

A Modal Analysis of Staged Computation

Rowan Davies
and
Frank Pfenning
Carnegie Mellon University

We show that a type system based on the intuitionistic modal logic S4 provides an expressive framework for specifying and analyzing computation stages in the context of typed λ -calculi and functional languages. We directly demonstrate the sense in which our $\lambda_e^{\rightarrow\Box}$ -calculus captures staging, and also give a conservative embedding of Nielson & Nielson's two-level functional language in our functional language Mini-ML $^{\Box}$, thus proving that binding-time correctness is equivalent to modal correctness on this fragment. In addition, Mini-ML $^{\Box}$ can also express immediate evaluation and sharing of code across multiple stages, thus supporting run-time code generation as well as partial evaluation.

Categories and Subject Descriptors: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—*lambda calculus and related systems, modal logic, proof theory*; F.3.3 [Logics and Meanings of Programs]: Studies of Program Constructs—*type structure*; F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages—*partial evaluation*; D.3.3 [Programming Languages]: Language Constructs and Features

General Terms: Languages, Theory

Additional Key Words and Phrases: Staged computation, binding times, run-time code generation

1. INTRODUCTION

Dividing a computation into separate stages is a common informal technique for the derivation of algorithms [Jørring and Scherlis 1986]. For example, instead of directly matching strings against a regular expression we may first compile the regular expression into a finite automaton and then execute the same automaton on different strings. Because significant efficiency gains can often be realized, there is

This is an extended and revised version of the conference paper [Davies and Pfenning 1996].

This work was sponsored in part by the National Science Foundation under grant CCR-9619832 and by the Advanced Research Projects Agency (ARPA) under Order No. 8313.

The first author was partly supported by a Hackett Studentship from the University of Western Australia. Part of this work was completed during a visit by the first author to BRICS (Basic Research in Computer Science, Centre of the Danish National Research Foundation). Address: 5000 Forbes Ave, Pittsburgh, Pennsylvania 15213-3891

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works, requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept, ACM Inc., 1515 Broadway, New York, NY 10036 USA, fax +1 (212) 869-0481, or permissions@acm.org.

a substantial body of work concerned with the automation of staged computation. *Partial evaluation* (see, for example, [Jones et al. 1993]) divides the computation into two stages based on the early availability of some function arguments. In practice this appears most successful when supported by *binding-time analysis* [Gomard and Jones 1991], which statically determines which parts of a computation may be carried out in the first phase, and which parts remain to be done in the second phase.

It often takes considerable ingenuity to write programs in such a way that they exhibit proper binding-time separation, that is, that the computation intended to occur when the early arguments become available can in fact be carried out. From a programmer's point of view it is therefore desirable to declare the expected binding-time separation and obtain constructive feedback when the computation may not be staged as expected. This suggests that the binding-time properties of a function should be expressed in a prescriptive type system, and that binding-time analysis should be a form of type checking. The work on two-level functional languages [Nielson and Nielson 1992] and some work on partial evaluation (for example, [Gomard and Jones 1991]) shows that this view is indeed possible.

Up to now these type systems have been motivated *algorithmically*, that is, they are explicitly designed to support specialization of a function to its early arguments. In this paper we show that they can also be motivated *logically*, and that the proper logical system for expressing computation stages is the intuitionistic variant of the modal logic S4 [Prawitz 1965]. This observation immediately gives rise to a natural generalization of standard binding-time analysis by allowing multiple computation stages, initiation of successor stages, and sharing of code across multiple stages. Such extensions are normally considered external issues. For example, Jones [Jones 1992] describes a typed framework for such concepts, but only at the level of operations on whole programs. Our framework instead provides these operations *within* the language of programs. This makes our approach particularly relevant to run-time code generation, where specialization takes place when the program is executed. Indeed, the authors and others have designed and implemented an extension of ML based on the type system described here which generates and executes abstract machine code at run time [Wickline et al. 1998; Wickline et al. 1998].

One of our conclusions is that when we extend the Curry-Howard isomorphism between proofs and programs from intuitionistic logic to the intuitionistic modal logic S4 we obtain a natural and logical explanation of computation stages. The isomorphism relates proofs in modal logic to functional programs which manipulate program fragments for later stages. Each world in the Kripke semantics of modal logic corresponds to a stage in the computation, and a term of type $\Box A$ corresponds to code to be executed in a future stage of the computation. The modal restrictions imposed on terms of type $\Box A$ guarantee that a function of type $B \rightarrow \Box A$ can carry out all computation concerned with its argument while generating the residual code of type A .

We begin by considering $\lambda_e^{\rightarrow\Box}$, a modal λ -calculus based on a natural-deduction formulation of intuitionistic modal S4. The presentation is new, but draws on ideas in [Bierman and de Paiva 1996; Pfenning and Wong 1995; Girard 1993].

We then construct a functional language Mini-ML $_e^\Box$ by augmenting $\lambda_e^{\rightarrow\Box}$ with a fixpoint operator, natural numbers, and pairs and endow it with a natural call-by-

value operational semantics along the lines of Mini-ML [Clément et al. 1986].

Mini-ML_e[□] can be somewhat awkward because it often requires a broad syntactic structuring of the program to directly reflect staging. This simplifies the study of staging properties of Mini-ML_e[□], but it also makes it difficult to directly relate it to previous work on staged languages, such as two-level languages [Nielson and Nielson 1992]. We thus consider a more implicit formulation of S4 motivated by its Kripke semantics following [Martini and Masini 1996; Pfenning and Wong 1995] and then augment it as before to form Mini-ML[□]. With some syntactic sugar, Mini-ML[□] is intended to serve as the basis for a conservative extension of ML with practical means to express and check staging of computation. The operational semantics of Mini-ML[□] is given by a type-preserving translation to Mini-ML_e[□] whose correctness is not entirely trivial. This translation also describes the first phase of a plausible compilation strategy for Mini-ML[□] for run-time code generation.

We then exhibit a simple full and faithful embedding of Nielson & Nielson’s two-level language [Nielson and Nielson 1992] in Mini-ML[□], providing further evidence that Mini-ML[□] provides an intuitively appealing, technically correct, and logically motivated view of staged computation.

2. A MODAL λ-CALCULUS

In this section we present the modal λ-calculus $\lambda_e^{\rightarrow\Box}$. We start by directly motivating the calculus in terms of “manipulation of code” and relate this to modal logic. We then present typing rules based on a natural deduction system for modal S4, give β and η rules for the modal \Box operator, and show that they satisfy subject reduction and expansion, respectively. We also demonstrate the relationship between $\lambda_e^{\rightarrow\Box}$ and computation staging via two theorems. We assume some familiarity with logical concepts such as modal logic, natural deduction, and the Curry-Howard isomorphism. We refer the reader to some of our recent work [Pfenning and Davies 2000] for a more basic discussion of these concepts for and to Chellas [Chellas 1980] for an introduction to modal logic.

2.1 Natural Deduction for Validity

A common feature of many forms of staged computation is the manipulation of code. Macro expanders and partial evaluators typically manipulate source expressions, run-time code generators typically manipulate object code or some form of intermediate code. To show how such manipulation of code may be accounted for in a typed framework, we start with a typed λ-calculus and introduce a new type constructor \Box , where $\Box A$ represents *code of type A*. This type remains abstract in the sense that we do not commit ourselves to a particular way of implementing it. In this way our type system can support diverse applications.

Next we have to decide which operations should be supported on code. First, we should be able to manipulate an arbitrary *closed* expression as code. This suggests a constructor **box** where **box** $E : \Box A$ if $E : A$ in the empty context. This is essentially the modal rule of necessitation. The second means of constructing code is by *substitution*: we can substitute code for a free variable appearing in code to obtain code. In a meaningful type system such substitution must be “hygienic” and rename bound variables if necessary to avoid capture. The restriction that we can only substitute *code* (and not arbitrary expressions) into code is reflected exactly

in one of Prawitz’s variants of the modal necessitation rule [Prawitz 1965]: we can infer that $\mathbf{box} E : \Box A$ from $E : A$ if all hypotheses of the latter derivation are of the form $x : \Box B$. This means that every free variable x in E must have a type of the form $\Box B$. Prawitz’s elimination rule allows us to infer A from $\Box A$. In terms of the functional interpretation, this suggest *evaluation*: we execute the *code* of type A to obtain a *value* of type A .

Unfortunately, the natural deduction formulation of modal logic based on these two rules does not obey subject reduction, as noticed by Prawitz himself and demonstrated by the following derivation:

$$\frac{\frac{\frac{\frac{\overline{\Box A \vdash \Box A}}{B, B \rightarrow \Box A, \Box A \vdash \Box \Box A}}{B, B \rightarrow \Box A \vdash \Box A \rightarrow \Box \Box A}}{B, B \rightarrow \Box A \vdash \Box \Box A}}{B, B \rightarrow \Box A \vdash \Box \Box A} \quad \frac{\frac{\overline{B, B \rightarrow \Box A \vdash B \rightarrow \Box A}}{B, B \rightarrow \Box A \vdash \Box A}}{B, B \rightarrow \Box A \vdash \Box A}}{B, B \rightarrow \Box A \vdash \Box \Box A}$$

This derivation contains a redex at the outermost level, due to the introduction and then elimination of \rightarrow . Attempting to reduce this redex yields the following, which violates the restriction on Prawitz’s necessitation rule:

$$\frac{\frac{\overline{B, B \rightarrow \Box A \vdash B \rightarrow \Box A} \quad \overline{B, B \rightarrow \Box A \vdash B}}{B, B \rightarrow \Box A \vdash \Box A}}{B, B \rightarrow \Box A \vdash \Box \Box A}$$

We can trace the difficulty to the global side-condition on the necessitation rule which requires assumptions to be of a particular form. If we express this condition directly on the level of the judgments, we are led to a different system which does satisfy subject reduction and other properties desirable for a system of natural deduction. To this end, we introduce two basic judgments on propositions, “ A is true” and “ A is valid”. We have hypotheses expressing that certain propositions are true and others are valid. We write

$$(A_1, \dots, A_m); (B_1, \dots, B_n) \vdash^e C$$

to express

Under the hypothesis that A_1, \dots, A_m are valid and B_1, \dots, B_n are true, C is true.

Since our main goal is the analysis of the Curry-Howard isomorphism between proofs and programs, we label the hypotheses and annotate C with a proof term E .

$$(u_1:A_1, \dots, u_m:A_m); (x_1:B_1, \dots, x_n:B_n) \vdash^e E : C$$

Here and throughout this paper, we presuppose that that all variables labelling hypotheses are distinct.

Taking the functional view for a moment, we think of u_1, \dots, u_m as variables ranging over code and x_1, \dots, x_n as variables ranging over values which may occur free in the expression E . Generally, we write Δ for a context $u_1:A_1, \dots, u_m:A_m$ declaring

modal variables u (also called *code variables*) and Γ for a context $x_1:B_1, \dots, x_n:B_n$ declaring *ordinary variables* x (also called *value variables*).

But how do we conclude that A is valid? In informal terms, A is valid if it is true under all possible interpretations. In other words, its derivation may not depend on any hypotheses about the truth of propositions. That is, we judge that C is valid under the hypothesis that A_1, \dots, A_m are valid if

$$(A_1, \dots, A_m); \cdot \vdash^e C$$

or, with proof terms,

$$(u_1:A_1, \dots, u_m:A_m); \cdot \vdash^e E : C$$

With respect to our functional interpretation, this means that E contains only free code variables, but no free value variables.

We now develop the inference rules characterizing the judgments and then introduce the logical connectives. First we have

$$\frac{x:A \text{ in } \Gamma}{\Delta; \Gamma \vdash^e x : A} \text{ovar}$$

since we can conclude that A is true from the hypothesis that A is true. But it is certainly also the case that A is true if A is valid.

$$\frac{u:A \text{ in } \Delta}{\Delta; \Gamma \vdash^e u : A} \text{mvar}$$

The transition from a judgment of validity to that of truth corresponds on the functional side to a transition from code to value. We will use this later to encode evaluation.

Second, we consider the substitution principles which are derived from the nature of the hypothetical judgments. In purely logical terms: if we have a derivation showing that C is true from the hypothesis that A is true, then we can substitute an actual derivation establishing the truth of A for all uses of the hypothesis. This results in a derivation for the truth of C which no longer depends on the hypothesis. With proof terms, the substitution principle for ordinary hypotheses reads:

Ordinary Substitution Principle

If $\Delta; \Gamma \vdash^e E_1 : A$ and $\Delta; (\Gamma, x:A, \Gamma') \vdash^e E_2 : B$
then $\Delta; (\Gamma, \Gamma') \vdash^e [E_1/x]E_2 : B$.

Similarly, we should be able to substitute a derivation demonstrating the validity of A for all uses of the hypothesis that A is valid.

Modal Substitution Principle

If $\Delta; \cdot \vdash^e E_1 : A$ and $(\Delta, u:A, \Delta'); \Gamma \vdash^e E_2 : B$
then $(\Delta, \Delta'); \Gamma \vdash^e [E_1/u]E_2 : B$.

It is critical here that A is valid and not just true, which should be obvious from what is said above. Therefore, we must require $\Delta; \cdot \vdash^e E_1 : A$ rather than just $\Delta; \Gamma \vdash^e E_1 : A$ (which would be unsound).

Eventually, when our system is complete, we have to *prove* the validity of the two substitution principles to verify that there is no mistake in the design of our

rules. Similar guiding properties of hypothetical judgments are *exchange* (the order of hypotheses should not matter), *weakening* (hypotheses need not be used) and *contraction* (hypotheses may be used more than once). All of these are proved in Section 2.4.

The next step is to introduce the logical connectives and operators. In natural deduction, these are characterized by introduction and elimination rules which must match in an appropriate way. One of the underlying principles of natural deduction is that connectives should be orthogonal to each other: each introduction or elimination rule should refer only to the connective whose meaning we define.

We first discuss this using the familiar implication (or function type, under the Curry-Howard correspondence). We want to express that $A \rightarrow B$ should be true if B is true under the hypothesis that A is true.

$$\frac{\Delta; (\Gamma, x:A) \vdash^e E : B}{\Delta; \Gamma \vdash^e \lambda x:A. E : A \rightarrow B} \rightarrow \text{I}$$

Note that Δ is not affected—validity does not enter the considerations for this connective. On proof terms, this rule explicitly introduces the function which maps proofs of A to proofs of B .

Conversely, if we know that $A \rightarrow B$ is true then B should be true under hypothesis A . So if we also know that A is true, we can conclude that B must be true.

$$\frac{\Delta; \Gamma \vdash^e E_2 : A \rightarrow B \quad \Delta; \Gamma \vdash^e E_1 : A}{\Delta; \Gamma \vdash^e E_2 E_1 : B} \rightarrow \text{E}$$

On proof terms, this applies the function E_2 which maps proofs of A to proofs of B to the given proof E_1 of A .

How do we know the introduction and elimination rules match and thus define a meaningful connective? We should verify two conditions: *local soundness* and *local completeness*. Local soundness ensures that we cannot gain information by introducing a connective and then immediately eliminating it—we must already be able to make the same judgment without the detour. This guarantees that the elimination rules are not too strong. Local completeness ensures that we can recover all information present in a connective: there is some way to apply the elimination rules so we can reconstitute a proof of the original proposition using its introduction rules. This guarantees that the elimination rules are not too weak.

On proof terms, local soundness and completeness are witnessed by local reduction and expansion, taking advantage of the substitution principles.

$$\frac{\frac{\Delta; \Gamma, x:A \vdash^e E_2 : B}{\Delta; \Gamma \vdash^e (\lambda x:A. E_2) : A \rightarrow B} \rightarrow \text{I} \quad \frac{\mathcal{D}_1}{\Delta; \Gamma \vdash^e E_1 : A}}{\Delta; \Gamma \vdash^e (\lambda x:A. E_2) E_1 : B} \rightarrow \text{E} \quad \Longrightarrow \quad \frac{\mathcal{D}'_2}{\Delta; \Gamma \vdash^e [E_1/x]E_2 : B} \rightarrow \text{E}$$

Here, \mathcal{D}'_2 is constructed by substitution of \mathcal{D}_1 into \mathcal{D}_2 , as indicated in the discussion of the (ordinary) substitution principle. On proof terms we have ordinary β -reduction. Local completeness is witnessed by η -expansion.

$$\Delta; \Gamma \vdash^e E : A \rightarrow B \xRightarrow{\mathcal{D}} \frac{\frac{\frac{\mathcal{D}' \quad \frac{\Delta; \Gamma, x:A \vdash^e E : A \rightarrow B \quad \frac{\Delta; \Gamma, x:A \vdash^e x : A}{\rightarrow E}}{\Delta; \Gamma, x:A \vdash^e E x : B}}{\Delta; \Gamma \vdash^e (\lambda x:A. E x) : A \rightarrow B}}{\rightarrow I}}{\Delta; \Gamma \vdash^e E : A \rightarrow B} \text{ovar}$$

Here, \mathcal{D}' is constructed by weakening from \mathcal{D} (we add the unused hypothesis $x:A$), which has no effect on the proof term E . On proof terms, therefore, we have ordinary η -expansion.

Next we consider the modal operator. $\Box A$ should be true if A is valid. Written as an inference rule:

$$\frac{\Delta; \cdot \vdash^e E : A}{\Delta; \Gamma \vdash^e \mathbf{box} E : \Box A} \Box I$$

Note how the premise enforces that A is valid by requiring the ordinary context to be empty. On proof terms this means that only modal variables from Δ can occur free in E .

The corresponding elimination is not straightforward. For example, Prawitz's rule from above which concludes A from $\Box A$ is locally sound but not complete. Intuitively, this should be clear because we are losing information when we make the step from “ $\Box A$ is true” to “ A is true”. An alternative rule which concludes “ A is valid” from “ $\Box A$ is true” is unsound, because the judgment that $\Box A$ is true may actually depend on hypotheses about the truth of other propositions.

Instead we reason as follows: if $\Box A$ is true under some hypotheses, then any judgment we make under the additional hypothesis that A is valid, must in fact be evident.

$$\frac{\Delta; \Gamma \vdash^e E_1 : \Box A \quad (\Delta, u:A); \Gamma \vdash^e E_2 : B}{\Delta; \Gamma \vdash^e \mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2 : B} \Box E$$

Thus the elimination rule for \Box introduces a modal hypothesis and the corresponding term construct has the form of a **let**. From the functional point of view, E_1 represents a value of type $\Box A$ containing some code. This code is accessible in E_2 with the name u . Local soundness and completeness with this construction are easily verified. Local soundness is guaranteed by the reduction

$$\frac{\frac{\frac{\mathcal{D}_1 \quad \Delta; \cdot \vdash^e E_1 : A}{\Delta; \Gamma \vdash^e \mathbf{box} E_1 : \Box A} \Box I \quad (\Delta, u:A); \Gamma \vdash^e E_2 : B}{\Delta; \Gamma \vdash^e \mathbf{let} \mathbf{box} u = \mathbf{box} E_1 \mathbf{in} E_2 : B} \Box E} \xRightarrow{\mathcal{D}_2} \Delta; \Gamma \vdash^e [E_1/u]E_2 : B \quad \mathcal{D}'_2$$

where \mathcal{D}'_2 is the derivation constructed by substitution as indicated in the modal substitution principle.

The expansion below demonstrates local completeness, since the result of the elimination rule applied to a derivation of $\Box A$ contains enough information to reconstitute a derivation of $\Box A$.

$$\frac{\Delta; \Gamma \vdash^{\mathcal{D}} E : \Box A \quad \frac{\frac{\text{mvar}}{(\Delta, u:A); \cdot \vdash^{\varepsilon} u : A} \Box I}{(\Delta, u:A); \Gamma \vdash^{\varepsilon} \mathbf{box} u : \Box A} \Box E}{\Delta; \Gamma \vdash^{\varepsilon} \mathbf{let} \mathbf{box} u = E \mathbf{in} \mathbf{box} u : \Box A} \Box E$$

Other standard logical connectives such as negation, conjunction, disjunction, universal and existential quantification can be defined by introduction and elimination rules in a similar manner to implication—they do not need to directly interact with the modal hypotheses. Since we are in the intuitionistic setting, the modal possibility operator $\Diamond A$ cannot be defined via negation. It can be characterized directly by introduction and elimination rules which are locally sound and complete, but only if we introduce a new judgment “ A is possibly true”, as we demonstrate in [Pfenning and Davies 2000]. We do not show the details here since it does not directly concern our main objective.

Our presentation simplifies that of the modal λ -calculus $\lambda_e^{\rightarrow \Box}$ from [Bierman and de Paiva 1996; Pfenning and Wong 1995] by eliminating the need for simultaneous substitution while preserving subject reduction. It is inspired by sequent calculi proposed by Andreoli [Andreoli 1992] for linear logic and by Girard [Girard 1993] for LU. Wadler [Wadler 1993] has formulated a linear λ -calculus with two contexts, which shares some features with our calculus. The methodology we followed is based on that of Martin-Löf [Martin-Löf 1996], and is described in more detail in [Pfenning and Davies 2000]. One novel feature of our methodology is its normative use of local soundness and completeness as witnessed by β -reduction and η -expansion. Note that only β -reduction has computational significance, while η -expansion internalizes an extensionality principle.

The elimination construct for \Box allows us to bind a variable u in Δ to code of type A , written as $\mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2$. Evaluation of code, certainly one of the most fundamental operations, is then definable by

$$eval \equiv (\lambda x:\Box A. \mathbf{let} \mathbf{box} u = x \mathbf{in} u) : \Box A \rightarrow A.$$

Here, and from now on, we associate \Box more strongly than \rightarrow to avoid excessive parentheses. Note that the opposite coercion, $\lambda x:A. \mathbf{box} x$, cannot be well-typed, since x is an arbitrary argument and will not necessarily be bound to code. Furthermore, it violates the concept of stage separation since x is an “early” argument which we refer to “late”, that is, inside \mathbf{box} . Here are two other examples of modal propositions and proofs from which the natural deductions can be easily reconstructed.

$$\begin{aligned} & \vdash^{\varepsilon} \lambda x:\Box(A \rightarrow B). \lambda y:\Box A. \mathbf{let} \mathbf{box} u = x \mathbf{in} \mathbf{let} \mathbf{box} v = y \mathbf{in} \mathbf{box} (uv) \\ & : \Box(A \rightarrow B) \rightarrow (\Box A \rightarrow \Box B) \end{aligned}$$

$$\begin{aligned} & \vdash^{\varepsilon} \lambda x:\Box A. \mathbf{let} \mathbf{box} u = x \mathbf{in} \mathbf{box} \mathbf{box} u \\ & : \Box A \rightarrow \Box \Box A \end{aligned}$$

Note that the first law holds in all modal logics, while the second and $eval$ above correspond to reflexivity and transitivity of the accessibility relation between worlds in the Kripke semantics [Kripke 1963] in axiomatic formulations of modal logics.

2.2 Syntax

We now summarize the system of natural deduction for S4 and its properties.

Types	$A ::= a \mid A_1 \rightarrow A_2 \mid \Box A$
Terms	$E ::= x \mid \lambda x:A. E \mid E_1 E_2 \mid$ $u \mid \mathbf{box} E \mid \mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2$
Ordinary Contexts	$\Gamma ::= \cdot \mid \Gamma, x:A$
Modal Contexts	$\Delta ::= \cdot \mid \Delta, u:A$

We use A, B for types, a for base types, x for ordinary variables and u for modal variables assuming that any variable can be declared at most once in a context. Bound variables may be renamed tacitly, and leading \cdot 's may be omitted from contexts. We write $[E'/x]E$ (and similarly for modal variables) for the result of substituting E' for x in E , renaming bound variables as necessary in order to avoid the capture of free variables in E' .

2.3 Typing Rules

$\Delta; \Gamma \vdash^e E : A$ E has type A in modal context Δ and ordinary context Γ .

Our system has the property that a valid term has a unique type and typing derivation, except for possibly unused hypotheses.

λ -calculus Fragment

$$\frac{x:A \text{ in } \Gamma}{\Delta; \Gamma \vdash^e x : A} \text{ovar}$$

$$\frac{\Delta; (\Gamma, x:A) \vdash^e E : B}{\Delta; \Gamma \vdash^e \lambda x:A. E : A \rightarrow B} \rightarrow I$$

$$\frac{\Delta; \Gamma \vdash^e E_1 : B \rightarrow A \quad \Delta; \Gamma \vdash^e E_2 : B}{\Delta; \Gamma \vdash^e E_1 E_2 : A} \rightarrow E$$

Modal Fragment

$$\frac{u:A \text{ in } \Delta}{\Delta; \Gamma \vdash^e u : A} \text{mvar}$$

$$\frac{\Delta; \cdot \vdash^e E : A}{\Delta; \Gamma \vdash^e \mathbf{box} E : \Box A} \Box I$$

$$\frac{\Delta; \Gamma \vdash^e E_1 : \Box A \quad (\Delta, u:A); \Gamma \vdash^e E_2 : B}{\Delta; \Gamma \vdash^e \mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2 : B} \Box E$$

2.4 Reduction and Expansion

The notions of β -reduction and η -expansion are fundamental to the λ -calculus. The preservation of types under β -reduction is the functional analog of local soundness

for rules of natural deduction; the preservation of types under η -expansion is the functional analog of local completeness. But first we need to verify the characteristic properties of hypothetical judgments: exchange, weakening, contraction, and substitution.

LEMMA 1 (STRUCTURAL PROPERTIES OF CONTEXTS).

- (1) If $(\Delta_1, u:A, v:B, \Delta_2); \Gamma \vdash^e E : C$ then $(\Delta_1, v:B, u:A, \Delta_2); \Gamma \vdash^e E : C$.
- (2) If $\Delta; (\Gamma_1, x:A, y:B, \Gamma_2) \vdash^e E : C$ then $\Delta; (\Gamma_1, y:B, x:A, \Gamma_2) \vdash^e E : C$.
- (3) If $\Delta; \Gamma \vdash^e E : C$ then $(\Delta, u:A); \Gamma \vdash^e E : C$.
- (4) If $\Delta; \Gamma \vdash^e E : C$ then $\Delta; (\Gamma, x:A) \vdash^e E : C$.
- (5) If $(\Delta, u:A, v:A); \Gamma \vdash^e E : C$ then $(\Delta, w:A); \Gamma \vdash^e [w/u][w/v]E : C$.
- (6) If $\Delta; (\Gamma, x:A, y:A) \vdash^e E : C$ then $\Delta; (\Gamma, z:A) \vdash^e [z/x][z/y]E : C$.

PROOF. By straightforward inductions over the structure of the given derivations. Recall the global assumption that each variable is declared at most once in a context, and that bound variables may be renamed tacitly. \square

LEMMA 2 (SUBSTITUTION).

- (1) If $\Delta; \Gamma \vdash^e E_1 : A$ and $\Delta; (\Gamma, x:A, \Gamma') \vdash^e E_2 : B$
then $\Delta; (\Gamma, \Gamma') \vdash^e [E_1/x]E_2 : B$.
- (2) If $\Delta; \cdot \vdash^e E_1 : A$ and $(\Delta, u:A, \Delta'); \Gamma \vdash^e E_2 : B$
then $(\Delta, \Delta'); \Gamma \vdash^e [E_1/u]E_2 : B$.

PROOF. By straightforward inductions on the typing derivations for E_2 . \square

The β -reductions and η -expansions on proof terms used in the preceding section to verify local soundness and completeness are summarized below.

$$\begin{aligned} (\lambda x:A. E_2) E_1 &\vdash^{\beta} [E_1/x]E_2 \\ \mathbf{let\ box\ } u = \mathbf{box\ } E_1 \mathbf{ in\ } E_2 &\vdash^{\beta} [E_1/u]E_2 \\ E : A \rightarrow B &\vdash^{\eta} \lambda x:A. E x \\ E : \Box A &\vdash^{\eta} \mathbf{let\ box\ } u = E \mathbf{ in\ box\ } u \end{aligned}$$

We now validate these rules by showing that they satisfy subject reduction.

THEOREM 3 (SUBJECT REDUCTION AND EXPANSION).

- (1) If $\Delta; \Gamma \vdash^e E : A$ and $E \vdash^{\beta} E'$ then $\Delta; \Gamma \vdash^e E' : A$.
- (2) If $\Delta; \Gamma \vdash^e E : A$ and $E : A \vdash^{\eta} E'$ then $\Delta; \Gamma \vdash^e E' : A$.

PROOF. In the case of a reduction we first apply inversion (i.e consider the possible final inferences in the derivation) and then use the substitution properties to obtain the result. In the case of an expansion we directly construct the typing derivation for the expanded term. \square

We will not discuss commuting conversions for the $\Box E$ rule here, since they are not particularly relevant to our intended application. Similarly, we will not present a formal proof of strong normalization, although this is easy to obtain by an embedding into the ordinary simply-typed λ -calculus.

2.5 Staged Computation

We now show the relationship between $\lambda_e^{\rightarrow\Box}$ and staged computation. It is our intention that those parts of a term enclosed by a **box** constructor should be considered “uninterpreted code”. Thus, when we construct a computational interpretation of $\lambda_e^{\rightarrow\Box}$ based on β -reduction, it is appropriate to omit the congruence rule for **box**. We have the judgment:

$$E \mapsto E' \quad E \text{ reduces to } E'$$

This judgment is defined by the following rules.

$$\frac{}{(\lambda x:A. E_2)E_1 \mapsto [E_1/x]E_2} \rightarrow\beta$$

$$\frac{}{\mathbf{let\ box\ } u = \mathbf{box\ } E_1 \mathbf{ in\ } E_2 \mapsto [E_1/u]E_2} \Box\beta$$

$$\frac{E \mapsto E'}{\lambda x:A. E \mapsto \lambda x:A. E'} \text{cong_lam}$$

$$\frac{E_1 \mapsto E'_1}{E_1 E_2 \mapsto E'_1 E_2} \text{cong_app1}$$

$$\frac{E_2 \mapsto E'_2}{E_1 E_2 \mapsto E_1 E'_2} \text{cong_app2}$$

$$\frac{E_1 \mapsto E'_1}{\mathbf{let\ box\ } u = E_1 \mathbf{ in\ } E_2 \mapsto \mathbf{let\ box\ } u = E'_1 \mathbf{ in\ } E_2} \text{cong_letbox1}$$

$$\frac{E_2 \mapsto E'_2}{\mathbf{let\ box\ } u = E_1 \mathbf{ in\ } E_2 \mapsto \mathbf{let\ box\ } u = E_1 \mathbf{ in\ } E'_2} \text{cong_letbox2}$$

We write \mapsto^* for the reflexive and transitive closure of \mapsto .

THEOREM 4 (SUBJECT REDUCTION WITH CONGRUENCES).
If $\Delta; \Gamma \vdash^e E : A$ and $E \mapsto^ E'$ then $\Delta; \Gamma \vdash^e E' : A$.*

PROOF. By a simple induction on the derivation of $E \mapsto^* E'$, using subject reduction (Theorem 3) for the base cases. \square

To demonstrate the relationship between this reduction relation and computation staging, we roughly follow the binding-time correctness criteria described by Palsberg [Palsberg 1993]. Palsberg presented a modular proof of correctness for binding-time analyses based on two-level languages, such as those studied in [Go-

mard and Jones 1991]. The first criterion is *consistency*, namely that static reduction of a well-annotated term cannot “go wrong”. In our case, well-annotated means well-typed, and the above subject reduction theorem corresponds roughly to the property of not “going wrong”. To make the correspondence more evident, we can simply note that a well-typed term cannot contain the “wrong” forms $(\mathbf{box} E'_1) E'_2$ and $\mathbf{let} \mathbf{box} u = (\lambda x:A. E'_1) \mathbf{in} E'_2$.

The second criterion for binding-time correctness is that when a stage is complete, no subterm occurrences that are marked as eliminable remain. In our case, the subterm occurrences in the scope of a \mathbf{box} constructor are code to be executed at a later stage and are therefore considered *persistent*; all other term occurrences are considered *eliminable*. Completing a stage means reducing a term until it can not be further reduced by the rules of the judgment $E \mapsto^* E'$. We call such terms *irreducible* to avoid confusion with subtly different notions such as head-normal form or weak head-normal form. Note that an irreducible term could still contain a “redex” in the traditional sense underneath a \mathbf{box} constructor. Since we only evaluate closed terms, the following theorem expresses that our language satisfies the second criterion for binding-time correctness.

THEOREM 5 (ELIMINABILITY). *If $\cdot; \cdot \vdash^e E : \square A$ and $E \mapsto^* E'$ and E' is irreducible, then E' contains no eliminable term occurrences.*

PROOF. By subject reduction, $\cdot; \cdot \vdash^e E' : \square A$. By inversion, and the fact that E' is irreducible, E' must have the form $\mathbf{box} E'_0$ for some E'_0 . Therefore all subterm occurrences in E' are in the scope of a \mathbf{box} constructor and hence persistent. \square

Thus, it appears that Palsberg’s two properties both follow relatively easily from subject reduction for $\lambda_e^{\rightarrow \square}$. However, there is still more to consider, because it is possible that eliminable subterms could reduce to persistent terms. This is ruled out syntactically in the two-level language studied by Palsberg, but in our case we need to show this explicitly. To argue about “images under reduction” we temporarily extend $\lambda_e^{\rightarrow \square}$ with labels.

Terms $E ::= \dots | E^l$

Labels have no impact on typing and can be reduced away.

$$\frac{\Delta; \Gamma \vdash^e E : A}{\Delta; \Gamma \vdash^e E^l : A} \text{LB}$$

$$\frac{}{E^l \mapsto E} \text{unlabel}$$

Recall that there is no congruence rule for \mathbf{box} so that the rule *unlabel* can not be applied to a label in the scope of a \mathbf{box} constructor.

Now suppose E_1 and E_2 differ only in that some subterms of E_1 have been labelled in E_2 . Then typing and reduction correspond between E_1 and E_2 . That is, $\Delta; \Gamma \vdash^e E_1 : A$ iff $\Delta; \Gamma \vdash^e E_2 : A$, and $E_1 \mapsto^* E'_1$ iff $E_2 \mapsto^* E'_2$ where E'_1

and E'_2 differ only in their labelling. This allows us to trace the “images under reduction” of eliminable parts of an unlabelled term by labelling all persistent subterm occurrences. The following theorem then expresses that only persistent parts of a term can yield persistent images. Here, for every subterm occurrence of the form E^l , both E and E^l are considered to be *labelled with l* . No other subterm occurrences are considered labelled.

THEOREM 6 (PERSISTENCE). *If $\Delta; \Gamma \vdash^e E : A$, all persistent subterm occurrences of E are labelled with l , and $E \mapsto E'$, then all persistent subterm occurrences of E' are labelled with l .*

PROOF. By induction on the derivation of $E \mapsto E'$.

Case:

$$\frac{}{(\lambda x:A. E_2)E_1 \mapsto [E_1/x]E_2} \rightarrow \beta$$

The result follows because the modal restriction in the typing rules does not allow x to appear in E_2 in a position enclosed by a **box** constructor. Formally, this case requires an auxiliary induction on E_2 .

Case:

$$\frac{}{\mathbf{let\ box\ } u = \mathbf{box\ } E_1 \mathbf{ in\ } E_2 \mapsto [E_1/u]E_2} \square\beta$$

The result follows because every subterm in E_1 is labelled with l .

Case:

$$\frac{E_2 \mapsto E'_2}{\mathbf{let\ box\ } u = E_1 \mathbf{ in\ } E_2 \mapsto \mathbf{let\ box\ } u = E_1 \mathbf{ in\ } E'_2} \text{cong_letbox2}$$

The result follows immediately by the induction hypothesis.

The other congruence cases are similar. If there were a congruence for **box**, that case would fail.

Case:

$$\frac{}{E^l \mapsto E} \text{unlabel}$$

The result is immediate, since E is not enclosed by **box** and therefore not persistent.

□

Interpreted as a statement about code manipulation during evaluation, this theorem says that we can never construct code from terms which were not originally code. This is one of the essential properties of $\lambda_e^{\rightarrow\Box}$ which makes it a suitable basis for languages allowing explicit code manipulation.

There is a dual property to persistence which is also enforced syntactically in the languages studied by Palsberg, namely that the eliminable parts of terms in the result of reduction only appear as the images of the eliminable parts of the original

term. It is an important aspect of $\lambda_e^{\rightarrow\Box}$ is that it *does not* have this property, as shown by the counterexample:

$$\mathbf{let\ box\ } u = \mathbf{box\ } E \mathbf{ in\ } u \mapsto E$$

E appears in a persistent (or code) position on the left, but in an eliminable (or value) position on the right. From the point of view of code manipulation it is easy to explain why this is allowed. In the languages we are interested in, the code representation for E can be *evaluated* to return a value for E , which stands in contrast to the languages studied by Palsberg.

Evaluation of code is expressed in $\lambda_e^{\rightarrow\Box}$ by occurrences of code variables u which are not enclosed by a **box** constructor. From a logical point of view, such instances correspond exactly to proofs which depend on the reflexivity of the Kripke reachability relation for modal S4. If we modify the rules to disallow such instances, we obtain the modal logic $K4$ and a corresponding modal λ -calculus which is somewhat closer to the two-level languages studied by Palsberg. If we also remove the transitivity of reachability from $K4$, we obtain the modal logic K and a corresponding modal λ -calculus which removes another feature of $\lambda_e^{\rightarrow\Box}$ that is not present in two-level languages, namely the ability to substitute code directly into code which is itself part of a code expression. So the following function from $\lambda_e^{\rightarrow\Box}$ would no longer be well-typed.

$$(\lambda x:\Box A. \mathbf{let\ box\ } u = \mathbf{box\ } E \mathbf{ in\ box\ box\ } u) : \Box A \rightarrow \Box\Box A$$

Allowing this feature in $\lambda_e^{\rightarrow\Box}$ seems reasonable and useful, though perhaps not as important as the inclusion of evaluation of code. We will come back to languages based on modal K later in Section 6, where we will briefly explain an exact correspondence between such languages and multi-level generalizations of two-level languages.

3. MODAL MINI-ML: EXPLICIT FORMULATION

This section presents Mini-ML $_e^{\Box}$, a language that combines some elements of Mini-ML [Clément et al. 1986] with the $\lambda_e^{\rightarrow\Box}$ -calculus of the previous section. For the sake of simplicity Mini-ML $_e^{\Box}$ is explicitly typed. ML-style or explicit polymorphism can also be added in a straightforward manner; we omit the details here in order to concentrate on the essential issues.

We present an operational semantics for the language, and demonstrate some basic properties such as type preservation. We also demonstrate the strong staging properties of the language. In the description of the operational semantics we choose the usual device of representing values (including code) by corresponding source expressions. This may be refined in different ways for lower-level semantics describing, for example, run-time code generation or partial evaluation.

3.1 Syntax

Types	$A ::= \mathbf{nat} \mid A_1 \rightarrow A_2 \mid A_1 \times A_2 \mid 1 \mid \Box A$
Terms	$E ::= x \mid \lambda x:A. E \mid E_1 E_2$ $\mid u \mid \mathbf{box} E \mid \mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2$ $\mid \langle E_1, E_2 \rangle \mid \mathbf{fst} E \mid \mathbf{snd} E$ $\mid \langle \rangle$ $\mid \mathbf{z} \mid \mathbf{s} E \mid (\mathbf{case} E_1 \mathbf{of} \mathbf{z} \Rightarrow E_2 \mid \mathbf{s} x \Rightarrow E_3)$ $\mid \mathbf{fix} x:A. E$
Ordinary Contexts	$\Gamma ::= \cdot \mid \Gamma, x:A$
Modal Contexts	$\Delta ::= \cdot \mid \Delta, u:A$

This language extends the one in the previous section, and we continue to use the conventions introduced in that section.

3.2 Typing Rules

Our typing rules for the Mini-ML fragment of the explicit language are completely standard, and we follow the previous section for the modal fragment.

$\Delta; \Gamma \vdash^e E : A$ E has type A in modal context Δ and ordinary context Γ .

Functions

$$\frac{x:A \text{ in } \Gamma}{\Delta; \Gamma \vdash^e x : A} \text{tpe_ovar} \qquad \frac{\Delta; (\Gamma, x:A) \vdash^e E : B}{\Delta; \Gamma \vdash^e \lambda x:A. E : A \rightarrow B} \text{tpe_lam}$$

$$\frac{\Delta; \Gamma \vdash^e E_1 : B \rightarrow A \quad \Delta; \Gamma \vdash^e E_2 : B}{\Delta; \Gamma \vdash^e E_1 E_2 : A} \text{tpe_app}$$

Code

$$\frac{u:A \text{ in } \Delta}{\Delta; \Gamma \vdash^e u : A} \text{tpe_mvar} \qquad \frac{\Delta; \cdot \vdash^e E : A}{\Delta; \Gamma \vdash^e \mathbf{box} E : \Box A} \text{tpe_box}$$

$$\frac{\Delta; \Gamma \vdash^e E_1 : \Box A \quad (\Delta, u:A); \Gamma \vdash^e E_2 : B}{\Delta; \Gamma \vdash^e \mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2 : B} \text{tpe_let_box}$$

Products

$$\frac{\Delta; \Gamma \vdash^e E_1 : A_1 \quad \Delta; \Gamma \vdash^e E_2 : A_2}{\Delta; \Gamma \vdash^e \langle E_1, E_2 \rangle : A_1 \times A_2} \text{tpe_pair}$$

$$\frac{\Delta; \Gamma \vdash^e E : A_1 \times A_2}{\Delta; \Gamma \vdash^e \mathbf{fst} E : A_1} \text{tpe_fst} \qquad \frac{\Delta; \Gamma \vdash^e E : A_1 \times A_2}{\Delta; \Gamma \vdash^e \mathbf{snd} E : A_2} \text{tpe_snd}$$

$$\frac{}{\Delta; \Gamma \vdash^e \langle \rangle : 1} \text{tpe_unit}$$

Natural Numbers

$$\begin{array}{c}
\frac{}{\Delta; \Gamma \vdash^e \mathbf{z} : \text{nat}} \text{tpe_z} \quad \frac{\Delta; \Gamma \vdash^e E : \text{nat}}{\Delta; \Gamma \vdash^e \mathbf{s} E : \text{nat}} \text{tpe_s} \\
\\
\frac{\Delta; \Gamma \vdash^e E_1 : \text{nat} \quad \Delta; \Gamma \vdash^e E_2 : A \quad \Delta; (\Gamma, x:\text{nat}) \vdash^e E_3 : A}{\Delta; \Gamma \vdash^e (\mathbf{case} E_1 \mathbf{of} \mathbf{z} \Rightarrow E_2 \mid \mathbf{s} x \Rightarrow E_3) : A} \text{tpe_case}
\end{array}$$

Recursion

$$\frac{\Delta; \Gamma, x:A \vdash^e E : A}{\Delta; \Gamma \vdash^e \mathbf{fix} x:A. E : A} \text{tpe_fix}$$

As before, there is only one rule which introduces variables into the modal context (here called `tpe_let_box`).

3.3 Operational Semantics

The Mini-ML fragment of our system has a standard call-by-value operational semantics. For the modal part, we represent code for E simply by `box E`, making the least commitment concerning lower-level implementations.

$$\text{Values } V ::= \lambda x:A. E \mid \langle V_1, V_2 \rangle \mid \langle \rangle \mid \mathbf{z} \mid \mathbf{s} V \mid \mathbf{box} E$$

We evaluate `let box x = E1 in E2` by substituting the code generated by evaluating E_1 for x in E_2 and then evaluating E_2 . The code generated by E_1 may then be evaluated during the evaluation of E_2 as necessary. On the λ -calculus and modal fragments our semantics corresponds to a reduction strategy for $\lambda_e^{\rightarrow\Box}$.

$E \hookrightarrow V$ Expression E evaluates to value V .

Functions

$$\begin{array}{c}
\frac{}{\lambda x:A. E \hookrightarrow \lambda x:A. E} \text{ev_lam} \\
\\
\frac{E_1 \hookrightarrow \lambda x. E'_1 \quad E_2 \hookrightarrow V_2 \quad [V_2/x]E'_1 \hookrightarrow V}{E_1 E_2 \hookrightarrow V} \text{ev_app}
\end{array}$$

Code

$$\begin{array}{c}
\frac{}{\mathbf{box} E \hookrightarrow \mathbf{box} E} \text{ev_box} \\
\\
\frac{E_1 \hookrightarrow \mathbf{box} E'_1 \quad [E'_1/u]E_2 \hookrightarrow V_2}{\mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2 \hookrightarrow V_2} \text{ev_let_box}
\end{array}$$

Products

$$\begin{array}{c}
 \frac{E_1 \hookrightarrow V_1 \quad E_2 \hookrightarrow V_2}{\langle E_1, E_2 \rangle \hookrightarrow \langle V_1, V_2 \rangle} \text{ev_pair} \\
 \frac{E \hookrightarrow \langle V_1, V_2 \rangle}{\mathbf{fst} E \hookrightarrow V_1} \text{ev_fst} \quad \frac{E \hookrightarrow \langle V_1, V_2 \rangle}{\mathbf{snd} E \hookrightarrow V_2} \text{ev_snd} \\
 \frac{}{\langle \rangle \hookrightarrow \langle \rangle} \text{ev_unit}
 \end{array}$$

Natural Numbers

$$\begin{array}{c}
 \frac{}{\mathbf{z} \hookrightarrow \mathbf{z}} \text{ev_z} \quad \frac{E \hookrightarrow V}{\mathbf{s} E \hookrightarrow \mathbf{s} V} \text{ev_s} \\
 \frac{E_1 \hookrightarrow \mathbf{z} \quad E_2 \hookrightarrow V}{(\mathbf{case} E_1 \text{ of } \mathbf{z} \Rightarrow E_2 \mid \mathbf{s} x \Rightarrow E_3) \hookrightarrow V} \text{ev_case_z} \\
 \frac{E_1 \hookrightarrow \mathbf{s} V'_1 \quad [V'_1/x]E_3 \hookrightarrow V}{(\mathbf{case} E_1 \text{ of } \mathbf{z} \Rightarrow E_2 \mid \mathbf{s} x \Rightarrow E_3) \hookrightarrow V} \text{ev_case_s}
 \end{array}$$

Recursion

$$\frac{[\mathbf{fix} x. E/x]E \hookrightarrow V}{\mathbf{fix} x. E \hookrightarrow V} \text{ev_fix}$$

The structural and substitution properties for $\lambda_e^{\rightarrow \square}$ extend to $\text{Mini-ML}_e^{\square}$ in a completely straightforward way, and we will make use of it below. We restate only the substitution lemma.

LEMMA 7 (SUBSTITUTION).

- (1) If $\Delta; \Gamma \vdash^e E_1 : A$ and $\Delta; (\Gamma, x:A, \Gamma') \vdash^e E_2 : B$
then $\Delta; (\Gamma, \Gamma') \vdash^e [E_1/x]E_2 : B$.
- (2) If $\Delta; \cdot \vdash^e E_1 : A$ and $(\Delta, u:A, \Delta'); \Gamma \vdash^e E_2 : B$
then $(\Delta, \Delta'); \Gamma \vdash^e [E_1/u]E_2 : B$.

PROOF. By straightforward inductions on the typing derivations for E_2 . \square

THEOREM 8 (DETERMINACY AND TYPE PRESERVATION).

- (1) If $E \hookrightarrow V$ then V is a value.
- (2) If $E \hookrightarrow V$ and $E \hookrightarrow \hat{V}$ then $V = \hat{V}$ (modulo renaming of bound variables).
- (3) If $E \hookrightarrow V$ and $\cdot; \cdot \vdash^e E : A$ then $\cdot; \cdot \vdash^e V : A$.

PROOF. By inductions over the structure of the derivation \mathcal{D} of $E \hookrightarrow V$. The cases for the non-modal part are completely standard. The cases for ev_box are trivial and those for ev_let_box are straightforward for value soundness and determinacy. We thus show only the ev_let_box case in the proof of type preservation.

Case:

$$\mathcal{D} = \frac{E_1 \hookrightarrow \mathbf{box} E'_1 \quad [E'_1/u]E_2 \hookrightarrow V_2}{\mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2 \hookrightarrow V_2} \text{ev_let_box}$$

$\cdot; \vdash^e E_1 : \Box B$ and
 $u:B; \cdot \vdash^e E_2 : A$
 $\cdot; \vdash^e \mathbf{box} E'_1 : \Box B$
 $\cdot; \vdash^e E'_1 : B$
 $\cdot; \vdash^e [E'_1/u]E_2 : A$
 $\cdot; \vdash^e V_2 : A$

by inversion
 by ind. hyp.
 by inversion
 by substitution lemma
 by ind. hyp.

□

Since our semantics for Mini-ML_e[□] is a the natural extension of a reduction strategy for $\lambda_e^{\rightarrow\Box}$, the staging correctness results in Section 2 carry over to Mini-ML_e[□]. We now briefly discuss the staging captured in Mini-ML_e[□] in informal terms.

Suppose that $\cdot; \vdash^e E : \Box A$ and $E \hookrightarrow V$. By value soundness and type preservation we have $V \equiv \mathbf{box} E'$. Thus the result consists only of residual code to be executed in the next stage. Further, by the modal restrictions, only terms enclosed by **box** constructors are ever substituted into other **box** constructors. As a result, the parts of the original program E not enclosed by any **box** constructor can be designated eliminable (static) since they will not appear in the residual code E' .

Further, the body of a **box** constructor can be considered persistent (dynamic) in the sense that we do not evaluate underneath the **box** constructor. The only way for evaluation to proceed to the body of the **box** constructor is by using the variable bound by a **let box** elimination construct to indicate where the delayed computation should be performed.

3.4 Examples

The following functions demonstrate the interaction between products and □. They correspond to proofs in a modal logic with conjunction, since products and conjunction are related by the Curry-Howard isomorphism.

$\lambda x: (\Box A) \times (\Box B). \mathbf{let} \mathbf{box} u_1 = \mathbf{fst} x \mathbf{in} \mathbf{let} \mathbf{box} u_2 = \mathbf{snd} x \mathbf{in} \mathbf{box} \langle u_1, u_2 \rangle$
 $: (\Box A) \times (\Box B) \rightarrow \Box(A \times B)$

$\lambda x: \Box(A \times B). \mathbf{let} \mathbf{box} u = x \mathbf{in} \langle \mathbf{box} \mathbf{fst} u, \mathbf{box} \mathbf{snd} u \rangle$
 $: \Box(A \times B) \rightarrow (\Box A) \times (\Box B)$

We now define the power function in Mini-ML_e[□] in such a way that it has type $\mathbf{nat} \rightarrow \Box(\mathbf{nat} \rightarrow \mathbf{nat})$, assuming a closed term $\mathit{times}:\mathbf{nat} \rightarrow \mathbf{nat} \rightarrow \mathbf{nat}$ (definable in the Mini-ML fragment in the standard way).

$$\begin{aligned} \mathit{power} &\equiv \mathbf{fix} p:\mathbf{nat} \rightarrow \Box(\mathbf{nat} \rightarrow \mathbf{nat}). \\ &\quad \lambda n:\mathbf{nat}. \mathbf{case} n \\ &\quad \mathbf{of} \mathbf{z} \Rightarrow \mathbf{box} (\lambda x:\mathbf{nat}. \mathbf{s} \mathbf{z}) \\ &\quad | \mathbf{s} m \Rightarrow \mathbf{let} \mathbf{box} q = p m \mathbf{in} \\ &\quad \quad \mathbf{box} (\lambda x:\mathbf{nat}. \mathit{times} x (q x)) \end{aligned}$$

The type $\text{nat} \rightarrow \Box(\text{nat} \rightarrow \text{nat})$ expresses that *power* evaluates everything that depends on the first argument of type nat (the exponent) and returns residual code of type $\Box(\text{nat} \rightarrow \text{nat})$. Indeed, we calculate with our operational semantics:

$$\begin{aligned} \text{power } \mathbf{z} &\hookrightarrow \mathbf{box} (\lambda x:\text{nat}. \mathbf{s } \mathbf{z}) \\ \text{power } (\mathbf{s } \mathbf{z}) &\hookrightarrow \mathbf{box} (\lambda x:\text{nat}. \text{times } x ((\lambda x:\text{nat}. \mathbf{s } \mathbf{z})x)) \\ \text{power } (\mathbf{s } (\mathbf{s } \mathbf{z})) &\hookrightarrow \mathbf{box} (\lambda x:\text{nat}. \text{times } x \\ &\quad ((\lambda x:\text{nat}. \text{times } x ((\lambda x:\text{nat}. \mathbf{s } \mathbf{z})x))x)) \end{aligned}$$

Modulo some trivial redices of variables for variables, this is the result we would expect from the partial evaluation of the power function.

3.5 Implementation Issues

The operational semantics of Mini-ML_e[□] may be implemented by a translation into pure Mini-ML, by the mapping:

$$\begin{aligned} \Box A &\mapsto 1 \rightarrow A \\ \mathbf{box } E &\mapsto \lambda x:1. E && \text{(where } x \text{ not free in } E) \\ \mathbf{let } \mathbf{box } u = E_1 \mathbf{in } E_2 &\mapsto (\lambda x:1 \rightarrow A. [x \langle \rangle / u] E_2) E_1 && \text{(where } x \text{ not in free } E_2) \end{aligned}$$

It may then appear that the modal fragment of Mini-ML_e[□] is redundant. Note, however, that the type $1 \rightarrow A$ does not express any binding-time properties, while $\Box A$ does. It is precisely this distinction which makes Mini-ML_e[□] interesting: The type checker will reject programs which may execute correctly, but for which the desired binding-time separation is violated. Without the modal operator, this property cannot be expressed and consequently not checked.

A more efficient implementation method would be to interpret $\Box A$ as a datatype representing code that calculates a value of type A . The representation must support substitution of one code fragment into another, as required by the *ev_let_box* rule. If the code is machine code, this naturally leads to the idea of templates, as used in run-time code generation (see [Keppel et al. 1993]). For many applications this code would instead be source expressions or some intermediate language, thus allowing optimization after code substitution, as in partial evaluation [Jones et al. 1993]. In our own experiments in run-time code generation [Wickline et al. 1998; Wickline et al. 1998], following ideas in [Leone and Lee 1994; Lee and Leone 1996], expressions of type $\Box A$ are compiled into *generating extensions* which emit machine code at run-time and then jump to it to effect evaluation. References to free code variables then represent calls from one generator to another.

4. A KRIPKE-STYLE MODAL λ -CALCULUS

The modal logic in Section 2 was motivated by the goal to capture validity. A valid proposition is one with a closed proof term, and closed proof terms correspond to code which can be explicitly manipulated and safely evaluated.

In this section we construct a modal logic based on Kripke's multiple-world interpretation of modal logic [Kripke 1963]. A world corresponds to a stage of computation during evaluation. A value computed at a given stage of computation is available in all *accessible* stages, according to the accessibility relation between worlds of the Kripke semantics. Subtly different modal logics arise, depending on the properties of the accessibility relation between worlds. They are captured by

structural rules built into the elimination rule for necessity ($\Box E$). In its most general form, we exactly capture validity and thereby the intuitionistic modal logic S4 presented in Section 2.

Our rules constitute a simplification of the system $\lambda^{\rightarrow\Box}$ in [Pfenning and Wong 1995] and [Davies and Pfenning 1996]. In particular we have replaced the structural rule **pop** by a more general form of elimination which can be motivated from the perspective of pure natural deduction.

We prove the correctness of our system by relating it to the natural deduction system for S4 presented in Section 2 via two translations between proof terms. In Section 5 we extend this formulation of modal logic to Mini-ML $^{\Box}$, which leads to a staged programming style akin to Lisp's **quasiquote**, **unquote**, and **eval**. Instead of giving a direct operational semantics for this language we present a type preserving compilation to the explicit language from Section 3. We give an embedding of a two-level language [Nielson and Nielson 1992] into our language in Section 6.

4.1 A Kripke-Style Natural Deduction System

In Kripke's interpretation of modal logic, the truth of a proposition is relativized to a world. Modal operators allow us to reason about the truth of a proposition in all worlds *accessible* from the current world. Imposing laws on the accessibility relation between worlds (such as reflexivity, transitivity, or symmetry) leads to different modal logics. A world in the sense of Kripke is represented by a context of hypotheses containing propositions we know to be true in this world. Based on this intuition, our main judgment has the form

$$\Gamma_1; \Gamma_2; \dots; \Gamma_n \vdash^i A$$

which expresses that A is true in the current world Γ_n . Furthermore, Γ_1 represents hypotheses true in some initial world, Γ_2 represents hypotheses true in an *arbitrary* world reachable from Γ_1 , and so on.

From the functional point of view, Γ_i binds variables available at stage i of the computation, where the proof term assigned to A may be executed at stage n . We refer to $\Gamma_1; \dots; \Gamma_n$ as the *world stack* or *context stack* since worlds are related to contexts by the Curry-Howard isomorphism. Note that there will always be at least one context in the context stack: the current world. We abbreviate (a possibly empty) context stack by Ψ .

As in Section 2 we now systematically develop a system of natural deduction for this judgment which includes proof terms. First, only the hypotheses in the current world are available to derive the conclusion.

$$\frac{x:A \text{ in } \Gamma}{\Psi; \Gamma \vdash^i x : A} \text{ var}$$

The substitution principle applies to arbitrary worlds, as long as we establish truth in the appropriate world.

Substitution Principle

If $\Psi; \Gamma \vdash^i M_1 : A$ and $\Psi; (\Gamma, x:A, \Gamma'); \Psi' \vdash^i M_2 : C$
 then $\Psi; (\Gamma, \Gamma'); \Psi' \vdash^i [M_1/x]M_2 : C$.

In the special case that Ψ' is empty, the current worlds in both given derivations coincide. We will formally demonstrate this property of the system later.

There are two kinds of structural properties. First, we have exchange, weakening and contraction within each world in the world stack. We will not formally restate this. Second, depending on the properties of the accessibility relation, we might have some structural properties on the world stack. In K we have none. If we add reflexivity of the accessibility relation, we reason as follows: If we have a judgment

$$\Psi; \Gamma; \Gamma'; \Psi' \vdash^i M : C$$

then Γ' contains hypotheses assumed to be true in an *arbitrary* world accessible from Γ . But Γ itself is accessible from Γ (by reflexivity), so C should still be true if we join Γ and Γ' .

$$\Psi; (\Gamma, \Gamma'); \Psi' \vdash^i M' : C$$

Whether M' is different from M depends on how much information is present in the term itself, as we shall see later. We refer to this as *modal fusion*.

For example, omitting proof terms, we read the judgment

$$\Box(A \rightarrow B); A \vdash^i B$$

as

If $\Box(A \rightarrow B)$ is true in some world w_1 and A is true in an arbitrary world w_2 reachable from w_1 , then B is true in w_2 .

If we assume reflexivity of the accessibility relation, we know that w_1 is accessible from w_1 , so we can infer from the above by replacing w_2 with w_1 :

If $\Box(A \rightarrow B)$ is true in w_1 and A is true in w_1 then B is true in w_1 .

In symbolic form, we write this as

$$\Box(A \rightarrow B), A \vdash^i B$$

which is exactly the result of fusion applied to the first judgment.

Next we add transitivity. Consider once again

$$\Psi; \Gamma; \Gamma'; \Psi' \vdash^i M : C.$$

Then if we add an *arbitrary* world accessible from Γ from which Γ' can be reached, the judgment should continue to hold, because Γ' is still accessible from Γ by transitivity.

$$\Psi; \Gamma; \cdot; \Gamma'; \Psi' \vdash^i M' : C$$

Again, M' may be identical to M or it can be directly created from M , depending how much information we represent in proof terms. We refer to this property as *modal weakening*. Note that by ordinary weakening, the new interposed world may also contain arbitrary assumptions without invalidating the judgment.

Now we define the connectives via their introduction and elimination rules. Implication only affects the current world and is similar to what we have presented in Section 2.

$$\frac{\Psi; (\Gamma, x:A) \vdash^i M : B}{\Psi; \Gamma \vdash^i \lambda x:A. M : A \rightarrow B} \rightarrow \text{I}$$

$$\frac{\Psi; \Gamma \vdash^i M : A \rightarrow B \quad \Psi; \Gamma \vdash^i N : A}{\Psi; \Gamma \vdash^i M N : B} \rightarrow \text{E}$$

Local soundness and completeness also works as before. The corresponding operations on proof terms are the familiar β -reduction and η -expansion.

Recall that $\Box A$ should be true in the current world if A is true in every reachable world. Since we have no information about the reachable worlds, we have no hypotheses about the truth of propositions in this world. Hence the introduction rule reads

$$\frac{\Psi; \Gamma; \cdot \vdash^i M : A}{\Psi; \Gamma \vdash^i \mathbf{box} M : \Box A} \Box \text{I}$$

The corresponding elimination rule states that if $\Box A$ is true in the current world, A must be true in every reachable world. Which worlds are reachable depends, of course, on the accessibility relation for the modal logic under consideration. In its most general form (S4), the elimination rule reads

$$\frac{\Psi; \Gamma \vdash^i M : \Box A}{\Psi; \Gamma; \Gamma_1; \dots; \Gamma_n \vdash^i \mathbf{unbox}_n M : A} \Box \text{E}$$

Note that Γ_1 is always accessible from Γ , so in K we only have \mathbf{unbox}_1 . The worlds $\Gamma_2, \dots, \Gamma_n$ are accessible from Γ only because of transitivity, so in modal logic with transitivity we also have \mathbf{unbox}_n for $n > 1$. Γ is accessible from itself in a modal logic whose accessibility relation is reflexive, so there we also have \mathbf{unbox}_0 .

Next we consider local soundness and completeness of the rules for the modal operator. Recall that soundness requires that an introduction rule followed by an elimination rule can be reduced to a direct derivation of the judgment.

$$\frac{\frac{\mathcal{D}}{\Psi; \Gamma; \cdot \vdash^i M : A} \Box \text{I}}{\Psi; \Gamma \vdash^i \mathbf{box} M : \Box A} \Box \text{E} \quad \Longrightarrow \quad \frac{\mathcal{D}'}{\Psi; \Gamma; \Gamma_1; \dots; \Gamma_n \vdash^i M' : A}$$

$$\frac{\Psi; \Gamma; \Gamma_1; \dots; \Gamma_n \vdash^i \mathbf{unbox}_n (\mathbf{box} M) : A}{\Psi; \Gamma; \Gamma_1; \dots; \Gamma_n \vdash^i M' : A} \Box \text{E}$$

where \mathcal{D}' and M' exist by the structural properties of world stacks (modal fusion in the case that $n = 0$, ordinary weakening in the case the $n = 1$, and ordinary weakening and $n - 1$ applications of modal weakening in the case that $n > 1$).

Local completeness is a bit simpler. We have the following η -expansion:

$$\Psi; \Gamma \vdash^i M : \Box A \stackrel{\mathcal{D}}{\Longrightarrow} \frac{\frac{\mathcal{D} \quad \Psi; \Gamma \vdash^i M : \Box A}{\Psi; \Gamma; \cdot \vdash^i \mathbf{unbox}_1 M : A} \Box E}{\Psi; \Gamma \vdash^i \mathbf{box} (\mathbf{unbox}_1 M) : \Box A} \Box I$$

We postpone a formal discussion of the rules for β -reduction and η -expansion on terms to Section 4.4.

There are two simple and consistent variations on this system.

The first arises from an analysis of local completeness: one can see that only \mathbf{unbox}_1 is necessary. The others (which are locally sound!) have been incorporated so that we need no explicit structural rules. However, \mathbf{unbox}_1 plus explicit rules for modal fusion and weakening also make a sensible system with the same derivable judgments. For our purposes, a formulation without explicit structural rules is preferable, since it allows more compact programs and a simpler meta-theory.

In the second variant we replace the constructor \mathbf{unbox}_n simply by \mathbf{unbox} . This would mean that $M' = M$ in the local reduction for \Box , and terms remain invariant under structural transformation of contexts. This more streamlined presentation of the calculus is not appropriate for our application, since the the index n in a term $\mathbf{unbox}_n M$ constructor determines the stage at which M is to be evaluated. Without the index, this would be ambiguous and depend on the typing derivation. In other words, the system would not be *coherent*.

4.2 Syntax

We summarize the syntax of the pure fragment.

Types	$A ::= a \mid A_1 \rightarrow A_2 \mid \Box A$
Terms	$M ::= x \mid \lambda x:A. M \mid M_1 M_2 \mid \mathbf{box} M \mid \mathbf{unbox}_n M$
Contexts	$\Gamma ::= \cdot \mid \Gamma, x:A$
Context Stacks	$\Psi ::= \odot \mid \Psi; \Gamma$

Here and subsequently we write \odot for the empty context stack to avoid possible confusion with the empty context.

4.3 Natural Deduction Judgment

We summarize the rules defining the main judgment, $\Psi; \Gamma \vdash^i M : A$ as motivated and developed in Section 4.1.

$$\begin{array}{c}
\frac{x:A \text{ in } \Gamma}{\Psi; \Gamma \vdash^i x : A} \text{var} \qquad \frac{\Psi; (\Gamma, x:A) \vdash^i M : B}{\Psi; \Gamma \vdash^i \lambda x:A. M : A \rightarrow B} \rightarrow \text{I} \\
\\
\frac{\Psi; \Gamma \vdash^i M : A \rightarrow B \quad \Psi; \Gamma \vdash^i N : A}{\Psi; \Gamma \vdash^i M N : B} \rightarrow \text{E} \\
\\
\frac{\Psi; \Gamma; \cdot \vdash^i M : A}{\Psi; \Gamma \vdash^i \mathbf{box} M : \Box A} \Box \text{I} \qquad \frac{\Psi; \Gamma \vdash^i M : \Box A}{\Psi; \Gamma; \Gamma_1; \dots; \Gamma_n \vdash^i \mathbf{unbox}_n M : A} \Box \text{E}
\end{array}$$

4.4 Properties of the Kripke-style λ -calculus

Structural transformations change the nature of the proof term by relabelling indices to the **unbox** constructor. Such a relabelling is also necessary to write out the rules for β -reduction and η -expansion. We define $\{n/m\}M$ inductively on the structure of M .

$$\begin{aligned}
\{n/m\}x &= x \\
\{n/m\}\lambda x:A. M &= \lambda x:A. \{n/m\}M \\
\{n/m\}M_1 M_2 &= (\{n/m\}M_1) (\{n/m\}M_2) \\
\{n/m\}\mathbf{box} M &= \mathbf{box} \{n/m+1\}M \\
\{n/m\}\mathbf{unbox}_p M &= \begin{cases} \mathbf{unbox}_p \{n/m-p\}M & \text{for } p < m \\ \mathbf{unbox}_{p+n-1} M & \text{for } p \geq m \end{cases}
\end{aligned}$$

This operation now allows us to state the substitution and structural transformation properties.

LEMMA 9 (MODAL STRUCTURAL TRANSFORMATION).

If $\Psi; \Gamma_0; \Delta_1; \dots; \Delta_m \vdash^i M : A$ then $\Psi; \Gamma_0; \dots; (\Gamma_n, \Delta_1); \dots; \Delta_m \vdash^i \{n/m\}M : A$.

PROOF. By induction over the structure of the given derivation \mathcal{D} of

$$\Psi; \Gamma_0; \Delta_1; \dots; \Delta_m \vdash^i M : A$$

In each case except for $\Box \text{E}$ we immediately apply the induction hypothesis and reconstruct an appropriate derivation from the results. In the case of $\Box \text{E}$ we distinguish two subcases.

Case:

$$\mathcal{D} = \frac{\mathcal{D}_1 \quad \Psi; \Gamma_0; \Delta_1; \dots; \Delta_{m-p} \vdash^i M_1 : \Box A}{\Psi; \Gamma_0; \Delta_1; \dots; \Delta_m \vdash^i \mathbf{unbox}_p M_1 : A} \Box \text{E}$$

for $p < m$. Then

$$\begin{array}{ll}
\Psi; \Gamma_0; \dots; (\Gamma_n, \Delta_1); \dots; \Delta_{m-p} \vdash^i \{n/m-p\}M_1 : \Box A & \text{by ind. hyp.} \\
\Psi; \Gamma_0; \dots; (\Gamma_n, \Delta_1); \dots; \Delta_m \vdash^i \mathbf{unbox}_p \{n/m-p\}M_1 : A & \text{by rule } \Box \text{E} \\
\Psi; \Gamma_0; \dots; (\Gamma_n, \Delta_1); \dots; \Delta_m \vdash^i \{n/m\}\mathbf{unbox}_p M_1 : A & \text{by definition of } \{n/m\}
\end{array}$$

Case:

$$\mathcal{D} = \frac{\mathcal{D}_1 \quad \Psi'; \Theta_{p-m} \vdash^i M_1 : \Box A}{\Psi'; \Theta_{p-m}; \dots; \Theta_0; \Delta_1; \dots \Delta_m \vdash^i \mathbf{unbox}_p M_1 : A} \Box E$$

for $p \geq m$ where $\Psi = \Psi'; \dots; \Theta_1$ and $\Theta_0 = \Gamma_0$. Then

$$\begin{aligned} & \Psi'; \Theta_{p-m}; \dots; \Theta_0; \Gamma_1; \dots; (\Gamma_n, \Delta_1); \dots \Delta_m \vdash^i \mathbf{unbox}_{p+n-1} M_1 : A \\ & \hspace{15em} \text{by rule } \Box E \text{ applied to } \mathcal{D}_1 \\ & \Psi; \Gamma_0; \dots; (\Gamma_n, \Delta_1); \dots; \Delta_m \vdash^i \{n/m\} \mathbf{unbox}_p M_1 : A \\ & \hspace{15em} \text{by definition of } \Psi, \Theta_0 \text{ and } \{n/m\}. \end{aligned}$$

□

The system also satisfies the usual structural properties of exchange, weakening and contraction in each of the contexts in the context stack. We only state the substitution property formally.

LEMMA 10 (SUBSTITUTION).

If $\Psi; \Gamma \vdash^i M_1 : A$ and $\Psi; (\Gamma, x:A, \Gamma'); \Psi' \vdash^i M_2 : C$ then $\Psi; (\Gamma, \Gamma'); \Psi' \vdash^i [M_1/x]M_2 : C$.

PROOF. By induction over the structure of the derivation of $\Psi; (\Gamma, x:A, \Gamma'); \Psi' \vdash^i M_2 : C$. □

Then we have the rules of β -reduction and η -expansion, corresponding to local reduction and expansion in natural deduction.

$$\begin{aligned} (\lambda x:A. M) N & \xrightarrow{\beta} [N/x]M \\ M : A_1 \rightarrow A_2 & \xrightarrow{\eta} \lambda x:A_1. M x \\ \mathbf{unbox}_n (\mathbf{box} M) & \xrightarrow{\beta} \{n/1\}M \\ M : \Box A & \xrightarrow{\eta} \mathbf{box} (\mathbf{unbox}_1 M) \end{aligned}$$

THEOREM 11 (SUBJECT REDUCTION AND EXPANSION).

- (1) If $\Psi; \Gamma \vdash^i M : A$ and $M \xrightarrow{\beta} M'$ then $\Psi; \Gamma \vdash^i M' : A$.
- (2) If $\Psi; \Gamma \vdash^i M : A$ and $M : A \xrightarrow{\eta} M'$ then $\Psi; \Gamma \vdash^i M' : A$

PROOF. In each case we apply inversion to the given typing derivation. For subject reduction we then either use modal structural transformation (Lemma 9) or substitution (Lemma 10). For subject expansion we directly construct a derivation of the conclusion, using weakening if necessary. □

4.5 Environments and Environment Stacks

In order to prove that the explicit (see Section 2) and implicit (Kripke style) formulations of S4 correspond, we need to develop some properties of environments and environment stacks. Roughly, an environment provides definitions for the modal variables available at a particular stage of the computation, while an environment stack extends this to all stages of a computation.

We define environments and stacks which bind patterns of the form $\mathbf{box} u$ to explicit terms E .

$$\begin{array}{l} \text{Environments} \quad \rho ::= \cdot \mid \rho, \mathbf{box} u = E \\ \text{Environment Stacks} \quad R ::= \odot \mid R; \rho \end{array}$$

The translations between the λ -calculi based on validity and multiple worlds require us to relate context pairs $\Delta; \Gamma$ to context stacks Ψ . This is achieved by the following typing judgments for environments and environment stacks. The latter ties in the context stacks of the implicit system. We use Θ to range over modal contexts.

$$\begin{array}{l} \Delta; \Gamma \vdash^e \rho : \Theta \quad \text{Environment } \rho \text{ matches } \Theta \text{ in contexts } \Delta \text{ and } \Gamma \\ \Psi \models^e R : \Delta \quad \text{Environment stack } R \text{ matches } \Delta \text{ in context stack } \Psi \end{array}$$

$$\frac{}{\Delta; \Gamma \vdash^e \cdot : \cdot} \text{tpv_empty}$$

$$\frac{\Delta; \Gamma \vdash^e \rho : \Theta \quad (\Delta, \Theta); \Gamma \vdash^e E : \Box A}{\Delta; \Gamma \vdash^e (\rho, \mathbf{box} u = E) : (\Theta, u:A)} \text{tpv_bind}$$

$$\frac{}{\Psi \models^e \odot : \cdot} \text{tpr_empty}$$

$$\frac{\Psi \models^e R : \Delta \quad \Delta; \Gamma \vdash^e \rho : \Theta}{\Psi; \Gamma \models^e (R; \rho) : (\Delta, \Theta)} \text{tpr_env}$$

These rules are equivalent to the following non-inductive ones, which are given here to aid the reader:

$$\frac{\forall 1 \leq i \leq n. \quad \Delta, u_1:A_1, \dots, u_{i-1}:A_{i-1}; \Gamma \vdash^e E_i : \Box A_i}{\Delta; \Gamma \vdash^e (\mathbf{box} u_1 = E_1, \dots, \mathbf{box} u_n = E_n) : (u_1:A_1, \dots, u_n:A_n)} \text{tpv_bind_n}$$

$$\frac{\forall 1 \leq i \leq n. \quad (\Theta_1, \dots, \Theta_{i-1}); \Gamma_i \vdash^e \rho_i : \Theta_i}{\Psi; \Gamma_1; \dots; \Gamma_n \models^e (\odot; \rho_1; \dots; \rho_n) : (\Theta_1, \dots, \Theta_n)} \text{tpr_env_n}$$

We will tacitly use weakening for typing of environment stacks, which directly follows from weakening on the typing judgment for the explicit modal λ -calculus. We also need to use the following property.

LEMMA 12 (ENVIRONMENT EXTENSION).

If $\Psi; \Gamma \models^e (R; \rho) : \Delta$ and $\Delta; \Gamma \vdash^e E : \Box A$ then $\Psi; \Gamma \models^e (R; \rho, \mathbf{box} u = E) : (\Delta, u:A)$.

PROOF. By inversion on the derivation for $(R; \rho)$ followed by a straightforward application of the typing rules for environments and environment stacks. \square

4.6 Translation from Explicit System

In this section we show that if A is true in the explicit system then A is also true in the implicit system. We show this by giving a translation on proof terms.

The difficulty in defining and proving the correctness of this translation lies in the relation between the modal and ordinary contexts on the explicit side and the context stack on the implicit side. This relationship can be maintained via the environment stacks defined in the preceding section.

$R; \rho \triangleright E \mapsto M$ Expression E translates to M in environment stack $R; \rho$

This judgment is defined by the following rules.

$$\begin{array}{c}
 \frac{}{R; \rho \triangleright x \mapsto x} \text{tx_ovar} \quad \frac{R; \rho \triangleright E \mapsto M}{R; \rho \triangleright \lambda x:A. E \mapsto \lambda x:A. M} \text{tx_lam} \\
 \\
 \frac{R; \rho \triangleright E_1 \mapsto M_1 \quad R; \rho \triangleright E_2 \mapsto M_2}{R; \rho \triangleright E_1 E_2 \mapsto M_1 M_2} \text{tx_app} \\
 \\
 \frac{R; \rho; \cdot \triangleright E \mapsto M}{R; \rho \triangleright \mathbf{box} E \mapsto \mathbf{box} M} \text{tx_box} \\
 \\
 \frac{R; \rho, \mathbf{box} u = E_1 \triangleright E_2 \mapsto M}{R; \rho \triangleright \mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2 \mapsto M} \text{tx_letbox} \\
 \\
 \frac{R; \rho'_n \triangleright E \mapsto M}{R; (\rho'_n, \mathbf{box} u = E, \rho''_n); \cdots; \rho_0 \triangleright u \mapsto \mathbf{unbox}_n M} \text{tx_mvar}
 \end{array}$$

THEOREM 13 (TRANSLATION FROM EXPLICIT SYSTEM).

Given $\Psi; \Gamma \Vdash^e (R; \rho) : \Delta$ and $\Delta; \Gamma \Vdash^e E : A$.

- (1) There is a unique M such that $R; \rho \triangleright E \mapsto M$.
- (2) Whenever $R; \rho \triangleright E \mapsto M$ then $\Psi; \Gamma \Vdash^i M : A$.

PROOF. Proposition 1 is proven by induction on the multiset extension of the subterm ordering of expressions in $R; \rho$ and E . For the case of a modal variable u we need to use the typing assumption to guarantee that `tx_mvar` applies, that is, that the environment stack contains an appropriate definition of u .

Proposition 2 is proven by induction on the structure of the derivation of $R; \rho \triangleright E \mapsto M$, applying inversion to the given typing derivations. In the case of modal variables u we have an auxiliary induction on the world index n .

We now show the proof of Proposition 2 in more detail. We assume we are given derivations

$$\Psi; \Gamma \Vdash^e (R; \rho) : \Delta \quad \text{and} \quad \Delta; \Gamma \Vdash^e E : A \quad \text{and} \quad R; \rho \triangleright E \mapsto M$$

We proceed by induction on the structure of \mathcal{T} , applying inversion to the typing derivations as needed in order to construct a derivation

$$\begin{array}{c}
 \mathcal{D} \\
 \Psi; \Gamma \Vdash^i M : A
 \end{array}$$

Case:

$$\mathcal{T} = \frac{}{R; \rho \triangleright x \mapsto x} \text{tx_ovar}$$

$\Delta; \Gamma \vdash^e x : A$ by assumption
 $x:A$ in Γ by inversion
 $\Psi; \Gamma \vdash^i x : A$ by rule var

Case:

$$\mathcal{T} = \frac{\frac{\mathcal{T}_2}{R; \rho \triangleright E_2 \mapsto M_2}}{R; \rho \triangleright \lambda x:A_1. E_2 \mapsto \lambda x:A_1. M_2} \text{tx_lam}$$

$\Delta; \Gamma \vdash^e \lambda x:A_1. E_2 : A$ by assumption
 $\Delta; \Gamma, x:A_1 \vdash^e E_2 : A_2$ and $A = A_1 \rightarrow A_2$ by inversion
 $\Psi; \Gamma, x:A_1 \models^e (R; \rho) : \Delta$ by assumption and weakening
 $\Psi; \Gamma, x:A_1 \vdash^i M_2 : A_2$ by ind. hyp.
 $\Psi; \Gamma \vdash^i \lambda x:A_1. M_2 : A_1 \rightarrow A_2$ by rule \rightarrow I

Case: tx_app is straightforward.

Case:

$$\mathcal{T} = \frac{\frac{\mathcal{T}_1}{R; \rho; \cdot \triangleright E_1 \mapsto M_1}}{R; \rho \triangleright \mathbf{box} E_1 \mapsto \mathbf{box} M_1} \text{tx_box}$$

$\Delta; \Gamma \vdash^e \mathbf{box} E_1 : A$ by assumption
 $\Delta; \cdot \vdash^e E_1 : A_1$ and $A = \Box A_1$ by inversion
 $\Psi; \Gamma \models^e (R; \rho) : \Delta$ by assumption
 $\Delta; \cdot \vdash^e \cdot : \cdot$ by tpv_empty
 $\Psi; \Gamma; \cdot \models^e (R; \rho; \cdot) : \Delta$ by tpr_env
 $\Psi; \Gamma; \cdot \vdash^i M_1 : A_1$ by ind. hyp.
 $\Psi; \Gamma \vdash^i \mathbf{box} M_1 : \Box A_1$ by \Box I

Case:

$$\mathcal{T} = \frac{\frac{\mathcal{T}_1}{R; \rho, \mathbf{box} u = E_1 \triangleright E_2 \mapsto M_2}}{R; \rho \triangleright \mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2 \mapsto M_2} \text{tx_letbox}$$

$\Delta; \Gamma \vdash^e \mathbf{let} \mathbf{box} u = E_1 \mathbf{in} E_2 : A$ by assumption
 $\Delta; \Gamma \vdash^e E_1 : \Box A_1$ and
 $(\Delta, u:A_1); \Gamma \vdash^e E_2 : A$ by inversion
 $\Psi; \Gamma \models^e (R; \rho) : \Delta$ by assumption
 $\Psi; \Gamma \models^e (R; \rho, \mathbf{box} u = E_1) : (\Delta, u:A_1)$ by Lemma 12
 $\Psi; \Gamma \vdash^i M_2 : A$ by ind. hyp.

Case:

$$\mathcal{T} = \frac{\frac{\mathcal{T}_1}{R_n; \rho'_n \triangleright E' \mapsto M'}}{R_n; (\rho'_n, \mathbf{box} u = E', \rho''_n); \cdots; \rho_0 \triangleright u \mapsto \mathbf{unbox}_n M'} \text{tx_mvar}$$

$\Delta; \Gamma \vdash^e u : A$	by assumption
$u : A \text{ in } \Delta$	by inversion
$\Psi_0; \Gamma_0 \Vdash^e R_n; (\rho'_n, \mathbf{box} u = E', \rho''_n); \dots; \rho_0 : \Delta$	
for $\Psi_0 = \Psi, \Gamma_0 = \Gamma$ and $\rho_0 = \rho$	by assumption
$\Delta = \Delta_0, \Theta_0,$	
$\Psi_0 \Vdash^e R_n; (\rho'_n, \mathbf{box} u = E', \rho''_n); \dots; \rho_1 : \Delta_0,$ and	
$\Delta_0; \Gamma_0 \vdash^e \rho_0 : \Theta_0$	by inversion
$\Psi_{i-1} = \Psi_i; \Gamma_i,$	
$\Delta_{i-1} = \Delta_i, \Theta_i,$	
$\Psi_i \Vdash^e R_n; (\rho'_n, \mathbf{box} u = E', \rho''_n); \dots; \rho_{i+1} : \Delta_i,$ and	
$\Delta_i; \Gamma_i \vdash^e \rho_i : \Theta_i$	by inversion for $i = 1, \dots, (n-1)$
$\Psi_{n-1} = \Psi_n; \Gamma_n,$	
$\Delta_{n-1} = \Delta_n, \Theta_n,$	
$\Psi_n \Vdash^e R_n : \Delta_n,$ and	
$\Delta_n; \Gamma_n \vdash^e (\rho'_n, \mathbf{box} u = E', \rho''_n) : \Theta_n$	by further inversion
$\Theta_n = \Theta'_n, u : A', \Theta''_n,$	
$\Delta_n; \Gamma_n \vdash^e \rho'_n : \Theta'_n$ and	
$(\Delta_n, \Theta'_n); \Gamma_n \vdash^e E' : \Box A'$	by further inversion
$A' = A$	since $u : A$ in $\Delta = \Delta_n, \Theta_n, \dots, \Theta_0$ is unique
$\Psi_n; \Gamma_n \Vdash^e (R_n; \rho'_n) : (\Delta_n, \Theta'_n)$	by rule <code>tpr.env</code>
$\Psi_n; \Gamma_n \vdash^i M' : \Box A$	by ind. hyp.
$\Psi_n; \Gamma_n; \dots; \Gamma_0 \vdash^i \mathbf{unbox}_n M' : A$	by rule <code>\Box E</code>

□

A special case of this theorem is the translation of a whole, well typed term: if $\cdot; \cdot \vdash^e E : A$ then there exists a unique M such that $\odot; \cdot \triangleright E \mapsto M$, and further we have that $\odot; \cdot \vdash^i M : A$.

4.7 Translation to Explicit System

To show that every proposition judged true in the implicit system is also true in the explicit system, we give another type-preserving translation on proof terms. This translation is the core of the compilation function we consider in Section 5.4. Again, the difficulty lies mainly in relating the context stack of the implicit system to the modal and ordinary contexts of the explicit systems.

The translation recursively extracts terms inside \mathbf{unbox}_n constructors and binds their translation to new variables, bound with a $\mathbf{let\ box}$ outside the n^{th} enclosing \mathbf{box} constructor. Variables thus bound occur exactly once.

We abstract over an environment by means of nested $\mathbf{let\ box}$ expressions.

$$\begin{aligned} \text{Let}(\cdot)(E) &= E \\ \text{Let}(\rho, \mathbf{box} u = E')(E) &= \text{Let}(\rho)(\mathbf{let\ box} u = E' \mathbf{in} E) \end{aligned}$$

We require a few straightforward properties of environments, but we explicitly state only the derived typing rule for environment abstractions.

$$\frac{\Delta; \Gamma \vdash^e \rho : \Theta \quad (\Delta, \Theta); \Gamma \vdash^e E : B}{\Delta; \Gamma \vdash^e \text{Let}(\rho)(E) : B} \text{tpe.env}$$

The merge operation $R_1 \mid R_2$ on environment stacks appends corresponding environments. We assume that the domains of the environments in R_1 and R_2 are disjoint so that the resulting environment stack is valid.

$$\begin{aligned} \odot \mid R_2 &= R_2 \\ R_1 \mid \odot &= R_1 \\ (R_1; \rho_1) \mid (R_2; \rho_2) &= (R_1 \mid R_2); (\rho_1, \rho_2) \end{aligned}$$

The translation is defined by the judgment

$$M \mapsto R \triangleright E \quad M \text{ compiles to term } E \text{ under stack } R$$

It is defined by the following rules.

$$\begin{array}{c} \frac{}{x \mapsto \odot \triangleright x} \text{tr_var} \quad \frac{M \mapsto R \triangleright E}{\lambda x:A. M \mapsto R \triangleright \lambda x:A. E} \text{tr_lam} \\ \\ \frac{M_1 \mapsto R_1 \triangleright E_1 \quad M_2 \mapsto R_2 \triangleright E_2}{M_1 M_2 \mapsto (R_1 \mid R_2) \triangleright E_1 E_2} \text{tr_app} \\ \\ \frac{M \mapsto \odot \triangleright E}{\mathbf{box} M \mapsto \odot \triangleright \mathbf{box} E} \text{tr_box0} \quad \frac{M \mapsto (R; \rho) \triangleright E}{\mathbf{box} M \mapsto R \triangleright \text{Let}(\rho)(\mathbf{box} E)} \text{tr_box1} \\ \\ \frac{M \mapsto R \triangleright E}{\mathbf{unbox}_0 M \mapsto R \triangleright \mathbf{let} \mathbf{box} u = E \mathbf{in} u} \text{tr_unbox0} \\ \\ \frac{M \mapsto R \triangleright E}{\mathbf{unbox}_{n+1} M \mapsto R; (\mathbf{box} u = E); \underbrace{\cdot; \dots; \cdot}_{n} \triangleright u} \text{tr_unbox1} \end{array}$$

The **tr_app** rule is restricted to context stacks R_1 and R_2 with disjoint domains. This can always be achieved by renaming of variables in the derivations of the two premisses.

THEOREM 14 (TRANSLATION TO EXPLICIT SYSTEM).

- (1) For any M there exist unique R and E such that $M \mapsto R \triangleright E$.
- (2) If $M \mapsto R \triangleright E$ and $\Psi; \Gamma \vdash^i M : A$ then for some Δ we have $\Psi \vdash^e R : \Delta$ and $\Delta; \Gamma \vdash^e E : A$.

PROOF. Proposition 1 is straightforward, since the translation is defined structurally on M with unique results (modulo renaming of bound variables, of course), except in the case of **box** M , where exactly one of **tr_box0** and **tr_box1** apply.

Proposition 2 follows by induction on the structure of the derivation \mathcal{T} of $M \mapsto R \triangleright E$. The proof requires a few simple lemmas such as weakening for \vdash^e and some immediate properties of $R_1 \mid R_2$ and $\text{Let}(\rho)(E)$ which we do not state here explicitly. We omit the cases for the non-modal constructors, which are straightforward.

Case:

$$\mathcal{T} = \frac{\mathcal{T}_1 \quad M_1 \mapsto \odot \triangleright E_1}{\mathbf{box} \ M_1 \mapsto \odot \triangleright \mathbf{box} \ E_1} \text{tr_box0}$$

$\Psi; \Gamma \vdash^i \mathbf{box} \ M_1 : A$ by assumption
 $\Psi; \Gamma; \cdot \vdash^i M_1 : A_1$ and $A = \square A_1$ by inversion
 $\Psi; \Gamma \models^e \odot : \Delta_1$ and by ind. hyp.
 $\Delta_1; \cdot \models^e E_1 : A_1$ for some Δ_1 by inversion
 $\Delta_1 = \cdot$ by tpr_empty
 $\Psi \models^e \odot : \cdot$ by \square
 $\cdot; \Gamma \models^e \mathbf{box} \ E_1 : \square A_1$

The last two lines are the desired conclusions for $\Delta = \cdot$.

Case:

$$\mathcal{T} = \frac{\mathcal{T}_1 \quad M_1 \mapsto (R; \rho) \triangleright E_1}{\mathbf{box} \ M_1 \mapsto R \triangleright \text{Let}(\rho)(\mathbf{box} \ E_1)} \text{tr_box1}$$

$\Psi; \Gamma \vdash^i \mathbf{box} \ M_1 : A$ by assumption
 $\Psi; \Gamma; \cdot \vdash^i M_1 : A_1$ and $A = \square A_1$ by inversion
 $\Psi; \Gamma \models^e (R; \rho) : \Delta_1$ and by ind. hyp.
 $\Delta_1; \cdot \models^e E_1 : A_1$ for some Δ_1 by inversion
 $\Psi \models^e R : \Delta'_1$ and by rule \square
 $\Delta'_1; \Gamma \models^e \rho : \Theta$ and $\Delta_1 = \Delta'_1, \Theta$ by rule tpi_env
 $(\Delta'_1, \Theta); \Gamma \models^e \mathbf{box} \ E_1 : \square A_1$
 $\Delta'_1; \Gamma \models^e \text{Let}(\rho)(\mathbf{box} \ E_1) : \square A_1$

Now we have the desired conclusions with $\Delta = \Delta'_1$.

Case:

$$\mathcal{T} = \frac{\mathcal{T}_1 \quad M_1 \mapsto R \triangleright E_1}{\mathbf{unbox}_0 \ M_1 \mapsto R \triangleright \mathbf{let} \ \mathbf{box} \ u = E_1 \ \mathbf{in} \ u} \text{tr_unbox0}$$

$\Psi; \Gamma \vdash^i \mathbf{unbox}_0 \ M_1 : A$ by assumption
 $\Psi; \Gamma \vdash^i M_1 : \square A$ by inversion
 $\Psi \models^e R : \Delta_1$ and by ind. hyp.
 $\Delta_1; \Gamma \models^e E_1 : \square A$ for some Δ_1 by rule tpe_mvar
 $(\Delta_1, u:A); \Gamma \models^e u : A$ by rule $\square E$
 $\Delta_1; \Gamma \models^e \mathbf{let} \ \mathbf{box} \ u = E_1 \ \mathbf{in} \ u : A$

Now we have the desired conclusions with $\Delta = \Delta_1$.

Case:

$$\mathcal{T} = \frac{\mathcal{T}_1 \quad M_1 \mapsto R \triangleright E_1}{\mathbf{unbox}_{n+1} \ M_1 \mapsto R; (\mathbf{box} \ u = E_1); \underbrace{\cdot; \dots; \cdot}_n \triangleright u} \text{tr_unbox1}$$

$\Psi; \Gamma \vdash^i \mathbf{unbox}_{n+1} M_1 : A$	by assumption
$\Psi'; \Gamma' \vdash^i M_1 : \Box A$ and $\Psi = \Psi'; \Gamma'; \Gamma_1; \dots; \Gamma_n$	by inversion
$\Psi' \Vdash^e R : \Delta_1$ and	
$\Delta_1; \Gamma' \Vdash^e E_1 : \Box A$ for some Δ_1	by ind. hyp.
$\Delta_1; \Gamma' \Vdash^e \cdot : \cdot$	by rule <code>tpv_empty</code>
$\Delta_1; \Gamma' \Vdash^e (\mathbf{box} u = E_1) : (u : A)$	by rule <code>tpv_bind</code>
$\Psi'; \Gamma' \Vdash^e (R; \mathbf{box} u = E_1) : (\Delta_1, u:A)$	by rule <code>tpr_env</code>
$\Psi'; \Gamma'; \Gamma_1; \dots; \Gamma_n \Vdash^e (R; \mathbf{box} u = E; \cdot; \dots; \cdot) : (\Delta_1, u:A)$	by n applications of <code>tpr_env</code>
$(\Delta_1, u:A); \Gamma \Vdash^e u : A$	by <code>tpe_mvar</code>

Now we have the desired conclusions for $\Delta = \Delta_1, u:A$. \square

A special case of this theorem is the translation of a whole, well typed term: if $\odot; \cdot \vdash^i M : A$ then there exists a unique E such that $M \mapsto \odot \triangleright E$, and further we have that $\cdot; \cdot \Vdash^e E : A$. Here, the emptiness of the context stack forces the environment stack to be empty.

5. MODAL MINI-ML: IMPLICIT FORMULATION

We now define Mini-ML[□], an “implicit” formulation of modal Mini-ML generalizing the λ -calculus core from the preceding section. The main advantage of this system over the explicit language is that altering the staging of a computation in a given program often only requires the insertion or deletion of modal constructors. In contrast, Mini-ML_e[□] requires that the structure of the program exactly mirror the staging, since the only way to refer to results from a previous stage is via code variables. Using **let** (a derived form in our fragment) to bind code variables we can still express staging more explicitly in Mini-ML[□] if we prefer; it is now a matter of style rather than a property enforced in the language.

Another motivation for Mini-ML[□] is that it can be directly related to the two-level λ -calculus (see Section 6) which would be much more difficult for Mini-ML_e[□] due to the different syntactic structuring required. Further, Mini-ML[□] is very similar to the quasi-quoting and eval mechanisms in LISP, which are relatively intuitive in practice. We believe that with some syntactic sugar along the lines of Scheme’s backquote and comma notation (as in the regular expression example in Section 7.3), Mini-ML[□] is a practical and theoretically well-founded basis for an extension of Standard ML. Indeed, experience with the two languages PML [Wickline et al. 1998] and Meta-ML [Taha and Sheard 1997; Taha et al. 1998; Moggi et al. 1999] indicates that such languages are indeed practical.

It may be helpful to consider the modal fragment of the implicit language to be a statically typed analog to the quasiquote mechanism in Scheme. Then **box** corresponds to **quasiquote** (‘) and **unbox**₁ to **unquote** (,). **unbox**₀ corresponds to **eval**. More generally, **unbox** _{n} corresponds to a generalized **unquote** which splices a quoted expression into a context with n levels of quasi-quoting. Note however that this analogy can also sometimes be misleading, and the actual behavior of code is closer to the quotations of a “semantically rationalized dialect” of Lisp called 2-Lisp [Smith 1984].

The operational semantics of the new system is given in terms of a type-preserving compilation to Mini-ML_e[□] which is a straightforward extension of the translation in

Section 4.7.

For some applications, such as emulating the two-level λ -calculus, weaker modal logics such as K are sufficient, as described in Section 6.4.

5.1 Syntax

We extend the logic to the core of a programming language as in Section 3.

$$\begin{array}{l}
\text{Types } A ::= \text{nat} \mid A_1 \rightarrow A_2 \mid A_1 \times A_2 \mid 1 \mid \Box A \\
\text{Terms } M ::= x \mid \lambda x:A. M \mid M_1 M_2 \\
\quad \mid \mathbf{box} M \mid \mathbf{unbox}_n M \\
\quad \mid \langle M_1, M_2 \rangle \mid \mathbf{fst} M \mid \mathbf{snd} M \\
\quad \mid \langle \rangle \\
\quad \mid \mathbf{z} \mid \mathbf{s} M \mid (\mathbf{case} M_1 \mathbf{of} \mathbf{z} \Rightarrow M_2 \mid \mathbf{s} x \Rightarrow M_3) \\
\quad \mid \mathbf{fix} x:A. M \\
\\
\text{Contexts } \Gamma ::= \cdot \mid \Gamma, x:A \\
\text{Context Stacks } \Psi ::= \odot \mid \Psi; \Gamma
\end{array}$$

5.2 Typing Rules

In this section we present typing rules for Mini-ML[□] using context stacks. The typing judgment has the form:

$\Psi; \Gamma \vdash^i M : A$ Term M has type A in local context Γ under stack Ψ .

Intuitively, each element Γ' of the context stack Ψ corresponds to a computation stage. The variables declared in Γ' are the ones whose values will be available during the corresponding evaluation phase. When we encounter a term $\mathbf{box} M$ during typing we enter a new evaluation stage, since M will be frozen during evaluation of the current stage. In this new phase, we are not allowed to refer to variables of the prior phases, since they may not be available when $\mathbf{box} M$ is unfrozen using \mathbf{unbox}_n . Thus, variables may only be looked up in the current context Γ (rule tpi_var) which is initialized as empty when we enter the body of a \mathbf{box} (rule tpi_box). However, *code* generated in the current or earlier stages may be used, which is represented by the rule tpi_unbox .

Functions

$$\begin{array}{c}
\frac{x:A \text{ in } \Gamma}{\Psi; \Gamma \vdash^i x : A} \text{tpi_var} \quad \frac{\Psi; (\Gamma, x:A) \vdash^i M : B}{\Psi; \Gamma \vdash^i \lambda x:A. M : A \rightarrow B} \text{tpi_lam} \\
\\
\frac{\Psi; \Gamma \vdash^i M : A \rightarrow B \quad \Psi; \Gamma \vdash^i N : A}{\Psi; \Gamma \vdash^i MN : B} \text{tpi_app}
\end{array}$$

Code

$$\frac{\Psi; \Gamma; \cdot \vdash^i M : A}{\Psi; \Gamma \vdash^i \mathbf{box} M : \Box A} \text{tpi_box} \quad \frac{\Psi; \Gamma \vdash^i M : \Box A}{\Psi; \Gamma; \Gamma_1; \dots; \Gamma_n \vdash^i \mathbf{unbox}_n M : A} \text{tpi_unbox}$$

Products

$$\frac{\Psi; \Gamma \vdash^i M_1 : A_1 \quad \Psi; \Gamma \vdash^i M_2 : A_2}{\Psi; \Gamma \vdash^i \langle M_1, M_2 \rangle : A_1 \times A_2} \text{tpi_pair}$$

$$\frac{\Psi; \Gamma \vdash^i M : A_1 \times A_2}{\Psi; \Gamma \vdash^i \mathbf{fst} M : A_1} \text{tpi_fst} \quad \frac{\Psi; \Gamma \vdash^i M : A_1 \times A_2}{\Psi; \Gamma \vdash^i \mathbf{snd} M : A_2} \text{tpi_snd}$$

$$\frac{}{\Psi; \Gamma \vdash^i \langle \rangle : 1} \text{tpi_unit}$$

Natural Numbers

$$\frac{}{\Psi; \Gamma \vdash^i \mathbf{z} : \text{nat}} \text{tpi_z} \quad \frac{\Psi; \Gamma \vdash^i M : \text{nat}}{\Psi; \Gamma \vdash^i \mathbf{s}M : \text{nat}} \text{tpi_s}$$

$$\frac{\Psi; \Gamma \vdash^i M_1 : \text{nat} \quad \Psi; \Gamma \vdash^i M_2 : A \quad \Psi; (\Gamma, x:\text{nat}) \vdash^i M_3 : A}{\Psi; \Gamma \vdash^i (\mathbf{case} M_1 \mathbf{of} \mathbf{z} \Rightarrow M_2 \mid \mathbf{s} x \Rightarrow M_3) : A} \text{tpi_case}$$

Recursion

$$\frac{\Psi; (\Gamma, x:A) \vdash^i M : A}{\Psi; \Gamma \vdash^i \mathbf{fix} x:A. M : A} \text{tpi_fix}$$

The reductions and expansions from Section 4 remain valid in this extended setting, as do the structural rules.

5.3 Examples in Implicit Form

We now show how we can define the power function in Mini-ML[□] with a different syntactic structure than in Mini-ML_e[□], though still with type $\text{nat} \rightarrow \square(\text{nat} \rightarrow \text{nat})$.

$$\begin{aligned} \text{power} &\equiv \mathbf{fix} p:\text{nat} \rightarrow \square(\text{nat} \rightarrow \text{nat}). \\ &\quad \lambda n:\text{nat}. \mathbf{case} n \\ &\quad \quad \mathbf{of} \mathbf{z} \Rightarrow \mathbf{box} (\lambda x:\text{nat}. \mathbf{s} \mathbf{z}) \\ &\quad \quad \mid \mathbf{s} m \Rightarrow \mathbf{box} (\lambda x:\text{nat}. \mathbf{times} x (\mathbf{unbox}_1 (p m) x)) \end{aligned}$$

As another example, we show how to define a function of type $\text{nat} \rightarrow \square \text{nat}$ that returns a **box**'ed copy of its argument:

$$\begin{aligned} \text{lift}_{\text{nat}} &\equiv \mathbf{fix} f:\text{nat} \rightarrow \square \text{nat}. \\ &\quad \lambda x:\text{nat}. \mathbf{case} x \\ &\quad \quad \mathbf{of} \mathbf{z} \Rightarrow \mathbf{box} \mathbf{z} \\ &\quad \quad \mid \mathbf{s} x' \Rightarrow \mathbf{box} (\mathbf{s} (\mathbf{unbox}_1 (f x'))) \end{aligned}$$

A similar term of type $A \rightarrow \square A$ that returns a **box**'ed copy of its argument exists exactly when every \rightarrow in A is enclosed by a \square . This justifies the inclusion of the *lift* primitive for base types in two-level languages such as in [Gomard and Jones

1991], and it seems natural to include such a primitive in a realistic extension of our language, as in [Wickline et al. 1998].

5.4 Compilation to Explicit Language

We do not define an operational semantics for Mini-ML[□] directly; instead we depend upon a translation to Mini-ML_e[□]. This extends the translation given in Section 4.7 in a straightforward way. We prefer this to a direct operational semantics on the implicit language since the translation should be identical to what a compiler would perform. We omit the obvious rules.

As an example of the compilation, it maps the definition of *power* from Section 5.3 to the one in Section 3.4. Note that the restructuring achieved by the compiler is similar to a staging transformation [Jørring and Scherlis 1986].

The operational semantics induced by the translation is different from some obvious ones defined directly on Mini-ML[□]. In [Martini and Masini 1996], for example, a simple reduction semantics is introduced for a system similar to the pure fragment of our implicit system. It does not reflect staging, and is instead used to prove a Church-Rosser theorem and strong normalization for a pure modal λ -calculus. Similarly, in [Pfenning and Wong 1995] an algorithm for converting pure modal λ -terms in implicit form to long normal form is given and proven correct. This algorithm bears no resemblance to the staged computation achieved via Mini-ML_e[□]. We also have constructed a direct operational semantics for Mini-ML[□] generalizing [Hatcliff 1995] that does capture staging, but prefer the compilation because it makes operational properties more evident. In particular, proving staging theorems for Mini-ML[□] directly would be much harder than taking advantage of the type-preserving compilation and proving the properties for Mini-ML_e[□].

6. A TWO-LEVEL LANGUAGE

In this section we define Mini-ML₂, a two-level functional language very close to the one described in [Nielson and Nielson 1992]. We then define a simple translation into Mini-ML[□] and prove that binding-time correctness in Mini-ML₂ is equivalent to modal correctness of the translation in Mini-ML[□].

A two-level language captures staging by explicitly annotating each occurrence of a term constructor as *compile-time* (often called *static*) or *run-time* (often called *dynamic*). Traditionally, expression constructors which can be evaluated at compile-time are overlined, those which cannot be evaluated until run-time are underlined. The process of annotating each term constructor in an expression is called *binding-time analysis*. Of course, not every possible annotation is valid. For example, the expression

$$\lambda x:\text{nat. } \overline{\text{case } \underline{x} \text{ of } \overline{z} \Rightarrow \overline{z} \mid \overline{s} \overline{y} \Rightarrow \overline{s} \overline{z}}$$

is not binding-time correct, since \underline{x} is not available until run-time, while the case statement is annotated to be executed at compile time, which is not possible.

We will not discuss binding-time analysis in this paper, only show how the resulting two-level terms are related to modal Mini-ML in its multiple-world formulation from Section 5.

Our language differs slightly from [Nielson and Nielson 1992] in that we inject *all* run-time types into compile-time types, instead of just function types. This

follows [Gomard and Jones 1991], where there is no such restriction. Also, we find it convenient to divide the variables and contexts into run-time and compile-time. All other differences to [Nielson and Nielson 1992] are due to minor differences between their underlying language and Mini-ML. Note that modal Mini-ML can accommodate arbitrary levels (not just two) and additional term operations (such as evaluation), so the two-level language we introduce in this section will be embedded into a relatively modest fragment of modal Mini-ML.

6.1 Syntax

Run-time Types	$\tau ::= \underline{\mathbf{nat}} \mid \tau_1 \Rightarrow \tau_2 \mid \tau_1 \times \tau_2 \mid \underline{\mathbf{1}}$
Compile-time Types	$\sigma ::= \overline{\mathbf{nat}} \mid \sigma_1 \Rightarrow \sigma_2 \mid \sigma_1 \times \sigma_2 \mid \overline{\mathbf{1}} \mid \tau$
Terms	
	$ \begin{aligned} e ::= & \underline{x} \mid \underline{\lambda x:\tau}. e \mid e_1 \underline{\textcircled{}} e_2 \\ & \mid \underline{\mathbf{fix}} \underline{x:\tau}. e \\ & \mid \underline{\langle e_1, e_2 \rangle} \mid \underline{\mathbf{fst}} e \mid \underline{\mathbf{snd}} e \\ & \mid \underline{\mathbf{z}} \mid \underline{\mathbf{s}} e \\ & \mid (\underline{\mathbf{case}} e_1 \underline{\mathbf{of}} \underline{\mathbf{z}} \Rightarrow e_2 \mid \underline{\mathbf{s}} \underline{x} \Rightarrow e_3) \\ & \mid \underline{\langle \rangle} \\ & \mid \underline{\overline{y}} \mid \underline{\lambda \overline{y}:\sigma}. e \mid e_1 \overline{\textcircled{}} e_2 \\ & \mid \underline{\mathbf{fix}} \overline{y:\sigma}. e \\ & \mid \overline{\langle e_1, e_2 \rangle} \mid \overline{\mathbf{fst}} e \mid \overline{\mathbf{snd}} e \\ & \mid \overline{\langle \rangle} \\ & \mid \overline{\mathbf{z}} \mid \overline{\mathbf{s}} e \\ & \mid (\overline{\mathbf{case}} e_1 \overline{\mathbf{of}} \overline{\mathbf{z}} \Rightarrow e_2 \mid \overline{\mathbf{s}} \overline{y} \Rightarrow e_3) \end{aligned} $
Run-time Contexts	$\Gamma ::= \cdot \mid \Gamma, \underline{x:\tau}$
Compile-time Contexts	$\Delta ::= \cdot \mid \Delta, \overline{y:\sigma}$

As a simple example we consider the two-level version of the power function.

$$\begin{aligned}
power \equiv & \underline{\mathbf{fix}} \overline{p} : \overline{\mathbf{nat}} \Rightarrow (\underline{\mathbf{nat}} \Rightarrow \underline{\mathbf{nat}}). \\
& \overline{\lambda \overline{n} : \overline{\mathbf{nat}}}. \overline{\mathbf{case}} \overline{n} \\
& \quad \overline{\mathbf{of}} \overline{\mathbf{z}} \Rightarrow (\underline{\lambda x:\mathbf{nat}}. \underline{\mathbf{s}} \underline{\mathbf{z}}) \\
& \quad \mid \overline{\mathbf{s}} \overline{m} \Rightarrow (\underline{\lambda x:\mathbf{nat}}. \underline{\mathbf{times}} \underline{\textcircled{}} \underline{x} \underline{\textcircled{}} ((\overline{p} \overline{\textcircled{}} \overline{m}) \underline{\textcircled{}} \underline{x}))
\end{aligned}$$

Recall that *times* is a curried function for multiplication represented as a closed term for simplicity. The type indicates that *power* takes a natural number as a compile-time argument and computes a residual run-time function from **nat** to **nat**. Otherwise the structure is very similar to the *power* function in its implicit formulation from Section 5.3. As we will see in Section 6.3 we can translate this to Mini-ML[□] by inserting a **box** constructor when an immediate subexpression of a compile-time term (overlined) is a run-time term (underlined). Conversely, when a compile-time term appears as an immediate subexpression of a run-time term we insert an **unbox**₁ constructor. It is easy to see that in this example we obtain the *power* function in implicit form, exactly as in Section 5.3:

$$\begin{aligned}
 \text{power} &\equiv \mathbf{fix} \ p:\text{nat} \rightarrow \square(\text{nat} \rightarrow \text{nat}). \\
 &\quad \lambda n:\text{nat}. \mathbf{case} \ n \\
 &\quad \quad \mathbf{of} \ \mathbf{z} \Rightarrow \mathbf{box} \ (\lambda x:\text{nat}. \mathbf{s} \ \mathbf{z}) \\
 &\quad \quad | \ \mathbf{s} \ m \Rightarrow \mathbf{box} \ (\lambda x:\text{nat}. \text{times} \ x \ (\mathbf{unbox}_1 \ (p \ m) \ x))
 \end{aligned}$$

6.2 Typing Rules

The typing rules of the two-level λ -calculus simultaneously verify staging and standard type-correctness, just as our explicit and implicit systems. We have two judgments:

$$\begin{array}{ll}
 \Delta; \Gamma \vdash^r e : \tau & \text{expression } e \text{ has run-time type } \tau \\
 \Delta \vdash^c e : \sigma & \text{expression } e \text{ has compile-time type } \sigma
 \end{array}$$

A compile-time expression can never depend on a run-time variable. Therefore, compile-time typing depends only on a compile-time context. A run-time expression may have embedded compile-time subexpressions and therefore carries compile-time variables (in Δ) as well as run-time variables in Γ .

Functions

$$\begin{array}{c}
 \frac{\underline{x}:\tau \text{ in } \Gamma}{\Delta; \Gamma \vdash^r \underline{x} : \tau} \text{tpr_var} \qquad \frac{\Delta; (\Gamma, \underline{x}:\tau_2) \vdash^r e : \tau}{\Delta; \Gamma \vdash^r \lambda \underline{x}:\tau_2. e : \tau_2 \rightarrow \tau} \text{tpr_lam} \\
 \\
 \frac{\Delta; \Gamma \vdash^r e_1 : \tau_2 \rightarrow \tau \quad \Delta; \Gamma \vdash^r e_2 : \tau_2}{\Delta; \Gamma \vdash^r e_1 \ @ \ e_2 : \tau} \text{tpr_app}
 \end{array}$$

Products

$$\begin{array}{c}
 \frac{\Delta; \Gamma \vdash^r e_1 : \tau_1 \quad \Delta; \Gamma \vdash^r e_2 : \tau_2}{\Delta; \Gamma \vdash^r \langle e_1, e_2 \rangle : \tau_1 \times \tau_2} \text{tpr_pair} \\
 \\
 \frac{\Delta; \Gamma \vdash^r e : \tau_1 \times \tau_2}{\Delta; \Gamma \vdash^r \mathbf{fst} \ e : \tau_1} \text{tpr_fst} \qquad \frac{\Delta; \Gamma \vdash^r e : \tau_1 \times \tau_2}{\Delta; \Gamma \vdash^r \mathbf{snd} \ e : \tau_2} \text{tpr_snd} \\
 \\
 \frac{}{\Delta; \Gamma \vdash^r \langle \rangle : \underline{1}} \text{tpr_unit}
 \end{array}$$

Natural Numbers

$$\begin{array}{c}
 \frac{}{\Delta; \Gamma \vdash^r \underline{\mathbf{z}} : \underline{\mathbf{nat}}} \text{tpr_z} \qquad \frac{\Delta; \Gamma \vdash^r e : \underline{\mathbf{nat}}}{\Delta; \Gamma \vdash^r \underline{\mathbf{s}} \ e : \underline{\mathbf{nat}}} \text{tpr_s} \\
 \\
 \frac{\Delta; \Gamma \vdash^r e_1 : \underline{\mathbf{nat}} \quad \Delta; \Gamma \vdash^r e_2 : \tau \quad \Delta; (\Gamma, \underline{x} : \underline{\mathbf{nat}}) \vdash^r e_3 : \tau}{\Delta; \Gamma \vdash^r (\mathbf{case} \ e_1 \ \mathbf{of} \ \underline{\mathbf{z}} \Rightarrow e_2 \ | \ \underline{\mathbf{s}} \ \underline{x} \Rightarrow e_3) : \tau} \text{tpr_case}
 \end{array}$$

Recursion

$$\frac{\Delta; (\Gamma, \underline{x}:\tau) \vdash^t e : \tau}{\Delta; \Gamma \vdash^t \mathbf{fix} \underline{x}:\tau. e : \tau} \text{tpr_fix}$$

Phase Transitions

$$\frac{\Delta \vdash^c e : \tau}{\Delta; \Gamma \vdash^t e : \tau} \text{down} \qquad \frac{\Delta; \cdot \vdash^t e : \tau}{\Delta \vdash^c e : \tau} \text{up}$$

Functions

$$\frac{\bar{y}:\sigma \text{ in } \Delta}{\Delta \vdash^c \bar{y} : \sigma} \text{tpc_var} \qquad \frac{\Delta, \bar{y}:\sigma_2 \vdash^c e : \sigma}{\Delta \vdash^c \bar{\lambda}\bar{y}:\sigma_2. e : \sigma_2 \Rightarrow \sigma} \text{tpc_lam}$$

$$\frac{\Delta \vdash^c e_1 : \sigma_2 \Rightarrow \sigma \quad \Delta \vdash^c e_2 : \sigma_2}{\Delta \vdash^c e_1 \bar{\otimes} e_2 : \sigma} \text{tpc_app}$$

Products

$$\frac{\Delta \vdash^c e_1 : \sigma_1 \quad \Delta \vdash^c e_2 : \sigma_2}{\Delta \vdash^c \langle \bar{e}_1, \bar{e}_2 \rangle : \sigma_1 \bar{\times} \sigma_2} \text{tpc_pair}$$

$$\frac{\Delta \vdash^c e : \sigma_1 \bar{\times} \sigma_2}{\Delta \vdash^c \mathbf{fst} e : \sigma_1} \text{tpc_fst} \qquad \frac{\Delta \vdash^c e : \sigma_1 \bar{\times} \sigma_2}{\Delta \vdash^c \mathbf{snd} e : \sigma_2} \text{tpc_snd}$$

$$\frac{}{\Delta \vdash^c \langle \bar{\cdot} \rangle : \bar{1}} \text{tpc_unit}$$

Natural Numbers

$$\frac{}{\Delta \vdash^c \bar{z} : \bar{\text{nat}}} \text{tpc_z} \qquad \frac{\Delta \vdash^c e : \bar{\text{nat}}}{\Delta \vdash^c \bar{s} e : \bar{\text{nat}}} \text{tpc_s}$$

$$\frac{\Delta \vdash^c e_1 : \bar{\text{nat}} \quad \Delta \vdash^c e_2 : \sigma \quad \Delta, \bar{y} : \bar{\text{nat}} \vdash^c e_3 : \sigma}{\Delta \vdash^c (\bar{\text{case}} e_1 \bar{\text{of}} \bar{z} \Rightarrow e_2 \mid \bar{s} \bar{y} \Rightarrow e_3) : \sigma} \text{tpc_case}$$

Recursion

$$\frac{\Delta, \bar{y}:\sigma \vdash^c e : \sigma}{\Delta \vdash^c \bar{\mathbf{fix}} \bar{y}:\sigma. e : \sigma} \text{tpc_fix}$$

Note that we remove run-time assumptions at the **down** rule, while in [Nielson and Nielson 1992] this is done later at the **up** rule, via a rule like:

$$\frac{\Delta; \cdot \vdash^r e : \tau}{\Delta; \Gamma \vdash^e e : \tau} \text{up}$$

Our change is justified since, by the structure of their rules, such assumptions can never be used in the compile-time deduction in between.

6.3 Translation to Implicit Language

The translation to Mini-ML[□] is now very simple. We translate both run-time and compile-time Mini-ML fragments directly, and insert \square , **box** and **unbox**₁ to represent the changes between phases. We define two mutually recursive functions to do this: $\|\cdot\|$ is the run-time translation and $|\cdot|$ is the compile-time translation. We overload this notation by using it for types, terms, and contexts. We write \underline{e} and \bar{e} to match any term whose top constructor matches the phase annotation.

Type Translation

$$\begin{array}{ll} \|\mathbf{nat}\| = \text{nat} & |\overline{\mathbf{nat}}| = \text{nat} \\ \|\tau_1 \rightarrow \tau_2\| = \|\tau_1\| \rightarrow \|\tau_2\| & |\sigma_1 \Rightarrow \sigma_2| = |\sigma_1| \rightarrow |\sigma_2| \\ \|\tau_1 \times \tau_2\| = \|\tau_1\| \times \|\tau_2\| & |\sigma_1 \times \sigma_2| = |\sigma_1| \times |\sigma_2| \\ \|\underline{1}\| = 1 & |\overline{1}| = 1 \\ & |\tau| = \square \|\tau\| \end{array}$$

Term Translation

$$\begin{array}{ll} \|\underline{x}\| = x & |\overline{y}| = y \\ \|\underline{\lambda x:\tau}. e\| = \lambda x:\|\tau\|. \|e\| & |\overline{\lambda \bar{y}:\sigma}. e| = \lambda y:|\sigma|. |e| \\ \|e_1 \underline{@} e_2\| = \|e_1\| \|e_2\| & |e_1 \overline{@} e_2| = |e_1| |e_2| \\ \|\underline{\mathbf{fix}} \underline{x}:\tau. e\| = \mathbf{fix} x:\|\tau\|. \|e\| & |\overline{\mathbf{fix}} \overline{y}:\sigma. e| = \mathbf{fix} y:|\sigma|. |e| \\ \|\langle e_1, e_2 \rangle\| = \langle \|e_1\|, \|e_2\| \rangle & |\langle e_1, e_2 \rangle| = \langle |e_1|, |e_2| \rangle \\ \|\underline{\mathbf{fst}} e\| = \mathbf{fst} \|e\| & |\overline{\mathbf{fst}} e| = \mathbf{fst} |e| \\ \|\underline{\mathbf{snd}} e\| = \mathbf{snd} \|e\| & |\overline{\mathbf{snd}} e| = \mathbf{snd} |e| \\ \|\langle \rangle\| = \langle \rangle & |\langle \rangle| = \langle \rangle \\ \|\underline{z}\| = z & |\overline{z}| = z \\ \|\underline{s} e\| = s \|e\| & |\overline{s} e| = s |e| \end{array}$$

$$\|\underline{\mathbf{case}} e_1 \underline{\mathbf{of}} \underline{z} \Rightarrow e_2 \mid \underline{s} \underline{x} \Rightarrow e_3\| = \mathbf{case} \|e_1\| \mathbf{of} z \Rightarrow \|e_2\| \mid s x \Rightarrow \|e_3\|$$

$$|\overline{\mathbf{case}} e_1 \overline{\mathbf{of}} \overline{z} \Rightarrow e_2 \mid \overline{s} \overline{y} \Rightarrow e_3| = \mathbf{case} |e_1| \mathbf{of} z \Rightarrow |e_2| \mid s y \Rightarrow |e_3|$$

$$\|\bar{e}\| = \mathbf{unbox}_1 |\bar{e}| \quad |e| = \mathbf{box} \|e\|$$

Context Translation

$$\begin{array}{ll} \|\cdot\| = \cdot & |\cdot| = \cdot \\ \|\Gamma, \underline{x}:\tau\| = \|\Gamma\|, x:\|\tau\| & |\Delta, \overline{y}:\sigma| = |\Delta|, y:|\sigma| \end{array}$$

6.4 Equivalence of Binding Time Correctness and Modal Correctness

In this section we show that binding-time correctness is equivalent to modal correctness of the translation to Mini-ML[□]. Note that even though we use Δ and Γ to denote contexts, the implicit language Mini-ML[□] employs context stacks, where $\odot; \Gamma_1; \dots; \Gamma_n$ is abbreviated as $\Gamma_1; \dots; \Gamma_n$.

THEOREM 15 (CONSERVATIVE EMBEDDING).

- (1) *If $\|e\| = M$ then:*
- (a) *if $\Delta; \Gamma \vdash^x e : \tau$ then $|\Delta|; \|\Gamma\| \vdash^i M : \|\tau\|$;*
 - (b) *if $|\Delta|; \|\Gamma\| \vdash^i M : A$ then $|\Delta|; \|\Gamma\| \vdash^x e : \tau$ with $\|\tau\| = A$.*
- (2) *If $|e| = M$ then:*
- (a) *if $\Delta \vdash^c e : \sigma$ then $|\Delta| \vdash^i M : |\sigma|$;*
 - (b) *if $|\Delta| \vdash^i M : A$ then $\Delta \vdash^c e : \sigma$ with $|\sigma| = A$.*

PROOF. By simultaneous induction on the definitions of $\|e\|$ and $|e|$. Note that we can take advantage of strong inversion properties, since we have exactly one typing rule for each term constructor in Mini-ML[□] and Mini-ML₂, plus the **up** and **down** rules to connect the \vdash^c and \vdash^x judgments.

We only show the two cases involving the **up** and **down** rules since all others are easy.

Case: $\|\bar{e}\| = \mathbf{unbox}_1 |\bar{e}|$, part 1a.

$$\begin{array}{ll} \Delta; \Gamma \vdash^x \bar{e} : \tau & \text{assumption} \\ \Delta \vdash^c \bar{e} : \tau & \text{by inversion (rule down)} \\ |\Delta| \vdash^i |\bar{e}| : \square \|\tau\| & \text{by i.h. 2a} \\ |\Delta|; \|\Gamma\| \vdash^i \mathbf{unbox}_1 |\bar{e}| : \|\tau\| & \text{by rule tpi_unbox} \end{array}$$

Case: $\|\bar{e}\| = \mathbf{unbox}_1 |\bar{e}|$, part 1b.

$$\begin{array}{ll} |\Delta|; \|\Gamma\| \vdash^i \mathbf{unbox}_1 |\bar{e}| : \|\tau\| & \text{assumption} \\ |\Delta| \vdash^i |\bar{e}| : \square \|\tau\| & \text{by inversion (rule tpi_unbox)} \\ \Delta \vdash^c \bar{e} : \tau & \text{by i.h. 2b} \\ \Delta; \Gamma \vdash^x \bar{e} : \tau & \text{by rule down} \end{array}$$

Case: $|e| = \mathbf{box} \|e\|$, part 2a.

$$\begin{array}{ll} \Delta \vdash^c \underline{e} : \tau & \text{assumption} \\ \Delta; \cdot \vdash^x \underline{e} : \tau & \text{by inversion (rule up)} \\ |\Delta|; \cdot \vdash^i \|\underline{e}\| : \|\tau\| & \text{by i.h. 1a} \\ |\Delta| \vdash^i \mathbf{box} \|\underline{e}\| : \square \|\tau\| & \text{by rule tpi_box} \end{array}$$

Case: $|e| = \mathbf{box} \|e\|$, part 2b.

$$\begin{array}{ll} |\Delta| \vdash^i \mathbf{box} \|\underline{e}\| : \square \|\tau\| & \text{assumption} \\ |\Delta|; \cdot \vdash^i \|\underline{e}\| : \|\tau\| & \text{by inversion (rule tpi_box)} \\ \Delta; \cdot \vdash^x \underline{e} : \tau & \text{by i.h. 1b} \\ \Delta \vdash^c \underline{e} : \tau & \text{by rule up} \end{array}$$

□

The translation and proof can be easily generalized from a two-level language to a B-level language [Nielson and Nielson 1992] with an infinite linear ordering. In this case the image of the translation on well-typed terms is exactly the fragment Mini-ML_K^\square , where \mathbf{unbox}_n is restricted to $n = 1$. This fragment corresponds to a weaker modal logic, K, in which we drop the assumption in S4 that the accessibility relation is reflexive and transitive [Martini and Masini 1996], and which we discussed briefly in Section 2. Thus a corollary of the generalized theorem is that Mini-ML_K^\square is equivalent to an infinite linear B-level language, since the translation is then a bijection which preserves correctness of typing. If we further restrict Mini-ML_K^\square so that rule $\mathbf{tpi_box}$ is only allowed when we have exactly one context in the context stack, then we obtain an exact correspondance with the two-level language, but this does not seem correspond to a previously studied logic.

7. EXAMPLES

We now present some standard examples from partial evaluation to illustrate the expressiveness of our language Mini-ML^\square . We use $\mathbf{let } x = E_1 \mathbf{in } E_2$ to introduce (non-polymorphic) top-level definitions; it may be considered syntactic sugar for $(\lambda x:A. E_2) E_1$.

7.1 Ackermann's Function

We now present a program for calculating Ackermann's function that specializes to the first argument. It is based on the following program:

```

fix acker:nat → nat → nat.
  λm:nat. case m
    of z ⇒ λn:nat. sn
    | s m' ⇒ λn:nat. case n
      of z ⇒ acker m' (s z)
      | s n' ⇒ acker m' (acker m n')

```

Now, if we attempt to directly insert the modal constructors to divide this program into two stages, we get the following:

```

fix acker:nat → □(nat → nat).
  λm:nat. case m
    of z ⇒ box (λn:nat. sn)
    | s m' ⇒ box (λn:nat. case n
      of z ⇒ (unbox1 (acker m')) (s z)
      | s n' ⇒ (unbox1 (acker m'))((unbox1 (acker m))n'))

```

Unfortunately, when applied to the first argument, this function generally will not terminate. This is a common problem in partial evaluation, and the usual solution is to employ memoization during specialization, which works for many programs. Here we will simply note that the problem in this case is a recursive call to $\mathit{acker } m$ while calculating $\mathit{acker } m$, which can be removed by adding an additional fixpoint as follows.

```

fix acker:nat → □(nat → nat).
λm:nat. case m
  of z ⇒ box (λn:nat. s n)
  | s m' ⇒ box (fix ackm. λn:nat.
    case n
      of z ⇒ (unbox1 (acker m')) (s z)
      | s n' ⇒ (unbox1 (acker m')) (ackm n'))

```

This function will always terminate. The recursive applications appearing inside **unbox**₁ constructors are evaluated when the first argument is given. The compilation of this function to Mini-ML_e[□] makes this more explicit:

```

fix acker:nat → □(nat → nat).
λm:nat. case m
  of z ⇒ box (λn:nat. s n)
  | s m' ⇒ let box f = acker m' in
    let box g = acker m' in
      box (fix ackm. λn:nat.
        case n of z ⇒ f (s z)
        | s n' ⇒ g (ackm n'))

```

Notice that *acker m'* is unnecessarily calculated twice. This would be avoided if memoization was employed during the compilation or if we had explicitly bound a variable to the result of this computation.

7.2 Inner Products

In [Glück and Jørgensen 1995] the calculation of inner products is given as an example of a program with more than two phases. We now show how this example can be coded in Mini-ML[□]. We assume a data type *vector* in the example, along with a function *sub*:nat → *vector* → nat to access the elements of a *vector*.

Then, the inner product example without staging is expressed in Mini-ML as follows:

```

fix ip:nat → vector → vector → nat.
λn:nat. case n
  of z ⇒ λv:vector. λw:vector. z
  | s n' ⇒ λv:vector. λw:vector.
    plus (times (sub n v) (sub n w))
    (ip n' v w)

```

We add in □, **box** and **unbox**_{*i*} to obtain a function with three computation stages which is shown in Figure 1. We assume a function *lift*_{nat} as defined earlier and a function *sub'*:nat → □(*vector* → nat) which is a specializing version of *sub*, that perhaps precomputes some pointer arithmetic based on the array index. We first define a staged version *times'* of *times* which avoids the multiplication in the specialization if the first argument is zero. This will speed up application of *iprod'* to its third argument, particularly in the case that the second argument is a sparse vector.

```

let times': $\Box$ (nat  $\rightarrow$   $\Box$ (nat  $\rightarrow$  nat)) =
  box ( $\lambda m$ :nat. case m
    of z  $\Rightarrow$  box ( $\lambda n$ :nat. z)
      | s m'  $\Rightarrow$  box ( $\lambda n$ :nat. times n (unbox1 (liftnat m))))
in let iprod' = fix ip:nat  $\rightarrow$   $\Box$ (vector  $\rightarrow$   $\Box$ (vector  $\rightarrow$  nat)).
   $\lambda n$ :nat. case n
    of z  $\Rightarrow$  box ( $\lambda v$ :vector. box ( $\lambda w$ :vector. z))
      | s n'  $\Rightarrow$  box ( $\lambda v$ :vector. box ( $\lambda w$ :vector.
        plus (unbox1 (unbox1 times' (unbox1 (sub' n) v))
          (unbox2 (sub' n) w))
          (unbox1 (unbox1 (ip n') v) w)))
in let iprod3 : vector  $\rightarrow$   $\Box$ (vector  $\rightarrow$  nat) = unbox0(iprod' 3)
in let iprod3a : vector  $\rightarrow$  nat = unbox0 (iprod3 [7, 0, 9])
in let iprod3b : vector  $\rightarrow$  nat = unbox0 (iprod3 [7, 8, 0])
in ...

```

Fig. 1. Staged code for inner product.

The last three lines show how to execute the result of a specialization using **unbox**₀ (corresponding to `eval` in Lisp). Also, the occurrence of **unbox**₂ indicates code used at the third stage but generated at the first. These two aspects could not be expressed within the multi-level language in [Glück and Jørgensen 1995].

Note the erasure of the **unbox**_{*i*} and **box** constructors in *iprod'* leaves the un-staged code, except that we used a different version of multiplication. The operational semantics of the two programs is of course quite different.

7.3 Regular Expression Matching

We now present a program for regular expression matching that specializes to a particular regular expression. We use the full Standard ML language, augmented with our modal constructors. Our program is based on the non-specializing one in Figure 2, which makes use of a continuation function that is called with the remaining input if the current matching succeeds. We assume the following datatype declaration:

```

datatype regexp
= Empty
| Plus of regexp * regexp
| Times of regexp * regexp
| Star of regexp
| Const of string

```

Note that there is a recursive call to `acc (Star(r))` in the case for `acc (Star(r))` which we can transform using a local definition, similar to the **fix** introduced in the Ackermann function example. This must be done so that specialization with respect to the regular expression terminates. The resulting code for this case is:

```

| acc (Star(r)) k s =
  let fun accStar k s =
      k s orelse
      acc r
      (fn ss => if s = ss then false

```

```

(* val acc : regexp -> (string list -> bool) -> (string list -> bool) *)
fun acc (Empty) k s = k s
  | acc (Plus(r1,r2)) k s = acc r1 k s orelse
                           acc r2 k s
  | acc (Times(r1,r2)) k s =
    acc r1 (fn ss => acc r2 k ss) s
  | acc (Star(r)) k s =
    k s orelse
    acc r (fn ss => if s = ss then false
                   else acc (Star(r)) k ss) s
  | acc (Const(str)) k (x::s) =
    (x = str) andalso k s
  | acc (Const(str)) k (nil) = false

(* val accept : regexp -> (string list -> bool) *)
fun accept r s =
  acc r (fn nil => true | (x::l) => false) s

```

Fig. 2. Unstaged regular expression matcher

```

                           else accStar k ss)
                           s
  in
    accStar k s
  end

```

Then, we can add in modal constructors to get the staged program in Figure 3 with the following types (using $[A]$ here to represent $\Box A$, following the syntax of PML [Wickline et al. 1998])

```

val acc2: regexp -> [(string list -> bool) -> (string list -> bool)]
val accept2: regexp -> [string list -> bool]

```

These types indicate that the required staging is achieved by the program. Inserting the modal constructors requires breaking up the function arguments, but is otherwise relatively straightforward. We use ‘ for **box** and ^ for **unbox**₁. More generally, we suggest using $\text{^}n$ for **unbox**_{*n*}.

We can now use our compilation to the explicit language Mini-ML_e[□] to get an equivalently staged program. We can then further translate to a program in pure Standard ML, which is staged in the same way, but without the modal annotations, as shown in Figure 4. It is unnecessary to replace $[A]$ by **unit** $\rightarrow A$ in this case, since the code constructor (‘) is only applied to values. We show this program only to demonstrate the staging described by the the modal annotated program. The program in Mini-ML_e[□] has the potential to be more efficient, since optimized code can be generated by a sophisticated implementation.

8. RELATED WORK

Our modal $\lambda_e^{\rightarrow\Box}$ -calculus is originally based on the modal λ -calculus presented by Bierman and de Paiva [Bierman and de Paiva 1996] and used by Pfenning and

```

(* val acc2 : regexp -> [(string list -> bool) -> (string list -> bool)] *)
fun acc2 (Empty) = ' fn k => fn s => k s
  | acc2 (Plus(r1,r2)) = ' fn k => fn s =>
    ^ (acc2 r1) k s orelse
    ^ (acc2 r2) k s
  | acc2 (Times(r1,r2)) = ' fn k => fn s =>
    ^ (acc2 r1) (fn ss => ^ (acc2 r2) k ss) s
  | acc2 (Star(r)) = ' fn k => fn s =>
    let fun acc2Star k s =
        k s orelse
        ^ (acc2 r)
          (fn ss => if s = ss then false
                    else acc2Star k ss)
    in
      s
    in
      acc2Star k s
    end
  | acc2 (Const(str)) = ' fn k =>
    (fn (x::ss) =>
      (x = ^(lift_string str))
      andalso k ss
    | nil => false)

(* val accept2 : regexp -> [string list -> bool] *)
fun accept2 r = ' fn s =>
  ^ (acc2 r) (fn nil => true | (x::l) => false) s

```

Fig. 3. Modally staged regular expression matcher

Wong [Pfenning and Wong 1995], who call it the “explicit system”. Our calculus avoids the use of simultaneous substitution by using both a modal context and an ordinary one, following the sequent calculi proposed by Andreoli [Andreoli 1992] for linear logic and by Girard [Girard 1993] for LU. The result is similar to the linear λ -calculus formulated by Wadler [Wadler 1993]. In separate work reported in [Pfenning and Davies 2000] we show how our calculus and associated natural deduction system can be carefully constructed following the methodology of Martin-Löf [Martin-Löf 1996] and also show how to account for modal possibility using two forms of conclusion.

The language Mini-ML_e[□] is constructed by combining $\lambda_e^{\rightarrow\Box}$ and Mini-ML [Clément et al. 1986]. The language Mini-ML[□] is based on the “implicit” modal λ -calculus presented in [Pfenning and Wong 1995], which uses a stack of ordinary contexts rather than two contexts. Mini-ML[□] avoids the **pop** structural rule of [Pfenning and Wong 1995], which is difficult to motivate from the point of view of natural deduction, by instead removing contexts from the stack at the $\Box E$ rule. The compilation from Mini-ML[□] to Mini-ML_e[□] is inspired by one direction of the proof of equivalence between the two calculi given in [Pfenning and Wong 1995]. Systems similar to the implicit modal λ -calculus of [Pfenning and Wong 1995] have

```

(* val acc3 : regexp -> (string list -> bool) -> (string list -> bool) *)
fun acc3 (Empty) = (fn k => fn s => k s)
  | acc3 (Plus(r1,r2)) =
    let val a1 = acc3 r1
        val a2 = acc3 r2
    in
      (fn k => fn s => a1 k s orelse a2 k s)
    end
  | acc3 (Times(r1,r2)) =
    let val a1 = acc3 r1
        val a2 = acc3 r2
    in
      (fn k => fn s => a1 (fn ss => a2 k ss) s)
    end
  | acc3 (Star(r1)) =
    let val a1 = acc3 r1
        fun acc3Star k s =
          k s orelse
            a1 (fn ss => if s = ss then false
                       else acc3Star k ss)
    in
      in
        (fn k => fn s => acc2 k s)
      end
    end
  | acc3 (Const(str)) =
    (fn k => (fn (x::s) => (x = str) andalso k s
                    | nil => false))

(* val accept3 : regexp -> (string list -> bool) *)
fun accept3 r =
  acc3 r (fn nil => true | (x::l) => false)

```

Fig. 4. Pure SML staged regular expression matcher

been proposed by Martini and Masini [Martini and Masini 1996], who introduce a simple reduction semantics, and Bourghuis [Borghuis 1994], who considers modal pure type systems. None of the prior work on modal λ -calculi has considered the relationship to computation staging.

Partly motivated by a previous version of the current paper [Davies and Pfenning 1996], Goubault-Larrecq [Goubault-Larrecq 1996; Goubault-Larrecq 1997] has proposed a formulation of modal λ -calculi using explicit substitutions. While this system has some interesting properties as a calculus, in particular giving a finer grained analysis of reduction and equality, it is unclear how this is relevant to the design of a programming language with staging primitives.

Despite some superficial similarities, our code types are quite different from Moggi's computational types based on monads [Moggi 1989; Moggi 1991] which only distinguish values from computations and do not allow expression of stage

separation. Moreover, our intended implementation of code is *intensional*, since we wish to allow refinements of our semantics to optimize code, while Moggi's computations are *extensional* with evaluation as the only operation. In other work [Pfenning and Davies 2000] we have been able to explain computational types cleanly in our framework via a combination of the intuitionistic possibility operator \diamond and necessity. This follows an earlier suggestion by Kobayashi [Kobayashi 1997] and a related investigation by Benton, Bierman and de Paiva [Benton et al. 1998] who establish a connection between the computational λ -calculus and lax logic [Fairtlough and Mendler 1997].

Another application of a modal λ -calculus to capture intensionality is reported in [Despeyroux et al. 1997]. This differs from our application in that it provides a primitive recursion construct for examining terms with types of the form $\Box A$. It also captures extensionality as given by the local expansions, thus obtaining a calculus suitable as a meta-logical framework, that is, a logical framework in which we can reason about the specified logics. The modal operator here serves to avoid the usual paradoxes which would arise if we incorrectly identify an expression with its denotation.

A further applicaiton of modal λ -calculi is to capture complexity classes such as polynomial-time computation [Hofmann 1998; Hofmann 2000] by restricting recursion to terms with types of the form $\Box A$.

We have shown how some standard examples of specialization can be expressed in Mini-ML $^\Box$. More complicated examples might require more advanced techniques to achieve the desired staging, such as the binding-time improvements used in partial evaluation. Memoizing when generating code is another useful technique used in partial evaluation, and [Wickline et al. 1998] shows how this technique can be programmed in a language with modal types. In general, we expect that any program that can be sucessfully partially evaluated should be able to be manually staged by the programmer in a language based on Mini-ML $^\Box$, although more experiments need to be conducted to determine how practical this is for large programs. See [Birkedal and Welinder 1993] for a description of a realistic partial evaluator for Standard ML and [Jones et al. 1993] for an overview of standard techniques and examples of partial evaluation.

One possible criticism of our languages is that they only manipulate *closed* code during execution, which restricts the staging that can be expressed compared to the two-level languages used in partial evaluation such as that proposed by Gomard and Jones [Gomard and Jones 1991]. This is solved in Mini-ML $^\Box$ by λ -abstracting code expressions over their free variables, and then later generating an application to the actual variables. This results in a number of variable for variable β -redices in the generated code in our examples. Of course, a lower level implementation could reduce these redices for efficiency, but this is not reflected in the language semantics. From a practical point of view, the reason for only treating closed code is that we need to be able to evaluate code without danger of encountering unbound variables. This is in contrast to the binding-time languages used in partial evaluation, which allow manipulation of code containing free variables, but do not support evaluation of code as a construct within the language. Instead, evaluation of the result of partial evaluation is an external operation applied only to whole programs, the properties of which have been studied separately by Jones [Jones 1992].

In other work [Davies 1996], one of the present authors has shown that the \bigcirc (“next”) operator from non-branching temporal logic exactly models the looser correctness criterion used in partial evaluation. Interestingly, the resulting languages are unsound when general references or value carrying exceptions are added, since these features allow a code expression with free variables to escape the binders for those variables. The language with \bigcirc corresponds closely to Mini-ML $^{\square}_K$ (see Section 6.4) except that it has the looser correctness criterion. Logically \square in Mini-ML $^{\square}_K$ exactly corresponds to \bigcirc in *branching* temporal logic, so the looser criterion is exactly captured by the move from branching to non-branching temporal logic.

Taha and Sheard [Taha and Sheard 1997] have directly constructed a language similar to Mini-ML $^{\square}$ which allows manipulation of code with free variables as well as type-safe evaluation, but their original design proved to be unsound in that free variables may be encountered during evaluation. This is fixed in [Taha et al. 1998], resulting in a language called Meta-ML which is sound in the absence of references and exceptions, but it seems more operationally, rather than logically motivated. More recent work on Meta-ML has concentrated on an idealized language [Moggi et al. 1999] and makes direct use of the results in the previous version of this paper [Davies and Pfenning 1996] as well as [Davies 1996].

Over the last few years there has been a lot of interest in run-time code generation in high-level languages. Engler, Hsieh and Kaashoek [Engler et al. 1996] describe an extension of the programming language C called ‘C (pronounced “tick C”)’ which uses similar mechanisms to Mini-ML $^{\square}$ to achieve computation staging. However, the type system lacks the modal restriction on variables, so it allows variables to be used when their values are not available, which may result in incorrect results or runtime errors. Consel and Noël [Consel and Noël 1996] describe a system called Tempo which allows both partial evaluation and run-time code generation for the C language. Standard binding-time analysis techniques are used, along with separate annotations to describe where run-time specialization should be done. These annotations roughly correspond to the **unbox** $_0$ construct in Mini-ML $^{\square}$, while the annotations resulting from the binding-time analysis roughly correspond to the **box** and **unbox** $_1$ constructs, although the restriction of Mini-ML $^{\square}$ to closed code means that this correspondence is not exact. Leone and Lee [Leone and Lee 1994; Lee and Leone 1996] describe a small ML-like language and a corresponding implementation called Fabius which treats curried functions as run-time code generators, using a new form of binding-time analysis. The staging achieved is quite different to that obtained using ordinary binding-time analysis, and one of the original motivations for the current work was to allow a formal characterization of this staging. Fabius also uses a very fast form of run-time code generation called deferred compilation.

Deferred compilation has also recently been used by Wickline, Lee and Pfenning [Wickline et al. 1998] as an implementation technique for a language based on Mini-ML $^{\square}_e$ called PML which includes most of core SML and performs run-time code generation based on modal types. The extension to core SML was very smooth even though the language includes polymorphism, datatypes, references and arrays. A simple compiler for this language has been completed, and work is continuing on improved implementations.

9. CONCLUSION AND FUTURE WORK

In this paper we have proposed a logical interpretation of binding times and staged computation in terms of the intuitionistic modal logic S4. We first presented the $\lambda_e^{\rightarrow\Box}$ -calculus, and formally demonstrated the sense in which it captures staging. We then extended this to the explicit language Mini-ML $^{\Box}_e$ (including recursion, natural numbers, and products) and presented its natural operational semantics. We continued by defining an implicit language Mini-ML $^{\Box}$ which might serve as the core for an extension of a language with the complexity of Standard ML, and which is syntactically similar to both Lisp's backquote and comma notation, as well as the languages used in partial evaluation. The operational semantics of Mini-ML $^{\Box}$ is given by a type-preserving compilation to Mini-ML $^{\Box}_e$. Further, Mini-ML $^{\Box}$ generalizes Nielson & Nielson's two-level functional language [Nielson and Nielson 1992] which is demonstrated by a conservative embedding theorem, an important technical result of this paper.

Our approach provides a general, logically motivated framework for staged computation that includes aspects of both partial evaluation and run-time code generation. As such it allows efficient code to be generated within a declarative style of programming, and provides an automatic check that the intended staging is achieved.

Our investigation remains at a relatively abstract level, thus providing a general framework in which various staging mechanisms may be studied from a new point of view. We implemented the original interpreter for Mini-ML $^{\Box}$ in the logic programming language Elf [Pfenning 1991], thus allowing us to perform small experiments at this abstract level. Concrete instances such as partial evaluation, run-time code generation, or macro expansion will require some additional considerations for their effective use and efficient implementation. The application to run-time code generation appears particularly promising and is described in more detail, including an extended example, in work by the current authors in conjunction with Wickline and Lee [Wickline et al. 1998]. We hope that future design and implementation work will lead to a practical full-scale programming language with computation staging based on modal types.

ACKNOWLEDGMENTS

The work described here has benefited from discussions with too many colleagues to acknowledge them here individually. We would like to particularly thank Olivier Danvy, Peter Lee, Philip Wickline and the anonymous referees for numerous fruitful discussions and suggestions.

REFERENCES

- ANDREOLI, J.-M. 1992. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation* 2, 3, 197–347.
- BENTON, P. N., BIERMAN, G. M., AND DE PAIVA, V. C. V. 1998. Computational types from a logical perspective. *Journal of Functional Programming* 8, 2 (March), 177–193.
- BIERMAN, G. AND DE PAIVA, V. 1996. Intuitionistic necessity revisited. Technical Report CSRP-96-10 (June), School of Computer Science, University of Birmingham.
- BIRKEDAL, L. AND WELINDER, M. 1993. Partial evaluation of Standard ML. Master's thesis, University of Copenhagen, Department of Computer Science. Available as Technical Report

- DIKU-report 93/22.
- BORGHUIS, T. 1994. *Coming to Terms with Modal Logic: On the Interpretation of Modalities in Typed λ -calculus*. Ph. D. thesis, Eindhoven University of Technology.
- CHELLAS, B. 1980. *Modal Logic: an Introduction*. Cambridge University Press, Cambridge and New York.
- CLÉMENT, D., DESPEYROUX, J., DESPEYROUX, T., AND KAHN, G. 1986. A simple applicative language: Mini-ML. In *Proceedings of the 1986 Conference on LISP and Functional Programming (LFP'86)* (1986), pp. 13–27. ACM Press.
- CONSEL, C. AND NOËL, F. 1996. A general approach for run-time specialization and its application to C. In *Proceedings of the 23rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'96)* (Jan. 1996), pp. 145–156.
- DAVIES, R. 1996. A temporal-logic approach to binding-time analysis. In E. CLARKE Ed., *Proceedings of the Eleventh Annual Symposium on Logic in Computer Science (LICS'96)* (July 1996), pp. 184–195.
- DAVIES, R. AND PFENNING, F. 1996. A modal analysis of staged computation. In *Proceedings of the 23rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'96)* (Jan. 1996), pp. 258–270.
- DESPEYROUX, J., PFENNING, F., AND SCHÜRMAN, C. 1997. Primitive recursion for higher-order abstract syntax. In R. HINDLEY Ed., *Proceedings of the Third International Conference on Typed Lambda Calculus and Applications (TLCA'97)* (Nancy, France, April 1997), pp. 147–163. Springer-Verlag LNCS 1210. An extended version is available as Technical Report CMU-CS-96-172, Carnegie Mellon University.
- ENGLER, D. R., HSIEH, W. C., AND KAASHOEK, M. F. 1996. ‘C: A language for high-level, efficient, and machine-independent dynamic code generation. In *Proceedings of the 23rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'96)* (Jan. 1996), pp. 131–144.
- FAIRTLOUGH, M. AND MENDLER, M. 1997. Propositional lax logic. *Information and Computation* 137, 1 (Aug.), 1–33.
- GIRARD, J.-Y. 1993. On the unity of logic. *Annals of Pure and Applied Logic* 59, 201–217.
- GLÜCK, R. AND JØRGENSEN, J. 1995. Efficient multi-level generating extensions for program specialization. In S. SWIERSTRA AND M. HERMENEGILDO Eds., *Programming Languages, Implementations, Logics and Programs (PLILP'95)* (Sept. 1995), pp. 259–278. Springer-Verlag LNCS 982.
- GOMARD, C. AND JONES, N. 1991. A partial evaluator for the untyped lambda-calculus. *Journal of Functional Programming* 1, 1 (January), 21–69.
- GOUBAULT-LARRECQ, J. 1996. On computational interpretations of the modal logic S4, parts I–III. Technical Reports 1996-33,34,35, Institut für Logik, Komplexität und Deduktionssysteme, Universität Karlsruhe, Karlsruhe, Germany.
- GOUBAULT-LARRECQ, J. 1997. On computational interpretations of the modal logic S4 - IIIb. confluence, termination of the $\lambda\text{ev}Q_H$ -calculus. Technical Report 3164 (May), INRIA, France.
- HATCLIFF, J. 1995. Mechanically verifying the correctness of an offline partial evaluator. In M. HERMENEGILDO AND S. SWIERSTRA Eds., *Proceedings of the 7th International Symposium on Programming Languages: Implementations, Logics and Programs (PLILP'95)* (Utrecht, The Netherlands, September 1995), pp. 279–298. Springer-Verlag LNCS 982.
- HOFMANN, M. 1998. A mixed modal/linear lambda calculus with applications to Bellantoni-Cook safe recursion. In M. NIELSEN AND W. THOMAS Eds., *Computer Science Logic (CSL'97): Annual Conference of the EACSL (Selected Papers)* (Aarhus, Denmark, Aug. 1998), pp. 275–294. Springer-Verlag LNCS 1414.
- HOFMANN, M. 2000. Programming languages capturing complexity classes. *ACM SIGACT News* 31, 1 (March), 31–42.
- JONES, N. D. 1992. Efficient algebraic operations on programs. In M. NIVAT, C. RATTRAY, T. RUS, AND G. SCOLLO Eds., *Algebraic Methodology and Software Technology (AMAST'91): Proc. of the Second International Conference*, pp. 393–420. London:

- Springer. A version appears as a chapter in [Jones et al. 1993].
- JONES, N. D., GOMARD, C., AND SESTOFT, P. 1993. *Partial Evaluation and Automatic Program Generation*. Prentice Hall International, International Series in Computer Science.
- JØRRING, U. AND SCHERLIS, W. L. 1986. Compilers and staging transformations. In *Conference Record of the Thirteenth Annual ACM Symposium on Principles of Programming Languages (POPL'86)* (St. Petersburg Beach, Florida, Jan. 1986), pp. 86–96.
- KEPPEL, D., EGGERS, S. J., AND HENRY, R. R. 1993. A case for runtime code generation. Technical Report TR 93-11-02 (Nov.), Department of Computer Science and Engineering, University of Washington.
- KOBAYASHI, S. 1997. Monad as modality. *Theoretical Computer Science* 175, 29–74.
- KRIPKE, S. A. 1963. Semantic analysis of modal logic. I: Normal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* 9, 67–96.
- LEE, P. AND LEONE, M. 1996. Optimizing ML with run-time code generation. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'96)* (Philadelphia, Pennsylvania, May 1996), pp. 137–148.
- LEONE, M. AND LEE, P. 1994. Lightweight run-time code generation. In *Proceedings of the Workshop on Partial Evaluation and Semantics-based Program Manipulation (PEPM'94)* (Orlando, Florida, June 1994). An earlier version appears as Carnegie Mellon School of Computer Science Technical Report CMU-CS-93-225, November 1993.
- MARTIN-LÖF, P. 1996. On the meanings of the logical constants and the justifications of the logical laws. *Nordic Journal of Philosophical Logic* 1, 1, 11–60.
- MARTINI, S. AND MASINI, A. 1996. A computational interpretation of modal proofs. In H. WANSING Ed., *Proof theory of Modal Logics*, pp. 213–241. Kluwer. Workshop proceedings.
- MOGGI, E. 1989. Computational lambda calculus and monads. In *Proceedings of the Fourth Symposium on Logic in Computer Science (LICS'89)* (Asilomar, California, June 1989), pp. 14–23. IEEE Computer Society Press.
- MOGGI, E. 1991. Notions of computation and monads. *Information and Computation* 93, 1, 55–92.
- MOGGI, E., TAHA, W., BENAÏSSA, Z.-E.-A., AND SHEARD, T. 1999. An idealized MetaML: Simpler, and more expressive. In *Proceedings of the European Symposium on Programming* (Amsterdam, March 1999), pp. 193–207.
- NIELSON, F. AND NIELSON, H. R. 1992. *Two-Level Functional Languages*. Cambridge University Press.
- PALSBERG, J. 1993. Correctness of binding time analysis. *Journal of Functional Programming* 3, 3 (July), 347–363.
- PFENNING, F. 1991. Logic programming in the LF logical framework. In G. HUET AND G. PLOTKIN Eds., *Logical Frameworks* (1991), pp. 149–181. Cambridge University Press.
- PFENNING, F. AND DAVIES, R. 2000. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science*. To appear.
- PFENNING, F. AND WONG, H.-C. 1995. On a modal λ -calculus for S4. In S. BROOKES AND M. MAIN Eds., *Proceedings of the Eleventh Conference on Mathematical Foundations of Programming Semantics (MFPS'95)* (New Orleans, Louisiana, March 1995). Available as *Electronic Notes in Theoretical Computer Science*, Volume 1, Elsevier.
- PRAWITZ, D. 1965. *Natural Deduction*. Almqvist & Wiksell, Stockholm.
- SMITH, B. C. 1984. Reflection and semantics in Lisp. In *Proceedings of the Eleventh Annual ACM Symposium on Principles of Programming Languages (POPL'84)* (Salt Lake City, Utah, January 1984), pp. 23–35. ACM Press.
- TAHA, W., BENAÏSSA, Z.-E.-A., AND SHEARD, T. 1998. Multi-stage programming: Axiomatization and type safety. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP'98)* (July 1998), pp. 918–929.
- TAHA, W. AND SHEARD, T. 1997. Multi-stage programming with explicit annotations. In *Proceedings of the ACM SIGPLAN Symposium on Partial Evaluation and Semantics Based Program Manipulation (PEPM'97)* (June 1997), pp. 203–217.

- WADLER, P. 1993. A taste of linear logic. In *Mathematical Foundations of Computing Science (MFCS'93)* (Gdansk, Aug. 1993), pp. 185–210. Springer-Verlag LNCS 711. Invited Talk.
- WICKLINE, P., LEE, P., AND PFENNING, F. 1998. Run-time code generation and modal-ML. In K. D. COOPER Ed., *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'98)* (Montreal, Canada, June 1998), pp. 224–235. ACM Press.
- WICKLINE, P., LEE, P., PFENNING, F., AND DAVIES, R. 1998. Modal types as staging specifications for run-time code generation. *ACM Computing Surveys* 30, 3es (Sept.).