | \text{ inr } x \Rightarrow N_2 \\
\hline
+ \text{E-nonlocal}
\end{array}$$

can be implemented as

```

let
  val mr = get[w'; 0]
      (case M of
        inl y => inl (hold[w'] y)
        | inr y => inr (hold[w'] y))
in
  case mr of
    inl x' => leta x = x'
              in N1
            end
  | inr x' => leta x = x'
              in N2
            end
end

```

Here, the type of value that we `get` is A at w' + B at w' , which is guaranteed to be mobile. Therefore, we retain all of the strength of the logical \vee connective

in IS5. The only downside is that this derivation requires us to enlist the `at` modality. The consequence is that the logical completeness relies on the presence of both `∨` and `at`, meaning that the two are not truly orthogonal in the language.

Like in Standard ML, labeled sums will actually be supported (along with recursive types) as part of the datatype mechanism.

Properties. The language presented here should have two main properties. First is type safety: Every well-typed program can take a step to another well-typed program, or is already a value. I have proved type safety in Twelf for a simplified but representative version of the calculus. Second is a connection to the logic: the modifications that we have made (particularly the use of `get` and locally-acting rules) should not have reduced the expressive power of the language. We prove for two related lambda calculi [29, 28] that this alternate formulation proves exactly the same theorems as the standard presentation of the logic. Both of these proofs have also been formalized in Twelf. Part of the thesis research will be extending this result to the pure subset of ML5 (with the `○` and `at` modalities) as well.

5 Code Examples

The following small examples illustrate how the constructs of ML5 can be used to build distributed applications.

Mobile code proxy. Many distributed programming languages provide a way of forming “proxies” for code that cannot itself move. These proxies can then be passed around, and used as stand-ins for the functions they mimic. In ML5, if the argument and return types of a function are mobile, then one can create a proxy for it. Here we’ll assume we have a variable $f:A \rightarrow B@w_0$ and will bind a valid variable $u \sim A \rightarrow B$ which is the mobile proxy for that function. This assumes that A and B are mobile.

```
putm w0a = localhost()
in putv u =
  (fn a =>
    putm v = a
    in
      get [w0; w0a] (f v)
    end)
  in
    ...
  end
end
```

We begin by computing the address of w_0 and making it valid, so that we can contact w_0 from any world. We then bind u to a function value, which is

the proxy. The proxy works by taking the argument (known to be mobile by assumption) and making it valid as well. The proxy then performs a `get` to apply the original function at w_0 and return the result.

Cords. Assuming some simple facilities for concurrency and the discovery of idle hosts, we can also emulate the features of Grid/ML in a statically-typed way. The Grid/ML client will be world h , which is the world that the program begins running on in ML5. Let's suppose that the following are bound in the context:

```
future::Type → Type,
future ~ ((unit → α) → α future),
now ~ (α future → α),
find_idle ~ (unit → ◇unit)
```

We assume Multilisp-like futures, which are a simple mechanism for fork-join parallelism. An α future is a computation, running concurrently, that produces a value of type α when synchronized on with `now`. Futures must allow the execution of arbitrary ML5 code, but needn't have any special features for distribution. We also expect a `find_idle` routine to give the address of an idle host, which is represented here as $\Diamond\text{unit}$. All of these primitives are valid hypotheses, so they can be used in any world.

With this, we can implement a "mobile spawn" function like the one that is provided by Grid/ML.

```
(* spawn ~ ∀α,ω.(□(α at ω) -> (α at ω) future) *)
putv ⟨α, ω⟩ spawn =
  (fn (code : □(α at ω)) =>
    future (fn () =>
      letd <w2, w2a, ()> = find_idle()
      in putm c = code
      in get[w2; w2a] (unbox c)
      end
    end))
in
  ...
end
```

This implementation of `spawn` takes as an argument a piece of suspended code suitable to run at any world. That suspended code must produce a result that is well-typed at the world ω . The function `spawn` is polymorphic in ω , so this is the caller's choice. It then produces a future that finds an idle host and runs the code there.

The following code shows how `spawn` could be used from the home world h :

```
val f = spawn(h, int)
      (box _,_.
```

```

    let fun nfactores (n : int) : int = ...
    in putm res = nfactores n
      in hold[h] res
      end
    end)

(* ... some computations in parallel ... *)
val answer : int = (leta a = now f in a end)

```

This program invokes `spawn` with some “cord” code. That cord computes the number of factors of some specific n , and coerces the resulting value (an integer, which is mobile) to the world h . There is no requirement that the result of the computation be mobile; if we like, it can be specifically tailored to the world that spawns it. Once we’ve spawned the cord f , we can later wait on its answer with `now`. Since it has type `int` at h and we are running at h , we can use `leta` to unencapsulate the result into a regular integer.

Although this requires more text than using the built-in primitives of Grid/ML, ML5 has the advantage of superior static checking. The type system ensures that the code given to `spawn` can run anywhere, and ensures that the result is well-formed at the world h . Unlike Grid/ML, we can also be more specific about the capabilities of the hosts on which our mobile code will run. Here we have represented the addresses of hosts as $\diamond\text{unit}$. If we wanted to express the idea that machines on the Grid have their own distinct set of capabilities, we need only amend the implementation of `find_idle` to return $\diamond\text{caps}$ for some appropriate description of its local capabilities.

Having toured the language’s features and looked at some simple examples, we now switch gears to a discussion of its implementation.

6 The ML5 Compiler

This section describes the planned implementation of a compiler for ML5. The early phases of the compiler have been studied in detail. The later phases have not; part of the thesis work will be to figure out the best course of action for these and take it.

6.1 Design Issues

Before walking through the phases of the compiler, it is important to recall some of the issues that will guide the design.

Data Representation. The chief thing that makes the implementation of ML5 different from a standard type-directed compiler for ML is that programs are able to manipulate data that are typed at different worlds. Because of our local-only elimination rules, the only way that these data can be used is in preparation of code or data also intended to be evaluated at other worlds. Therefore,

the implementation has considerable flexibility in the way these abstract remote values are represented—the choices can be guided almost entirely by the marshalling strategy.

Marshalling. Because ML5 programs run on a network, we require the ability to marshal code and data into raw bytes to be sent over that network. The compiler will support the runtime in this task by providing marshalling and unmarshalling functions for each abstract type. (Recall that any ML5 value can be marshalled, even if the type is not mobile.)

Certification. In order to use ML5 in our trustless Grid computing project, the compiler must produce certificates that demonstrate the safety of the mobile code. I will do this by building a type-directed compiler that preserves and translates types throughout the phases of compilation. The compiler will output typed assembly code for the TALT system [12].

The phases of compilation are as follows. First, the external language (ML5) is elaborated into a simple internal language (IL) (Section 6.2). Second, we convert this IL to continuation passing style (CPS) (Section 6.3). Next, we perform closure conversion to implement lexically scoped closures with closed functions (Section 6.4). Each of these phases is standard [26] apart from the new features of ML5, so the description concentrates on those.

6.2 Elaboration

The purpose of elaboration is to replace features of convenience in ML5 with the simpler type theoretic constructs of the IL. For the purpose of this proposal, those features are the \square and \diamond modalities and world polymorphism.

World Quantification. In the IL, type inference is no longer an issue. We can therefore abandon ML5’s restriction to prenex polymorphism, which was expressed in terms of a syntactic class of “polytypes.” Instead we simply have the type constructor $\forall\omega.A$:

$$\frac{\Gamma, \omega \text{ world} \vdash v : A@w}{\Gamma \vdash \Lambda\omega.v : \forall\omega.A@w} \forall I \quad \frac{\Gamma \vdash M : \forall\omega.A@w}{\Gamma \vdash M\langle w' \rangle : [w'/\omega]A@w} \forall E$$

The IL supports a “type-erasure semantics” [26] where types are purely static entities, and the term constructors that manipulate them can be erased without changing the dynamic meaning of the program. To achieve this, the constructor Λ is restricted to values. This is all we need for the elaboration of ML5, where polymorphism is also restricted to values.

The IL also has a first-class existential quantifier, which will be used for

closure conversion and the elaboration of \diamond :

$$\frac{\Gamma \vdash w' \text{ world} \quad \Gamma \vdash v : [w'/\omega]A@w}{\Gamma \vdash \text{pack}[w']v \text{ as } \exists\omega.A : \exists\omega.A@w} \exists I \quad \frac{\Gamma \vdash M : \exists\omega.A@w \quad \Gamma, \omega \text{ world}, x:A \vdash N : C@w}{\Gamma \vdash \text{unpack } M \text{ as } \omega, x \text{ in } N : C@w} \exists E$$

The Elaboration Relation. Elaboration is given as a function on ML5 typing derivations. As a shorthand, we write this as a function on expressions, $\llbracket \cdot \rrbracket_E^{A@w}$, parameterized by the current world w and the type A of the expression being translated.² We also give a function that translates ML5 types $\llbracket \cdot \rrbracket_E$. The cases for eliminating the \square and \diamond modalities are:

$$\begin{aligned} \llbracket \square A \rrbracket_E &= \forall\omega.((\text{unit} \rightarrow \llbracket A \rrbracket_E) \text{ at } \omega) \\ \llbracket \diamond A \rrbracket_E &= \exists\omega.(\llbracket A \rrbracket_E \text{ at } \omega \times \omega \text{ addr}) \\ \llbracket \text{box } \omega'.M \rrbracket_E^{\square A@w} &= \Lambda\omega'.\text{hold}_{\omega'}\lambda().\llbracket M \rrbracket_E^{A@w'} \\ \llbracket \text{unbox } M \rrbracket_E^{A@w} &= \text{leta } x = \llbracket M \rrbracket_E^{\square A@w}(\omega) \text{ in } x() \text{ end} \\ \llbracket \text{here } M \rrbracket_E^{\diamond A@w} &= \text{let } x = \llbracket M \rrbracket_E^{A@w} \text{ in} \\ &\quad \text{let } a = \text{localhost}() \text{ in} \\ &\quad \text{pack}[w](x, a) \text{ as } \exists\omega.(\llbracket A \rrbracket_E \text{ at } \omega \times \omega \text{ addr}) \\ \llbracket \text{letd } \langle\omega, a, x\rangle = M \text{ in } N \rrbracket_E^{C@w} &= \text{unpack } \llbracket M \rrbracket_E^{\diamond A@w} \text{ as } \omega, y \text{ in} \\ &\quad \text{let } x = \text{fst } y \text{ in} \\ &\quad \text{let } a = \text{snd } y \text{ in } \llbracket N \rrbracket_E^{C@w} \end{aligned}$$

World polymorphism in ML5 is translated into uses of the $\forall\omega$ quantifier in the IL. Because polymorphic types are a different syntactic class in ML5, this requires another translation $\llbracket \cdot \rrbracket_E^\forall$ for polytypes.

$$\begin{aligned} \llbracket \forall\alpha_1, \dots, \alpha_n, \omega_1, \dots, \omega_m. A \rrbracket_E^\forall &= \forall\alpha_1, \dots, \forall\alpha_n. \forall\omega_1, \dots, \forall\omega_m. \llbracket A \rrbracket_E \\ \llbracket x \langle A_1, \dots, A_n, w_1, \dots, w_m \rangle \rrbracket_E^{A@w} &= x \langle A_1 \rangle \dots \langle A_n \rangle \langle w_1 \rangle \dots \langle w_m \rangle \\ \llbracket \text{let val } \langle\alpha_1, \dots, \alpha_n, \omega_1, \dots, \omega_n\rangle x = v \text{ in } N \rrbracket_E^{C@w} &= \text{let val } x = \Lambda\alpha_1, \dots, \Lambda\alpha_n. \Lambda\omega_1, \dots, \Lambda\omega_n. \llbracket v \rrbracket_E^{\omega@w} \\ &\quad \text{in } \llbracket N \rrbracket_E^{C@w} \end{aligned}$$

Type Polymorphism. One final phase is needed to eliminate the need for types at run-time. The program must be able to marshal values of any type, including the abstract type variables bound by Λ . I will achieve this by using dictionary-passing to make available a marshalling and unmarshalling function for each type variable. This is a standard implementation technique for type classes [44]. Section 6.5 explains marshalling in more detail.

²In some recursive invocations of the function, we must come up with the type of a subexpression, where that type is not evident from the expression itself. This is why elaboration works on typing derivations (where all of the types of subterms are present) rather than directly on the expressions. In those recursive calls we write $_$ for the missing type.

exps	c	::=	$\text{leta } x = v \text{ in } c$ $\text{letg } u = v \text{ in } c$ $\text{lift } u = v \text{ in } c$ $\text{let } x = \text{fst } v \text{ in } c$ $\text{let } x = \text{snd } v \text{ in } c$ $\text{let } x = \text{mkpair } v_1, v_2 \text{ in } c$ $\text{let } x = \text{localhost}() \text{ in } c$ \dots $\text{go}[w; v_a] c$ $\text{call } v_f(v_x)$ halt
values	v	::=	$\langle v_1, v_2 \rangle$ $\text{hold}_w v$ $\lambda x. c$ $\Lambda \omega. v$ $\Lambda \alpha. v$ $\text{cval } v$ \bar{w} u x \dots
types	A, B	::=	$A \text{ cont} \mid A \times B \mid A \text{ at } w$ $\exists A \mid \forall \omega. A \mid \exists \omega. A$ $w \text{ addr} \mid \dots$

Figure 3: A representative sample of the CPS language

6.3 Continuation Passing Style

After elaborating into the IL, the goal of the compiler is to repeatedly simplify the code by removing the use of constructs that have no direct analogue in machine code. The first such translation serializes the evaluation of expressions and makes control flow explicit by putting the program in continuation passing style [4].

Other than the translation of valid expressions and `get`, CPS conversion is standard. Figure 3 contains a representative sample of the CPS language (the dynamic semantics for this fragment are given in Appendix B). A value of type $A \text{ cont}$ is a continuation expecting a value of type A . Its canonical form is $\lambda x. c$. The type $\exists A$ is a new modality that is used to compile the $\circ A$ modality. A value of type $\exists A$ is an encapsulated value that is well-typed anywhere. It is an internalization of the $u \sim A$ hypothesis just as the `at` modality is an internalization of the $x:A@w$ hypothesis. The canonical form of type $\exists A$ is `cval` v . The value \bar{w} represents an address for the world w . Such values are introduced by the `localhost` operation.

$$\begin{aligned} \llbracket A \rightarrow B \rrbracket_c &= (\llbracket A \rrbracket_c \times (\llbracket B \rrbracket_c \text{ cont})) \text{ cont} \\ \llbracket \odot A \rrbracket_c &= (\odot \llbracket A \rrbracket_c) \text{ cont cont} \end{aligned}$$

Figure 4: Selected cases for the CPS type translation

We use the judgment

$$\Gamma \vdash c@w$$

to type-check CPS expressions, because they do not return (and so they cannot be distinguished by return type). We use

$$\Gamma \vdash v : A@w$$

to type-check values. The interesting typing rules are as follows:

$$\begin{aligned} &\frac{\Gamma \vdash v_a : w' \text{ addr}@w \quad \Gamma \vdash c@w'}{\Gamma \vdash \text{go}[w'; v_a] c@w} \text{ go} \\ &\frac{\Gamma \vdash v_f : A \text{ cont}@w \quad \Gamma \vdash v_x : A@w}{\Gamma \vdash \text{call } v_f(v_x)@w} \text{ cont E} \\ &\frac{A \text{ cmobile} \quad \Gamma \vdash v : A@w \quad \Gamma, u \sim A \vdash c@w}{\Gamma \vdash \text{lift } u = v \text{ in } c@w} \text{ lift} \\ &\frac{\Gamma, \omega \text{ world} \vdash v : A@w}{\Gamma \vdash \text{cval } v : \odot A@w} \odot \text{I} \\ &\frac{\Gamma \vdash v : \odot A@w \quad \Gamma, u \sim A \vdash c@w}{\Gamma \vdash \text{letg } u = v \text{ in } c@w} \odot \text{E} \end{aligned}$$

In the CPS language, `go` replaces ML5's `get`. The `go` construct invokes a continuation at a remote world. In contrast to the IL's `get`, `go` is not restricted to mobile types, because it is only a control-flow construct. The mobility restriction shows up in `lift`, which is also used in the translation of `get` below.

The `call` primitive is used to pass a value to a continuation. It simply requires that the argument is the same type that the continuation expects, and that both are typed at the same world. The `lift` construct binds a valid variable to a value whose type is `cmobile`. The `cmobile` judgment is analogous to the `mobile` judgment from ML5 and its definition appears in Figure 5. The `cval` construct, which is the introduction form for the \odot modality, requires a value. The typing rule checks that the value is well-typed in any world by checking it in a fresh hypothetical world. The elimination form, `letg`, simply binds the valid variable in the continuation.

There are four translations to convert from the IL to CPS, which translate IL types, expressions, values, and valid expressions. The translation of IL types

$$\begin{array}{c}
\frac{}{b \text{ cmobile}} \quad \frac{A \text{ cmobile} \quad B \text{ cmobile}}{A \times B \text{ cmobile}} \\
\frac{\alpha \text{ cmobile}}{\mu\alpha.A \text{ cmobile}} \quad \frac{A \text{ cmobile} \quad B \text{ cmobile}}{A + B \text{ cmobile}} \\
\frac{A \text{ cmobile}}{\forall\alpha.A \text{ cmobile}} \quad \frac{A \text{ cmobile}}{\forall\omega.A \text{ cmobile}} \\
\frac{A \text{ cmobile}}{\exists\alpha.A \text{ cmobile}} \quad \frac{A \text{ cmobile}}{\exists\omega.A \text{ cmobile}} \\
\frac{}{w \text{ addr cmobile}}
\end{array}$$

Figure 5: The definition of the cmobile judgment

to CPS types is given by a function $\llbracket \cdot \rrbracket_c$. The interesting cases are given in Figure 4. An important lemma states that this translation preserves our notion of mobility:

Lemma 1 (Preservation of Mobility)

If A mobile then $\llbracket A \rrbracket_c$ cmobile.

Proof is straightforward by induction on the derivation of A mobile. \square

The translation of IL terms to CPS terms is given in terms of a higher-order function `convert`. It takes as arguments the current world, an IL expression, and a meta-level continuation to which the result of converting the IL expression is passed. For example, the translation for tuple construction is as follows:

$$\begin{array}{l}
\text{convert } w \langle M_1, M_2 \rangle \mathcal{K} = \\
\text{convert } w M_1 \mathcal{K}_1 \\
\text{where } \mathcal{K}_1(v_1) = \\
\text{convert } w M_2 \mathcal{K}_2 \\
\text{where } \mathcal{K}_2(v_2) = \\
\text{let } x = \text{mkpair } v_1, v_2 \\
\text{in } \mathcal{K}(x)
\end{array}$$

The translation of IL values is given by a second function, $\llbracket \cdot \rrbracket_{cv}$, which produces

a CPS value. This is standard; we give only a few cases:

$$\begin{aligned}
\llbracket n \rrbracket_{\text{cv}}^w &= n \\
\llbracket \lambda x. M \rrbracket_{\text{cv}}^w &= \lambda y. \text{let } x = \text{fst } y \text{ in} \\
&\quad \text{let } k = \text{snd } y \text{ in} \\
&\quad \text{convert } w \ M \ \mathcal{K} \\
&\quad \text{where } \mathcal{K} \ v = \text{call } k(v) \\
&\quad \vdots \\
\llbracket \text{cir } S \rrbracket_{\text{cv}}^w &= \llbracket S \rrbracket_{\text{cs}}^w
\end{aligned}$$

The case for `cir` makes use of the third translation function, $\llbracket \cdot \rrbracket_{\text{cs}}^w$. It takes a valid expression of type A and the current world, and produces a CPS value of type $(\wp[A]_c)$ cont cont. This function is defined as follows:

$$\begin{aligned}
\llbracket \text{valid } v \rrbracket_{\text{cs}}^w &= \lambda k. \text{call } k(\text{cval } v) \\
\llbracket \text{mobile } M \rrbracket_{\text{cs}}^w &= \lambda k. \text{convert } w \ M \ \mathcal{K} \\
&\quad \text{where } \mathcal{K}(v) = \\
&\quad \quad \text{lift } u = v \ \text{in} \\
&\quad \quad \text{call } k(\text{cval } u)
\end{aligned}$$

We can now give the interesting cases of `convert`. The elimination form for \bigcirc is translated into a use of `letg`:

$$\begin{aligned}
\text{convert } w \ (\text{letv } u = M \ \text{in } N) \ \mathcal{K} &= \\
\text{convert } w \ M \ \mathcal{K}_1 & \\
\text{where } \mathcal{K}_1(v) &= \text{call } v(\lambda x. \text{letg } u = x \ \text{in} \\
&\quad \text{convert } w \ N \ \mathcal{K})
\end{aligned}$$

The result of converting M (of type $\bigcirc A$) is a value of type $(\wp[A]_c)$ cont cont. We call this continuation with a continuation that unpacks its argument into a validity variable and proceeds with the translated body of the `letv`.

The other interesting case is the translation of `get`:

$$\begin{aligned}
\text{convert } w \ (\text{get}[w'; M_a]M_r) \ \mathcal{K} &= \\
\text{convert } w \ M_a \ \mathcal{K}_1 & \\
\text{where } \mathcal{K}_1(v_{a'}) &= \\
\quad \text{let } a = \text{localhost} \ \text{in} & \\
\quad \text{lift } u_a = a \ \text{in} & \\
\quad \text{go}[w'; v_{a'}] \ \text{convert } N \ w' \ \mathcal{K}_r & \\
\text{where } \mathcal{K}_r(v_r) &= \\
\quad \text{lift } u_r = v_r \ \text{in} & \\
\quad \text{go}[w; u_a] \ \mathcal{K}(u_r) &
\end{aligned}$$

The translation consists of two uses of `go`; one to transfer control to the remote world, and one to come back. We start by evaluating M_a to the value $v_{a'}$, the address for the remote world. We then compute the local address a , which will be used to get back. Because we need to use this address from the remote

world, we create a valid hypothesis u_a from it with `lift` (addresses are mobile). We then transfer control to the world w' with `go`. While in the world w' , the translation of N is evaluated. Since we want to use the result v_r back at the original world w , it is made into a valid hypothesis using `lift` again. The typing condition for the IL `get` construct requires A to be `mobile`, so by Lemma 1 the type of this value is `cmobile` as required. The code then returns to the world w to continue with the final result.

Properties. The CPS translation should satisfy the following properties, where $\llbracket \Gamma \rrbracket_c$ is the natural extension of $\llbracket A \rrbracket_c$ to contexts:

Property 1 (Type soundness of CPS translation)

- (a) If $\Gamma \vdash v : A@w$ (in the IL)
then $\llbracket \Gamma \rrbracket_c \vdash \llbracket v \rrbracket_{cv}^w : \llbracket A \rrbracket_c@w$ (in CPS).
- (b) If $\Gamma \vdash S : A@w$ (in the IL)
then $\llbracket \Gamma \rrbracket_c \vdash \llbracket S \rrbracket_{cs}^w : (\ominus \llbracket A \rrbracket_c) \text{ cont cont } @w$ (in CPS).
- (c) If $\Gamma \vdash M : A@w$ (in the IL)
and for all v such that $\llbracket \Gamma \rrbracket_c \vdash v : \llbracket A \rrbracket_c@w$, $\llbracket \Gamma \rrbracket_c \vdash \mathcal{K}(v)@w$ (in CPS)
then $\llbracket \Gamma \rrbracket_c \vdash (\text{convert } w M \mathcal{K})@w$ (in CPS).

This property has been proven and formalized in Twelf for a similar language without the validity judgment, \circ , and \ominus modalities.

6.4 Closure Conversion

The next important step in compilation is closure conversion. Although low-level machine languages have the ability to represent higher-order code (by passing addresses of code), there is no direct support for nesting these code blocks within one another with lexical scope. Closure conversion implements lexical scope by constructing an explicit *environment* for each λ , so that each function is actually closed.

Once every function is closed, we can hoist the definitions to the top level, and refer to them by label rather than variable. In our distributed setting, the idea is that each label globally identifies a code block, so that we need only to pass around labels and environments in order to implement mobile code. To do this, every host must have the code for every label that it might be asked to execute; one conservative approximation simply gives every host access to all code in the program.

Closure conversion in this setting is conceptually standard, but is complicated by the presence of valid hypotheses and hypotheses at worlds other than the current world. For instance, the following (IL) function at world w_0 has a free variable at w_1 :

$$a:w_1 \text{ addr}@w_0, x:\text{int}@w_1 \vdash \lambda y.y + \text{get}[w_1;a]x : \text{int} \rightarrow \text{int}@w_0$$

To generate an environment for this function, we need to reify the remote hypothesis as a value. Because it does not make sense here, we do this by using the `at` modality.

The destination language of closure conversion is the same as for CPS, except for the typing rule for the `go` construct. For the sake of economy, we reuse the CPS language and add a new construct `go_cc` which we use in place of the old.

The closure conversion algorithm is given by three functions. One converts CPS types to CPS types. The only types that are affected by the translation are function (continuation) types. The type of a converted continuation will be an existential:

$$\llbracket A \text{ cont} \rrbracket_{\text{cc}} = \exists \alpha_{\text{env}}. \alpha_{\text{env}} \times ((\llbracket A \rrbracket_{\text{cc}} \times \alpha_{\text{env}}) \text{ cont})$$

The other two functions translate CPS expressions and values, respectively. For the translation of values, only lambdas are affected:³

$$\llbracket \lambda x. c \rrbracket_{\text{cc}} = \text{pack } \langle B_{\text{env}}, \langle v_{\text{env}}, \lambda p. c' \rangle \rangle \\ \text{as } \exists \alpha_{\text{env}}. \alpha_{\text{env}} \times ((\llbracket A \rrbracket_{\text{cc}} \times \alpha_{\text{env}}) \text{ cont})$$

where...

$$c' = \text{let } x = \pi_2 p \text{ in} \\ \text{let } e = \pi_1 p \text{ in} \\ \text{leta } \text{FV}(c)_1 = \pi_1 e \text{ in} \\ \vdots \\ \text{leta } \text{FV}(c)_n = \pi_n e \text{ in} \\ \text{letg } \text{FSV}(c)_1 = \pi_{n+1} e \text{ in} \\ \vdots \\ \text{letg } \text{FSV}(c)_m = \pi_{n+m} e \text{ in} \\ \llbracket c \rrbracket_{\text{cc}}$$

$$v_{\text{env}} = \langle \text{hold}_{\text{FVW}(c)_1}(\text{FV}(c)_1), \dots, \text{hold}_{\text{FVW}(c)_n}(\text{FV}(c)_n), \\ \text{cval}(\text{FSV}(c)_1), \dots, \text{cval}(\text{FSV}(c)_m) \rangle$$

$$B_{\text{env}} = \llbracket \text{FVT}(c)_1 \rrbracket_{\text{cc}} \text{ at } \text{FVW}(c)_1 \times \dots \times \llbracket \text{FVT}(c)_n \rrbracket_{\text{cc}} \text{ at } \text{FVW}(c)_n \times \\ \wp(\text{FSVT}(c)_1) \times \dots \times \wp(\text{FSVT}(c)_m)$$

n = Number of free regular variables in c .

$\text{FV}(c)_i$ = The i^{th} free regular variable of c .

$\text{FVT}(c)_i$ = The type of the i^{th} free regular variable of c .

$\text{FVW}(c)_i$ = The world of the i^{th} free regular variable of c .

m = Number of free valid variables in c .

$\text{FSV}(c)_i$ = The i^{th} free valid variable of c .

$\text{FSVT}(c)_i$ = The type of the i^{th} free valid variable of c .

Note that the translation of $\lambda x. c$, a value, is also a value. Because we have lasting restrictions on the bodies of *e.g.* `hold`, it is important to preserve the syntactic classes of values and expressions in these translations.

³Here we assume the language has n -ary products with projections π_i .

The translation of expressions only affects `call` and `go`:

$$\begin{aligned} \llbracket \text{call } v_f(v_x) \rrbracket_{\text{cc}} &= \text{unpack } \llbracket v_f \rrbracket_{\text{cc}} \text{ as } \langle \alpha_{\text{env}}, p \rangle \text{ in} \\ &\quad \text{let } f = \text{fst } p \text{ in} \\ &\quad \text{let } e = \text{snd } p \text{ in} \\ &\quad \text{let } a = \text{mkpair} \llbracket v_x \rrbracket_{\text{cc}}, e \text{ in} \\ &\quad \text{call } f(a) \\ \llbracket \text{go}[w; v_a]c \rrbracket_{\text{cc}} &= \text{go_cc}[w, \llbracket v_a \rrbracket_{\text{cc}}] \llbracket \lambda \cdot c \rrbracket_{\text{cc}} \end{aligned}$$

The translation of `call` is a standard closure call. We unpack the closure, build the argument to the function (which consists of the original argument along with the function's environment) and call it. The translation of `go` is necessary in order to explicitly represent the continuation as a closure, which is what is transmitted over the network in the implementation. The typing rule for `go_cc` is:

$$\frac{\Gamma \vdash v_a : w' \text{ addr@w} \quad \Gamma \vdash v_k : \exists \alpha_{\text{env}}. \alpha_{\text{env}} \times ((\langle \rangle \times \alpha_{\text{env}}) \text{ cont})@w'}{\Gamma \vdash \text{go_cc}[w'; v_a] v_k @w} \text{ go_cc}$$

Properties. Closure conversion should admit the following properties:

Property 2 (Soundness of closure conversion)

- (a) If $\Gamma \vdash v : A@w$ (in CPS)
then $\llbracket \Gamma \rrbracket_{\text{cc}} \vdash \llbracket v \rrbracket_{\text{cc}} : \llbracket A \rrbracket_{\text{cc}}@w$ (in CPS).
- (b) If $\Gamma \vdash c@w$ (in CPS)
then $\llbracket \Gamma \rrbracket_{\text{cc}} \vdash \llbracket c \rrbracket_{\text{cc}}@w$ (in CPS).
- (c) If $\Gamma \vdash v : A@w$ (in CPS)
then for every subterm v' of the form $\lambda x.c$ within $\llbracket v' \rrbracket_{\text{cc}}$,
 $\text{FV}(v') = \text{FSV}(v') = \emptyset$.
- (d) If $\Gamma \vdash c@w$ (in CPS)
then for every subterm v of the form $\lambda x.c'$ within $\llbracket c \rrbracket_{\text{cc}}$,
 $\text{FV}(v) = \text{FSV}(v) = \emptyset$.

Because of the relative difficulty of formalizing context operations like `FV` in `Twelf`, this proof has only been done on paper. Like Property 1, the proof was done for a similar language without a validity judgment or \exists modality.

6.5 Backend

The backend of the compiler transforms the low-level compiled program into typed assembly language that can be executed by the machine. For the `ML5` compiler, I will target the `TALT` system in active development by Crary *et al.* at Carnegie Mellon [12].

I have not yet designed the backend. Vanderwaart's resource-bounded Popcorn compiler [?] has a code generator that produces `TALT` code from a low-level typed intermediate language. I plan to try to reuse this code with minor modifications. Even if this turns out to be impossible, generation of typed

assembly language is a solved problem; there are only a few new challenges in this setting:

Marshalling. ML5 programs must be able to marshal arbitrary code and data into raw bytes for transmission on the network. By using a dictionary-passing scheme [44] we can arrange for marshalling and unmarshalling functions to be available whenever they are needed. Writing these marshalling functions in a type-safe way will take special care, but the problem has been studied before [14, 13].

Binding to Local Resources. More interesting will be the problem of identifying and binding to local resources. One of the primary motivations of the work is the controlled access to local resources (for example, a value `lineprinter : string → unit@w8108`). During unmarshalling, code that makes use of a local resource must have these open references bound to those local resources. The problem is slightly different from marshalling, because in general, the host that prepares the marshalled code can only name those resources—it cannot make sense of them in the same way it can make sense of common data like tuples and portable code.

I solved a very similar problem in the Hemlock compiler with an easy ad hoc solution. Generalizing that technique to this case would mean providing a uniform table on each host with one entry for every local resource that is used in the program. The table is initialized by the runtime. On the host where the resource is present, the table entry contains a tagged reference to the resource, and accesses to it succeed. Accesses to non-existent resources fail with a runtime error. Although this scheme would work for the ML5 implementation as well, it is worth investigating how to avoid this run-time check. However, note that even if the implementation incurs the cost of a run-time check, the type system guarantees that such a check will never fail.

Allocation and Valuability. Allocation conversion [26] is a pass that makes the allocation and initialization of memory explicit. The standard way of allocating the tuple $\langle 830, 1 \rangle$, for example, is to generate code such as the following:⁴

```
let x = malloc(8, int * int)
let _ = x[0] := 830
let _ = x[1] := 1
```

Several constructs in the internal language, such as `hold` and `cval`, have a syntactic restriction to values that is necessary for safety. Many values, however, require allocation. This makes those values into effectful expressions, such that those programs can no longer be typed in our language! One solution is to relax the value restriction to one of *valuability*. A “valuable” expression may have

⁴The actual code would be more complicated, mostly due to type systems for initialization, but that is not at issue here.

effects, but it can safely be evaluated by one world on behalf of another. I have only sketched such a system on paper.

By postponing allocation conversion until after translation into a language like TALT that is not modally typed, we can sidestep this issue completely. Still, it is interesting to figure out what the right thing to do in this case is, even if I do not actually implement it.

7 The ML5 Runtime

Related to the backend of the compiler is the language runtime. The runtime is the software that enables a compiled program to run by mediating between it and the environment. In the case of a distributed programming language, this includes forming and maintaining the network on which the program runs.

The ML5 runtime will be similar to the ConCert Conductor [27], and I expect to be able to reuse much of the network code developed for that software. In many ways it will be simpler, because it does not need to support features like automatic failure recovery and peer-to-peer node discovery. However, there are again some special concerns for ML5:

Concurrency. A logical account of concurrency is firmly outside the scope of this thesis. However, some minimal support for concurrency will likely be necessary to develop a compelling distributed application. A simple fork/join mechanism like futures (as used in Section 5) may suffice. Typed CML-style channels [36] might be easily implementable using the same marshalling mechanisms already planned.

Any of these should integrate cleanly with the spatial distribution features of ML5, the implementation burden being the main cost. The demands of the application should dictate which features I include.

Garbage Collection. In order to support long-running programs, the runtime must reclaim abandoned storage. The TALT runtime already includes a conservative garbage collector for locally allocated data. However, as references to local data spread (in the environments of mobile code) to other places in the network, we cannot safely reclaim these data without ensuring that all global references have been discarded. Distributed garbage collection is a problem that has been studied, but the solutions are quite complicated [35]. At a minimum, it may be feasible to simply never reclaim data that may have non-local references. If not, then the most complicated scheme I would attempt would be very naïve, stopping the entire world in order to conduct a scan for data with no global references.

8 Application

I plan to demonstrate the effectiveness of ML5 and its implementation by building an example distributed application. The application will be chosen to show off the features that I have time to implement, so I have not committed to a specific application yet. Optimally, the application I choose would have the following characteristics. Most of its complexity should come from the way it is distributed, not from the particular computations it does at each host. It should manage location-dependent resources and effects in a non-trivial way, and should make use of the higher-order programming techniques enabled by ML5's type system. Finally, a similar application should have an implementation in at least one related system, so that the two can be compared.

9 Related Work

Distributed computing is a huge area of computer science, suggesting a large body of related work. Most of this work is only superficially relevant to this proposal, and the relationship can be summarized in terms of a few major themes. A few more closely related languages are described in detail below.

A Language, not a Tool. ML5 is an abstract language designed for spatially distributed computing on a foundation of logic. In contrast, many systems used for distributed computing are ad hoc extensions of existing languages. In many languages used in industry, the mechanisms for distributed computing are provided in the form of libraries and stub generators (*e.g.* CORBA [46, 41], DCOM [7], .Net [31], Web Services [43, 42]).

By building ML5 from first principles, and by using the Curry-Howard Isomorphism to guide its design, we are able to isolate key concepts of spatial distribution and codify them in a natural way within the type system. The result is a language that is more elegant.

Marshalling and Data Mobility. A common shortcoming of distributed languages and tools is that they place restrictions on the types of data that can be marshalled and sent to other hosts. CORBA, whose basic primitive for distributed programming is the remote procedure call, limits the types of RPC arguments to a few simple base types. In Java RMI [15], objects matching the `Serializable` interface can be marshalled, but this does not include references to local resources such as file handles. I contend that this restriction stems from a mistaken conflation of two separate notions. First is the implementation technique of marshalling, which is simply a way of representing a piece of data in a format suitable for transmission on a network. The second is the potential mobility of data, that is, the fact that some data make sense at more than one world. In most languages these ideas are conflated because distributed applications essentially consist of a collection of separate programs, each of which

can only understand data from its own perspective. When a Java program performs a remote procedure call, the function that is called can't help but require that all its arguments make sense to it, because the only notion of "making sense" is one of "making sense locally." Therefore, every remote procedure call requires that the arguments shift from making sense to the caller to making sense at the receiver. Because this shift always occurs at the same time as marshalling, they seem like the same operation.

The modal type system of ML5 makes the distinction easy to see, because it endows programs with the ability to simultaneously reason about the well-formedness of values from multiple perspectives. An ML5 program may observe that a value it holds also makes sense at another world, and therefore can coerce it to a value well-typed at that world. If it then performs a remote procedure call, it may pass along that value, knowing it makes sense in the called procedure's world. A program can also pass along a value that only makes sense to it—in which case the called procedure knows that the value it has received is only well-typed at the caller's world. In either case, marshalling is simply an implementation technique; any value can be marshalled for transmission, even if it does not make sense on one end of the transaction. This permits more flexibility in the way programs are written, particularly when programming with higher-order functions.

The Location of Effects. Another feature of ML5's modal typing discipline is that it is able to directly control access to local resources in a statically checked way. I believe that no other distributed programming language features this ability. For instance, in Java RMI almost all checks happen dynamically, which means that more bugs can survive until run-time. These run-time checks include the dynamic serializability of arguments or return types of methods that are invoked remotely, as well as any checks that the Java program must make to ensure that local resources are present when it needs them. For an example of the latter: when implementing a client/server database system, it is possible for the client to instantiate a local instance of the server object

```
Server serv = new Server();
```

instead of getting a remote instance via RMI

```
Server serv =  
    (Server)Naming.lookup("rmi://server.org/dbserv");
```

Unless special precautions are taken, this compiles and works until the server code attempts to find the local database on the client computer, which fails. Similarly, the client can call any public helper code or methods on objects that are intended for use only by the server. If the programmer wishes to exclude such erroneous programs, it is often possible to use Java mechanisms like `private` fields and inner classes to control access to location-dependent resources. To do so, one must identify the concept of location with the concept of class membership (the only notion of principal in the language). In contrast,

ML5 provides a type structure that allows the programmer to directly express these location dependencies. I claim that this is exactly the relevant concept for spatially distributed programs.

9.1 Distributed ML-Like Languages

Several lines of research have extended ML-like languages with explicit support for distributed computing. Some present a very different abstraction of the network; for instance, D’Caml’s model [45] is a single parallel computer with a unified store. This model solves the problem of controlling access to localized resources by making every resource available everywhere through the automatic creation of proxies. This has the disadvantage of obscuring the performance characteristics of the code, because any read, write, or even copy of a memory location can require a request over the network. ML5 is lower level in that it attempts to describe the spatial distribution of data with a type structure rather than abstract away from it.

Acute. Acute [39] is a new ML-like language for distributed computing. The model it supports is more difficult than ours: a dynamic linking setting where interacting programs across hosts may be upgraded to different versions while the program is running. The general mechanism in Acute for controlling locality is its dynamic rebinding system. This allows the programmer to mark libraries so that references to them are rebound to local versions when code arrives at a site. In ML5, every variable or label refers to a single binding site, be it bound in the program or in the initial context that describes the local resources in the network. This improves the programmer’s ability to reason statically about programs. On the other hand, dynamic binding is a reality of many distributed applications. It may be possible to explicitly implement dynamic binding on top of ML5’s static type system, which would be interesting future work.

Like several other languages in this category, Acute allows for the marshalling of arbitrary values, and because of the dynamic framework must take special care to preserve abstraction between distributed modules and module versions. In comparison, because we check the entire program during a single static link, we are able to treat abstraction as a high-level issue that is simply not present when running machine code.

Alice. Alice [37, 1] is a project to build a new SML-like language with many additional features, including distribution. In Alice, marshalling is called *pickling*, and the story is similar to Java: Stateful values such as arrays are freshly copied, and pickling fails at runtime for “sited” values (those that contain site-specific resources like file handles or thread IDs). As in Java, sitedness is not evident from the type of a value. Hosts have the ability to publish picklable values for receipt by other hosts, and to spawn new threads of control. As

usual, there is no typing support of the kind we offer, in fact, “you must be careful not to use sited values in a service!” [2] or run-time errors may occur.

9.2 Modal Logic in Distributed Computing

Other lines of research have used modal logic (in various forms) for distributed computing. Borghuis and Feijs’s Modal Type System for Networks [6] provides a logic and operational semantics⁵ for network tasks with stationary services and mobile data. They use \Box , 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 . With no way of internalizing mobility as a proposition, the calculus limits mobile data to base types. Services are similarly restricted to depth-one arrow types. By using \Box for mobile code and \Diamond or at for stationary resources, we believe our resulting calculus is both simpler and more general.

Cardelli and Gordon [10] provide an early example of using modal logic for reasoning about programs spatially, later refined by Caires and Cardelli [8, 9]. They do not take a propositions-as-types view of their logic; instead, they start from a process calculus, mobile ambients [10], and develop a classical logic for reasoning about their behaviors. Therefore, their modal logic is very different from intuitionistic S5 and includes connectives for stating temporal properties, security properties, and properties of parallel compositions. In contrast, the proposed work gives a computational interpretation of the logic as a type system in the tradition of ML.

Hennessey et al. [18] developed a distributed version of the π -calculus and impose a complex static type system in order to constrain and describe behavior. Similarly, Schmitt and Stefani [38] develop a distributed, higher-order version of the Join Calculus with a complex behavioral type system. In comparison, our system is much simpler, eliminating the complexities of concurrency, access control, and related considerations. By basing our system on the Curry-Howard correspondence, we have a purely logical analysis and, furthermore, propose a straightforward path for integration into a full-scale functional language for realistic programs.

Moody [24] gives a system based on the constructive modal logic S4 due to Pfenning and Davies [34]. In contrast to our explicit world formalism, this language is based on judgments A true (here), A poss (somewhere), and A valid (everywhere). 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 model. From the standpoint of a multiple world semantics, the accessibility relation of S4 satisfies only reflexivity and transitivity, not symmetry. From the computational point of view, accessibility in his work describes process interdependence rather than connections between actual network locations. Programs are therefore somewhat higher-level and express *potential mobility* instead of explicit code motion as in the `get` construct. In particular, due to the lack of symmetry it is not

⁵By way of compilation into shell scripts!

possible to go back to a source world after a potentially remote procedure call except by returning a value.

Jia and Walker [20] give a judgmental account of an explicit world S5-based logic that is very similar to ours. The main difference is in the envisioned operational semantics. Theirs is a process calculus admitting “action-at-a-distance,” where a primitive notion is the passive synchronization of two processes: one that waits for the value of a label, and one that computes it. In comparison, we have designed a semantics that we believe is more amenable to standard functional programming style. Our decomposition of the rules for the modalities means that we have only one primitive that induces communication between nodes, and it has a simple, active, operational behavior. To our knowledge, there is no implementation of their calculus as a programming language.

10 Conclusion

I have presented a plan for the design and construction of a language for spatially distributed computing. The language is designed from first principles on a foundation of logic, using the modal logic Intuitionistic S5 for its ability to represent spatial reasoning. Through its modal type system, ML5 is able to statically control access to location-sensitive resources. The proposed work is novel, feasible, and has both theory and practice components. I conclude with a rough timeline for the completion of the thesis.

Timeline. This project consists of a substantial body of theory (a new language design and associated logics), and implementation (a compiler and distributed runtime). For the purposes of the thesis, these parts may be of varying quality. For instance, a high-performance ML compiler alone can require dozens of man-years of work, but is not required to support the thesis. For the work to be done, I adopt a taxonomy that prioritizes tasks into the following categories:

- **Primary**, meaning that this is a central contribution of original work, deserving maximal attention
- **Secondary**, meaning an interesting and original contribution that is not necessary, but that would strengthen the thesis. Many of these items will become “future work”
- **Tarpit**, meaning a notoriously difficult or involved problem, and so I will only do the bare minimum to make the system usable, leaving the problem open for “future work”
- **Solved**, meaning a task that is well-understood and standard, and so I will do only the bare minimum to make the system usable (or re-use existing technology)

The following is a plan for completing the proposed work, annotated with a basic strategy according to the above. Rather than give time estimates, I give approximate fractions of the total time.

Theory: Language and Logic Design	
<p>Complete Twelf proofs and add other types. <i>Primary.</i> What needs to be done here is to work the validity judgment through the phases of the formalized compiler (elaboration, CPS conversion). Since I have worked these cases out on paper, I do not foresee any problems. Adding sums and recursive types (etc.) to the language and proofs will be standard.</p>	6%
<p>Network signatures and separate compilation. <i>Secondary.</i> The language as presented does not provide a way of specifying the initial configuration of resources in the network. The language should be extended to give a way to describe the local resources at other hosts, and to compile and link against this interface to them. My recent work on a separate compilation system for Standard ML [?] should help provide guidance in designing this.</p>	6%?
Implementation: Compiler	
<p>Parser, elaboration, CPS, closure conversion. <i>Primary.</i> I can re-use some of the code from Hemlock (for instance, the pattern matcher) here.</p>	19%
<p>Standard IL Optimizations. <i>Solved.</i> I will spend a small amount of time implementing the optimizations that make the biggest difference for the least effort (dead code elimination), but this is only to make the compiler suitable for reasonable applications.</p>	2%
<p>Optimizations of mobile primitives. <i>Secondary.</i> It may be important to do special optimizations for our domain (closure representation, dead and constant gets, etc.). Time permitting, I would like to work on this, since it is likely that there is much low-hanging fruit.</p>	2%?

<p>Certifying back-end. <i>Solved.</i> I would like to re-use Joe Vanderwaart’s popcorn TALT compiler as the back-end, if possible. Although I will need to make some changes to support primitives necessary for ML5-specific features (marshalling), I hope that these will be relatively minor, and that the bulk of boilerplate code will be taken care of.</p>	13%
Implementation: Runtime	
<p>Basic Infrastructure. <i>Primary.</i> Some basic infrastructure is required to run ML5 programs. This would allow a program to arrive on the target machine, be certified, and executed. Although we cannot use the ConCert v2 system for this, much of that code can be repurposed to build the ML5 runtime. (This runtime is also much simpler, barring the points below, because it does not need to implement work-stealing, P2P node discovery, a distributed hash table, or fault tolerance.) Still, this will be a significant implementation effort.</p>	13%
<p>Concurrency support. <i>Tarpit.</i> Good support for concurrency opens up the system to a large number of design and implementation issues which cannot reasonably fit within the time frame, and which are anyway outside the scope of the thesis. Some basic concurrency support may however be necessary to write a compelling example application.</p>	∞%
<p>Failure tolerance. <i>Tarpit.</i> We have designed a fault-tolerant network with the ConCert v2 system. However, the ideas are largely incompatible here because we support effects; fault tolerance would need to be able to log, roll-back, and replay effects. While there are some known techniques for this (such as message logging [21]), they would make the implementation substantially more complex than it needs to be to support the thesis.⁶</p>	∞%

⁶Some very simple support for enabling application authors to handle failure may be possible by simply raising a handleable exception when a `get` does not succeed?

Garbage collection.	6%
<i>Tarpit.</i> This is a difficult problem, but a basic solution will be necessary to make the system usable. I will begin with the simplest possible collector and, when designing the runtime, make sure not to preclude more advanced collectors should the need for them arise.	
Application	
Application.	6%
<i>Primary.</i> The primary difficulty will be finding an interesting problem to solve that does not require sophisticated concurrency or fault tolerance, since I want to avoid those.	
Dissertation	
Writing.	26%

References

- [1] Alice web site, 2005. URL: <http://www.ps.uni-sb.de/alice/>.
- [2] Alice web site: Distribution, 2005. URL: <http://www.ps.uni-sb.de/alice/manual/distribution.html>.
- [3] David P. Anderson. BOINC: A system for public-resource computing and storage. In *Proceedings of the 5th IEEE/ACM International Workshop on Grid Computing*, Pittsburgh, USA, November 2004.
- [4] Andrew Appel. *Compiling With Continuations*. Cambridge University Press, Cambridge, 1992.
- [5] Robert D. Blumofe and Philip A. Lisiecki. Adaptive and reliable parallel computing on networks of workstations. In *USENIX 1997 Annual Technical Conference on UNIX and Advanced Computing Systems*, pages 133–147, Anaheim, California, 1997.
- [6] 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.
- [7] D. Box. *Understanding COM*. Addison-Wesley, Reading, MA, 1997.
- [8] Luís Caires and Luca Cardelli. A spatial logic for concurrency (part I). In *Theoretical Aspects of Computer Software (TACS)*, pages 1–37. Springer-Verlag LNCS 2215, October 2001.
- [9] Luís Caires and Luca Cardelli. A spatial logic for concurrency (part II). In *Proceedings of the 13th International Conference on Concurrency Theory (CONCUR)*, pages 209–225, Brno, Czech Republic, August 2002. Springer-Verlag LNCS 2421.
- [10] Luca Cardelli and Andrew D. Gordon. Anytime, anywhere. modal logics for mobile ambients. In *Proceedings of the 27th Symposium on Principles of Programming Languages (POPL)*, pages 365–377. ACM Press, 2000.

- [11] B. Chang, K. Crary, M. DeLap, R. Harper, J. Liszka, T. Murphy VII, and F. Pfenning. Trustless grid computing in ConCert. In M. Parashar, editor, *Grid Computing – Grid 2002 Third International Workshop*, pages 112–125, Berlin, November 2002. Springer-Verlag.
- [12] Karl Crary. Toward a foundational typed assembly language (expanded version). Technical Report CMU-CS-02-196, Carnegie Mellon University, 2002.
- [13] Karl Crary, Michael Hicks, and Stephanie Weirich. Safe and flexible dynamic linking of native code. In *2000 ACM SIGPLAN Workshop on Types in Compilation*, September 2000.
- [14] Karl Crary and Stephanie Weirich. Flexible type analysis. In *International Conference on Functional Programming*, pages 233–248, 1999.
- [15] Jim Farley. *Java Distributed Computing*. O’Reilly, January 1998.
- [16] Ian Foster and Carl Kesselman. The Globus project: a status report. *Future Generation Computer Systems*, 15(5–6):607–621, 1999.
- [17] Ian Foster and Carl Kesselman, editors. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, San Francisco, California, 1999.
- [18] Matthew Hennessy, Julian Rathke, and Nobuko Yoshida. SafeDPi: A language for controlling mobile code. Report 02/2003, Department of Computer Science, University of Sussex, October 2003.
- [19] Hybrid logics bibliography, 2005. URL: <http://hylo.loria.fr/content/papers.php>.
- [20] Limin Jia and David Walker. Modal proofs as distributed programs (extended abstract). *European Symposium on Programming*, 2004.
- [21] David B. Johnson and Willy Zwaenepoel. Recovery in distributed systems using optimistic message logging and checkpointing. In *Proc. 7th Annual ACM Symp. on Principles of Distributed Computing*, pages 171–181, Toronto (Canada), 1988.
- [22] Clarence Irving Lewis. *A Survey of Symbolic Logic*. University of California Press, 1918.
- [23] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, Cambridge, Massachusetts, 1997.
- [24] Jonathan Moody. Modal logic as a basis for distributed computation. Technical Report CMU-CS-03-194, Carnegie Mellon University, October 2003.
- [25] Greg Morrisett. *Compiling with Types*. PhD thesis, Carnegie Mellon University, December 1995. Available as CMU Technical Report CMU–CS–95–226.
- [26] Greg Morrisett, David Walker, Karl Crary, and Neal Glew. From System F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3):527–568, May 1999.
- [27] Tom Murphy, VII. Functional grid programming with ConCert, 2004. Unpublished draft; <http://tom7.org/papers/>.
- [28] Tom Murphy, VII, Karl Crary, and Robert Harper. Distributed control flow with classical modal logic. In Luke Ong, editor, *14th Annual Conference of the European Association for Computer Science Logic (CSL 2005)*, Lecture Notes in Computer Science. Springer, August 2005.
- [29] Tom Murphy, VII, Karl Crary, Robert Harper, and Frank Pfenning. A symmetric modal lambda calculus for distributed computing. In *Proceedings of the 19th IEEE Symposium on Logic in Computer Science (LICS 2004)*. IEEE Press, July 2004.

- [30] George C. Necula. Proof-carrying code. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '97)*, pages 106–119, Paris, January 1997.
- [31] .Net web site, 2005. URL: <http://msdn.microsoft.com/netframework/>.
- [32] Sungwoo Park. A modal language for the safety of mobile values. Technical Report CMU-CS-05-124, Carnegie Mellon, Pittsburgh, Pennsylvania, USA, May 2005.
- [33] Leaf Petersen. *Certifying Compilation for Standard ML in a Type Analysis Framework*. PhD thesis, Carnegie Mellon University, May 2005. Available as CMU Technical Report CMU-CS-05-135.
- [34] 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.
- [35] David Plainfossé and Marc Shapiro. A survey of distributed collection techniques. Technical report, BROADCAST, 1994.
- [36] John Reppy. *Concurrent Programming in ML*. Cambridge University Press, 1999.
- [37] Andreas Rossberg, Didier Le Botlan, Guido Tack, Thorsten Brunklau, and Gert Smolka. *Alice Through the Looking Glass*, volume 5 of *Trends in Functional Programming*. Intellect, Munich, Germany, 2004.
- [38] Alan Schmitt and Jean-Bernard Stefani. The M-calculus: A higher-order distributed process calculus. In *Conference Record of the 30th Symposium on Principles of Programming Languages*, pages 50–61, New Orleans, Louisiana, January 2003. ACM Press.
- [39] Peter Sewell, James J. Leifer, Keith Wansbrough, Francesco Zappa Nardelli, Mair Allen-Williams, Pierre Habouzit, and Viktor Vafeiadis. Acute: high-level programming language design for distributed computation. In *International Conference on Functional Programming (ICFP) 2005*, Tallinn, Estonia, September 2005.
- [40] Alex Simpson. *The Proof Theory and Semantics of Intuitionistic Modal Logic*. PhD thesis, University of Edinburgh, 1994.
- [41] Steve Vinoski. CORBA: integrating diverse applications within distributed heterogeneous environments. *IEEE Communications Magazine*, 14(2), 1997.
- [42] W3C Web Service Description Working Group, David Booth and Canyang Kevin Liu, eds. Web services description language (WSDL) version 2.0 part 0: primer. Technical Report WD-wsdl20-primer-20050510, W3C, URL: <http://www.w3.org/TR/wsdl20-primer>, May 2005.
- [43] W3C web services working group, 2002. URL: <http://www.w3.org/2002/ws/>.
- [44] P. Wadler and S. Blott. How to make ad-hoc polymorphism less ad-hoc. In *Conference Record of the 16th Annual ACM Symposium on Principles of Programming Languages*, pages 60–76. ACM, January 1989.
- [45] Ken Wakita, Takashi Asano, and Masataka Sassa. D'Caml: A native distributed ML compiler for heterogeneous environment. In Patrick Amestoy, Philippe Berger, Michel J. Daydé, Iain S. Duff, Valérie Frayssé, Luc Giraud, and Daniel Ruiz, editors, *Euro-Par 1999 Parallel Processing, 5th International Euro-Par*

Conference, volume 1685 of *Lecture Notes in Computer Science*, pages 914–924. Springer, 1999.

- [46] Zhonghua Yang and Keith Duddy. CORBA: A platform for distributed object computing (a state-of-the-art report on OMG/CORBA). *Operating Systems Review*, 30(2):4–31, 1996.

A Intuitionistic S5

(a)	
$\frac{\Gamma, \omega \text{ world} \vdash M : A@w}{\Gamma \vdash \text{box } \omega.M : \Box A@w} \Box I$	$\frac{\Gamma \vdash M : \Box A@w}{\Gamma \vdash \text{unbox } M : A@w} \Box E$
$\frac{\Gamma \vdash A \text{ true}@w'}{\Gamma \vdash \Diamond A \text{ true}@w} \Diamond I$	$\frac{\Gamma \vdash \Diamond A \text{ true}@w' \quad \Gamma, \omega \text{ world}, A \text{ true}@w \vdash C \text{ true}@w}{\Gamma \vdash C \text{ true}@w} \Diamond E$
$\frac{\Gamma, A \text{ true}@w \vdash B \text{ true}@w}{\Gamma \vdash A \supset B \text{ true}@w} \supset I$	$\frac{\Gamma \vdash A \supset B \text{ true}@w \quad \Gamma \vdash A \text{ true}@w}{\Gamma \vdash B \text{ true}@w} \supset E$
$\frac{\Gamma \vdash A \text{ true}@w \quad \Gamma \vdash B \text{ true}@w}{\Gamma \vdash A \wedge B \text{ true}@w} \wedge I$	$\frac{\Gamma \vdash A \wedge B \text{ true}@w}{\Gamma \vdash A \text{ true}@w} \wedge E_1$
$\frac{\Gamma \vdash A \wedge B \text{ true}@w}{\Gamma \vdash B \text{ true}@w} \wedge E_2$	$\frac{\Gamma \vdash A \text{ true}@w}{\Gamma \vdash A \vee B \text{ true}@w} \vee I_1$
$\frac{\Gamma \vdash B \text{ true}@w}{\Gamma \vdash A \vee B \text{ true}@w} \vee I_2$	$\frac{\Gamma \vdash A \vee B \text{ true}@w' \quad \Gamma, A \text{ true}@w' \vdash C@w \quad \Gamma, B \text{ true}@w' \vdash C@w}{\Gamma \vdash C \text{ true}@w} \vee E$
(b)	
$\frac{\Gamma \vdash A \text{ true}@w'}{\Gamma \vdash A \text{ at } w' \text{ true}@w} \text{ at } I$	$\frac{\Gamma \vdash A \text{ at } w \text{ true}@w'}{\Gamma \vdash A \text{ true}@w} \text{ at } E$

Figure 6: (a) The standard presentation of Intuitionistic S5, as used by *e.g.* Simpson [40]. (b) The extension to add the at modality, as in *e.g.* Jia and Walker [20].

B CPS Language

$\mathbf{w}::\text{leta } x = \text{hold}_w v \text{ in } c$	\mapsto	$\mathbf{w}::[v/x]c$
$\mathbf{w}::\text{letg } u = \text{cval } v \text{ in } c$	\mapsto	$\mathbf{w}::[v/u]c$
$\mathbf{w}::\text{lift } u = v \text{ in } c$	\mapsto	$\mathbf{w}::[v/u]c$
$\mathbf{w}::\text{let } x = \text{fst } \langle v_1, v_2 \rangle \text{ in } c$	\mapsto	$\mathbf{w}::[v_1/x]c$
$\mathbf{w}::\text{let } x = \text{snd } \langle v_1, v_2 \rangle \text{ in } c$	\mapsto	$\mathbf{w}::[v_2/x]c$
$\mathbf{w}::\text{let } x = \text{mkpair } v_1, v_2 \text{ in } c$	\mapsto	$\mathbf{w}::[\langle v_1, v_2 \rangle / x]c$
$\mathbf{w}::\text{let } x = \text{localhost}() \text{ in } c$	\mapsto	$\mathbf{w}::[\bar{\mathbf{w}}/x]c$
$\mathbf{w}::\text{go}[\mathbf{w}_1; \bar{\mathbf{w}}_1] c$	\mapsto	$\mathbf{w}_1::c$
$\mathbf{w}::\text{call } (\lambda x. c)(v)$	\mapsto	$\mathbf{w}::[v/x]c$
$\mathbf{w}::\text{halt}$	\mapsto	$\mathbf{w}::\text{halt}$

Figure 7: The dynamic semantics for the CPS language. It is given in terms of a stepping relation \mapsto between pairs of CPS machines of the form $\mathbf{w}::c$, where \mathbf{w} is the current world of evaluation and c is the current continuation expression. We give only the rules for the constructs that appear in Figure 3. The expression `halt` has a self-transition for the ease of stating type safety (Property 3)

Property 3 (Type Safety of CPS Language)

- (a) *If $\cdot \vdash c @ w$ (in CPS)
then there exists w', c' such that
 $w::c \mapsto w'::c'$.*
- (b) *If $\cdot \vdash c @ w$ (in CPS)
and $w::c \mapsto w'::c'$
then $\cdot \vdash c' @ w'$.*