

Contextual Modal Type Theory

ALEKSANDAR NANEVSKI

Harvard University

and

FRANK PFENNING

Carnegie Mellon University

and

BRIGITTE PIENKA

McGill University

The intuitionistic modal logic of necessity is based on the judgmental notion of categorical truth. In this paper we investigate the consequences of relativizing these concepts to explicitly specified contexts. We obtain contextual modal logic and its type-theoretic analogue. Contextual modal type theory provides an elegant, uniform foundation for understanding meta-variables and explicit substitutions. We sketch some applications in functional programming and logical frameworks.

Categories and Subject Descriptors: F.4.1 [**Theory of Computation**]: Mathematical Logic and Formal Languages—*Modal Logic*; D.3.3 [**Software**]: Language Constructs and Features—*Frameworks*

General Terms: Design, Theory

Additional Key Words and Phrases: Type theory, logical frameworks, intuitionistic modal logic

1. INTRODUCTION

The dictionary¹ defines *context* as “*the parts of a discourse that surround a word or passage and can throw light on its meaning*” and also as “*the interrelated conditions in which something exists or occurs*”. One can see from these definitions that the notion of context, with somewhat differentiated meanings, is fundamental in linguistics, artificial intelligence, and logic. Narrowing our scope to logic, the notion of context enters almost immediately because the truth of a proposition is not absolute, but depends on the context we consider it in.

Taking the next step, namely allowing the propositions in a logic itself to speak about contexts, is the subject of modal logic and fraught with many difficulties and paradoxes. But it is also very important. The study of contexts in logic has ram-

¹Merriam-Webster Online

This material is based upon work supported by the National Science Foundation under Grant No. 0325808.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2005 ACM 1529-3785/2005/0700-0001 \$5.00

ifications for topics such as modularity, reflection, spatial and temporal reasoning, and meta-reasoning. Because of the close relationship between constructive logic and programming languages, contexts are also directly relevant to computational phenomena, particularly data abstraction, run-time code generation, distributed computation, partial evaluation, and meta-programming.

Narrowing our scope a bit further, in this paper we think of a context as consisting of the hypotheses we make while trying to establish a conclusion. Contexts are therefore intrinsically tied to hypothetical judgments. We investigate how to internalize reasoning about contexts within the logic while retaining an intuitionistic and predicative viewpoint regarding the meaning of the logical connectives. In this manner we can give our core calculus several interesting interpretations. One, in the realm of functional programming, relates contexts to intensional expressions and run-time code generation in the presence of free variables. Another, in the realm of logic programming, explains explicit substitutions and meta-variables. In each case, the strength of the underlying logical foundations provides new insights for language design which has so far been largely motivated operationally. We return to these applications in Sections 6 and 7.

There seem to be two main approaches to intuitionistic modal logic. One, which we call *nominal*, assigns names to contexts and includes explicit judgments to relate them. This can be seen as the intuitionistic analogue of Kripke’s classical multiple-world semantics [Kripke 1959] and has been investigated in some depth by Simpson [1994] and applied to partial evaluation [Davies 1996] and distributed computation [VII et al. 2004]. We will not pursue the nominal approach here. The other, which we call *structural*, identifies the context with the propositions it contains. Pfenning and Davies [2001] present the first exploration in this direction, singling out the empty context which represents reasoning without truth assumptions. This provides a constructive meaning explanation for necessity, possibility, and lax truth. In this paper we carry the structural approach significantly further, allowing arbitrary contexts to be captured within a proposition. This sheds new light on some important ideas. For example, *explicit substitutions* [Cartmell 1986; Abadi et al. 1990] are inevitable as evidence that we can reach one context from another. Another example is *meta-variables* which arise as we introduce hypotheses about the truth of propositions in given contexts.

We start our investigation in the simplest case, namely a propositional logic with a contextual modality. We develop both natural deduction (Section 2) and the sequent calculus (Section 3). In view of applications in computer science, we then generalize to a type theory in Section 4, which differs from the logic in that it carries explicit proof terms. We then generalize further to add dependent types (Section 5) and show applications to staged functional programming (Section 6) and logical frameworks (Section 7). We close with a discussion of related and future work.

2. INTUITIONISTIC CONTEXTUAL MODAL LOGIC

The philosophical foundation of our development is Martin-Löf’s approach of separating judgments from propositions [Martin-Löf 1996]. The most basic judgment is the *truth* of a *proposition*, written as A true. We explain the meaning of a propo-

sition by presenting the means of inferring its truth via introduction rules, and for exploiting the knowledge of its truth via elimination rules.

2.1 Hypothetical Judgments

In order to explain the meaning of implication and later of the modal operators in this manner, we need the notion of a *hypothetical judgment*. We write

$$x_1:A_1 \text{ true}, \dots, x_n:A_n \text{ true} \vdash A \text{ true}$$

to express that A is true whenever all hypotheses A_1, \dots, A_n are true. We always label hypotheses with distinct variables x_i in order to avoid any ambiguity in proofs. The properties of hypothetical judgments are captured by the hypothesis rule and the substitution principle. We abbreviate a collection of assumptions by Γ .

Hypothesis Rule.

$$\frac{}{\Gamma, x:A \text{ true}, \Gamma' \vdash A \text{ true}} \text{hyp}_x$$

Note that without distinct labels for assumptions, this rule would be ambiguous.

Substitution Principle

If $\Gamma \vdash A \text{ true}$ and $\Gamma, x:A \text{ true}, \Gamma' \vdash C \text{ true}$ then $\Gamma, \Gamma' \vdash C \text{ true}$.

While the hypothesis rule is primitive, the substitution principle should always be *admissible*: whenever we have a proof of $A \text{ true}$ then we can substitute that proof for all uses of an assumption $A \text{ true}$ in another proof. We will not further elaborate here on the standard notion of hypothetical judgments.

Now we can explain the meaning of implication.

Implication Introduction. $A \rightarrow B$ is true if B is true under the hypothesis that A is true.

$$\frac{\Gamma, x:A \text{ true} \vdash B \text{ true}}{\Gamma \vdash A \rightarrow B \text{ true}} \rightarrow^I_x$$

By our convention regarding well-formed contexts, x must not already be declared in Γ .

Implication Elimination. Conversely, if we know that $A \rightarrow B$ is true, then B is true under hypothesis that A is true. By the substitution principle we should therefore be able to obtain a proof of B from a proof of A .

$$\frac{\Gamma \vdash A \rightarrow B \text{ true} \quad \Gamma \vdash A \text{ true}}{\Gamma \vdash B \text{ true}} \rightarrow^E$$

In order to check that the elimination rules are not too strong, we should verify that any introduction of a connective immediately followed by its elimination can be reduced. This means the elimination rule cannot extract more knowledge from a proposition than contributed by its proof. In other words, it is *sound*. Local reductions also give rise to computation in the setting of functional programming.

Local Reduction for Implication.

$$\frac{\frac{\mathcal{D}}{\Gamma, x:A \text{ true} \vdash B \text{ true}} \rightarrow |^x \quad \frac{\mathcal{E}}{\Gamma \vdash A \text{ true}} \rightarrow E}{\Gamma \vdash B \text{ true}} \rightarrow E \quad \Longrightarrow_R \quad \frac{\mathcal{D}'}{\Gamma \vdash B \text{ true}}$$

Here, \mathcal{D}' can be obtained from \mathcal{D} and \mathcal{E} by application of the substitution principle. Computationally, this corresponds to a β -reduction.

In order to check that the elimination rules are not too weak, we should verify that we can recover the knowledge that contributed to the proof of a proposition. In other words, starting from any proof of a proposition, we can apply the elimination rules in such a way that we can reintroduce the proposition from the results. Local expansions correspond to extensionality in functional programming.

Local Expansion for Implication.

$$\frac{\frac{\mathcal{D}}{\Gamma \vdash A \rightarrow B \text{ true}} \Longrightarrow_E \quad \frac{\frac{\mathcal{D}'}{\Gamma, x:A \text{ true} \vdash A \rightarrow B \text{ true}} \rightarrow |^x \quad \frac{\text{hyp}_x}{\Gamma, x:A \text{ true} \vdash A \text{ true}}}{\Gamma, x:A \text{ true} \vdash B \text{ true}} \rightarrow E}{\Gamma \vdash A \rightarrow B \text{ true}} \Longrightarrow_E$$

Here, \mathcal{D}' can be obtained from \mathcal{D} by *weakening*, that is, adjoining the hypothesis $x:A \text{ true}$ to every judgment in \mathcal{D} . In functional programming, local expansion corresponds to an η -expansion.

2.2 Categorical Judgments

In this section we review the critical notion of a *categorical judgment* [Pfenning and Davies 2001]. We say that A is *categorically true* or *valid* if its truth does not depend on any hypotheses about the truth of other propositions. However, it may depend on the validity of other proposition, because their validity may not depend on truth assumptions.

We write Δ for a labeled collection $u_1::A_1 \text{ valid}, \dots, u_k::A_k \text{ valid}$ of assumptions about the validity of propositions, using “.” for an empty collection. With this notation we separate assumptions about validity and truth and write

$$\Delta; \Gamma \vdash J$$

where J is either $A \text{ true}$ or $A \text{ valid}$.

Definition of Validity.

$$\frac{\Delta; \cdot \vdash A \text{ true}}{\Delta; \Gamma \vdash A \text{ valid}}$$

Here the premise represents a truth-categorical judgment. Conversely, if A is valid then we can conclude it is true.

$$\frac{}{(\Delta, u::A \text{ valid}, \Delta'); \Gamma \vdash A \text{ true}} \text{vldhyp}_u$$

Another way to think about this definition is in terms of a multiple-world semantics. We say that A is valid if A is true in every possible world. Since we can never circumscribe all possible worlds, we instead require that A must be true in a world about which we assume nothing ($\Gamma = \cdot$), except for the propositions we assumed to be always true (Δ).

Note that by the definition of validity we can immediately reduce the goal of proving A valid to A true. We can always apply this silently and thereby remove the judgment A valid from consideration as the conclusion of a hypothetical judgment. No such trick applies to assumptions of the form A valid, so we require a new substitution principle.

Substitution Principle for Validity

If $\Delta; \cdot \vdash A$ true and $(\Delta, u::A \text{ valid}, \Delta'); \Gamma \vdash C$ true then $(\Delta, \Delta'); \Gamma \vdash C$ true.

Note that the first assumption expresses $\Delta; \Gamma \vdash A$ valid, thereby justifying the substitution principle directly from the nature of hypothetical judgments in general.

Using the categorical judgment, it is now possible to introduce a modal operator for necessary truth, $\Box A$. Since we need a more general operator we forego this analysis here; it can be found in the paper by Pfenning and Davies [2001]. The result is intuitionistic S4 which captures reflexivity and transitivity of accessibility between worlds in its nominal formulation [Simpson 1994].

2.3 Contextual Validity

The judgment A valid expresses that A is true in any world. In this section we relativize this judgment: let Ψ be a collection of assumptions $y_1:B_1$ true, \dots , $y_m:B_m$ true. We say that A is valid relative Ψ (written as A valid $[\Psi]$) if A is true in every world in which B_1, \dots, B_m are true. We similarly generalize assumptions about validity so that Δ now has the form $u_1::A_1$ valid $[\Psi_1], \dots, u_k::A_k$ valid $[\Psi_k]$.

This generalization does not directly add expressive power, since a judgment of the general form A valid $[y_1:B_1$ true, $\dots, y_n:B_m$ true] holds iff $(B_1 \rightarrow \dots (B_m \rightarrow A))$ valid. However, there is an important difference in the proof theory, which we exploit in several applications to obtain context calculi for staged computation and meta-variables.

As common in programming languages, we refer to Ψ as a *context* and to the judgment A valid $[\Psi]$ as *contextual validity*.

Definition of Contextual Validity. In order to prove A valid $[\Psi]$ we have to prove A using only Ψ and without any other current hypotheses about truth. Assumptions about contextual validity, however, carry over as before.

$$\frac{\Delta; \Psi \vdash A \text{ true}}{\Delta; \Gamma \vdash A \text{ valid}[\Psi]}$$

As in the case of ordinary validity in Section 2.2, proving $\Delta; \Gamma \vdash A$ valid $[\Psi]$ can always be immediately reduced to $\Delta; \Psi \vdash A$ true. Therefore we are justified in streamlining our analysis by not considering validity on the right-hand side of a hypothetical judgment

Contextual Entailment. We can exploit the knowledge that $A \text{ valid}[\Psi]$ to conclude that $A \text{ true}$, but only if we can show that all propositions in Ψ follow from the current assumptions. We write

$$\Delta; \Gamma \vdash \Psi$$

to express that all the propositions in Ψ are true using the assumptions from Γ .

$$\frac{\Delta; \Gamma \vdash B_1 \text{ true} \quad \dots \quad \Delta; \Gamma \vdash B_m \text{ true}}{\Delta; \Gamma \vdash y_1:B_1 \text{ true}, \dots, y_m:B_m \text{ true}} \text{ ctx}$$

Contextual Hypothesis Rule. Now we can state the contextual hypothesis rule.

$$\frac{(\Delta, u::A \text{ valid}[\Psi], \Delta'); \Gamma \vdash \Psi}{(\Delta, u::A \text{ valid}[\Psi], \Delta'); \Gamma \vdash A \text{ true}} \text{ ctxhyp}_u$$

In order to use the assumption that A is valid relative Ψ , we must show that our current hypotheses in Γ are strong enough to establish Ψ .

With the contextual hypothesis rule in place, it is easy to verify that the judgment $A \text{ valid}$ from the previous section is recovered as $A \text{ valid}[\cdot]$. Of course, this is a special case of contextual validity where the relevant context is taken to be empty. For example, the instance of the contextual hypothesis rule

$$\frac{\frac{}{(\Delta, u::A \text{ valid}[\cdot], \Delta'); \Gamma \vdash \cdot} \text{ ctx}}{(\Delta, u::A \text{ valid}[\cdot], \Delta'); \Gamma \vdash A \text{ true}} \text{ ctxhyp}_u$$

has no premises because the empty context “.” has no members and so becomes the modal hypothesis rule vldhyp_u .

Contextual Substitution Principle. The substitution principle generalizes in a straightforward way.

If $\Delta; \Psi \vdash A \text{ true}$ and $(\Delta, u::A \text{ valid}[\Psi], \Delta'); \Gamma \vdash C \text{ true}$ then $(\Delta, \Delta'); \Gamma \vdash C \text{ true}$.

The first condition expresses $\Delta; \Gamma \vdash A \text{ valid}[\Psi]$. The principle is therefore justified by the general substitution property of hypothetical judgments. The associated substitution operation substitutes the proof of A under Ψ for uses of the assumption that A is valid under Ψ . A more precise description is possible when we add proof terms in Section 4.

Contextual Identity Principle. One particular instance of contextual entailment combined with the ordinary hypothesis rule is the contextual identity principle.

$$\Delta; \Psi \vdash \Psi$$

This holds because each conclusion $B_j \text{ true}$ is proved in one step from the corresponding assumption $y_j:B_j \text{ true}$ by the hypothesis rule applied to y_j . Like the substitution principle, this is not a primitive rule, but every instance of it is derivable in our system.

2.4 Contextual Modal Necessity

A modal logic arises from the judgmental definition of contextual validity if we internalize the judgment as a modal propositional operator. We write $[\Psi]A$ for the *proposition* asserting that A is valid in context Ψ . The difference is of course that a proposition can appear as part of other propositions, while judgments cannot be combined with propositional connectives.

Necessity Introduction. The introduction rule is straightforward, simply unfolding the definition of validity.

$$\frac{\Delta; \Psi \vdash A \text{ true}}{\Delta; \Gamma \vdash [\Psi]A \text{ true}} \quad \square I$$

Necessity Elimination. The elimination rule is based on the substitution principle: if we can prove $[\Psi]A$ then we are justified in assuming $A \text{ valid}[\Psi]$.

$$\frac{\Delta; \Gamma \vdash [\Psi]A \text{ true} \quad (\Delta, u::A \text{ valid}[\Psi]); \Gamma \vdash C \text{ true}}{\Delta; \Gamma \vdash C \text{ true}} \quad \square E^u$$

As usual, we assume that u is not already declared in Δ so that all hypotheses have distinct labels.

Local Reduction. We have to show that an introduction of necessity followed immediately by its elimination can be reduced, avoiding the detour.

$$\frac{\frac{\frac{\mathcal{D}}{\Delta; \Psi \vdash A \text{ true}}}{\Delta; \Gamma \vdash [\Psi]A \text{ true}} \quad \square I \quad \frac{\mathcal{E}}{(\Delta, u::A \text{ valid}[\Psi]); \Gamma \vdash C \text{ true}}}{\Delta; \Gamma \vdash C \text{ true}} \quad \square E^u}{\Delta; \Gamma \vdash C \text{ true}} \quad \Longrightarrow_R \quad \frac{\mathcal{E}'}{\Delta; \Gamma \vdash C \text{ true}}$$

Here, \mathcal{E}' is obtained from \mathcal{E} application of the contextual substitution principle, replacing uses of the assumption u with the proof \mathcal{D} .

Local Expansion. We have to show that we can apply the elimination rule in such a way that we can reconstitute a proof of $[\Psi]A$ from the result.

$$\frac{\frac{\mathcal{D}}{\Delta; \Gamma \vdash [\Psi]A \text{ true}} \quad \frac{\frac{\frac{\mathcal{E}}{(\Delta, u::A \text{ valid}[\Psi]); \Psi \vdash \Psi}}{(\Delta, u::A \text{ valid}[\Psi]); \Psi \vdash A \text{ true}} \text{ctxhyp}_u}{(\Delta, u::A \text{ valid}[\Psi]); \Gamma \vdash [\Psi]A \text{ true}} \quad \square I}{\Delta; \Gamma \vdash [\Psi]A \text{ true}} \quad \square E^u}{\Delta; \Gamma \vdash [\Psi]A \text{ true}} \quad \Longrightarrow_E$$

Here the derivation \mathcal{E} exists by the contextual identity principle. This is analogous to an appeal to the contextual substitution principle during local reduction.

We call the resulting logic intuitionistic contextual modal logic (ICML). The rules of natural deduction are summarized in Figure 1.

$$\begin{array}{c}
\frac{}{\Delta; (\Gamma, x:A \text{ true}, \Gamma') \vdash A \text{ true}} \text{hyp}_x \quad \frac{(\Delta, u::A \text{ valid}[\Psi], \Delta'); \Gamma \vdash \Psi}{(\Delta, u::A \text{ valid}[\Psi], \Delta'); \Gamma \vdash A \text{ true}} \text{ctxhyp}_u \\
\frac{\Delta; \Gamma, x:A \text{ true} \vdash B \text{ true}}{\Delta; \Gamma \vdash A \rightarrow B \text{ true}} \rightarrow\text{I}^x \quad \frac{\Delta; \Gamma \vdash A \rightarrow B \text{ true} \quad \Delta; \Gamma \vdash A \text{ true}}{\Delta; \Gamma \vdash B \text{ true}} \rightarrow\text{E} \\
\frac{\Delta; \Psi \vdash A \text{ true}}{\Delta; \Gamma \vdash [\Psi]A \text{ true}} \Box\text{I} \quad \frac{\Delta; \Gamma \vdash [\Psi]A \text{ true} \quad \Delta, u::A \text{ valid}[\Psi]; \Gamma \vdash C \text{ true}}{\Delta; \Gamma \vdash C \text{ true}} \Box\text{E}^u \\
\frac{\Delta; \Gamma \vdash B_1 \text{ true} \quad \dots \quad \Delta; \Gamma \vdash B_m \text{ true}}{\Delta; \Gamma \vdash y_1:B_1 \text{ true}, \dots, y_m:B_m \text{ true}} \text{ctx}
\end{array}$$

Fig. 1. Natural Deduction for ICML

2.5 Examples

Before presenting more of the theory of intuitionistic contextual modal logic, we give some example derivations. We follow the usual convention that the modal operator (here $[\Psi]$) binds more tightly than the other logical connectives and that “ \rightarrow ” associates to the right. Where the names of the variables in contexts Ψ do not matter we abbreviate $x:A \text{ true}$ simply by A .

- (1) $\vdash [C]A \rightarrow [C, D]A \text{ true}$
- (2) $\vdash [C, C]A \rightarrow [C]A \text{ true}$
- (3) $\vdash [A]A \text{ true}$
- (4) $\vdash [A]B \rightarrow [A][B]C \rightarrow [A]C \text{ true}$
- (5) $\vdash [\cdot]A \rightarrow A \text{ true}$
- (6) $\vdash [C]A \rightarrow [D][C]A \text{ true}$
- (7) $\vdash [C](A \rightarrow B) \rightarrow ([D]A \rightarrow [C, D]B) \text{ true}$
- (8) $\vdash [A](A \rightarrow B) \rightarrow ([B]C \rightarrow [A]C) \text{ true}$
- (9) $\not\vdash A \rightarrow [\cdot]A \text{ true}$
- (10) $\not\vdash A \rightarrow [B]A \text{ true}$
- (11) $\not\vdash (A \rightarrow B) \rightarrow [B]C \rightarrow [A]C \text{ true}$

The derivation of (1) closely mirrors the local expansion for the type $[\Psi]A$, as presented in the previous section.

$$\begin{array}{c}
\frac{}{x:[C]A \text{ true} \vdash [C]A \text{ true}} \text{hyp}_x \quad \frac{\frac{\frac{u::A \text{ valid}[C]; x:C \text{ true}, y:D \text{ true} \vdash C \text{ true}}{u::A \text{ valid}[C]; x:C \text{ true}, y:D \text{ true} \vdash C} \text{ctx}}{u::A \text{ valid}[C]; x:C \text{ true}, y:D \text{ true} \vdash A \text{ true}} \text{ctxhyp}_u}{u::A \text{ valid}[C]; \cdot \vdash [C, D]A \text{ true}} \Box\text{I} \\
\frac{x:[C]A \text{ true} \vdash [C, D]A \text{ true}}{\vdash [C]A \rightarrow [C, D]A \text{ true}} \rightarrow\text{I}^x \quad \Box\text{E}^u
\end{array}$$

We now turn to the derivation of (5). The proof of (5) shows that a proposition derived relative to an empty context is actually unconditionally true, and can thus be obtained in an arbitrary context.

$$\frac{\frac{\frac{}{x:[\cdot]A \text{ true} \vdash [\cdot]A \text{ true}}{\text{hyp}_x} \quad \frac{\frac{}{u::A \text{ valid}[\cdot]; \cdot \vdash \cdot}{}{\text{ctx}}}{u::A \text{ valid}[\cdot]; \cdot \vdash A \text{ true}}{\text{ctxhyp}_u}}{\frac{}{x:[\cdot]A \text{ true} \vdash A \text{ true}}{\square E^u}}{\vdash [\cdot]A \rightarrow A \text{ true}} \rightarrow I^x$$

We omit the remaining derivations. On the other hand, the last three are not derivable in general. We do not yet have the tools to prove this, but we may attempt the proof of (9) from the bottom up.

$$\frac{\frac{\frac{}{\vdash A \text{ true}}{}{\square I}}{x:A \text{ true} \vdash [\cdot]A \text{ true}}{\vdash A \rightarrow [\cdot]A \text{ true}} \rightarrow I^x$$

As can be seen, this approach leaves us with the generally unattainable subgoal of proving $A \text{ true}$ in an empty context. The fact that $A \rightarrow [\cdot]A \text{ true}$ is indeed unprovable will be an immediate consequence of the completeness of the cut-free sequent calculus introduced in Section 3.

2.6 Elementary Properties

In this section we formally prove some properties of intuitionistic contextual modal logic on the connectives we have introduced so far. The main normal form theorem is postponed to Section 4.5; here we concentrate on other structural properties.

Since we have a separate judgment for contextual entailment, we need to state several properties for a conclusion $A \text{ true}$ but also for the conclusion Ψ . We use J to stand for either of these two judgments. We begin with weakening and two simple properties of contextual entailment. We will often use the fundamental properties tacitly.

THEOREM 2.1 WEAKENING, IDENTITY, AND EXTENSION.

- (1) (**Weakening**) If $\Delta; \Gamma \vdash J$ then $\Delta; (\Gamma, x:A \text{ true}, \Gamma') \vdash J$
and $(\Delta, u::A \text{ valid}[\Psi], \Delta'); \Gamma \vdash J$.
- (2) (**Identity**) $\Delta; \Psi \vdash \Psi$.
- (3) (**Extension**) If $\Delta; \Gamma \vdash \Psi$ then $\Delta; (\Gamma, x:A \text{ true}) \vdash (\Psi, x:A \text{ true})$.

PROOF. By straightforward inductions. \square

Next, the substitution properties, some of which we already stated as design principles above.

THEOREM 2.2 SUBSTITUTION.

- (1) (**Substitution**) If $\Delta; \Gamma \vdash A \text{ true}$ and $\Delta; (\Gamma, x:A \text{ true}, \Gamma') \vdash J$
then $\Delta; (\Gamma, \Gamma') \vdash J$.

- (2) (*Contextual Substitution*) If $\Delta; \Psi \vdash A$ and $(\Delta, u::A \text{ valid}[\Psi], \Delta'); \Gamma \vdash J$ then $(\Delta, \Delta'); \Gamma \vdash J$.
- (3) (*Simultaneous Substitution*) If $\Delta; \Gamma \vdash \Psi$ and $\Delta; \Psi \vdash J$ then $\Delta; \Gamma \vdash J$.

PROOF. For substitution and contextual substitution by induction over the second given derivation. For simultaneous substitution by induction over the second given derivation where we augment the first given derivation using extension whenever necessary to apply the induction hypothesis. \square

The only property we have not yet discussed is simultaneous substitution, so named because it eliminates all assumptions in the context Ψ in one operation. If J is itself a context it corresponds to composition of two substitutions, as will become clear when we make proof terms explicit. Otherwise it is simply the application of a simultaneous substitution to a term.

3. VERIFICATIONS

The system of natural deduction presented so far provides a meaning explanation of the connectives in terms of rules of proof. As a preliminary check we have ascertained that, locally, the introduction of a connective followed by its elimination can be reduced. The corresponding global property means that we can always understand the meaning of a proposition by examining only its constituents. More specifically, to accept the meaning explanations as sound we require each true proposition to have a *verification* which only refers to constituent propositions. In general, a proof may not have this property, and the existence of local reductions by themselves is insufficient to guarantee it.

We therefore develop a sequent calculus which is based on two judgments: A has a verification (written $A \text{ verif}$) and A is a hypothesis (written $A \text{ hyp}$). Ignoring, for the moment, validity, our form of hypothetical judgment is

$$x_1:A_1 \text{ hyp}, \dots, x_n:A_n \text{ hyp} \Longrightarrow C \text{ verif}$$

where, as usual, all variables labeling hypotheses must be distinct. We refer to this judgment form as a *sequent*. The derivation of a sequent will serve as its *verification* and must therefore mention only propositions already in the sequent.

These two new judgments are connected explicitly via an *init* rule and two principles which will always be admissible: *identity* and *cut*. We give these three first in their simplest form, ignoring modalities.

Initial Sequents. First we state that we are willing to accept a hypothesis as a verification when the proposition is atomic and can not be decomposed further. We write P for atomic propositions.

$$\frac{}{\Gamma, x:P \text{ hyp}, \Gamma' \Longrightarrow P \text{ verif}} \text{init}_x$$

Note that because judgments about P on the left- and right-hand side are different, this is not a consequence of the general definition of hypothetical judgments, but is an explicit rule relating two judgments.

Identity Principle. There is a verification of any proposition from itself as an assumption.

$$\Gamma, x:A \text{ hyp}, \Gamma' \Longrightarrow A \text{ verif for any proposition } A.$$

Identity is not a rule in our system, since a verification of A should always analyze its structure.

Cut Principle. The principle of cut states that if we have a verification of A then we are justified in making the assumption A hyp.

$$\text{If } \Gamma \Longrightarrow A \text{ verif and } \Gamma, x:A \text{ hyp}, \Gamma' \Longrightarrow C \text{ verif then } \Gamma, \Gamma' \Longrightarrow C \text{ verif.}$$

Cut is not a rule in our system, since it would destroy the very nature of verifications: the proposition A need not occur in the conclusion, but would appear in both premises.

3.1 Sequent Calculus for Contextual Modal Logic

Given the system of natural deduction, it is easy to generalize the ideas above to a sequent calculus that incorporates verifications both for truth and validity. Following standard convention, we do not explicitly mention judgments and labels of hypotheses, identifying them instead by their position in sequents. Sequent then have the form

$$\Delta; \Gamma \Longrightarrow C$$

asserting that C has a verification, and

$$\Delta; \Gamma \Longrightarrow \Psi$$

asserting that each proposition in Ψ has a verification. Moreover, we take the liberty of writing Γ, A and $\Delta, A[\Psi]$ if A or $A[\Psi]$ occurs anywhere in Γ or Δ , respectively. These conventions should be understood only as a short-hand notation and do not introduce any essential ambiguity into the system.

The inference rules for sequents can be derived systematically from the rules for natural deduction. For each connective there are right and left rules, corresponding to the introduction and elimination rules of natural deduction. We have arranged the system so that no explicit rules are required for weakening and contraction, which simplifies the proof of the admissibility of cut and the relationship to natural deduction. We forego the straightforward development and just summarize the resulting system in Figure 2.

These rules are designed to correspond most directly to natural deduction, which means that some of the premises are redundant with respect to provability. Specifically, $A \rightarrow B$ in the right premise of $\rightarrow\text{L}$ and $[\Psi]A$ in the premise of $\Box\text{L}$ are not needed if we are only concerned with provability, although we do need them if we want to model all normal natural deductions directly.

First, we show that cut is admissible in this system. Recall that we follow the general conventions of the sequent calculus and freely allow exchange between assumptions.

THEOREM 3.1 ADMISSIBILITY OF CUT.

(1) (**Weakening**) If $\Delta; \Gamma \Longrightarrow C$ then $\Delta, A[\Psi]; \Gamma \Longrightarrow C$ and $\Delta; \Gamma, A \Longrightarrow C$

$$\begin{array}{c}
\frac{}{\Delta; \Gamma, P \Longrightarrow P} \text{init} \qquad \frac{\Delta, A[\Psi]; \Gamma \Longrightarrow \Psi \quad \Delta, A[\Psi]; \Gamma, A \Longrightarrow C}{\Delta, A[\Psi]; \Gamma \Longrightarrow C} \text{reflect} \\
\frac{\Delta; \Gamma, A \Longrightarrow B}{\Delta; \Gamma \Longrightarrow A \rightarrow B} \rightarrow R \qquad \frac{\Delta; \Gamma, A \rightarrow B \Longrightarrow A \quad \Delta; \Gamma, A \rightarrow B, B \Longrightarrow C}{\Delta; \Gamma, A \rightarrow B \Longrightarrow C} \rightarrow L \\
\frac{\Delta; \Psi \Longrightarrow A}{\Delta; \Gamma \Longrightarrow [\Psi]A} \Box R \qquad \frac{\Delta, A[\Psi]; \Gamma, [\Psi]A \Longrightarrow C}{\Delta; \Gamma, [\Psi]A \Longrightarrow C} \Box L \\
\frac{\Delta; \Gamma \Longrightarrow B_1 \quad \dots \quad \Delta; \Gamma \Longrightarrow B_m}{\Delta; \Gamma \Longrightarrow B_1, \dots, B_m} \text{ctx}
\end{array}$$

Fig. 2. Sequent Calculus for ICML

- (2) (**Identity**) $\Delta; \Gamma, A \Longrightarrow A$ for any proposition A .
(3) (**Contextual Identity**) $\Delta; \Psi \Longrightarrow \Psi$ for any context Ψ .
(4) (**Cut**) If $\Delta; \Gamma \Longrightarrow A$ and $\Delta; \Gamma, A \Longrightarrow C$ then $\Delta; \Gamma \Longrightarrow C$
(5) (**Contextual Cut**) If $\Delta; \Psi \Longrightarrow A$ and $\Delta, A[\Psi]; \Gamma \Longrightarrow C$ then $\Delta; \Gamma \Longrightarrow C$

PROOF. Weakening is a trivial induction over the given derivation. Note that the size or structure of a derivation does not change when weakening is applied. We can therefore use it tacitly in the proof of the other properties.

The identity principles are proved by induction on the structure of A and Ψ , respectively.

Admissibility of cut and contextual cut is proved by a nested induction, first on the structure of A and $A[\Psi]$, and second on the structure of the two derivations that are cut. \square

Now it is easy to establish that every derivation has a verification and vice versa. We only state the main property along these lines. We commit a slight abuse of notation by writing Γ and Δ on both sides, even though the actual judgments in hypotheses for natural deductions and sequents are different (A true vs. A hyp).

THEOREM 3.2 PROOFS AND VERIFICATIONS. $\Delta; \Gamma \vdash A$ true iff $\Delta; \Gamma \Longrightarrow A$.

PROOF. From left to right by induction over the structure of the given proof, using admissibility of identity for hypotheses and cut for elimination rules.

From right to left by induction over the structure of the given verification, using substitution properties. \square

By examining the relationships between proofs and verifications we can see that the admissibility of cut is the global version of local soundness, and that the admissibility of identity is the global version of local completeness. A more formal analysis of this relationship is beyond the scope of the present paper.

3.2 Examples

We revisit only two of the examples from Section 2.5, giving the sequent derivation of (5) and showing that (9) is not provable. We use the admissible **Identity** instead of the **init** rule to emphasize that A need not be atomic.

$$\begin{array}{c}
\frac{}{A[\cdot]; [\cdot]A \Longrightarrow \cdot} \text{ctx} \quad \frac{}{A[\cdot]; [\cdot]A, A \Longrightarrow A} \text{Identity} \\
\hline
\frac{}{A[\cdot]; [\cdot]A \Longrightarrow A} \text{reflect} \\
\frac{}{[\cdot]A \Longrightarrow A} \square L \\
\frac{}{\Longrightarrow [\cdot]A \rightarrow A} \rightarrow R
\end{array}$$

We can now show that (5) is not derivable for an arbitrary A . In each sequent there is at most one applicable rule in the bottom-up construction of a potential derivation, leaving us with the subgoal of verifying A in an empty context of hypotheses. This subgoal does not have a verification in general.

$$\frac{\frac{\cdot \Longrightarrow A}{A \Longrightarrow [\cdot]A}}{\Longrightarrow A \rightarrow [\cdot]A}$$

4. SIMPLE CONTEXTUAL MODAL TYPE THEORY

In this section we develop a type theory based on contextual modal logic. We postpone the treatment of dependencies to the next section and cover functions and the contextual modality we have introduced. This could serve as the foundation for a functional language, although there are some difficulties of extending it to a fully dependent type theory, as we see in Section 5.

4.1 Proof Term Assignment

In the terminology of Martin-Löf [1996], the truth judgment for ICML is *synthetic* in that we have to provide a separate proof as evidence for the judgment. By annotating the judgment with proof terms we obtain two benefits: first, the judgment is now *analytic* in the sense that it contains its evidence, and second, it provides us with a notation for programs under the Curry-Howard interpretation of proofs as programs.

When assigning proof terms we have to be careful to respect variable scope and allow consistent renaming of bound variables, because some operations are applied to full contexts and hence to all variables in a term simultaneously. We therefore chose an abstract syntax where binding is made explicit, using a dot (“.”) to separate a variable or context on the left from its scope on the right.

This proof term assignment uses the following proof term language, where x stands for ordinary variables and u for contextual modal variables. Keeping in mind one of our applications, we will refer to u as standing for *meta-variables*.

$$\begin{array}{l}
\text{Terms } M, N ::= x \mid \text{lam}(x:A. M) \mid \text{app}(M, N) \\
\quad \quad \quad \mid \text{clo}(u, \sigma) \mid \text{box}(\Psi. M) \mid \text{letbox}(M, u. N) \\
\text{Substitutions } \sigma, \tau ::= \cdot \mid \sigma, M/x \\
\text{Contexts } \Gamma, \Psi ::= \cdot \mid \Gamma, x:A \\
\text{Modal Contexts } \Delta ::= \cdot \mid \Delta, u::A[\Psi]
\end{array}$$

$$\begin{array}{c}
\frac{}{\Delta; (\Gamma, x:A, \Gamma') \vdash x : A} \text{hyp} \qquad \frac{(\Delta, u::A[\Psi], \Delta'); \Gamma \vdash \sigma : \Psi}{(\Delta, u::A[\Psi], \Delta'); \Gamma \vdash \text{clo}(u, \sigma) : A} \text{ctxhyp} \\
\frac{\Delta; \Gamma, x:A \vdash M : B}{\Delta; \Gamma \vdash \text{lam}(x:A. M) : A \rightarrow B} \rightarrow I \qquad \frac{\Delta; \Gamma \vdash M : A \rightarrow B \quad \Delta; \Gamma \vdash N : A}{\Delta; \Gamma \vdash \text{app}(M, N) : B} \rightarrow E \\
\frac{\Delta; \Psi \vdash M : A}{\Delta; \Gamma \vdash \text{box}(\Psi. M) : [\Psi]A} \Box I \qquad \frac{\Delta; \Gamma \vdash M : [\Psi]A \quad \Delta, u::A[\Psi]; \Gamma \vdash N : C}{\Delta; \Gamma \vdash \text{letbox}(M, u. N) : C} \Box E \\
\frac{\Delta; \Gamma \vdash N_1 : B_1 \quad \dots \quad \Delta; \Gamma \vdash N_m : B_m}{\Delta; \Gamma \vdash (N_1/y_1, \dots, N_m/y_m) : (y_1:B_1, \dots, y_m:B_m)}
\end{array}$$

Fig. 3. Proof Term Assignment for ICML

Recall that all variables declared in a context or modal context must be distinct.

Since the binding structure is somewhat non-standard, we explicitly define the set of free variables for terms, $\text{FV}(M)$, and substitutions, $\text{FV}(\sigma)$.

$$\begin{array}{ll}
\text{FV}(x) & = \{x\} \\
\text{FV}(\text{lam}(x:A. M)) & = \text{FV}(M) - \{x\} \\
\text{FV}(\text{app}(M, N)) & = \text{FV}(M) \cup \text{FV}(N) \\
\text{FV}(\text{clo}(u, \sigma)) & = \text{FV}(\sigma) \\
\text{FV}(\text{box}(\Psi. M)) & = \{\} \\
\text{FV}(\text{letbox}(M, u. N)) & = \text{FV}(M) \cup \text{FV}(N) \\
\text{FV}(\cdot) & = \{\} \\
\text{FV}(\sigma, M/x) & = \text{FV}(\sigma) \cup \text{FV}(M)
\end{array}$$

The term constructor $\text{box}(\Psi. M)$ binds all the variables from the context Ψ simultaneously. Consequently, the renaming of bound variables or α -conversion now also includes terms of the form $\text{box}(\Psi. M)$ where variables in Ψ may be renamed, and we tacitly apply α -conversion as necessary. Furthermore, according to the typing rule $\Box I$ from Figure 3, Ψ is required to contain all the free variables of M in $\text{box}(\Psi. M)$. Thus, $\text{FV}(\text{box}(\Psi. M))$ is appropriately defined to be empty. M may still contain free meta-variables from the modal context Δ , so we next define the free meta-variables in a term and substitution, $\text{FMV}(M)$ and $\text{FMV}(\sigma)$.

$$\begin{array}{ll}
\text{FMV}(x) & = \{\} \\
\text{FMV}(\text{lam}(x:A. M)) & = \text{FMV}(M) \\
\text{FMV}(\text{app}(M, N)) & = \text{FMV}(M) \cup \text{FMV}(N) \\
\text{FMV}(\text{clo}(u, \sigma)) & = \{u\} \cup \text{FMV}(\sigma) \\
\text{FMV}(\text{box}(\Psi. M)) & = \text{FMV}(M) \\
\text{FMV}(\text{letbox}(M, u. N)) & = \text{FMV}(M) \cup (\text{FMV}(N) - \{u\}) \\
\text{FMV}(\cdot) & = \{\} \\
\text{FMV}(\sigma, M/x) & = \text{FMV}(\sigma) \cup \text{FMV}(M)
\end{array}$$

The context Ψ in $\text{box}(\Psi. M)$ binds ordinary variables, but does not bind any meta-variables, so the free meta-variables of $\text{box}(\Psi. M)$ are those of M . On the other hand, the free meta-variables of $\text{letbox}(M, u. N)$ include those of M and N , but u

has to be excluded because it is bound in $\text{letbox}(M, u. N)$ with the scope extending through N . Of course, bound meta-variables are subject to α -conversion, too.

4.2 Examples

In this section we present the ICML proof terms corresponding to the propositions from the example in Section 2.5. For brevity, we omit the types from the proof terms.

$$\begin{aligned}
M_1 & : [x:C]A \rightarrow [y_1:C, y_2:D]A = \\
& \quad \text{lam}(z. \text{letbox}(z, u. \text{box}(y_1, y_2. \text{clo}(u, [y_1/x]))) \\
M_2 & : [x_1:C, x_2:C]A \rightarrow [y:C]A = \\
& \quad \text{lam}(z. \text{letbox}(z, u. \text{box}(y. \text{clo}(u, [y/x_1, y/x_2]))) \\
M_3 & : [x:A]A = \\
& \quad \text{box}(x. x) \\
M_4 & : [x:A]B \rightarrow [y_1:A][y_2:B]C \rightarrow [z:A]C = \\
& \quad \text{lam}(x_1. \text{lam}(x_2. \\
& \quad \quad \text{letbox}(x_1, u. \text{letbox}(x_2, v. \text{box}(z. \text{letbox}(\text{clo}(v, [z/y_1]), \\
& \quad \quad \quad w. \text{clo}(w, [\text{clo}(u, [z/x])/y_2]))) \\
M_5 & : [\cdot]A \rightarrow A = \\
& \quad \text{lam}(x. \text{letbox}(x, u. \text{clo}(u, [\cdot]))) \\
M_6 & : [x:C]A \rightarrow [y_1:D][y_2:C]A = \\
& \quad \text{lam}(x_1. \text{letbox}(x_1, u. \text{box}(y_1. \text{box}(y_2. \text{clo}(u, [y_2/x]))) \\
M_7 & : [x:C](A \rightarrow B) \rightarrow [y:D]A \rightarrow [z_1:C, z_2:D]B = \\
& \quad \text{lam}(x_1. \text{lam}(x_2. \\
& \quad \quad \text{letbox}(x_1, u. \text{letbox}(x_2, v. \text{box}(z_1, z_2. \text{app}(\text{clo}(u, [z_1/x]), \text{clo}(v, [z_2/y]))) \\
M_8 & : [x:A](A \rightarrow B) \rightarrow [y:B]C \rightarrow [z:A]C = \\
& \quad \text{lam}(x_1. \text{lam}(x_2. \\
& \quad \quad \text{letbox}(x_1, u. \text{letbox}(x_2, v. \text{box}(z:A. \text{clo}(v, [\text{app}(\text{clo}(u, [z/x]), z)/y])))
\end{aligned}$$

4.3 Substitution on Terms

In this section we define the operations of substitution, simultaneous substitution, and substitution for meta-variables, both into a term and substitution. These will all be total operations since any side condition can be satisfied by α -conversion. First, ordinary capture-avoiding substitution on a single variable, $[M/x]N$ and $[M/x]\sigma$.

$$\begin{aligned}
[M/x](x) &= M \\
[M/x](y) &= y \quad \text{if } y \neq x \\
[M/x](\text{lam}(y:B. N)) &= \text{lam}(y:B. [M/x]N) \quad \text{provided } y \notin \text{FV}(M) \text{ and } y \neq x \\
[M/x](\text{app}(N_1, N_2)) &= \text{app}([M/x]N_1, [M/x]N_2) \\
[M/x](\text{clo}(u, \sigma)) &= \text{clo}(u, [M/x]\sigma) \\
[M/x](\text{box}(\Psi. N)) &= \text{box}(\Psi. N) \\
[M/x](\text{letbox}(N_1, u. N_2)) &= \text{letbox}([M/x]N_1, u. [M/x]N_2) \\
[M/x](\cdot) &= \cdot \\
[M/x](\sigma, N/y) &= [M/x]\sigma, ([M/x]N)/y
\end{aligned}$$

The substitution commutes with the constructors in all the cases, except in the case for $\text{box}(\Psi. M)$. As already discussed, $\text{box}(\Psi. M)$ does not contain any free variables, and in particular, it does not contain the free variable x . Thus, $\text{box}(\Psi. M)$ is not changed when $[M/x]$ is applied to it.

This substitution satisfies the following versions of the substitution principle, annotated with proof terms.

THEOREM 4.1 SUBSTITUTION ON TERMS.

- (1) If $\Delta; \Gamma \vdash M : A$ and $\Delta; \Gamma, x:A, \Gamma' \vdash N : C$ then $\Delta; \Gamma, \Gamma' \vdash [M/x]N : C$.
- (2) If $\Delta; \Gamma \vdash M : A$ and $\Delta; \Gamma, x:A, \Gamma' \vdash \sigma : \Psi$ then $\Delta; \Gamma, \Gamma' \vdash [M/x]\sigma : \Psi$.

PROOF. By induction on the structure of the second given derivation. \square

Next we define simultaneous substitution $[\sigma]M$ and $[\sigma]\tau$. It is only total when the substitution σ is defined on all free variables in M and τ , respectively. This will be satisfied, because simultaneous substitution is only applied when the assumptions of the theorem following this definition are satisfied. Simultaneous substitutions commute with the term constructors, as one would expect. The only exception are the terms of the form $\text{box}(\Psi. M)$; they do not contain any free variables and thus remain unchanged by the substitution.

$$\begin{aligned}
[\sigma_1, M/x, \sigma_2](x) &= M \\
[\sigma](\text{lam}(y:B. N)) &= \text{lam}(y:B. [\sigma, y/y]N) \quad \text{if } y \notin \text{FV}(\sigma) \text{ and } y \notin \text{dom}(\sigma) \\
[\sigma](\text{app}(N_1, N_2)) &= \text{app}([\sigma]N_1, [\sigma]N_2) \\
[\sigma](\text{clo}(u, \tau)) &= \text{clo}(u, [\sigma]\tau) \\
[\sigma](\text{box}(\Psi. N)) &= \text{box}(\Psi. N) \\
[\sigma](\text{letbox}(N_1, u. N_2)) &= \text{letbox}([\sigma]N_1, u. [\sigma]N_2) \\
[\sigma](\cdot) &= \cdot \\
[\sigma](\tau, N/y) &= [\sigma]\tau, ([\sigma]N)/y
\end{aligned}$$

Sometimes we need to rename the domain of a substitution to match a given context. When $\sigma = (M_1/x_1, \dots, M_n/x_n)$ and $\Psi = (y_1:A_1, \dots, y_n:A_n)$ then $\sigma/\Psi = (M_1/y_1, \dots, M_n/y_n)$.

Simultaneous substitutions satisfy the simultaneous substitution principle, annotated with proof terms. The second property amounts to composition of the substitutions τ and σ .

THEOREM 4.2 SIMULTANEOUS SUBSTITUTION ON TERMS.

- (1) If $\Delta; \Gamma \vdash \sigma : \Psi$ and $\Delta; \Psi \vdash N : C$ then $\Delta; \Gamma \vdash [\sigma]N : C$.
(2) If $\Delta; \Gamma \vdash \sigma : \Psi$ and $\Delta; \Psi \vdash \tau : \Theta$ then $\Delta; \Gamma \vdash [\sigma]\tau : \Theta$.

PROOF. By induction on the structure of the second given derivation. \square

Substitutions for meta-variables u are a little more difficult. One complication is that simultaneous substitutions σ are not first class: from the logical point of view it makes sense to restrict them to $\text{clo}(u, \sigma)$. So we cannot just substitute a term M to obtain $\text{clo}(M, \sigma)$ because the latter is not well-formed. The solution to this problem becomes clear when we recall that u should stand for meta-variables. The closure $\text{clo}(u, \sigma)$ is a postponed substitution. As soon as we know which term u should stand for we can apply σ to it.

Moreover, because of α -conversion, the variables that are substituted at different occurrences of u may be different. As a result, substitution for a meta-variable must carry a context, written as $[\Psi.M/u]N$ and $[\Psi.M/u]\sigma$ where Ψ binds all free variables in M . This complication can be eliminated in an implementation of our calculus based on de Bruijn indexes.

$$\begin{aligned}
\llbracket \Psi.M/u \rrbracket(x) &= x \\
\llbracket \Psi.M/u \rrbracket(\text{lam}(y:B. N)) &= \text{lam}(y:B. \llbracket \Psi.M/u \rrbracket N) \\
\llbracket \Psi.M/u \rrbracket(\text{app}(N_1, N_2)) &= \text{app}(\llbracket \Psi.M/u \rrbracket N_1, \llbracket \Psi.M/u \rrbracket N_2) \\
\llbracket \Psi.M/u \rrbracket(\text{clo}(u, \tau)) &= \llbracket \llbracket \Psi.M/u \rrbracket \tau / \Psi \rrbracket M \\
\llbracket \Psi.M/u \rrbracket(\text{clo}(v, \tau)) &= \text{clo}(v, \llbracket \Psi.M/u \rrbracket \tau) \quad \text{provided } v \neq u \\
\llbracket \Psi.M/u \rrbracket(\text{box}(\Gamma. N)) &= \text{box}(\Gamma. \llbracket \Psi.M/u \rrbracket N) \\
\llbracket \Psi.M/u \rrbracket(\text{letbox}(N_1, v. N_2)) &= \text{letbox}(\llbracket \Psi.M/u \rrbracket N_1, v. \llbracket \Psi.M/u \rrbracket N_2) \\
&\quad \text{provided } v \notin \text{FMV}(M) \text{ and } v \neq u \\
\llbracket \Psi.M/u \rrbracket(\cdot) &= \cdot \\
\llbracket \Psi.M/u \rrbracket(\tau, N/y) &= \llbracket \Psi.M/u \rrbracket \tau, (\llbracket \Psi.M/u \rrbracket N)/y
\end{aligned}$$

Applying $\llbracket \Psi.M/u \rrbracket$ to the closure $\text{clo}(u, \tau)$ first obtains the simultaneous substitution $\tau' = \llbracket \Psi.M/u \rrbracket \tau$, but instead of returning $\text{clo}(M, \tau')$, it proceeds to eagerly apply τ' to M . Before τ' can be carried out, however, its domain must be renamed to match the variables in Ψ , denoted by τ'/Ψ .

While the definition of the discussed case may seem circular at first, it is actually well-founded. The computation of τ' recursively invokes $\llbracket \Psi.M/u \rrbracket$ on τ , a constituent of $\text{clo}(u, \tau)$. Then τ'/Ψ is applied to M , but applying simultaneous substitutions has already been defined without appeal to meta-variable substitution.

As an illustration of this operation, consider the expression

$$\llbracket [x:A \rightarrow A, y:A.\text{app}(x, y)/u]\text{clo}(u, [\text{lam}(z:A. z)/x, t/y]) \rrbracket$$

Here we assume that $u::B[x:A \rightarrow A, y:A]$ and $t:A$ (and t does not depend on u), so that all the terms involved are well typed. We obtain $\text{app}(\text{lam}(z:A. z), t)$, because upon substituting $\text{app}(x, y)$ for u , the local variables x and y are immediately replaced with $\text{lam}(z:A. z)$ and t , respectively.

Another property worth emphasizing here is that the side conditions in the definition of meta-variable substitution only involve meta-variables, but no checks involving ordinary variables are needed. This property will be exploited in Section 7 to efficiently implement meta-variables in a logical framework.

Substitution of a meta-variable satisfies the following contextual substitution property. Again, this is the contextual substitution property from Section 2.6 annotated with proof terms.

THEOREM 4.3 CONTEXTUAL SUBSTITUTION ON TERMS.

- (1) If $\Delta; \Psi \vdash M : A$ and $(\Delta, u::A[\Psi], \Delta'); \Gamma \vdash N : C$
then $(\Delta, \Delta'); \Gamma \vdash \llbracket \Psi.M/u \rrbracket N : C$.
- (2) If $\Delta; \Psi \vdash M : A$ and $(\Delta, u::A[\Psi], \Delta'); \Gamma \vdash \tau : \Theta$
then $(\Delta, \Delta'); \Gamma \vdash \llbracket \Psi.M/u \rrbracket \tau : \Theta$.

PROOF. By simple inductions on the second given derivation, appealing to Theorem 4.2 in the case for meta-variables. \square

4.4 Proof Reductions and Expansions

In this section we present the local reductions and expansions, employing the substitution operations defined above. First, the two local reductions.

$$\begin{aligned} \text{app}(\text{lam}(x:A. N), M) &\Longrightarrow_R [M/x]N \\ \text{letbox}(\text{box}(\Psi. M), u. N) &\Longrightarrow_R \llbracket \Psi.M/u \rrbracket N \end{aligned}$$

In order to formulate the local expansion for the contextual modality, we need to construct identity substitutions. The corresponding operation id , is defined for each context Ψ as follows.

$$\begin{aligned} \text{id}_{(\cdot)} &= (\cdot) \\ \text{id}_{(\Psi, x:A)} &= (\text{id}_\Psi, x/x) \end{aligned}$$

Obviously, id_Ψ stands for different substitutions, depending on the particular instance of Ψ . In other words, id is a meta-level operation just like substitution, and not a proper part of the proof term calculus. We easily verify that it is indeed an identity substitution.

THEOREM 4.4 IDENTITY SUBSTITUTION.

- (1) $\Delta; \Psi \vdash \text{id}_\Psi : \Psi$ for any Ψ .
- (2) $[\text{id}_\Psi]M = M$ provided $\Delta; \Psi \vdash M : A$ for some Δ and A .
- (3) $[\text{id}_\Psi]\sigma = \sigma$ provided $\Delta; \Psi \vdash \sigma : \Gamma$ for some Δ and Γ .

PROOF. By easy inductions. \square

Now we can write out the two local expansions.

$$\begin{aligned} M : A \rightarrow B &\Longrightarrow_E \text{lam}(x:A. \text{app}(M, x)) \quad \text{for } x \notin \text{FV}(M) \\ M : [\Psi]A &\Longrightarrow_E \text{letbox}(M, u. \text{clo}(u, \text{id}_\Psi)) \end{aligned}$$

Due to the presence of the `letbox` elimination, the calculus also requires two commuting reductions to reach a normal form in which no types are introduced and later eliminated. An example to illustrate the need for the commuting reduction is

$$\begin{aligned} \text{lam}(x:[\cdot]A. \text{app}(\text{letbox}(x, u. \text{lam}(z:[\cdot]A \rightarrow [\cdot]A. \text{app}(z, \text{box}(\cdot. \text{clo}(u, \cdot))))), \\ \text{lam}(w:[\cdot]A. w))) \end{aligned}$$

In a more natural mathematical syntax:

$$\lambda x.(\text{let box } u = x \text{ in } \lambda z.z (\text{box } u) \text{ end}) (\lambda w.w)$$

This term has no redex because the `letbox` hides the application of $\lambda z.z (\text{box } u)$ to $\lambda w.w$. This leads to the following commuting reductions:

$$\begin{aligned} \text{app}(\text{letbox}(N_1, v. N_2), M) &\Longrightarrow_C \text{letbox}(N_1, v. \text{app}(N_2, M)) \\ &\quad \text{provided } v \notin \text{FMV}(M) \\ \text{letbox}(\text{letbox}(N_1, v. N_2), u. M) &\Longrightarrow_C \text{letbox}(N_1, v. \text{letbox}(N_2, u. M)) \\ &\quad \text{provided } v \notin \text{FMV}(u.M) \end{aligned}$$

It is easy to see from the substitution properties that local reductions, local expansions, and commuting reductions preserve types.

THEOREM 4.5 SUBJECT REDUCTION AND EXPANSION.

- (1) If $\Delta; \Gamma \vdash M : A$ and $M \Longrightarrow_R M'$ then $\Delta; \Gamma \vdash M' : A$.
- (2) If $\Delta; \Gamma \vdash M : A$ and $M : A \Longrightarrow_E M'$ then $\Delta; \Gamma \vdash M' : A$.
- (3) If $\Delta; \Gamma \vdash M : A$ and $M \Longrightarrow_C M'$ then $\Delta; \Gamma \vdash M' : A$.

PROOF. For subject reduction by inversion on the given typing derivation and appeal to one of the substitution properties. For subject expansion the proof is direct. For commuting reductions by inversion on the giving typing derivation and weakening. \square

4.5 Strong Normalization

In order to prove strong normalization, we can exploit strong normalization for the simply-typed λ -calculus with sums under permutation conversions [de Groote 2002]. The result is somewhat weaker than the corresponding property for the usual calculi of explicit substitution, because we treat application of substitutions as a single-step meta-level operation. For a reduction relation \Longrightarrow , we write \Longrightarrow^* for zero or more reductions and \Longrightarrow^+ for one or more reductions. We omit the straightforward congruence rules for all reductions and note that subject reduction (Theorem 4.5) continues to hold.

For the target of our translation, we use the following fragment of the simply-typed λ -calculus, where a stands for base types.

$$\begin{aligned} \text{Types } \alpha, \beta, \gamma &::= a \mid \alpha \rightarrow \beta \mid \alpha + \beta \\ \text{Terms } s, t, r &::= x \mid \lambda x.t \mid st \mid \iota_1(t) \mid \iota_2(t) \mid \delta(s, x.t, y.r) \end{aligned}$$

λ -abstraction and application are as usual, ι_1 and ι_2 are left and right injections into a sum, and δ represents a case construct. We continue to use our convention that $x.t$ binds x with scope t . We omit the standard typing rules but show the reductions. De Groote calls the last two *permutation conversions* even though they

are applied as reductions.

$$\begin{aligned}
(\lambda x.t) s &\Longrightarrow_D [s/x]t \\
\delta(\iota_1(s), x.t, y.r) &\Longrightarrow_D [s/x]t \\
\delta(\iota_2(s), x.t, y.r) &\Longrightarrow_D [s/y]r \\
\delta(s, x.t, y.r) p &\Longrightarrow_D \delta(s, x.t p, y.r p) \\
\delta(\delta(s, x.t, y.r), u.p, v.q) &\Longrightarrow_D \delta(s, x.\delta(t, u.p, v.q), y.\delta(r, u.p, v.q)) \\
&\quad \text{provided } x \notin \text{FV}(u.p, v.q), y \notin \text{FV}(u.p, v.q)
\end{aligned}$$

We use three auxiliary operations on terms and types in the target language to form iterated function types, abstract over a context, and apply a term to a sequence of arguments.

$$\begin{aligned}
\Pi(x_1:\beta_1, \dots, x_m:\beta_m)(\alpha) &= \beta_1 \rightarrow \dots \rightarrow \beta_m \rightarrow \alpha \\
\Lambda(x_1:\beta_1, \dots, x_m:\beta_m)(s) &= \lambda x_1 \dots \lambda x_m.s \\
\text{App}(s)(t_1, \dots, t_m) &= s t_1 \dots t_m
\end{aligned}$$

The interpretation of types is a simple compositional translation $(A)^\circ$, where $[y_1:B_1, \dots, y_m:B_m]A$ is mapped $\alpha + \alpha$, where $\alpha = B_1^\circ \rightarrow \dots \rightarrow B_m^\circ \rightarrow A^\circ$. The sum is necessary because the target of the translation does not have a **let** construct. We also extend this translation to ordinary hypotheses and contextual validity assumptions.

$$\begin{aligned}
(a)^\circ &= a \\
(A \rightarrow B)^\circ &= A^\circ \rightarrow B^\circ \\
([\Psi]A)^\circ &= (\Pi(\Psi^\circ)(A^\circ)) + (\Pi(\Psi^\circ)(A^\circ)) \\
(\cdot)^\circ &= \cdot \\
(\Gamma, x:A)^\circ &= \Gamma^\circ, x:A^\circ \\
(\cdot)^\circ &= \cdot \\
(\Delta, u::A[\Psi])^\circ &= \Delta^\circ, u:\Pi(\Psi^\circ)(A^\circ)
\end{aligned}$$

In the translation of terms $(M)^\circ$, meta-variables are mapped to ordinary variables that take additional arguments, and closures $\text{clo}(u, (M_1/y_1, \dots, M_m/y_m))$ are mapped to corresponding sequences of applications $((u M_1) \dots M_m)$. The **letbox** construct is translated to a case construct in order to preserve the commuting reductions.

$$\begin{aligned}
(x)^\circ &= x \\
(\text{lam}(x:A.M))^\circ &= \lambda x.M^\circ \\
(\text{app}(M, N))^\circ &= M^\circ N^\circ \\
(\text{clo}(u, \sigma))^\circ &= \text{App}(u)(\sigma^\circ) \\
(\text{box}(\Psi.M))^\circ &= \iota_1(\Lambda(\Psi^\circ)(M^\circ)) \\
(\text{letbox}(M, u.N))^\circ &= \delta(M^\circ, u.N^\circ, u.N^\circ) \\
(\cdot)^\circ &= \cdot \\
(\sigma, M/x)^\circ &= \sigma^\circ, M^\circ
\end{aligned}$$

The choice of the first or second injection in the translation of **box** is arbitrary.

Under these translations, the following theorem is easy to show. We omit some properties pertaining to simultaneous substitutions that follow directly from the properties of terms.

THEOREM 4.6 INTERPRETATION.

- (1) If $\Delta; \Gamma \vdash M : A$ then $\Delta^\circ, \Gamma^\circ \vdash M^\circ : A^\circ$.
- (2) If $\Delta; \Gamma \vdash M : A$ and $\Delta; \Gamma, x:A, \Gamma' \vdash N : B$ then $([M/x]N)^\circ = [M^\circ/x]N^\circ$.
- (3) If $\Delta; \Gamma \vdash \sigma : \Psi$ and $\Delta; \Psi \vdash M : A$ then $\mathbf{App}(\Lambda(\Psi^\circ)(M^\circ))(\sigma^\circ) \Longrightarrow_D^* ([\sigma]M)^\circ$.
- (4) If $\Delta; \Psi \vdash M : A$ and $(\Delta, u::A[\Psi], \Delta'); \Gamma \vdash N : C$ then $[\Lambda(\Psi^\circ)(M^\circ)/u]N^\circ \Longrightarrow_D^* ([\Psi.M/u]N)^\circ$.

PROOF. By simple structural inductions. The typing assumptions are necessary for the substitution properties because $[M/x](\mathbf{box}(\Psi.N)) = \mathbf{box}(\Psi.N)$. The reductions are necessary in the last part to eliminate potential redexes after substituting a λ -abstraction for u . \square

THEOREM 4.7 STRONG NORMALIZATION.

- (1) If $\Delta; \Gamma \vdash M : A$ and $M \Longrightarrow_R N$ or $M \Longrightarrow_C N$ then $M^\circ \Longrightarrow_D^+ N^\circ$.
- (2) The combination of \Longrightarrow_R and \Longrightarrow_C is strongly normalizing on well-typed terms.

PROOF. Part (1) follows by definition of \Longrightarrow_R and \Longrightarrow_C , inversion on the typing derivation, and appeal to the preceding theorem. Part (2) follows from Part (1) and the strong normalization theorem for the simply-typed λ -calculus with sums [de Groote 2002]. \square

5. DEPENDENT CONTEXTUAL MODAL TYPE THEORY

As shown in the previous section, a type theory with the contextual modal operator requires commuting reductions and therefore lacks unique canonical forms. In the context of functional programming this is acceptable, because commuting reductions never arise when evaluating a closed term because we do not normalize under variable binding operators. For a logical framework such as LF [Harper et al. 1993], however, object-language expressions and deductions are represented by canonical forms in the type theory underlying the logical framework. Therefore commuting conversions are to be avoided: they complicate the definitional equality and make it difficult to obtain adequate encodings due to the lack of uniqueness of canonical forms.

In this section we define a dependent contextual modal type theory in the LF family. In order to avoid commuting conversions, we internalize contextual validity not via $[\Psi]A$, but instead we internalize the hypothetical validity judgment. A similar technique has been employed in the definition of linear logical frameworks in order to retain canonical forms [Hodas and Miller 1994; Cervesato and Pfenning 2002]. We exploit here a recent presentation technique for logical frameworks due to Watkins et al. [2002] in which only canonical forms are well-typed. In order to achieve this we divide the term calculus into *atomic objects* R and *normal objects* M . We reuse the same letters as before for various syntactic categories; the reader should therefore keep in mind that they are appropriately restricted in this section. In particular, all references to M and N stand for normal objects, and all references to A , B , or C stand for normal types. Furthermore, contexts Γ and Ψ contain only declarations $x:A$ where A is normal, all terms occurring in substitutions σ are either normal (in N/x) or atomic (in $R//x$), and so on. Finally, while the syntax

only guarantees that terms N are normal (that is, contain no β -redexes), the typing rules will in addition guarantee that all well-typed terms are fully η -expanded.

Normal Kinds	$K ::= \text{type} \mid \Pi x:A.K \mid \Pi u::A[\Psi].K$
Atomic Types	$P, Q ::= a \mid \text{app}(P, N) \mid \text{mapp}(P, \hat{\Psi}.N)$
Normal Types	$A, B, C ::= P \mid \Pi x:A.B \mid \Pi u::A[\Psi].B$
Atomic Objects	$R ::= x \mid \text{app}(R, N) \mid \text{clo}(u, \sigma) \mid \text{mapp}(R, \hat{\Psi}.N)$
Normal Objects	$M, N ::= \text{lam}(x. M) \mid \text{mlam}(u. M) \mid R$
Contexts	$\Gamma, \Psi ::= \cdot \mid \Gamma, x:A$
Substitutions	$\sigma ::= \cdot \mid \sigma, N/x \mid \sigma, R//x$
Modal Contexts	$\Delta ::= \cdot \mid \Delta, u::A[\Psi]$
Signatures	$\Sigma ::= \cdot \mid \Sigma, a:K \mid \Sigma, c:A$

Since we do not need types inside objects, we write $\hat{\Psi}$ for a list of variables x_1, \dots, x_n which we think of as a context Ψ without types.

Signatures declare global constants and never change in the course of a typing derivation. We therefore suppress the signatures throughout. It is important that the restriction regarding validity does not affect occurrences of constants, since they are assumed to be meaningful in all possible worlds.

At the heart of the presentation technique are two observations. The first is that we can characterize canonical forms via bi-directional type-checking. The second is that we can formulate normalization as a primitive recursive functional, exploiting the structure of the types and objects. To see the need for the latter, consider the standard rule for application

$$\frac{\Gamma \vdash M : \Pi x:A.B \quad \Gamma \vdash N : A}{\Gamma \vdash M N : [N/x]B} \text{PIE}$$

This is not correct here because $[N/x]B$ is in general not a canonical form. Instead we need to define a form of substitution $[N/x]_A^a(B)$ which is guaranteed to return a canonical form, given that N is canonical at type A , and B is a canonical type. We postpone the somewhat tedious definition of this operator to Section 5.1, where one can also find the related operations $[[\hat{\Psi}.M/u]_{A[\Psi]}^a(B)$, $[\sigma]_{\Psi}^a(B)$, and auxiliary definitions. We refer to these operations collectively as *hereditary substitutions*.

Typing at the level of objects is divided into three judgments:

$$\begin{aligned} \Delta; \Gamma \vdash M \Leftarrow A & \text{ Check normal object } M \text{ against canonical } A \\ \Delta; \Gamma \vdash R \Rightarrow A & \text{ Synthesize canonical } A \text{ for atomic object } R \\ \Delta; \Gamma \vdash \sigma \Leftarrow \Psi & \text{ Check } \sigma \text{ against } \Psi \end{aligned}$$

We always assume that Δ and Γ and the subject (M , R , or σ) are given, and that the contexts Δ and Γ contain only canonical types. For checking $M \Leftarrow A$ we also assume A is given and canonical, and similarly for checking $\sigma \Leftarrow \Psi$ we assume Ψ is given and all types in it are canonical. For synthesis $R \Rightarrow A$ we assume R is given and we generate a canonical A . Similarly, at the level of types we have

$$\begin{aligned} \Delta; \Gamma \vdash A \Leftarrow \text{type} & \text{ Check normal type } A \\ \Delta; \Gamma \vdash P \Rightarrow K & \text{ Synthesize kind } K \text{ for atomic type family } P \end{aligned}$$

with corresponding assumptions on the constituents. We omit the judgments for well-formed kinds; the ones for contexts are given later in this section.

Judgmental Rules. When checking a normal object that happens to be atomic (that is, has the form R) against a type A we have to synthesize the type for R and compare it with A . Since all synthesized types are canonical, this comparison is simple α -conversion. Moreover, we need to enforce that A is atomic. Otherwise our terms would be β -normal but not necessarily η -long.

$$\frac{\Delta; \Gamma \vdash R \Rightarrow P' \quad P' = P}{\Delta; \Gamma \vdash R \Leftarrow P} \Rightarrow \Leftarrow$$

The rules for ordinary variables and constants are as in the simply typed case. For meta-variables we need to be careful about directions and dependencies. While $\text{clo}(u, \sigma)$ synthesizes a type, we need the type of u , namely $A[\Psi]$ so we can check σ against Ψ . Some renaming is left implicit here, as the variables in the domain of σ should match the variables declared in Ψ . Moreover, we need to apply σ to transport A from Ψ (upon which it may depend) to Γ . The hereditary simultaneous substitution $[\sigma]_{\Psi}^{\alpha}(A)$ always returns a canonical type.

$$\frac{}{\Delta; \Gamma, x:A, \Gamma' \vdash x \Rightarrow A} \text{ var} \quad \frac{c:A \in \Sigma}{\Delta; \Gamma \vdash c \Rightarrow A} \text{ con}$$

$$\frac{\Delta, u::A[\Psi], \Delta'; \Gamma \vdash \sigma \Leftarrow \Psi}{\Delta, u::A[\Psi], \Delta'; \Gamma \vdash \text{clo}(u, \sigma) \Rightarrow [\sigma]_{\Psi}^{\alpha}(A)} \text{ mvar}$$

Dependent Functions. Now that our type theory is dependent, we have to be concerned about the well-formedness of types. The formation rule for dependent function types is familiar and straightforward.

$$\frac{\Delta; \Gamma \vdash A \Leftarrow \text{type} \quad \Delta; \Gamma, x:A \vdash B \Leftarrow \text{type}}{\Delta; \Gamma \vdash \Pi x:A. B \Leftarrow \text{type}} \text{ \Pi F}$$

This formation rule entails some restrictions on the occurrences of x in B , since x is an ordinary variable rather than a meta-variable. For example, x cannot occur (free) in $\hat{\Psi}.M$ in $\text{mapp}(R, \hat{\Psi}.M)$ explained below. This is important for the consistency of the type theory.

In general, introduction forms for a type constructor break down a type when read from the conclusion to the premise. This means if the type in the conclusion is given, we can extract the type for the premise, and therefore introduction forms should be checked against a type. Conversely, elimination forms break down the type when read from premise to conclusion. This means if the type in the premise can be synthesized, we can extract the component type for the conclusion, and therefore elimination forms should synthesize their type. Together with the consideration that the types of variables should always be known, this yields the following rule for the dependent function types $\Pi x:A.B$.

$$\frac{\Delta; \Gamma, x:A \vdash M \Leftarrow B}{\Delta; \Gamma \vdash \text{lam}(x. M) \Leftarrow \Pi x:A. B} \text{ \Pi I} \quad \frac{\Delta; \Gamma \vdash R \Rightarrow \Pi x:A. B \quad \Delta; \Gamma \vdash M \Leftarrow A}{\Delta; \Gamma \vdash \text{app}(R, M) \Rightarrow [M/x]_A^{\alpha}(B)} \text{ \Pi E}$$

Contextual Modal Functions. Contextual modal functions can also be dependent, but the corresponding constructor binds a meta-variable instead of an ordinary variable.

$$\frac{\Delta \vdash \Psi \text{ ctx} \quad \Delta; \Psi \vdash A \Leftarrow \text{type} \quad \Delta, u::A[\Psi]; \Gamma \vdash B \Leftarrow \text{type}}{\Delta; \Gamma \vdash \Pi u::A[\Psi].B \Leftarrow \text{type}} \Pi^{\square} \text{F}$$

The first premise shows that for a declaration $u::A[\Psi]$, A must be a valid type in Ψ . In particular, Ψ can not mention any ordinary variables in Γ from the surrounding context. This restriction on the formation of dependent function types is again essential in order to obtain a clean language for meta-variables.

The rule for meta-variables should now be easy to follow, given the foregoing development. The introduction rule for forming a modal function is straightforward. The argument to a modal function must bind its free variables to avoid unexpected capture and related problems. However, the types of these free variables are not needed, since they can be obtained from the type synthesized by R . We therefore write here $\hat{\Psi}$ for the result of erasing types from the context Ψ . Some renaming of bound variables may be necessary to apply this rule to bring the variables in $A[\Psi]$ in accordance with the variables in $\hat{\Psi}.M$.

$$\frac{\Delta, u::A[\Psi]; \Gamma \vdash M \Leftarrow B}{\Delta; \Gamma \vdash \text{mlam}(u.M) \Leftarrow \Pi u::A[\Psi].B} \Pi^{\square} \text{I}$$

$$\frac{\Delta; \Gamma \vdash R \Rightarrow \Pi u::A[\Psi].B \quad \Delta; \Psi \vdash M \Leftarrow A}{\Delta; \Gamma \vdash \text{mapp}(R, \hat{\Psi}.M) \Rightarrow \llbracket \hat{\Psi}.M/u \rrbracket_{A[\Psi]}^a(B)} \Pi^{\square} \text{E}$$

Simultaneous Substitutions. Simultaneous substitutions are checked against a context which prescribes types for the terms in the substitution. The principal complication here are dependencies. When we check $\Delta; \Gamma \vdash (\sigma, M/x) \Leftarrow (\Psi, x:A)$ then A will be canonical in Ψ , while the substitution itself is checked in Γ . So we need to apply σ to A before checking M against the result. However, this substitution must return a canonical type, so we once again need hereditary substitution.

$$\frac{\Delta; \Gamma \vdash (\cdot) \Leftarrow (\cdot)}{\Delta; \Gamma \vdash \sigma \Leftarrow \Psi \quad \Delta; \Gamma \vdash M \Leftarrow [\sigma]_{\Psi}^a(A)} \frac{\Delta; \Gamma \vdash (\sigma, M/x) \Leftarrow (\Psi, x:A)}{\Delta; \Gamma \vdash \sigma \Leftarrow \Psi \quad \Delta; \Gamma \vdash R \Rightarrow A' \quad A' = [\sigma]_{\Psi}^a(A)} \frac{\Delta; \Gamma \vdash (\sigma, R//x) \Leftarrow (\Psi, x:A)}$$

Besides M/x for canonical M , there is a second way to construct a substitution to replace a variable by an atomic term R , written $R//x$. This is justified from the nature of hypothetical judgments, since an assumption $x:A$ represents $x \Rightarrow A$ so we can substitute R for x if $R \Rightarrow A$. Again, dependencies slightly complicate the form of the rule which is the third one listed above.

Substitutions $R//x$ are necessary so that we can extend a given substitution with $x//x$ when traversing a binding operator in a type-free way. We could not extend substitutions with x/x , since x is not a canonical term unless it is of atomic type. Identity substitutions can now have the form $x_1//x_1, \dots, x_n//x_n$.

Atomic Types. Atomic types mirror the structure of atomic terms, except that the head of an atomic type is always a constant and never a variable.

$$\frac{a:K \in \Sigma}{\Delta; \Gamma \vdash a \Rightarrow K} \text{acon} \quad \frac{\Delta; \Gamma \vdash P \Rightarrow K \quad K = \text{type}}{\Delta; \Gamma \vdash P \Leftarrow \text{type}} \Rightarrow \Leftarrow$$

$$\frac{\Delta; \Gamma \vdash P \Rightarrow \Pi x:A.K \quad \Delta; \Gamma \vdash N \Leftarrow A}{\Delta; \Gamma \vdash \text{app}(P, N) \Rightarrow [N/x]_A^k(K)} \Pi E$$

$$\frac{\Delta; \Gamma \vdash P \Rightarrow \Pi u::A[\Psi].K \quad \Delta; \Psi \vdash N \Leftarrow A}{\Delta; \Gamma \vdash \text{mapp}(P, \hat{\Psi}.N) \Rightarrow [\hat{\Psi}.N/u]_{A[\Psi]}^k(K)} \Pi^\square E$$

Well-Formed Contexts. The context formation rules are foreshadowed by the formation rules for types. We have the judgments $\Delta \vdash \Gamma \text{ ctx}$ and $\vdash \Delta \text{ mctx}$. We omit the straightforward judgments defining the well-formedness of signatures.

$$\frac{}{\Delta \vdash \cdot \text{ ctx}} \quad \frac{\Delta \vdash \Gamma \text{ ctx} \quad \Delta; \Gamma \vdash A \Leftarrow \text{type}}{\Delta \vdash (\Gamma, x:A) \text{ ctx}}$$

$$\frac{}{\vdash \cdot \text{ mctx}} \quad \frac{\vdash \Delta \text{ mctx} \quad \Delta \vdash \Psi \text{ ctx} \quad \Delta; \Psi \vdash A \Leftarrow \text{type}}{\vdash (\Delta, u::A[\Psi]) \text{ mctx}}$$

We show these rules here to emphasize that the context Δ of meta-variables is ordered from left to right: meta-variables declared earlier can appear in the types of later meta-variables, but not vice versa.

5.1 Hereditary Substitution Operations

The substitution functions we need must construct canonical terms, since those are the only ones that are well-formed (and the only ones of interest in a logical framework). Hence, in places where the ordinary substitution would construct a redex $(\lambda y. M) N$ we must continue, substituting N for y in M . But that could again create a redex, so any redex created by this embedded substitution operation must also continue, and so on. We therefore call this operation *hereditary substitution*.

The difficulty is that hereditary substitution could easily fail to terminate. Consider, for example, $[(\lambda x. x x)/y]^h(y y)$ which arises if one tries to compute the normal form of $(\lambda y. y y) (\lambda x. x x)$. Clearly, on well-typed terms this should not occur, but since hereditary substitution is part of the typing judgment we have a mutual dependency. Fortunately, hereditary substitution can be defined as a primitive recursive functional if we pass in the type of the variable we substitute for. For example, if we hereditarily substitute $[\lambda y. M/x](x N)$ then, if everything is well-typed, $x:A_1 \rightarrow A_2$ for some A_1 and A_2 . So we would write $[\lambda y. M/x]_{A_1 \rightarrow A_2}(x N)$, indexing the operation with the type of x . If the substitution is to continue hereditarily we must further substitute N for y in M , again hereditarily. But notice that if the

original substitution is well-typed then $y:A_1$, so we invoke $[N/y]_{A_1}M$. Even though we switched from substituting into $x N$ to substituting into M (which may be more complex than $x N$), we have reduced the type of the variable we are substituting for from $A_1 \rightarrow A_2$ to A_1 .

In the formal development it is simpler if we can stick to the structure of the example above and use only non-dependent types in hereditary substitutions. We therefore first define type approximations α and an erasure operation $()^-$ that removes dependencies. Before applying any hereditary substitution $[M/x]_A^a(B)$ we first erase dependencies to obtain $\alpha = A^-$ and then carry out the hereditary substitution formally as $[M/x]_\alpha^a(B)$. A similar convention applies to the other forms of hereditary substitutions.

$$\begin{aligned} \text{Type approximation } \alpha, \beta &::= a \mid \alpha \rightarrow \beta \mid \alpha[\gamma] \Rightarrow \beta \\ \text{Context approximation } \gamma, \psi &::= \cdot \mid \gamma, x:\alpha \mid \gamma, x:- \end{aligned}$$

The last form of context approximation, $x:-$ is needed when the approximate type of x is not available.² It does not arise directly from erasure.

Types and contexts relate to type and context approximations via an erasure operation $()^-$ which we overload to work on types and contexts.

$$\begin{aligned} (a)^- &= a \\ (\text{app}(P, N))^- &= P^- \\ (\text{mapp}(P, \Psi.N))^- &= P^- \\ (\Pi x:A. B)^- &= A^- \rightarrow B^- \\ (\Pi u::A[\Psi]. B)^- &= A^-[\Psi^-] \Rightarrow B^- \\ (\cdot)^- &= \cdot \\ (\Psi, x:A)^- &= \Psi^-, x:A^- \end{aligned}$$

We can now define $[M/x]_\alpha^n(N)$ and $[M/x]_\alpha^r(R)$ by nested induction, first on the structure of the type approximation α and second on the structure of the objects N and R . In other words, we either go to a smaller type approximation (in which case the objects can become larger), or the type approximation remains the same and the objects become smaller. We write $\alpha \leq \beta$ and $\alpha < \beta$ if α occurs in β (as a proper subexpression in the latter case). Such occurrences can be inside a context approximation ψ in the modal approximation $\beta_1[\psi] \Rightarrow \beta_2$, so we also write $\alpha < \psi$ if $\alpha \leq \beta$ for some $y:\beta$ in ψ , and we write $\alpha < \beta[\psi]$ if $\alpha \leq \beta$ or $\alpha < \psi$.

If the original term is not well-typed, a hereditary substitution, though terminating, cannot always return a meaningful term. We formalize this as failure to return a result. Later we show that on well-typed terms, hereditary substitution will always return well-typed terms.

As a convention when discussing expressions that may fail to be defined, whenever we state an equality $T_1 = T_2$, we imply that T_1 and T_2 are both defined and equal.

The LF Fragment. We start with the LF fragment to highlight the basic ideas in the simpler case before adding meta-variables. We have the following operations,

²See the definition of $[\sigma]_\psi^n(\text{lam}(y. M))$ in the electronic appendix.

where the superscript indicates the domain of the operation.³

$$\begin{array}{ll}
[M/x]_{\alpha}^n(N) = N' & \text{Hereditary substitution into } N \\
[M/x]_{\alpha}^r(R) = R' \quad \text{or} \quad M' : \alpha' & \text{Hereditary substitution into } R \\
[M/x]_{\alpha}^a(A) = A' & \text{Hereditary substitution into } A \\
[M/x]_{\alpha}^p(P) = P' & \text{Hereditary substitution into } P
\end{array}$$

Any of these functions can implicitly fail, which is propagated to the top level. Substitution $[M/x]_{\alpha}^r$ into an atomic term could return either an atomic term or a normal term, depending on the variable at the head of the term. If the variable at the head is different from x it returns an atomic term. If it is equal to x it must call substitution hereditarily, in which case the result is normal. In the latter case we also return a type approximation $\alpha' \leq \alpha$. One further remark: where hereditary substitution $[M/x]_{\alpha}^*$ is invoked in the typing judgment, we have $M \Leftarrow A$ for some A such that $\alpha = (A)^-$. As we recurse we maintain this invariant.

The cases for substitution into a normal term are just compositional.

$$\begin{array}{ll}
[M/x]_{\alpha}^n(\text{lam}(y. N)) = \text{lam}(y. N') & \text{where } N' = [M/x]_{\alpha}^n(N) \\
& \text{choosing } y \notin \text{FV}(M) \text{ and } y \neq x \\
[M/x]_{\alpha}^n(R) & = M' \quad \text{if } [M/x]_{\alpha}^r(R) = M' : \alpha' \\
[M/x]_{\alpha}^n(R) & = R' \quad \text{if } [M/x]_{\alpha}^r(R) = R' \\
[M/x]_{\alpha}^n(N) & \text{fails} \quad \text{otherwise}
\end{array}$$

Because of compositionality, the hereditary substitution may fail to be defined only if the substitutions into the subterms are undefined. The side conditions $y \notin \text{FV}(M)$ and $y \neq x$ do not cause failure, because they can always be satisfied by appropriately renaming y .

Substitution into an atomic term may fail to be defined depending on what is returned as a result of substituting into the function part of an application.

$$\begin{array}{ll}
[M/x]_{\alpha}^r(x) & = M : \alpha \\
[M/x]_{\alpha}^r(y) & = y \quad \text{for } y \neq x \\
[M/x]_{\alpha}^r(\text{app}(R, N)) = \text{app}(R', N') & \text{if } [M/x]_{\alpha}^r(R) = R' \text{ and } N' = [M/x]_{\alpha}^n(N) \\
[M/x]_{\alpha}^r(\text{app}(R, N)) = M_2 : \alpha_2 & \text{if } [M/x]_{\alpha}^r(R) = \text{lam}(y. M') : \alpha_1 \rightarrow \alpha_2 \\
& \text{where } \alpha_1 \rightarrow \alpha_2 \leq \alpha \text{ and } N' = [M/x]_{\alpha}^n(N) \\
& \text{and } M_2 = [N'/y]_{\alpha_1}^n(M') \\
[M/x]_{\alpha}^r(\text{app}(R, N)) & \text{fails} \quad \text{otherwise}
\end{array}$$

For example, if $[M/x]_{\alpha}^r(\text{app}(R, N))$ does not return an atomic term or a lambda expression, or if the returned approximation $\alpha_1 \rightarrow \alpha_2$ is not a subexpression of α , then the substitution into $\text{app}(R, N)$ is undefined. Shortly, we will prove that $\alpha_1 \rightarrow \alpha_2$ will always be a subexpression of α so that the check is superfluous. We nevertheless include the check here to make it obvious that the hereditary substitution is well-founded, because recursive appeals to substitutions take place on smaller terms with equal approximation, or on a smaller approximation.

³We use n for normal object, r for atomic object, a for normal type, p for atomic type, and letter s for substitution; k for kind is used earlier but its straightforward definition is omitted. We also use γ for contexts, δ for modal contexts. In general, we write $*$ in the substitution $[M/x]_{\alpha}^*$ to denote n, r, a , etc.

Since variables only range over objects, hereditary substitution for types A and atomic types P is trivially compositional. We only show these cases here and omit them in the rest of the development.

$$\begin{aligned}
[M/x]_\alpha^a(\Pi y:B.C) &= \Pi y:B'.C' && \text{where } B' = [M/x]_\alpha^a(B) \text{ and } C' = [M/x]_\alpha^a(C) \\
&&& \text{choosing } y \notin \text{FV}(M) \text{ and } y \neq x \\
[M/x]_\alpha^a(P) &= P' && \text{where } P' = [M/x]_\alpha^p(P) \\
[M/x]_\alpha^p(a) &= a \\
[M/x]_\alpha^p(\text{app}(P, N)) &= \text{app}(P', N') && \text{where } P' = [M/x]_\alpha^p(P) \text{ and } N' = [M/x]_\alpha^n(N)
\end{aligned}$$

An even more straightforward definition applies for the substitution into kinds.

The Full Language. Following the above idea and the definition of non-hereditary substitution, it is not difficult to extend the definitions to the full language. We present the complete definition of hereditary substitution in the electronic appendix, using the notation

$$\begin{aligned}
&[M/x]_\alpha^*(-) \\
&\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^*(-) \\
&[\sigma]_\psi^*(-)
\end{aligned}$$

for $*$ = $n, r, s, a, p, k, \gamma, \delta$ and $-$ = $N, R, \sigma, A, P, K, \Gamma, \Delta$, respectively.

To illustrate the extension, we present here a characteristic fragment of the definition of hereditary modal substitution, that is, the substitution for a meta-variable. We will concentrate on the case where we apply a hereditary modal substitution to an atomic term $\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R)$. The result of this substitution is either an atomic term R' or a normal term $M' : \alpha'$ with a type approximation.

In the definition of this substitution, we require that $\hat{\Psi}$ lists exactly the variables from the context approximation ψ which appears as a bounding index of the substitution. Because $\hat{\Psi}$ in $\hat{\Psi}.M$ is a binder and its variables can be renamed, this requirement really expresses that the number of bound variables in $\hat{\Psi}.M$ matches the number of variables in the context approximation ψ .

We highlight the most interesting cases, namely when substituting into a meta-variable closure $\text{clo}(u, \sigma)$ and a modal application $\text{mapp}(R, \hat{\Gamma}.N)$.

First, the cases where we have reached a closure. If the meta-variable at the head is the same as we are substituting for, we need to invoke a hereditary simultaneous substitution, similar to the non-hereditary meta-variable substitution in Section 4.3.

$$\begin{aligned}
\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\text{clo}(u, \tau)) &= M' : \alpha && \text{where } \tau' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^s(\tau) \\
&&& \text{and } M' = [\tau'/\hat{\Psi}]_\psi^n(M) \\
\llbracket \hat{\Psi}.M/u \rrbracket_{A[\Psi]}^r(\text{clo}(v, \tau)) &= \text{clo}(v, \tau') && \text{if } v \neq u \\
&&& \text{with } \tau' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^s(\tau) \\
\llbracket \hat{\Psi}.M/u \rrbracket_{A[\Psi]}^r(\text{clo}(v, \tau)) &\text{ fails} && \text{otherwise}
\end{aligned}$$

The failure case only applies when one of the appropriate recursive calls fails.

For an application, the result differs depending on whether the head of R is equal to the meta-variable we are substituting for or not. If so, then substitution needs to proceed hereditarily and we return a normal term together with it approximate

type. If not, we just substitute recursively and re-build an application.

$$\begin{aligned}
\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\text{mapp}(R, \hat{\Gamma}.N)) &= M_2 : \alpha_2 && \text{if } \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R) = \\
&&& \text{mlam}(v.M') : \alpha_1[\gamma] \Rightarrow \alpha_2 \\
&&& \text{where } \alpha_1[\gamma] \Rightarrow \alpha_2 \leq \alpha[\psi] \\
&&& \text{and } \hat{\gamma} = \hat{\Gamma} \\
&&& \text{and } N' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(N) \\
&&& \text{and } M_2 = \llbracket \hat{\Gamma}.N'/v \rrbracket_{\alpha_1[\gamma]}^n(M') \\
\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\text{mapp}(R, \hat{\Gamma}.N)) &= \text{mapp}(R', \hat{\Gamma}.N') && \text{if } \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R) = R' \\
&&& \text{and } N' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(N) \\
\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\text{mapp}(R, \hat{\Gamma}.N)) &\text{ fails} && \text{otherwise}
\end{aligned}$$

In the above cases we do not substitute into $\hat{\Gamma}$, as $\hat{\Gamma}$ is just a list of ordinary variables without any type information. In the first case, we need to verify that $\hat{\gamma} = \hat{\Gamma}$ in order to establish our invariant to allow the final computation of M_2 . Since $\hat{\Gamma}.N$ binds the variable in $\hat{\Gamma}$ and therefore allows their renaming, this requires only that γ and $\hat{\Gamma}$ have the same length. If not, hereditary substitution will fail, as it will in the case where the result of substituting into R returns a normal term which is not a meta-variable abstraction.

5.2 Properties of Hereditary Substitutions

The first property states that the hereditary substitution operations terminate, independently of whether the terms involved are well-typed or not.

THEOREM 5.1 TERMINATION.

- (1) If $\llbracket M/x \rrbracket_{\alpha}^r(R) = M' : \alpha'$ then $\alpha' \leq \alpha$
- (2) $\llbracket M/x \rrbracket_{\alpha}^*(\cdot)$, $\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^*(\cdot)$ and $\llbracket \sigma \rrbracket_{\psi}^*(\cdot)$ terminate, either by returning a result or failing after a finite number of steps.

PROOF. The first part is verified by induction on the definition of all operations. The second part follows by a nested induction, first on the structure of α , $\alpha[\psi]$ and ψ and second on the structure of the term we apply hereditary substitution to. Note that type approximations contain context approximations and vice versa, so the structural ordering on these is meaningful and well-founded. \square

Termination is the key property allows us to disentangle the mutual dependency between typing and normalization which is the most difficult obstacle to overcome in the meta-theory of dependent type theories.

THEOREM 5.2 DECIDABILITY OF TYPE CHECKING.

All judgments in the dependent contextual modal type theory are decidable.

PROOF. The typing judgments are syntax-directed and therefore clearly decidable. Hereditary substitution always terminates, giving us a decision procedure for dependent typing. \square

It is remarkable that we can prove decidability even without proving that hereditary substitution of well-typed terms into well-typed terms again yields well-typed

terms. Nonetheless, such a theorem is critical for the logical framework methodology, and to relate our formulation of the type theories to others in the literature. We state here only the version of the theorem needed to validate the operations that arise during type-checking; more details and appropriate generalizations can be found in the electronic appendix.

THEOREM 5.3 HEREDITARY SUBSTITUTION PRINCIPLES FOR TYPES.

Assume the contexts in the given judgments below are well-formed.

- (1) *If $\Delta; \Gamma \vdash M \Leftarrow A$ and $\Delta; \Gamma, x:A \vdash B \Leftarrow \text{type}$ then $\Delta; \Gamma \vdash [M/x]_{A^-}^a(B) \Leftarrow \text{type}$.*
- (2) *If $\Delta; \Psi \vdash M \Leftarrow A$ and $\Delta, u::A[\Psi]; \Gamma \vdash B \Leftarrow \text{type}$
then $\Delta; [\hat{\Psi}.M/u]_{A^-[\Psi^-]}^a(\Gamma) \vdash [\hat{\Psi}.M/u]_{A^-[\Psi^-]}^a(B) \Leftarrow \text{type}$*
- (3) *If $\Delta; \Gamma \vdash \sigma \Leftarrow \Psi$ and $\Delta; \Psi \vdash B \Leftarrow \text{type}$ then $\Delta; \Gamma \vdash [\sigma]_{\Psi^-}^a(B) \Leftarrow \text{type}$.*

PROOF. Direct consequence of the hereditary substitution principles given in the electronic appendix. In brief, we generalize the judgments over all syntactic categories and then proceed by nested induction on the index of the hereditary substitution (A^- , $A^-[\Psi^-]$, and Ψ^- , respectively) and the structure of the derivation we substitute into. \square

For more detail on the technique of hereditary substitutions and discussion of related issues we refer the reader to the technical report by Watkins et al. [2002]. Simultaneous and modal substitution add some bulk to the development, but do not require any essentially new ideas not already introduced above.

6. STAGED FUNCTIONAL COMPUTATION

In this section we show how contextual modal type theory can avoid generating spurious redexes in staged computation and run-time code generation, producing simpler and more efficient code than is possible with validity alone.

Davies and Pfenning [2001] have proposed the use of the modal necessity operator \Box in order to provide type-theoretic support for staged computation and, more specifically, run-time code generation. Abstractly, values of type $\Box A$ have the form $\text{box}(M)$, where M is an (unevaluated) source expression. In other words, the box constructor functions like a quotation operator. The introduction rule

$$\frac{\Delta; \cdot \vdash M : A}{\Delta; \Gamma \vdash \text{box}(M) : \Box A} \Box I$$

enforces that M does not contain any free ordinary variables x , only other meta-variables in Δ . In the implementation of this calculus for run-time code generation [Wickline et al. 1998], a value $\text{box}(M)$ of type $\Box A$ is actually not maintained as a source expression, but compiled to a generator which will produce compiled code for M at run-time. The modal restriction means that programs are correctly staged, that is, all the information to generate efficient code from M is indeed available when the generator is invoked. Abstractly, the elimination rule

$$\frac{\Delta; \Gamma \vdash M : \Box A \quad \Delta, u::A; \Gamma \vdash N : C}{\Delta; \Gamma \vdash \text{let box } u = M \text{ in } N \text{ end} : C} \Box E$$

binds u to the source expression that M evaluates to and substitutes it in N . Variables $u::A$ then refer to source expressions.

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

In the implementation, an occurrence of u in N inside a box operator corresponds to the inner generator invoking the outer one during its operation, while an occurrence of u in N not underneath a box will invoke the generator and then jump to it, executing the generated code.

As mentioned earlier, in the syntax of this paper, $\Box A$ would be written as $[\cdot]A$. Proof terms are also simple notational variants: $\text{box}(M)$ would be $\text{box}((\cdot).M)$ and $\text{let box } u = M \text{ in } N \text{ end}$ would be $\text{letbox}(M, u, N)$, and variables u become $\text{clo}(u, \cdot)$.

As a simple standard example, consider the exponentiation function $\text{exp } n \ b$, which computes b^n . We present the examples in a syntax reminiscent of Standard ML, including integers, some integer operations, and recursion. We also do not provide a formal operational semantics, but it is a straightforward call-by-value functional language based on the reductions \Longrightarrow_R from Section 4.4. In this setting box is akin to a quotation operator and we do not evaluate in its scope.

```
(* val exp : int -> (int -> int) *)
fun exp 0 = fn b => 1
  | exp n = fn b => b * exp (n-1) b
```

This function will always return a functional object (that is, a closure in the implementation) after one step.

We can stage the computation in such a way that when given the exponent n it generates a function to compute b^n . Using modal types we can not only enforce this staging as shown below, we can also implement it so that it returns a source expression (or its corresponding generator) rather than just a closure.

```
(* val exp : int -> [](int -> int) *)
fun exp 0 = box (fn b => 1)
  | exp n = let box u = exp (n-1) in box (fn b => b * u b) end
```

We then have (using some variable renaming for readability)

```
exp 2 ==> box (fn b2 => b2 * (fn b1 => b1 * (fn b0 => 1) b1) b2)
```

If we carry out some variable-for-variable reductions, we can see that the output represents a somewhat less direct version of

```
box (fn b2 => b2 * (b2 * 1))
```

which would avoid unnecessary function calls and is therefore the ultimately desired residual expression.

The logical view of staged computation proposed above is simple, elegant, and effective [Wickline et al. 1998]. When used as a mechanism for meta-programming, however, it has the frequently cited drawback [Davies 1996; Taha and Sheard 1997] that the expressions that are manipulated must be closed. The example of `exp 2` shown above is one example where the residual code could be improved further if one could manipulate expressions with free variables as in work by Davies [1996]

and recent research on MetaML [Calcagno et al. 2004] which are based on temporal logic instead of the modal logic of validity. Moving to temporal logic, however, comes at a price, because we still need to track closed expressions if we want to be able to evaluate them explicitly, which accounts for a good deal of the complexity of the MetaML type system.

Instead of using temporal type theory we can overcome these difficulties by generalizing from modal logic to contextual modal logic. Recall from Sections 2.4 and 4.1 that contextual modal logic is an extension of the modal logic proposed above where $\Box A$ corresponds exactly to $[\cdot]A$. We can now reformulate the function of type $\text{int} \rightarrow [\cdot](\text{int} \rightarrow \text{int})$ into a function of type $\text{int} \rightarrow [\mathbf{b}:\text{int}]\text{int}$, where the output is closed *except* that it can mention one variable (named \mathbf{b} in the type).

In the concrete syntax, the `box` operator now is itself a binding operator (instead of taking a function as an argument), where we write `box b1, ..., bn => M` for `box(b1, ..., bn. M)`. We write a substitution as a tuple $[M_1, \dots, M_n]$ which is applied to a meta-variable u by position, rather than by name in the form $u[M_1, \dots, M_n]$. Below is the improved exponentiation function.

```
(* val exp : int -> [b:int]int *)
fun exp 0 = box b => 1
  | exp n = let box u = exp (n-1) in box b => b * u[b] end
```

When substituting an expression M for u in $u b$ in the previous version, we generate a residual function application $M b$. Here we substitute $[(b_2.M)/u](u [b_1/b_2])$ which directly applies the substitution $[b_1/b_2]$ to M . In this manner, we directly obtain the desired residual expression.

```
exp 2 ==> box (b => b * (b * 1))
```

Other approaches, such as Taha and Nielsen's environment classifiers [Taha and Nielsen 2003], may be more parsimonious on some examples and perhaps more amenable to type inference (which we do not consider here), but they come at a high price in syntactic and conceptual complexity. One of the consequences of this complexity is that it is difficult to see how one might formally reason about the correctness of staged programs. Our type-theoretic setting could provide an advantage in that respect since it is consistent with dependent types. A brief discussion of this possibility can be found among the future work in Section 9.2.

We close this section with the observation that we can convert between the two forms of abstraction.

```
(* val f : [] (A -> B) -> [y:A]B *)
fun f x = let box u = x in box y => u[] y end

(* val g : [y:A]B -> [] (A -> B) *)
fun g z = let box v = z in box (fn y => v[y]) end
```

The function to evaluate an expression still requires its argument to be closed, when the language is pure.

```
(* val eval : []A -> A *)
fun eval x = let box u = x in u[] end
```


However, when the language has non-termination, exceptions, or other expression of arbitrary type, we can write, for example,

```
(* val eval' : [y:B]A -> A *)
fun eval' x = let box u = x in u[raise exn] end
```

Then `eval'` (`box y => M`) raises the exception `exn` if the evaluation of M would refer to y .

7. IMPLEMENTATION OF META-VARIABLES IN LOGICAL FRAMEWORKS

In this section, we sketch how dependent contextual modal type theory as presented in Section 5 can be employed for the efficient implementation of meta-variables in a logical framework. This is a different form of application than given in the previous section, not only because we rely on canonical forms, but also because explicit quantification over meta-variables is at present not available to the user, only in the implementation. More details on this application in the context of the Twelf system [Pfenning and Schürmann 1999] are given by Pientka [2003] and Pientka and Pfenning [2003]. For the non-dependent case, similar observations (without any logical foundations) have been made by Dowek et al. [1995; 1996] and later generalized by Muñoz [2001].

An implementation of a logical framework must handle meta-variables, which are used both during type reconstruction and during proof search. These meta-variables are placeholders for terms in the language under consideration, to be determined or constrained by unification, proof search, or perhaps even directly by the user. Since dependent type theories reify proofs as objects, unsolved subgoals can also be represented as meta-variables.

When the value of a meta-variable is determined we say that the meta-variable is *instantiated* in order to distinguish this from the process of substitution of a term for an ordinary variable. To avoid misunderstandings, we will refer to ordinary variables as *parameters*. In contrast to meta-variables, these should be interpreted universally, that is, they are not subject to instantiation.

There are four principal questions in the handling of meta-variables:

- (1) Which parameters (that is, bound or universally quantified variables) are allowed to occur in a term that instantiates a meta-variable?
- (2) What constraints does the type theory impose on occurrences of meta-variables in the contexts or types of other meta-variables?
- (3) How do we implement meta-variables and the instantiation operation?
- (4) Which algorithm do we use for unification or constraint simplification?

The first three questions apply, regardless of whether the meta-variables stand for open subgoals during proof search, or if they are placeholders for terms to be determined by unification.

The last one is beyond the scope of this paper, although Dowek et al. [1995; 1996] and Pientka [Pientka 2003] provide answers that are consistent with our approach. For the first three questions, dependent contextual modal type theory provides clean and simple answers that have been experimentally validated in Twelf.

Parameter Occurrences. The question which parameters may occur in the term that instantiates a meta-variable is answered quite directly by our type theory. A meta-variable $u::A[\Psi]$ can depend exactly on the parameters in Ψ . All occurrences of u in a term have the form $\text{clo}(u, \sigma)$ with the typing rule

$$\frac{\Delta, u::A[\Psi], \Delta'; \Gamma \vdash \sigma \Leftarrow \Psi}{\Delta, u::A[\Psi], \Delta'; \Gamma \vdash \text{clo}(u, \sigma) \Rightarrow [\sigma]_{\Psi}^a(A)} \text{ mvar}$$

This means that σ is a mediating substitution, providing a translation from the context Ψ to the ambient context Γ in effect where u occurs. Different occurrences of u may be under different mediating substitutions σ , avoiding possible confusions of variable names. While this means of achieving consistency may appear to be somewhat heavy, it is in fact quite efficiently implementable by using explicit substitutions as sketched below.

Meta-Variable Occurrences. Recall the rules to verify that modal contexts are well-formed.

$$\frac{}{\vdash \cdot \text{mctx}} \quad \frac{\vdash \Delta \text{ mctx} \quad \Delta \vdash \Psi \text{ ctx} \quad \Delta; \Psi \vdash A \Leftarrow \text{type}}{\vdash (\Delta, u::A[\Psi]) \text{ mctx}}$$

This means that we must be able to present all meta-variables that appear in a given term in a linear order so that the types and contexts of meta-variables further to the right may mention meta-variables further to the left. Since the set of currently available meta-variables is usually not made explicit in an implementation, we had missed this invariant in an earlier implementation of unification for the LF logical framework. Fortunately, in the case of higher-order patterns in the sense of Miller [1991], this restriction does not place any additional burden on the implementation since the steps of the algorithm implicitly maintain this invariant [Pientka 2003]. However, when constraints outside the pattern fragment are generated and maintained, additional work during an occurs-check may be necessary to enforce this invariant.

There is a related question: which occurrences of *meta-variables* in the substitution term for another meta-variable are sound. In the first-order, non-dependent case this reduces to a simple occurs-check. In this situation, the question is most easily answered by analyzing simultaneous substitution for meta-variables. This can easily be done in the framework set out here. The details are beyond the scope of this paper and can be found in Pientka's Ph.D. thesis [2003].

Implementation Considerations. The type theory suggests certain implementation choices which we briefly discuss here. The first is dictated by the fact that closures $\text{clo}(u, \sigma)$ are part of the syntax of canonical forms. Consequently, the implementation should support substitutions as explicit syntactic objects. In order to avoid the excessive renaming inherent in the substitutions $\llbracket \hat{\Psi}.M/u \rrbracket_{A[\Psi]}^*$ we are led to a representation using de Bruijn indices, since multiple renamings of a context Ψ are identical. Moreover, the type-free context $\hat{\Psi}$ in $\llbracket \hat{\Psi}.M/u \rrbracket$ becomes unnecessary, since it only records a bound on the de Bruijn index of variables free in M . But these can be obtained from the type $u::A[\Psi]$, so they do not need to be maintained during substitution.

The next critical observation is that in the absence of abstraction over meta-variables, the substitution operation (hereditary or otherwise) for a meta-variable does not need to check *any* condition related to bound variables (please see the electronic appendix). This would suggest to represent meta-variables directly by pointers, and to instantiate meta-variables by updating these pointers, called *grafting* by Dowek et al. [1995; 1996]. Clearly, since each variable occurs as part of a closure $\text{clo}(u, \sigma)$ with an appropriate mediating substitution σ , updating u in place does not destroy any typing or variable naming invariant.

However, there is a catch: our system only treats canonical forms (the only terms of interest in a logical framework), the result of replacing u by a term M would no longer be canonical and not even be syntactically meaningful. The solution is to adopt a notion of *postponed hereditary substitution* directly modeled after the notion of *explicit substitution* [Abadi et al. 1990]. However, instead of giving postponed hereditary substitutions first-class status and investigate properties such as strong normalization or the Church-Rosser property, we view them only as an implementation technique. That means we should never directly analyze terms such as $\text{clo}(M, \sigma)$ unless M is a meta-variable, and all operations should propagate the substitution inwards. One can think of this as a lazy implementation of the hereditary substitution operations detailed in Section 5. With this change, grafting via pointer-update is a correct implementation of instantiation as given by $[\hat{\Psi}.M/u]$.

This notion of lazy hereditary substitution provides a clean and simple organizing principle for many operations on terms in the Twelf implementation. These operations come in two forms: one operating on a weak head-normal form, and one operating on a term with a possible postponed hereditary substitution which is pushed downwards far enough to expose a term in weak head-normal form. As we traverse the term, we therefore expose the canonical term layer by layer, carrying out only as much of the hereditary substitution as necessary. A similar technique has proved quite effective in λ Prolog, although both the formal basis and the implementation differs in some important aspects [Liang and Nadathur 2002]. A further discussion is beyond the scope of this paper, but some additional material can be found in a related paper [Pientka and Pfenning 2003].

8. RELATED WORK

We roughly divide the related work into research on modal logic, functional programming, and meta-variables in theorem proving.

8.1 Modal logics of context

The formal notion of context has been extensively researched for various applications in artificial intelligence. One of the early proposals is by Weyhrauch [1979], motivated by Weyhrauch's work in interactive theorem proving. Later, McCarthy [1987; 1993], proposed a logic of context as a solution to the problem of generality in common-sense reasoning. McCarthy's premise is that logical axioms are usually not general, but crucially depend on the context in which they are asserted. Thus, a formalization of the notion of context is a necessary first step in formulating general axioms for common-sense knowledge, because it facilitates a precise expression of the different meanings that a logical statement can take in different contexts. This line of research eventually led to the formulation of the *propositional logic of*

contexts (PLC) by Buvač et al. [Buvač et al. 1995], and an extension to first-order logic [Buvač 1996]. In these theories, contexts are first class objects (i.e., the theory provides terms that denote contexts), and are modeled by sets of propositions. Each context can make statements about the truth of formulas in other contexts, by using the modal operator *ist*. For example, in a context c_1 we can assert the formula $ist(c_2, A)$, establishing that, when viewed from c_1 , the proposition A is true in the context c_2 .

The work of Giunchiglia on contextual reasoning [Giunchiglia 1993] is motivated by an observation that reasoning about a certain problem usually involves only a small subset of the overall available knowledge. Thus, Giunchiglia proposes that the knowledge be organized into contexts – modular, but related units, that formulate *local theories*, which any particular reasoning agent may or may not employ. As a realization of this idea, we mention the work on *multi-context systems* [Giunchiglia and Serafini 1994; Serafini and Giunchiglia 2002]. In multi-context systems, contexts represent the propositions associated with the knowledge or belief of a particular agent. Each context is therefore labeled by a unique identifier of the agent it represents. The multi-context systems are multi-lingual, meaning that facts from different contexts can be represented using different languages, and there is an apparatus for specifying translations between the languages. Facts that span multiple contexts are specified via so-called *bridge rules*, which are inference rules whose premises and conclusions belong to different languages. Technically, multi-context systems are founded in the multi-modal version of the modal logic K. For a formal relationship between PLC and multi-context systems, we refer to a paper by Serafini and Bouquet [2004].

Attardi and Simi [1995; 1994] propose the notion of a viewpoint as means of expressing relativized truth. A viewpoint is a set of propositions representing assumptions of a theory of interest. For example, the sentence $in('A', vp)$ states that the proposition A is true in the context of a viewpoint vp . The operator *in* for entailment in a viewpoint does not modify the proposition A , but rather, expects the syntactic representation of A (here denoted as $'A'$). The syntactic nature of entailment in a viewpoint is exploited in the representation of other modal operators like knowledge and belief, which are viewed as first-order predicates over *syntactic sentences*, rather than as operators on propositions. The process of entering a new viewpoint is called reification, and exporting a knowledge from one viewpoint into another is called reflection. A proposition can be reified only if it is derived in a restricted way, without using reflection. This restriction is motivated by the inconsistency that arises in formal systems capable of both representing their own syntax, and also reasoning by reflection [Montague 1963].

A type theoretic approach to the formulation of the logic of contexts has been advocated by Thomason [1999; 2003]. In addition to the types of individuals, truth-values, and functions, this type theory features the type w of *possible worlds*. There are no constants or variables of type w , so there is no explicit lambda abstraction over worlds. Instead, the type theory contains operators for forming *intensions*, which are functions over possible worlds. For example, if $e:A$ is an expression, then \hat{e} is an intensional expression and has type $w \rightarrow A$. Dually, if e is an intensional expression of type $w \rightarrow A$, then \check{e} is its *extension*, and has type A . The type

theory further contains the type of *indices*, whose elements denote differently in different contexts. Typical examples are the expression “I” and “here”, which, obviously, have different meanings depending on the context of the person that interprets them. Indices can be built into more complex types, like the type of context-dependent propositions, or the type of context-dependent intensions.

Finally, we mention the work of de Paiva [2003], who studies the relationship between several known propositional logics of contexts, like the PLC logic and multi-context systems mentioned above, and proposes a constructive version of the multi-modal K as the common foundation. The paper further discusses the proof term assignment for this logic, and the corresponding Curry-Howard isomorphism, along the lines of Alechina et al. [2001]. This proof term assignments is different from ours in that it does not separate the ordinary from the meta-variables. As a consequence, it has to employ explicit substitutions even when the modalities are not indexed by contexts.

8.2 Context in functional programming

In functional programming, various formulations of context as a primitive programming construct have been considered [Sato et al. 2001; Sato et al. 2002; Mason 1999; Hashimoto and Ohori 2001], with applications in explaining binding structure and dynamic binding [Dami 1996; 1998], and in incremental program construction [Lee and Friedman 1996].

For example, the $\lambda\kappa\epsilon$ -calculus of Sato et al. [2002], allows a simultaneous abstraction over a set of variables. The expression $\kappa\{\Psi\}.M$ abstracts the variables listed in Ψ from the expression $M : A$. The type of $\kappa\{\Psi\}.M$ is A^Ψ , similar to our type $[\Psi]A$. There are many distinctions, however, between $\lambda\kappa\epsilon$ and the proof terms for ICML, arising mostly because $\lambda\kappa\epsilon$ is not based on modal logic. For example, the context in $\lambda\kappa\epsilon$ associates variables with types, but not with the context that they depend on. This leads to a somewhat complicated formulation, where each variable must be assigned an integer level, and the typing rules and the operation of substitution must perform arithmetic over levels. Another difference is that $\lambda\kappa\epsilon$ -calculus admits outside free variables to appear in the body of a context abstraction. This is prevented in ICML by the typing rules for $[\Psi]A$, and is essential for the modeling and implementation of meta-variables. As explained in Section 7, substitutions for meta-variables do not need to check any side conditions regarding bound variables, allowing meta-variables to be represented by pointers. Without the modal restriction, this aspect of meta-variables would not have been captured.

The $\lambda\epsilon$ -calculus of Sato et al. [2001] is a precursor to $\lambda\kappa\epsilon$. The $\lambda\epsilon$ -calculus provides explicit substitution of terms for variables, but a variable may be used only if it is bound by the explicit substitution.

Mason [1999] extends the untyped λ -calculus with a primitive notion of context, and the related operations for declaring and filling context holes. Holes are similar to our meta-variables, in the sense that each hole is decorated with its corresponding substitution. Mason considers operations of strong and weak application of variable substitutions. Both operations propagate the substitution down to the holes, but differ in their action at the holes. Strong application composes the substitution with the hole’s substitutions (possibly changing the domain of the hole’s substitution), while weak application only distributes over the terms of the hole’s substitution.

Mason’s operation of hole filling correspond to our modal substitution, in the sense that the substitution decorating each hole is applied (weakly or strongly) over the term being placed into the hole. The paper does not consider abstraction over holes, nor is there a term constructor similar to our `box` that introduces a modal distinction between terms.

Hashimoto and Ogori [2001] present a typed calculus which internalizes the notion of computation in context via a type of functions from contexts to values. The calculus distinguishes between ordinary variables and hole variables. Hole variables correspond to our meta-variables; they are typed in a separate hole context Δ which associates with every hole variable u a type A , an interface Ψ (roughly corresponding to our context), but also an explicit substitution σ that specifies the bindings of the hole. Explicit substitutions only rename variables with other variables. Associating holes and substitutions in Δ complicates the system significantly and reduces its expressiveness. For example, each hole variable can only be associated with one substitution, and the typing rules must non-trivially manipulate the hole context. Even more severely, the hole context must be linear, that is, hole variables can only be used once, and ordinary variables can be referenced only when the hole context is empty.

The mentioned works are computationally motivated, and typically do not explore the logical aspects of contexts. A modal logic of contexts as a foundation for functional programming with various kinds of effects has recently been proposed by Nanevski [2003; 2004]. In that work, modal validity classifies computations that are effectful in the sense that the computation results depend on the run-time environment, but execution does not change the run-time environment itself. Examples include programs that read (but do not write) from memory, or programs with control-flow effects. Computations that may change the run-time environment (for example, a computation that writes into memory) are ascribed a type corresponding to modal possibility. Modal types for effectful computations are indexed by sets of *names*, where each name corresponds to a particular effect instance. In this respect, sets of names are similar to context from the current paper. There are significant differences between the two, however; the principal difference being that names have global identity. One and the same name can appear in the index of many different modal operators. In contrast, a variable bound in a context has local scope and is thus necessarily different from any other bound variable in the program.

A modal λ -calculus has been proposed as a foundation for staged computation and metaprogramming by Davies and Pfenning [2001], and was subsequently used as a language for run-time code generation by Wickline et al. [1998; 1998]. In this work, the modal constructor $\Box A$ is used to classify *closed* code of type A . Closed code contains no free ordinary variables, and is thus executable at run time. The applicability of closed code is sometimes restrictive, so alternative approaches for typing *open* code were developed; most significantly, λ° [Davies 1996], and MetaML [Taha and Sheard 1997]. Both of these calculi are based on a proof term system for a variant of temporal logic with a modality expressing truth at the next time moment. From the computational perspective, the modal type classifies open code, but the code expressions cannot be evaluated at run time unless further guarantees

are provided. For some time, the focus of the work on staged computation and metaprogramming has been in combining these two approaches. Taha and Nielsen [2003] and Calcagno et al. [2004] achieve this by relying on environment classifiers, which stand for an unnamed set of variables that a term may be open in. These calculi contain the types of code open in a set of environment classifiers, but also the operations of universal quantification over classifiers. Another approach to typing open code is the ν^\square -calculus developed by Nanevski and Pfenning [2002], where names stand for free variables, and there is also universal quantification over sets of names. The ν^\square -calculus is similar to the calculus presented here, except that in ν^\square modalities are indexed by sets of names rather than by contexts.

8.3 Context in theorem proving

In general we can find two distinct motivations for context calculi in theorem proving. As discussed in the previous section, context calculi provide a foundation for meta-variables which need to be instantiated during proof search via higher-order unification. The concise description of higher-order unification requires a clear distinction between bound variables and meta-variables and efficient implementation techniques rely on instantiation of meta-variables which allows capture. One of the first approaches to allow first-order replacement and avoid the aforementioned problems in the simply-typed setting was developed by Dowek et al. [1995; 1996]. It relies on de Bruijn indices and explicit substitutions and was later extended to dependently typed theories by Muñoz [2001; 2000]. In this approach explicit substitutions can be associated with any term, not only with meta-variables. First-order replacement is achieved in two steps. First, terms are pre-cooked in such a way that scoping constraints are preserved. The main idea is, if some meta-variable X occurs under some λ -abstractions, then substituting M for X needs some processing called lifting to avoid capture of bound variables in M . Second, we can use first-order instantiation, called grafting. Although the use of de Bruijn indices leads to a simple formal system and it is very useful in an implementation, nameless representation of variables via de Bruijn indices are usually hard to read and critical principles are obfuscated by the technical notation.

Independently, Strecker [1999] developed a calculus with meta-variables as first-class objects for a variant of the Extended Calculus of Construction. Similar to our approach, meta-variables are associated with a substitution to eliminate the need for raising and β -reduction during instantiation of meta-variables. However there are several differences. First of all, there is no explicit context for the meta-variables. As a consequence, it is difficult and it may be expensive to verify that an instantiation is valid, since meta-variable dependencies need to be computed afresh. Second, the substitutions attached to meta-variables are not simultaneous substitutions. As a result, the decidability of the typing rules is not directly obvious. Finally, the calculus lacks a logical foundation.

Context calculi not only play an important role in concisely describing meta-variables during higher-order unification, they also are an integral part of a foundation for interactive theorem proving where we need to reason about incomplete proofs or proof with holes [Magnusson 1995; Geuvers and Jojgov 2002; Jojgov 2003; Bogner and de Vrijer 2001]. The status of holes in proofs which may be filled later is essentially the same as meta-variables. Therefore, not surprisingly,

many approaches build directly or indirectly on the treatment of meta-variables in higher-order unification and suffer the same problems and difficulties. For example, in Geuvers and Jojgov [2002], open terms are represented via a kind of meta-level Skolem function. However, in general reduction and instantiation of meta-variables (or holes) do not commute. This problem also arises in Bognar and de Vrijer [2001]. Our work resolves many of these aforementioned problems, since reduction and instantiation naturally commute and require no special treatment. Hence this work may also be viewed as a foundation for treating incomplete proofs.

Recently, Sato et al. [Sato et al. 2003] proposed a calculus of meta-variables. Their approach is fundamentally different from the one presented in this paper. Each variable is given a level, which classifies the variables into bound variables (level 0), meta-variables (level 1) and meta-meta-variables (level 2) etc. Then the authors define a “textual” substitution which allows capture. There are two main obstacles with using this approach. First, since textual substitutions are not capture-avoiding, we will lose confluence. The second problem is that some reductions may get stuck. To address these problems the authors suggest to define reductions in such a way that it takes into account the different levels and keep track of arities of functions. This leads to a carefully engineered system which is confluent and strongly normalizing, although not very intuitive. Moreover, a formulation of algorithms such as higher-order unification with this calculus seems cumbersome.

9. EXTENSIONS AND FUTURE WORK

We briefly sketch three main avenues of future work: adding contextual possibility, defining a dependent necessity for reasoning about staged computation, and allowing explicit quantification over substitutions.

9.1 Contextual possibility

The contextual modal type theory can be extended to include a relativized version of modal possibility, and we have already presented some preliminary steps in that direction [Nanevski et al. 2003]. The idea is to introduce a new judgment $A \text{ poss} \langle \Psi \rangle$ which expresses existential quantification over possible worlds. The judgment $A \text{ poss} \langle \Psi \rangle$ holds if there *exists* a world in which Ψ and A are simultaneously true. This is dual to relativized validity $A \text{ valid} [\Psi]$ which holds if A is true in *every* world in which Ψ is true. The judgment $A \text{ poss} \langle \Psi \rangle$ is internalized by means of a type operator $\langle \Psi \rangle A$, which degenerates into the customary operator for possibility $\diamond A$ when Ψ is the empty context. The typing rules for relativized possibility are an indexed variant of the rules presented by Pfenning and Davies [2001]. We show here only the simply typed variant of the rules.

$$\begin{array}{c}
\frac{\Delta; \Gamma \vdash \sigma : \Psi \quad \Delta; \Gamma \vdash M : A \text{ true}}{\Delta; \Gamma \vdash \langle \sigma, M \rangle : A \text{ poss}(\Psi)} \text{ poss} \\
\frac{\Delta; \Gamma \vdash E : A \text{ poss}(\Psi)}{\Delta; \Gamma \vdash \text{dia}(E) : \langle \Psi \rangle A \text{ true}} \diamond I \\
\frac{\Delta; \Gamma \vdash M : \langle \Psi \rangle A \text{ true} \quad \Delta; \Psi, x:A \text{ true} \vdash E : B \text{ poss}(\Theta)}{\Delta; \Gamma \vdash \text{letdia}(M, \langle \Psi, x \rangle. E) : B \text{ poss}(\Theta)} \diamond E \\
\frac{\Delta; \Gamma \vdash M : [\Psi] A \text{ true} \quad \Delta, u::A \text{ valid}[\Psi]; \Gamma \vdash E : B \text{ poss}(\Theta)}{\Delta; \Gamma \vdash \text{letbox}(M, u. E) : B \text{ poss}(\Theta)} \square E^p
\end{array}$$

The additional rule $\square E^p$ is not related to possibility per se, but eliminates the validity modality into the possibility judgment. It is needed so that the system satisfies the subformula property [Pfenning and Davies 2001]. The rest of ICML can easily be extended to the new judgment. For example, contextual substitution simply commutes with all the new constructors. Extending simultaneous substitutions to the new cases is slightly more involved, as presented below.

$$\begin{array}{ll}
[\sigma](\text{dia}(E)) & = \text{dia}([\sigma]E) \\
[\sigma]\langle \tau, M \rangle & = \langle [\sigma]\tau, [\sigma]M \rangle \\
[\sigma](\text{letdia}(M, \langle \Psi, x \rangle. E)) & = \text{letdia}([\sigma]M, \langle \Psi, x \rangle. E) \\
[\sigma](\text{letbox}(M, u. E)) & = \text{letbox}([\sigma]M, u. [\sigma]E)
\end{array}$$

Observe that in the `letdia` clause of the definition, the substitution σ is not applied to E ; the free variables of E are all bound by Ψ, x , and thus cannot be influenced by σ .

Just as Pfenning and Davies [2001], we require a new substitution operation $\langle\langle F / \langle \Psi, x \rangle \rangle\rangle E$, where F and E are expressions witnessing relativized possibility. If $F : A \text{ poss}(\Psi)$ and $\Psi, x:A \text{ true} \vdash E : B \text{ poss}(\Theta)$ then $\langle\langle F / \langle \Psi, x \rangle \rangle\rangle E$ is a witness of relativized possibility $B \text{ poss}(\Theta)$. The operation is defined by induction on the structure of F :

$$\begin{array}{l}
\langle\langle \langle \sigma, M \rangle / \langle \Psi, x \rangle \rangle\rangle E = [\sigma / \Psi, M / x]E \\
\langle\langle \text{letdia}(M, \langle \Gamma, y \rangle. F') / \langle \Psi, x \rangle \rangle\rangle E = \text{letdia}(M, \langle \Gamma, y \rangle. \langle\langle F' / \langle \Psi, x \rangle \rangle\rangle E) \\
\langle\langle \text{letbox}(M, u. F') / \langle \Psi, x \rangle \rangle\rangle E = \text{letbox}(M, u. \langle\langle F' / \langle \Psi, x \rangle \rangle\rangle E)
\end{array}$$

where we assume Γ, y and Ψ, x are variable disjoint. The local reduction and expansion for contextual possibility can then be defined as:

$$\begin{array}{ll}
\text{letdia}(\text{dia}(F), \langle \Psi, x \rangle. E) & \Longrightarrow_R \langle\langle F / \langle \Psi, x \rangle \rangle\rangle E \\
\text{letbox}(\text{box}(\Psi. M), u. E) & \Longrightarrow_R \llbracket \Psi. M / u \rrbracket E \\
M : \langle \Psi \rangle A & \Longrightarrow_E \text{dia}(\text{letdia}(M, \langle \Psi, x \rangle. \langle \text{id}_\Psi, x \rangle))
\end{array}$$

We currently do not have any convincing applications of contextual possibility, although the relationship to monadic programming [Pfenning and Davies 2001], mobile computation [Moody 2003], and, more generally, interactions of programs with their environment [Nanevski 2004] provide some avenues for further investigation.

9.2 Dependent necessity

Another direction for exploration is the dependent typing for the modal operator \Box and its indexed variant. Currently, \Box is defined only for simple types. In the dependent case we have a somewhat simplified operator $\Pi u::A[\Psi].B$, which internalizes the dependence on the meta-variable u , but does not internalize the validity judgment itself. As we explained in Section 5, this preserves the existence and properties of canonical forms of the underlying logical framework. However, if we are not concerned with canonical forms, the dependently typed rules for \Box may be formulated as follows.

$$\frac{\Delta; \Psi \vdash M : A}{\Delta; \Gamma \vdash \text{box}(\Psi. M) : [\Psi]A} \Box I$$

$$\frac{\Delta; \Gamma, x:[\Psi]A \vdash C(x) : \text{type} \quad \Delta; \Gamma \vdash M : [\Psi]A \quad \Delta, u::A[\Psi]; \Gamma \vdash N : C(\text{box } u)}{\Delta; \Gamma \vdash \text{letbox}(M, u. N) : C(M)} \Box E$$

A complete dependent type theory with Π , Σ , equality and the full set of contextual modal operators is likely to be a suitable framework for reasoning about staged programs. For example, in such a type theory we could declare a staged, dependently typed version of the exponentiation function from Section 6 as follows,

$$\text{exp} : \Pi x:\text{int}. \Sigma e:[b:\text{int}]\text{int}. \Pi n:\text{int}. \text{letbox}(e, u. \text{clo}(u, n/b)) = n^x$$

Here, the scope of the Σ plays the role of a correctness proof for e . It should be possible to eliminate this correctness proof with standard program extraction techniques or proof irrelevance [Pfenning 2001] to obtain the function presented earlier of type $\text{int} \rightarrow [b:\text{int}]\text{int}$.

9.3 Internalizing explicit substitutions

A further extension that we intend to pursue concerns internalizing simultaneous substitutions and giving them a first class status. This would allow us to pass substitutions as function arguments, and to quantify over substitution variables. Such an extension may potentially have interesting applications, especially in the modularity of encoding of object theories in a logical framework and perhaps in capturing operations on contexts such as closure conversion for a functional language in a more immediate way than possible in LF. The proper setting for such an extension might be the ∇ -calculus [Schürmann et al. 2005].

10. CONCLUSION

In this paper we present an intuitionistic contextual modal logic and its type theory (in both simple and dependently typed setting), as a suitable foundation for formalizing contexts in theorem proving and functional and logic programming. From the logical standpoint, our system is a relativized version of an intuitionistic modal logic S4. From the computational standpoint, it is a logically justified calculus of explicit simultaneous substitutions, with very pleasing theoretical properties.

As we outlined in the section on related work, various contextual lambda calculi have been proposed before, and various modal logics of contexts have been proposed

before. However, we believe that we are the first to unify the two disparate views and application domains, and provide a type theoretic treatment for both.

We illustrate our logic and its type theory by applying it to two well-known problems in the areas of staged computation and logical frameworks.

In functional staged computation, one of the challenges has been to design a type system that can differentiate between code that is manipulated as data (object code), and code that performs the manipulation (meta code). Such a type system should satisfy two further requirements: (1) it should be possible to coerce object code with no free variables into meta code and execute it when required, and (2) it should be possible to compose object code with free variables into larger object code, while imposing capture of free variables. The two requirements are difficult to reconcile because the type system has to recognize when object code is closed in order to provide for (1), but also when object code is not closed and what its free variables are in order to provide for (2). Our type system provides for both by using the modal type constructor $[\Psi]A$ to classify object code with free variables Ψ ; when Ψ is empty, that object code can be converted into meta code and executed.

In logical frameworks, meta-variables are place-holders for as yet unknown terms, to be instantiated via unification or via proof search. Each meta-variable admits terms that may depend only on a particular subset of parameters, and a potential unifier that does not respect this restriction has to be rejected. Tracking the correspondence between the meta-variables and their respective sets of admissible parameters is obviously very important for the soundness of the logical framework and can also have significant consequences for the framework efficiency.

Our proposal for meta-variables centers on the dependently typed version of intuitionistic contextual modal logic, where meta-variables correspond to hypotheses of contextual validity. For example, the meta-variable $u::A \text{ valid}[\Psi]$ admits only terms of type A that depend on the parameters listed in the local context Ψ . The modal formulation explains and justifies the optimization strategies like grafting, lowering and raising, that have previously been employed in implementations of logical frameworks and theorem provers, but have only been justified algorithmically, rather than logically. We also believe that ICML and its type theory is the first proposal that uncovers the modal nature of explicit substitutions in logical frameworks.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library by visiting the following URL: <http://www.acm.org/pubs/citations/journals/tocl/2005-V-N/p1-URLend>.

ACKNOWLEDGMENTS

We would like to thank Rowan Davies and Kevin Watkins for discussions relating to the subject of this paper, and Yuki Yoshi Kameyama for comments on an earlier draft.

REFERENCES

- ABADI, M., CARDELLI, L., CURIEN, P.-L., AND LÈVY, J.-J. 1990. Explicit substitutions. In *Conference Record of the Seventeenth Annual ACM Symposium on Principles of Programming Languages*. ACM Press, New York, NY, 1-14.
- ACM Transactions on Computational Logic, Vol. V, No. N, September 2005.

- Languages, San Francisco, California.* ACM, 31–46.
- ALECHINA, N., MENDLER, M., DE PAIVA, V., AND RITTER, E. 2001. Categorical and Kripke semantics for Constructive S4 modal logic. In *International Workshop on Computer Science Logic, CSL'01*, L. Fribourg, Ed. Lecture Notes in Computer Science, vol. 2142. Springer, Paris, France, 292–307.
- ATTARDI, G. AND SIMI, M. 1994. Proofs in context. In *Principles of Knowledge Representation and Reasoning*, J. Doyle, E. Sandewall, and P. Torasso, Eds. 16–26.
- ATTARDI, G. AND SIMI, M. 1995. A formalization of viewpoints. *Fundamenta Informaticae* 23, 3, 149–173.
- BJØRNER, N. AND MUÑOZ, C. 2000. Absolute explicit unification. In *Proceedings of the 11th International Conference on Rewriting Techniques and Applications (RTA'00), Norwich, UK*. LNCS, vol. 1833. 31–46. <http://heory.stanford.edu/people/nikolaj/rta.ps>.
- BOGNAR, M. AND DE VRIJER, R. 2001. A calculus of lambda calculus context. *Journal of Automated Reasoning* 27, 1, 29–59.
- BUVAČ, S. 1996. Quantificational logic of context. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*. 600–606.
- BUVAČ, S., BUVAČ, V., AND MASON, I. 1995. Metamathematics of contexts. *Fundamenta Informaticae* 23, 3, 412–419.
- CALCAGNO, C., MOGGI, E., AND TAHA, W. 2004. ML-like inference for classifiers. In *Proceedings of the 13th European Symposium on Programming (ESOP'04)*, D. Schmidt, Ed. Springer-Verlag LNCS 2986, Barcelona, Spain, 79–93.
- CARTMELL, J. 1986. Generalized algebraic theories and contextual categories. *Annals of Pure and Applied Logic* 32, 209–243.
- CERVESATO, I. AND PFENNING, F. 2002. A linear logical framework. *Information & Computation* 179, 1 (Nov.), 19–75.
- DAMI, L. 1996. Functional programming with dynamic binding. In *Object Applications*, D. Tsichritzis, Ed. Technical Report, University of Geneva, 155–172.
- DAMI, L. 1998. A lambda-calculus for dynamic binding. *Theoretical Computer Science* 192, 2, 201–231.
- DAVIES, R. 1996. A temporal logic approach to binding-time analysis. In *Proceedings of the Eleventh Annual Symposium on Logic in Computer Science*, E. Clarke, Ed. IEEE Computer Society Press, New Brunswick, New Jersey, 184–195.
- DAVIES, R. AND PFENNING, F. 2001. A modal analysis of staged computation. *Journal of the ACM* 48, 3 (May), 555–604.
- DE GROOTE, P. 2002. On the strong normalisation of intuitionistic natural deduction with permutation-conversions. *Information and Computation* 178, 2 (Nov.), 441–464.
- DE PAIVA, V. 2003. Natural deduction and context as (constructive) modality. In *Modelling and Using Context*, P. Blackburn, C. Ghidini, R. M. Turner, and F. Giunchiglia, Eds. Lecture Notes in Artificial Intelligence, vol. 2680. Springer, 116–129.
- DOWEK, G., HARDIN, T., AND KIRCHNER, C. 1995. Higher-order unification via explicit substitutions. In *Proceedings of the Tenth Annual Symposium on Logic in Computer Science*, D. Kozen, Ed. IEEE Computer Society Press, San Diego, California, 366–374.
- DOWEK, G., HARDIN, T., KIRCHNER, C., AND PFENNING, F. 1996. Unification via explicit substitutions: The case of higher-order patterns. In *Proceedings of the Joint International Conference and Symposium on Logic Programming*, M. Maher, Ed. MIT Press, Bonn, Germany, 259–273.
- GEUVERS, H. AND JOJGOV, G. I. 2002. Open proofs and open terms: a basis for interactive logic. In *Proceedings of the 16th International on Computer Science Logic (CSL'02) Edinburgh, Scotland, UK, September 22-25*, J. C. Bradfield, Ed. Lecture Notes in Computer Science (LNCS 2471). Springer, 537–552.
- GIUNCHIGLIA, F. 1993. Contextual reasoning. *Epistemologica* 16, 345–364.
- GIUNCHIGLIA, F. AND SERAFINI, L. 1994. Multilanguage hierarchical logics, or: How to do without modal logics. *Artificial Intelligence* 65, 1, 29–70.
- ACM Transactions on Computational Logic, Vol. V, No. N, September 2005.

- HARPER, R., HONSELL, F., AND PLOTKIN, G. 1993. A framework for defining logics. *Journal of the Association for Computing Machinery* 40, 1 (January), 143–184.
- HASHIMOTO, M. AND OHORI, A. 2001. A typed context calculus. *Theoretical Computer Science* 266, 1–2, 249–272.
- HODAS, J. AND MILLER, D. 1994. Logic programming in a fragment of intuitionistic linear logic. *Information and Computation* 110, 2, 327–365.
- JOJGOV, G. I. 2003. Holes with binding power. In *Types for Proofs and Programs, Second International Workshop, TYPES 2002, Berg en Dal, The Netherlands, April 24–28, 2002, Selected Papers*, H. Geuvers and F. Wiedijk, Eds. Lecture Notes in Computer Science (LNCS 2646). Springer, 162–181.
- KRIPKE, S. 1959. A completeness theorem in modal logic. *Journal of Symbolic Logic* 24, 1–14.
- LEE, S.-D. AND FRIEDMAN, D. P. 1996. Enriching the lambda calculus with contexts: toward a theory of incremental program construction. In *International Conference on Functional Programming, ICFP'96*. 239–250.
- LIANG, C. AND NADATHUR, G. 2002. Tradeoffs in the intensional representation of lambda terms. In *Proceedings of the 13th International Conference on Rewriting Techniques and Applications*, S. Tison, Ed. Springer-Verlag LNCS 2378, Copenhagen, Denmark, 192–206.
- MAGNUSSON, L. 1995. The implementation of ALF – a proof editor based on martin-löf’s monomorphic type theory with explicit substitutions. Ph.D. thesis, Chalmers University of Technology and Göteborg University.
- 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.
- MASON, I. A. 1999. Computing with contexts. *Higher-Order and Symbolic Computation* 12, 2, 171–201.
- MCCARTHY, J. 1987. Generality in artificial intelligence. *Communications of the ACM* 30, 12, 1030–1035.
- MCCARTHY, J. 1993. Notes on formalizing contexts. In *International Joint Conference on Artificial Intelligence*. Chambery, France, 555–560.
- MILLER, D. 1991. A logic programming language with lambda-abstraction, function variables, and simple unification. *Journal of Logic and Computation* 1, 4, 497–536.
- MONTAGUE, R. 1963. Syntactical treatment of modalities, with corollaries on reflexion principles and finite axiomatizability. *Acta Philosophica Fennica* 16, 153–167.
- MOODY, J. 2003. Modal logic as a basis for distributed computation. Tech. Rep. CMU-CS-03-194, Carnegie Mellon University. Oct.
- MUÑOZ, C. 2001. Dependent types and explicit substitutions. *Mathematical Structures in Computer Science* 11, 1 (February), 91–129. It also appears as report NASA/CR-1999-209722 ICASE No. 99-43.
- NANEVSKI, A. 2003. From dynamic binding to state via modal possibility. In *International Conference on Principles and Practice of Declarative Programming, PPDP'03*. Uppsala, Sweden, 207–218.
- NANEVSKI, A. 2004. Functional programming with names and necessity. Ph.D. thesis, Computer Science Department, Carnegie Mellon University.
- NANEVSKI, A. AND PFENNING, F. 2002. Staged computation with names and necessity. *Journal of Functional Programming*. To appear.
- NANEVSKI, A., PIENKA, B., AND PFENNING, F. 2003. A modal foundation for meta variables. In *Proceedings of MERLIN 2003*. Uppsala, Sweden.
- PFENNING, F. 2001. Intensionality, extensionality, and proof irrelevance in modal type theory. In *Proceedings of the 16th Annual Symposium on Logic in Computer Science (LICS'01)*, J. Halpern, Ed. IEEE Computer Society Press, Boston, Massachusetts, 221–230.
- PFENNING, F. AND DAVIES, R. 2001. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science* 11, 4, 511–540.
- PFENNING, F. AND SCHÜRMAN, C. 1999. System description: Twelf - a meta-logical framework for deductive systems. In *International Conference on Automated Deduction, CADE'99*, ACM Transactions on Computational Logic, Vol. V, No. N, September 2005.

- H. Ganzinger, Ed. Lecture Notes in Artificial Intelligence, vol. 1632. Springer-Verlag, Trento, Italy, 202–206.
- PIENTKA, B. 2003. Tabled higher-order logic programming. Ph.D. thesis, Computer Science Department, Carnegie Mellon University.
- PIENTKA, B. AND PFENNING, F. 2003. Optimizing higher-order pattern unification. In *19th International Conference on Automated Deduction, Miami, USA*, F. Baader, Ed. Lecture Notes in Computer Science (LNAI 2741). Springer, 473–487.
- SATO, M., SAKURAI, T., AND BURSTALL, R. 2001. Explicit environments. *Fundamenta Informaticae* 45, 1-2, 79–115.
- SATO, M., SAKURAI, T., AND KAMEYAMA, Y. 2002. A simply typed context calculus with first-class environments. *Journal of Functional and Logic Programming* 2002, 4 (March).
- SATO, M., SAKURAI, T., KAMEYAMA, Y., AND IGARASHI, A. 2003. Calculi of meta-variables. In *Proceedings of the 17th International on Computer Science Logic (CSL'03) Vienna, Austria, August 25-30*, M. Baaz and J. A. Makowsky, Eds. Lecture Notes in Computer Science (LNCS 2803). Springer, 484–497.
- SCHÜRMAN, C., POSWOLSKY, A., AND SARNAT, J. 2005. The ∇ -calculus: Functional programming with higher-order encodings. In *Proceedings of the 7th International Conference on Typed Lambda Calculi and Applications (TLCA'05)*, P. Urzyczyn, Ed. Springer-Verlag LNCS 3461, Nara, Japan, 339–353.
- SERAFINI, L. AND BOUQUET, P. 2004. Comparing formal theories of context in AI. *Artificial Intelligence* 155, 1-2, 41–67.
- SERAFINI, L. AND GIUNCHIGLIA, F. 2002. ML system: A proof theory for contexts. *Journal of Logic, Language and Information* 11, 4, 471–518.
- SIMPSON, A. K. 1994. The proof theory and semantics of intuitionistic modal logic. Ph.D. thesis, University of Edinburgh.
- STRECKER, M. 1999. Construction and deduction in type theories. Ph.D. thesis, Universität Ulm.
- TAHA, W. AND NIELSEN, M. F. 2003. Environment classifiers. In *Conference Record of the 30th Annual Symposium on Principles of Programming Languages (POPL'03)*, G. Morrisett, Ed. ACM Press, New Orleans, Louisiana, 26–37.
- TAHA, W. AND SHEARD, T. 1997. Multi-stage programming with explicit annotations. In *Workshop on Partial Evaluation and Semantics-Based Program Manipulation, PEPM'97*. Amsterdam, 203–217.
- THOMASON, R. H. 1999. Type theoretic foundations for context, part 1: Contexts as complex type-theoretic objects. In *Proceedings of the Second International and Interdisciplinary Conference on Modeling and Using Contexts, CONTEXT'99*, P. Bouquet, L. Serafini, P. Brézillon, M. Benerecetti, and F. Castellani, Eds. Trento, Italy, 352–374.
- THOMASON, R. H. 2003. Dynamic contextual intensional logic: Logical foundations and an application. In *Modeling and Using Context*, P. Blackburn, C. Ghidini, R. M. Turner, and F. Giunchiglia, Eds. Lecture Notes in Artificial Intelligence, vol. 2680. Springer, 328–341.
- VII, T. M., CRARY, K., HARPER, R., AND PFENNING, F. 2004. A symmetric modal lambda calculus for distributed computing. In *Symposium on Logic in Computer Science, LICS'04*, H. Ganzinger, Ed. Turku, Finland, 286–295.
- WATKINS, K., CERVESATO, I., PFENNING, F., AND WALKER, D. 2002. A concurrent logical framework I: Judgments and properties. Tech. Rep. CMU-CS-02-101, Department of Computer Science, Carnegie Mellon University. Revised May 2003.
- WEYHRAUCH, R. W. 1979. Prolegomena to a theory of mechanized formal reasoning. *Artificial Intelligence* 13, 1-2, 133–170.
- WICKLINE, P., LEE, P., AND PFENNING, F. 1998. Run-time code generation and Modal-ML. In *Conference on Programming Language Design and Implementation, PLDI'98*, K. D. Cooper, Ed. ACM Press, Montreal, Canada, 224–235.
- 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.
- ACM Transactions on Computational Logic, Vol. V, No. N, September 2005.

Received *month year*; revised *month year*; accepted *month year*

THIS DOCUMENT IS THE ONLINE-ONLY APPENDIX TO:

Contextual Modal Type Theory

ALEKSANDAR NANEVSKI

Harvard University

and

FRANK PFENNING

Carnegie Mellon University

and

BRIGITTE PIENTKA

McGill University

ACM Transactions on Computational Logic, Vol. V, No. N, September 2005, Pages 1–47.

A. DEFINITION OF HEREDITARY SUBSTITUTIONS

A.1 Input and Output of Hereditary Substitutions

$$\begin{aligned}
 [M/x]_{\alpha}^n(N) &= N' \\
 [M/x]_{\alpha}^r(R) &= R' \quad \text{or} \quad M' : \alpha' \\
 [M/x]_{\alpha}^s(\tau) &= \tau' \\
 [M/x]_{\alpha}^a(B) &= B' \\
 [M/x]_{\alpha}^p(P) &= P' \\
 [M/x]_{\alpha}^k(K) &= K' \\
 [M/x]_{\alpha}^{\gamma}(\Gamma) &= \Gamma'
 \end{aligned}$$

$$\begin{aligned}
 [\hat{\Psi}.M/u]_{\alpha[\psi]}^n(N) &= N' \\
 [\hat{\Psi}.M/u]_{\alpha[\psi]}^r(R) &= R' \quad \text{or} \quad M' : \alpha' \\
 [\hat{\Psi}.M/u]_{\alpha[\psi]}^s(\tau) &= \tau' \\
 [\hat{\Psi}.M/u]_{\alpha[\psi]}^a(B) &= B' \\
 [\hat{\Psi}.M/u]_{\alpha[\psi]}^p(P) &= P' \\
 [\hat{\Psi}.M/u]_{\alpha[\psi]}^k(K) &= K' \\
 [\hat{\Psi}.M/u]_{\alpha[\psi]}^{\gamma}(\Gamma) &= \Gamma' \\
 [\hat{\Psi}.M/u]_{\alpha[\psi]}^{\delta}(\Delta) &= \Delta'
 \end{aligned}$$

$$\begin{aligned}
 [\sigma]_{\psi}^n(N) &= N' \\
 [\sigma]_{\psi}^r(R) &= R' \quad \text{or} \quad M' : \alpha' \\
 [\sigma]_{\psi}^s(\tau) &= \tau' \\
 [\sigma]_{\psi}^a(B) &= B' \\
 [\sigma]_{\psi}^p(P) &= P'
 \end{aligned}$$

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2005 ACM 1529-3785/2005/0700-0001 \$5.00

A.2 Hereditary Substitution for a Variable

$[M/x]_{\alpha}^r(x)$	$= M : \alpha$	
$[M/x]_{\alpha}^r(y)$	$= y$	if $y \neq x$
$[M/x]_{\alpha}^r(\text{app}(R, N))$	$= \text{app}(R', N')$	where $R' = [M/x]_{\alpha}^r(R)$ and $N' = [M/x]_{\alpha}^n(N)$
$[M/x]_{\alpha}^r(\text{app}(R, N))$	$= M_2 : \alpha_2$	if $[M/x]_{\alpha}^r(R) = \text{lam}(y. M') : \alpha_1 \rightarrow \alpha_2$, where $\alpha_1 \rightarrow \alpha_2 \leq \alpha$ and $N' = [M/x]_{\alpha}^n(N)$ and $M_2 = [N'/y]_{\alpha_1}^n(M')$
$[M/x]_{\alpha}^r(\text{clo}(v, \tau))$	$= \text{clo}(v, \tau')$	where $\tau' = [M/x]_{\alpha}^s(\tau)$
$[M/x]_{\alpha}^r(\text{mapp}(R, \hat{\Gamma}.N))$	$= \text{mapp}(R', \hat{\Gamma}.N)$	where $R' = [M/x]_{\alpha}^r(R)$
$[M/x]_{\alpha}^r(\text{mapp}(R, \hat{\Gamma}.N))$	$= M_2 : \alpha_2$	if $[M/x]_{\alpha}^r(R) = \text{mlam}(u. M') : \alpha_1[\gamma] \Rightarrow \alpha_2$ and $\alpha_1[\gamma] \Rightarrow \alpha_2 \leq \alpha$ and $\hat{\gamma} = \hat{\Gamma}$ and $M_2 =$ $[[\hat{\Gamma}.N/u]_{\alpha_1[\gamma]}^n(M')$
$[M/x]_{\alpha}^r(R)$	fails	otherwise
$[M/x]_{\alpha}^n(\text{lam}(y. N))$	$= \text{lam}(y. N')$	where $N' = [M/x]_{\alpha}^n(N)$ choosing $y \notin \text{FV}(M)$ and $y \neq x$
$[M/x]_{\alpha}^n(\text{mlam}(u. N))$	$= \text{mlam}(u. N')$	where $N' = [M/x]_{\alpha}^n(N)$ choosing $u \notin \text{FMV}(M)$
$[M/x]_{\alpha}^n(R)$	$= M'$	if $[M/x]_{\alpha}^r(R) = M' : \alpha'$
$[M/x]_{\alpha}^n(R)$	$= R'$	if $[M/x]_{\alpha}^r(R) = R'$
$[M/x]_{\alpha}^n(N)$	fails	otherwise
$[M/x]_{\alpha}^s(\cdot)$	$= \cdot$	
$[M/x]_{\alpha}^s(\tau, N/y)$	$= \tau', N'/y$	where $\tau' = [M/x]_{\alpha}^s(\tau)$ and $N' = [M/x]_{\alpha}^n(N)$
$[M/x]_{\alpha}^s(\tau, R//y)$	$= \tau', R'//y$	if $[M/x]_{\alpha}^r(R) = R'$ and $\tau' = [M/x]_{\alpha}^s(\tau)$
$[M/x]_{\alpha}^s(\tau, R//y)$	$= \tau', M'/y$	if $[M/x]_{\alpha}^r(R) = M' : \alpha'$ and $\tau' = [M/x]_{\alpha}^s(\tau)$
$[M/x]_{\alpha}^s(\tau)$	fails	otherwise
$[M/x]_{\alpha}^p(a)$	$= a$	
$[M/x]_{\alpha}^p(\text{app}(P, N))$	$= \text{app}(P', N')$	where $P' = [M/x]_{\alpha}^p(P)$ and $N' = [M/x]_{\alpha}^n(N)$
$[M/x]_{\alpha}^p(\text{mapp}(P, \hat{\Gamma}.N))$	$= \text{mapp}(P', \hat{\Gamma}.N)$	where $P' = [M/x]_{\alpha}^p(P)$
$[M/x]_{\alpha}^p(P)$	fails	otherwise
$[M/x]_{\alpha}^a(P)$	$= P'$	where $P' = [M/x]_{\alpha}^p(P)$
$[M/x]_{\alpha}^a(\Pi y:A. B)$	$= \Pi y:A'. B'$	if $A' = [M/x]_{\alpha}^a(A)$ and $B' = [M/x]_{\alpha}^a(B)$ choosing $y \notin \text{FV}(M)$ and $y \neq x$
$[M/x]_{\alpha}^a(\Pi u::A[\Psi]. B)$	$= \Pi u::A[\Psi]. B'$	if $B' = [M/x]_{\alpha}^a(B)$ choosing $u \notin \text{FMV}(M)$
$[M/x]_{\alpha}^a(A)$	fails	otherwise

$[M/x]_{\alpha}^k(\mathbf{type})$	$= \mathbf{type}$	
$[M/x]_{\alpha}^k(\Pi y:A. K)$	$= \Pi y:A'. K'$	if $A' = [M/x]_{\alpha}^a(A)$ and $K' = [M/x]_{\alpha}^k(K)$ choosing $y \notin \mathbf{FV}(M)$ and $y \neq x$
$[M/x]_{\alpha}^k(\Pi u::A[\Psi]. K)$	$= \Pi u::A[\Psi]. K'$	if $K' = [M/x]_{\alpha}^a(K)$ choosing $u \notin \mathbf{FMV}(M)$
$[M/x]_{\alpha}^k(K)$	fails	otherwise
$[M/x]_{\alpha}^{\gamma}(\cdot)$	$= \cdot$	
$[M/x]_{\alpha}^{\gamma}(\Gamma, y:A)$	$= \Gamma', y:A'$	if $\Gamma' = [M/x]_{\alpha}^{\gamma}(\Gamma)$ and $A' = [M/x]_{\alpha}^a(A)$
$[M/x]_{\alpha}^{\gamma}(\Gamma)$	fails	otherwise

A.3 Hereditary Substitution for a Meta-Variable

$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(x)$	$= x$	
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\mathbf{app}(R, N))$	$= \mathbf{app}(R', N')$	if $\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R) = R'$ with $N' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(N)$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\mathbf{app}(R, N))$	$= M_2 : \alpha_2$	if $\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R) = \mathbf{lam}(y. M')$: $\alpha_1 \rightarrow \alpha_2$ for $\alpha_1 \rightarrow \alpha_2 \leq \alpha[\psi]$ and $N' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(N)$ and $M_2 = [N'/y]_{\alpha_1}^n(M')$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\mathbf{clo}(u, \tau))$	$= M' : \alpha$	where $\tau' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^s(\tau)$ and $M' = [\tau']_{\psi}^n(M)$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\mathbf{clo}(v, \tau))$	$= \mathbf{clo}(v, \tau')$	if $v \neq u$ with $\tau' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^s(\tau)$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\mathbf{mapp}(R, \hat{\Gamma}.N))$	$= \mathbf{mapp}(R', \hat{\Gamma}.N')$	if $\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R) = R'$ and $N' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(N)$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(\mathbf{mapp}(R, \hat{\Gamma}.N))$	$= M_2 : \alpha_2$	if $\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R) = \mathbf{mlam}(v. M')$: $\alpha_1[\gamma] \Rightarrow \alpha_2$ where $\alpha_1[\gamma] \Rightarrow \alpha_2 \leq \alpha[\psi]$ and $\hat{\gamma} = \hat{\Gamma}$ and $N' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(N)$ and $M_2 = \llbracket \hat{\Gamma}.N'/v \rrbracket_{\alpha_1[\gamma]}^n(M')$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R)$	fails	otherwise
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(\mathbf{lam}(y. N))$	$= \mathbf{lam}(y. N')$	where $N' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(N)$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(\mathbf{mlam}(v. N))$	$= \mathbf{mlam}(v. N')$	where $N' = \llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(N)$ choosing $v \notin \mathbf{FMV}(M)$ and $v \neq u$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(R)$	$= R'$	if $\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R) = R'$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(R)$	$= M'$	if $\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^r(R) = M' : \alpha'$
$\llbracket \hat{\Psi}.M/u \rrbracket_{\alpha[\psi]}^n(N)$	fails	otherwise

$[\hat{\Psi}.M/u]_{\alpha[\psi]}^s(\cdot)$	$= \cdot$	
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^s(\tau, N/y)$	$= \tau', N'/y$	where $\tau' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^s(\tau)$ and $N' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^n(N)$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^s(\tau, R//y)$	$= \tau', R'//y$	if $[\hat{\Psi}.M/u]_{\alpha[\psi]}^r(R) = R'$ with $\tau' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^s(\tau)$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^s(\tau, R//y)$	$= \tau', M'/y$	if $[\hat{\Psi}.M/u]_{\alpha[\psi]}^r(R) = M' : \alpha'$ with $\tau' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^s(\tau)$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^s(\tau)$	fails	otherwise
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^p(a)$	$= a$	
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^p(\mathbf{app}(P, N))$	$= \mathbf{app}(P', N')$	if $P' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^p(P)$ and $N' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^n(N)$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^p(\mathbf{mapp}(P, \hat{\Gamma}.N))$	$= \mathbf{mapp}(P', \hat{\Gamma}.N')$	if $P' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^p(P)$ and $N' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^n(N)$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^p(P)$	fails	otherwise
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^a(P)$	$= P'$	where $P' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^p(P)$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^a(\Pi x:A. B)$	$= \Pi x:A'. B'$	where $A' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^a(A)$ and $B' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^a(B)$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^a(\Pi v::A[\Gamma]. B)$	$= \Pi v::A'[\Gamma']. B'$	where $A' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^a(A)$ and $\Gamma' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^\gamma(\Gamma)$ and $B' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^a(B)$ choosing $v \notin \text{FMV}(M)$ and $v \neq u$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^a(A)$	fails	otherwise
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^k(\mathbf{type})$	$= \mathbf{type}$	
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^k(\Pi x:A. K)$	$= \Pi x:A'. K'$	where $A' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^a(A)$ and $K' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^k(K)$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^k(\Pi v::A[\Gamma]. K)$	$= \Pi v::A'[\Gamma']. K'$	where $A' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^a(A)$ and $\Gamma' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^\gamma(\Gamma)$ and $K' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^k(K)$ choosing $v \notin \text{FMV}(M)$ and $v \neq u$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^k(K)$	fails	otherwise
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^\gamma(\cdot)$	$= \cdot$	
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^\gamma(\Gamma, x:A)$	$= \Gamma', x:A'$	where $\Gamma' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^\gamma(\Gamma)$ and $A' = [\hat{\Psi}.M/u]_{\alpha[\psi]}^a(A)$
$[\hat{\Psi}.M/u]_{\alpha[\psi]}^\gamma(\Gamma)$	fails	otherwise

$$\begin{aligned}
[[\hat{\Psi}.M/u]_{\alpha[\psi]}^{\delta}(\cdot)] &= \cdot \\
[[\hat{\Psi}.M/u]_{\alpha[\psi]}^{\delta}(\Delta, v::A[\Gamma])] &= \Delta', v::A'[\Gamma'] \text{ where } \Delta' = [[\hat{\Psi}.M/y]_{\alpha[\psi]}^{\delta}(\Delta)] \\
&\text{and } A' = [[\hat{\Psi}.M/y]_{\alpha[\psi]}^{\alpha}(A)] \\
&\text{and } \Gamma' = [[\hat{\Psi}.M/y]_{\alpha[\psi]}^{\gamma}(\Gamma)] \\
&\text{choosing } v \neq u \\
[[\hat{\Psi}.M/u]_{\alpha[\psi]}^{\delta}(\Delta)] &\text{ fails} \quad \text{otherwise}
\end{aligned}$$

A.4 Hereditary Simultaneous Substitution

$$\begin{aligned}
[\sigma]_{\psi}^r(x) &= M : \alpha && \text{if } M/x \in \sigma/\psi \text{ and } x:\alpha \in \psi \\
[\sigma]_{\psi}^r(x) &= R && \text{if } R//x \in \sigma/\psi \\
[\sigma]_{\psi}^r(\text{app}(R, N)) &= \text{app}(R', N') && \text{where } [\sigma]_{\psi}^r(R) = R' \text{ and } [\sigma]_{\psi}^n(N) = N' \\
[\sigma]_{\psi}^r(\text{app}(R, N)) &= M_2 : \alpha_2 && \text{if } [\sigma]_{\psi}^r(R) = \text{lam}(y.M') : \alpha_1 \rightarrow \alpha_2 \text{ for } \alpha_1 \rightarrow \\
&&& \alpha_2 \leq \psi \text{ with } N' = [\sigma]_{\psi}^n(N) \\
&&& \text{and } M_2 = [N'/y]_{\alpha_1}^n(M') \\
[\sigma]_{\psi}^r(\text{clo}(v, \tau)) &= \text{clo}(v, \tau') && \text{where } \tau' = [\sigma]_{\psi}^s(\tau) \\
[\sigma]_{\psi}^r(\text{mapp}(R, \hat{\Gamma}.N)) &= \text{mapp}(R', \hat{\Gamma}.N) && \text{if } [\sigma]_{\psi}^r(R) = R' \\
[\sigma]_{\psi}^r(\text{mapp}(R, \hat{\Gamma}.N)) &= M_2 : \alpha_2 && \text{if } [\sigma]_{\psi}^r(R) = \text{mlam}(u.M) : \alpha_1[\gamma] \Rightarrow \alpha_2 \text{ for } \\
&&& \alpha_1[\gamma] \Rightarrow \alpha_2 \leq \psi \text{ and } \hat{\gamma} = \hat{\Gamma} \text{ and } M_2 = \\
&&& [[\hat{\Gamma}.N/u]_{\alpha_1[\gamma]}^n(M)] \\
[\sigma]_{\psi}^r(R) &\text{ fails} && \text{otherwise} \\
[\sigma]_{\psi}^n(\text{lam}(y.N)) &= \text{lam}(y.N') && \text{where } N' = [\sigma, y//y]_{\psi, y, \cdot}^n(N) \\
&&& \text{choosing } y \notin \text{FV}(\sigma), \text{dom}(\sigma) \\
[\sigma]_{\psi}^n(\text{mlam}(u.N)) &= \text{mlam}(u.N') && \text{where } N' = [\sigma]_{\psi}^n(N), \text{choosing } u \notin \text{FMV}(\sigma) \\
[\sigma]_{\psi}^n(R) &= M' && \text{if } [\sigma]_{\psi}^n(R) = M' : \alpha' \\
[\sigma]_{\psi}^n(R) &= R' && \text{if } [\sigma]_{\psi}^n(R) = R' \\
[\sigma]_{\psi}^n(N) &\text{ fails} && \text{otherwise} \\
[\sigma]_{\psi}^s(\cdot) &= \cdot \\
[\sigma]_{\psi}^s(\tau, N/y) &= \tau', N'/y && \text{where } \tau' = [\sigma]_{\psi}^s(\tau) \text{ and } N' = [\sigma]_{\psi}^n(N) \\
[\sigma]_{\psi}^s(\tau, R//y) &= \tau', R//y && \text{if } [\sigma]_{\psi}^n(R) = R' \text{ with } \tau' = [\sigma]_{\psi}^s(\tau) \\
[\sigma]_{\psi}^s(\tau, R//y) &= \tau', M'/y && \text{if } [\sigma]_{\psi}^n(R) = M' : \alpha' \text{ with } \tau' = [\sigma]_{\psi}^s(\tau) \\
[\sigma]_{\psi}^s(\tau) &\text{ fails} && \text{otherwise} \\
[\sigma]_{\psi}^p(a) &= a \\
[\sigma]_{\psi}^p(\text{app}(P, N)) &= \text{app}(P', N') && \text{where } P' = [\sigma]_{\psi}^p(P) \text{ and } N' = [\sigma]_{\psi}^n(N) \\
[\sigma]_{\psi}^p(\text{mapp}(P, \hat{\Gamma}.N)) &= \text{mapp}(P', \hat{\Gamma}.N) && \text{where } P' = [\sigma]_{\psi}^p(P) \\
[\sigma]_{\psi}^p(P) &\text{ fails} && \text{otherwise}
\end{aligned}$$

$$\begin{array}{lll}
[\sigma]_{\psi}^a(P) & = P' & \text{where } P' = [\sigma]_{\psi}^p(P) \\
[\sigma]_{\psi}^a(\Pi x:A. B) & = \Pi x:A'. B' & \text{if } A' = [\sigma]_{\psi}^a(A) \text{ and } B' = [\sigma, x//x]_{\psi, x:\alpha}^a(B) \\
& & \text{choosing } x \notin \text{FV}(\sigma), \text{dom}(\sigma) \\
[\sigma]_{\psi}^a(\Pi u::A[\Gamma]. B) & = \Pi u::A[\Gamma]. B' & \text{if } B' = [\sigma]_{\psi}^a(B) \\
& & \text{choosing } u \notin \text{FMV}(\sigma) \\
[\sigma]_{\psi}^a(A) & \text{fails} & \text{otherwise}
\end{array}$$

B. SUBSTITUTION PRINCIPLES

We begin by defining the *head* of an atomic term as the variable or meta-variable at its beginning.

$$\begin{aligned}
\text{head}(x) &= x \\
\text{head}(\text{clo}(u, \tau)) &= u \\
\text{head}(\text{app}(R, M)) &= \text{head}(R) \\
\text{head}(\text{mapp}(R, \Psi.M)) &= \text{head}(R)
\end{aligned}$$

Whether the result of a hereditary substitution into an atomic term R is an atomic term R' or a normal term $M' : \alpha'$ depends solely on the head variable of R . This property is needed in several subsequent proofs, so we state it formally.

LEMMA B.1 HEREDITARY SUBSTITUTIONS AND HEADS.

- (1) If $[M/x]_{\alpha}^r(R)$ exists, then
 - (a) $[M/x]_{\alpha}^r(R) = R'$ is atomic iff $\text{head}(R) \neq x$.
 - (b) $[M/x]_{\alpha}^r(R) = M' : \alpha'$ is normal iff $\text{head}(R) = x$
- (2) If $[\hat{\Psi}.M/u]_{\alpha[\psi]}^r(R)$ exists, then
 - (a) $[\hat{\Psi}.M/u]_{\alpha[\psi]}^r(R) = R'$ is atomic iff $\text{head}(R) \neq u$
 - (b) $[\hat{\Psi}.M/u]_{\alpha[\psi]}^r(R) = M' : \alpha'$ is normal iff $\text{head}(R) = u$
- (3) If $[\sigma]_{\psi}^r(R)$ exists, then
 - (a) $[\sigma]_{\psi}^r(R) = R'$ is atomic iff $\text{head}(R) = x$ and $R''//x \in \sigma/\psi$.
 - (b) $[\sigma]_{\psi}^r(R) = M' : \alpha'$ is normal iff $\text{head}(R) = x$ and $M''/x \in \sigma/\psi$.

PROOF. By straightforward induction on the structure of R . \square

Next we need to show that hereditary substitutions are defined and do not change the term if a variable does not appear in the term we substitute into.

LEMMA B.2 TRIVIAL HEREDITARY SUBSTITUTIONS. *Let T range over expressions of any syntactic category (i.e., kinds, atomic types, normal types, atomic terms, normal terms, substitutions, contexts and modal contexts), and let $*$ \in $\{k, p, a, r, n, s, \gamma, \delta\}$, correspondingly.*

- (1) If $x \notin \text{FV}(T)$, then $[M/x]_{\alpha}^*(T) = T$.
- (2) If $u \notin \text{FMV}(T)$, then $[\hat{\Psi}.M/u]_{\alpha[\psi]}^*(T) = T$.
- (3) If $x \notin \text{FV}(T)$, then $[\sigma, R//x]_{\psi, x:\alpha}^*(T) = [\sigma, M/x]_{\psi, x:\alpha}^*(T) = [\sigma]_{\psi}^*(T)$.

PROOF. By straightforward induction on the structure of T . \square

When investigating properties of ordinary substitution, we usually have to verify that $[M/x]([N/y]O) = ([M/x]N/y)([M/x]O)$, assuming that y is not in the free variables of M . Similar principles apply to hereditary substitutions, although the fact that hereditary substitutions are partial operations on ill-typed terms makes their formulation more tedious.

LEMMA B.3 COMPOSITION OF HEREDITARY SUBSTITUTIONS. *Suppose that T ranges over expressions of any syntactic category (i.e., kinds, atomic types, normal types, atomic terms, normal terms, substitutions, contexts and modal contexts), and let $*$ \in $\{k, p, a, r, n, s, \gamma, \delta\}$, correspondingly.*

- (1) *If $y \notin \text{FV}(M_0)$, and $[M_0/x]_\alpha^*(T) = T_0$, $[M_1/y]_\beta^*(T) = T_1$ and $[M_0/x]_\alpha^n(M_1)$ exist, then $[M_0/x]_\alpha^*(T_1) = [[M_0/x]_\alpha^n(M_1)/y]_\beta^*(T_0)$.*
- (2) *If $v \notin \text{FMV}(M_0)$, and $[M_0/x]_\alpha^*(T) = T_0$, $[\hat{\Psi}_1.M_1/v]_{\beta[\psi_1]}^*(T) = T_1$, then $[M_0/x]_\alpha^*(T_1) = [[\hat{\Psi}_1.M_1/v]_{\beta[\psi_1]}^*(T_0)]$.*
- (3) *If $[\sigma_1]_{\psi_1}^*(T) = T_1$ and $[M_0/x]_\alpha^s(\sigma)$ exists, then $[M_0/x]_\alpha^*(T_1) = [[M_0/x]_\alpha^s(\sigma_1)]_{\psi_1}^*(T)$.*
- (4) *If $[\hat{\Psi}_0.M_0/u]_{\alpha[\psi_0]}^*(T) = T_0$ and $[M_1/y]_\beta^*(T) = T_1$ and $[\hat{\Psi}_0.M_0/u]_{\alpha[\psi_0]}^n(M_1) = M'_1$ exist, then $[\hat{\Psi}_0.M_0/u]_{\alpha[\psi_0]}^*(T_1) = [M'_1/y]_\beta^*(T_0)$.*
- (5) *If $v \notin \text{FMV}(M_0)$, and $[\hat{\Psi}_0.M_0/u]_{\alpha[\psi_0]}^*(T) = T_0$ and $[\hat{\Psi}_1.M_1/v]_{\beta[\psi_1]}^*(T) = T_1$ and $[\hat{\Psi}_0.M_0/u]_{\alpha[\psi_0]}^n(M_1) = M'_1$ exist, then $[\hat{\Psi}_0.M_0/u]_{\alpha[\psi_0]}^*(T_1) = [[\hat{\Psi}_1.M'_1/v]_{\beta[\psi_1]}^*(T_0)]$.*
- (6) *If $[\hat{\Psi}_0.M_0/u]_{\alpha[\psi_0]}^*(T) = T_0$ and $[\sigma_1]_{\psi_1}^*(T) = T_1$ and $[\hat{\Psi}_0.M_0/u]_{\alpha[\psi_0]}^s(\sigma_1) = \sigma'_1$ exist, then $[\hat{\Psi}_0.M_0/u]_{\alpha[\psi_0]}^*(T_1) = [\sigma'_1]_{\psi_1}^*(T_0)$.*
- (7) *If $y \notin \text{dom}(\sigma_0)$, $\text{FV}(\sigma_0)$, and $[\sigma_0, y//y]_{\psi_0, y, _}^*(T) = T_0$ and $[M_1/y]_\beta^*(T) = T_1$, and $[\sigma_0]_{\psi_0}^n(M)$ exists, then $[\sigma_0]_{\psi_0}^*(T_1) = [[\sigma_0]_{\psi_0}^n(M_1)/y]_\beta^*(T_0)$.*
- (8) *If $v \notin \text{FMV}(\sigma_0)$, and $[\sigma_0]_{\psi_0}^*(T) = T_0$ and $[\hat{\Psi}_1.M_1/v]_{\beta[\psi_1]}^*(T) = T_1$, then $[\sigma_0]_{\psi_0}^*(T_1) = [[\hat{\Psi}_1.M_1/v]_{\beta[\psi_1]}^*(T_0)]$.*
- (9) *If $[\tau_1]_{\psi_1}^*(T)$ and $[\tau_0]_{\psi_0}^s(\tau_1)$ exist, then $[\tau_0]_{\psi_0}^*([\tau_1]_{\psi_1}^*(T)) = [[\tau_0]_{\psi_0}^s(\tau_1)]_{\psi_1}^*(T)$.*

PROOF. By nested induction, first on approximate types α , β and contexts ψ_0 , ψ_1 , and then on the structure of the expression being substituted into. \square

Before we can state the hereditary substitution principles we have to overcome one further technical obstacle. When applying a simultaneous hereditary substitution to an abstraction, we do not have a type or approximate type for the abstracted variable available. For example,

$$[\sigma]_\psi^n(\text{lam}(y. N)) = \text{lam}(y. N') \quad \text{where } N' = [\sigma, y//y]_{\psi, y, _}^n(N)$$

where we choose $y \notin \text{FV}(\sigma)$, $\text{dom}(\sigma)$. That means that we can not maintain the invariant that $\psi = \Psi^-$ whenever we apply $[\sigma]_\psi$ for $\sigma \Leftarrow \Psi$. We say that ψ *underapproximates* Ψ with respect to σ if ψ differs from Ψ^- only in that some declarations $x:\alpha$ may be replaced by $x:_$ if $R//x$ is in σ . Fortunately, any underapproximation is still sufficient to guarantee termination and the desired hereditary substitution properties.

Now we are prepared to state the hereditary substitution principles. We assume that all contexts in given judgments are well-formed. We also exploit our prior convention that all dependencies in an index to a hereditary substitution are erased before the substitution is applied, writing, for example, $[M/x]_A^r(R)$ instead of $[M/x]_{A^-}^r(R)$.

THEOREM B.4 HEREDITARY SUBSTITUTION PRINCIPLES.

Assume all contexts in the judgments given below are well-formed.

- (1) *If $\Delta; \Gamma \vdash M \Leftarrow A$ and $\Delta; \Gamma, x:A, \Gamma_1 \vdash R \Rightarrow C$ and the context $\Gamma'_1 = [M/x]_A^\gamma(\Gamma_1)$ exists and is well-formed (i.e., $\Delta \vdash \Gamma, \Gamma'_1 \text{ ctx}$), then $[M/x]_A^r(R)$ exists and the type $C' = [M/x]_A^\alpha(C)$ exists and is well-formed (i.e., $\Delta; \Gamma, \Gamma'_1 \vdash C' \Leftarrow \text{type}$) and*
 - (a) *if $[M/x]_A^r(R) = R'$ is atomic, then $\Delta; \Gamma, \Gamma'_1 \vdash R' \Rightarrow C'$.*
 - (b) *if $[M/x]_A^r(R) = M' : \alpha'$ is normal, then $\Delta; \Gamma, \Gamma'_1 \vdash M' \Leftarrow C'$ and $\alpha' = C^-$.*
- (2) *If $\Delta; \Gamma \vdash M \Leftarrow A$ and $\Delta; \Gamma, x:A, \Gamma_1 \vdash N \Leftarrow C$ and the context $\Gamma'_1 = [M/x]_A^\gamma(\Gamma_1)$ and type $C' = [M/x]_A^\alpha(C)$ exist and are well-formed (i.e., $\Delta \vdash \Gamma, \Gamma'_1 \text{ ctx}$ and $\Delta; \Gamma, \Gamma'_1 \vdash C' \Leftarrow \text{type}$), then $\Delta; \Gamma, \Gamma'_1 \vdash [M/x]_A^n(N) \Leftarrow C'$.*
- (3) *If $\Delta; \Psi \vdash M \Leftarrow A$ and $\Delta, u::A[\Psi], \Delta_1; \Gamma \vdash R \Rightarrow C$, and the modal context $\Delta'_1 = [\hat{\Psi}.M/u]_{A[\Psi]}^\delta(\Delta_1)$ and the context $\Gamma' = [[\hat{\Psi}.M/u]_{A[\Psi]}^\gamma(\Gamma)]$ exist and are well-formed (i.e., $\vdash \Delta, \Delta'_1 \text{ mctx}$, and $\Delta, \Delta'_1 \vdash \Gamma' \text{ ctx}$), then $[[\hat{\Psi}.M/u]_{A[\Psi]}^r(R)]$ exists and the type $C' = [[\hat{\Psi}.M/u]_{A[\Psi]}^\alpha(C)]$ exists and is well-formed (i.e. $\Delta, \Delta'; \Gamma' \vdash C' \Leftarrow \text{type}$) and*
 - (a) *if $[[\hat{\Psi}.M/u]_{A[\Psi]}^r(R)] = R'$ is atomic, then $\Delta, \Delta'_1; \Gamma' \vdash R' \Rightarrow C'$*
 - (b) *if $[[\hat{\Psi}.M/u]_{A[\Psi]}^r(R)] = M' : \alpha'$ is normal, then $\Delta, \Delta'_1; \Gamma' \vdash M' \Leftarrow C'$, and $\alpha' = C^-$.*
- (4) *If $\Delta; \Psi \vdash M \Leftarrow A$ and $\Delta, u::A[\Psi], \Delta_1; \Gamma \vdash N \Leftarrow C$, and the modal context $\Delta'_1 = [\hat{\Psi}.M/u]_{A[\Psi]}^\delta(\Delta_1)$ and the context $\Gamma' = [[\hat{\Psi}.M/u]_{A[\Psi]}^\gamma(\Gamma)]$ and the type $C' = [[\hat{\Psi}.M/u]_{A[\Psi]}^\alpha(C)]$ exist and are well-formed (i.e., $\vdash \Delta, \Delta'_1 \text{ mctx}$, and $\Delta, \Delta'_1 \vdash \Gamma' \text{ ctx}$ and $\Delta, \Delta'_1; \Gamma' \vdash C' \Leftarrow \text{type}$), then $\Delta, \Delta'_1; \Gamma' \vdash [[\hat{\Psi}.M/u]_{A[\Psi]}^n(N)] \Leftarrow C'$.*
- (5) *If $\Delta; \Gamma \vdash \sigma \Leftarrow \Psi$ and $\Delta; \Psi \vdash R \Rightarrow C$ and ψ is any underapproximation of Ψ with respect to σ , then $[\sigma]_\psi^r(R)$ exists and the type $C' = [\sigma]_\psi^\alpha(C)$ exists and is well-formed (i.e., $\Delta; \Gamma \vdash C' \Leftarrow \text{type}$) and*
 - (a) *if $[\sigma]_\psi^r(R) = R'$ is atomic, then $\Delta; \Gamma \vdash R' \Rightarrow C'$*
 - (b) *if $[\sigma]_\psi^r(R) = M' : \alpha'$ is normal, then $\Delta; \Gamma \vdash M' \Leftarrow C'$, and $\alpha' = C^-$.*
- (6) *If $\Delta; \Gamma \vdash \sigma \Leftarrow \Psi$ and $\Delta; \Psi \vdash N \Leftarrow C$ and ψ is any underapproximation of Ψ with respect to σ such that the type $C' = [\sigma]_\psi^\alpha(C)$ exists and is well-formed (i.e., $\Delta; \Gamma \vdash C' \Leftarrow \text{type}$), then $\Delta; \Gamma \vdash [\sigma]_\psi^n(N) \Leftarrow C'$.*

PROOF. We generalize this property and then proceed by nested induction, first on the structure of the approximation A^- and underapproximation ψ to the index type A and index context Ψ , and then on the structure of the second derivation in each of the cases, using the lemmas on trivial hereditary substitutions and composition of hereditary substitutions (Lemmas B.2 and B.3). We also use that the erasure of a type is invariant under substitution (for example, $([M/x]_A^\alpha(B))^- = B^-$) without explicitly stating this as lemma.

The generalization involves stating the similar principles for hereditary substitutions into other syntactic categories, that is, kinds, atomic types, normal types, contexts and modal contexts. We omit these cases here because hereditary substitutions into the mentioned syntactic categories are mostly compositional, and their substitution principles do not contribute any new insights. \square