

A SYNTACTIC APPROACH TO TYPE SOUNDNESS

Andrew K. Wright*

Matthias Felleisen*

Department of Computer Science
Rice University
Houston, TX 77251-1892

June 18, 1992

Rice Technical Report TR91-160

To appear in: *Information and Computation*

Abstract

We present a new approach to proving type soundness for Hindley/Milner-style polymorphic type systems. The keys to our approach are (1) an adaptation of subject reduction theorems from combinatory logic to programming languages, and (2) the use of rewriting techniques for the specification of the language semantics. The approach easily extends from polymorphic functional languages to imperative languages that provide references, exceptions, continuations, and similar features. We illustrate the technique with a type soundness theorem for the core of STANDARD ML, which includes the first type soundness proof for polymorphic exceptions and continuations.

1 Type Soundness

Static type systems for programming languages attempt to prevent the occurrence of *type errors* during execution. A definition of type error depends on a specific language and type system, but always includes the use of a function on arguments for which it is not defined, and the attempted application of a non-function. A static type system is *sound* if well-typed programs cannot cause type errors; a programming language with a sound static type system is *strongly typed* [7].

While it is easy to design a sound type system for an explicitly typed monomorphic language, the formulation of a sound type system for a language based on Hindley/Milner-style polymorphism and type inference is delicate [6, 8, 15, 17, 20, 37]. Although the treatment of purely functional languages is relatively well understood in this framework, the incorporation of imperative features such as references and exceptions requires extreme care, and the soundness of such systems is not obvious. A formal proof of soundness is

*This research was supported in part by the United States Department of Defense under a National Defense Science and Engineering Graduate Fellowship, by NSF grant CCR 89-17022, and by Texas Grant ARP/ATP-91 #003604-014.

required to lend credibility to the claim that such a system prevents type errors, and may be a crucial tool in the design of the type system.

We view a static type system for an implicitly typed language as a filter that selects well-typed programs from a larger universe of untyped programs. A partial function $eval : Programs \rightarrow Answers \cup \{WRONG\}$ defines the semantics of untyped programs. The result `WRONG` is returned for programs whose evaluation causes a type error; $eval$ is undefined for programs whose evaluation does not terminate. Let $\triangleright e : \tau$ mean that the type system assigns program e the type τ , *i.e.*, e is well-typed. The simplest soundness property states that well-typed programs do not yield `WRONG` [20].

Definition (Weak Soundness) *If $\triangleright e : \tau$ then $eval(e) \neq \text{WRONG}$.*

While weak soundness establishes that a static type system achieves its primary goal of preventing type errors, it is often possible to demonstrate a stronger property that relates the answer produced to the type of the program. If we view each type τ as denoting different subsets V^τ of the set of all answers V , then strong soundness states that an answer v produced by a terminating program of type τ is an element of the subset V^τ .

Definition (Strong Soundness) *If $\triangleright e : \tau$ and $eval(e) = v$ then $v \in V^\tau$.*

Strong soundness permits an implementation of the language to associate the representation of a value with its type, and thereby omit the representation tags required by dynamically typed languages. Weak soundness follows from strong soundness since `WRONG` is not a member of V^τ for any type τ .

Significant effort has been invested in proving type soundness for Hindley/Milner-style type systems, and their practical realization in the programming language STANDARD ML [22, 23]. Soundness proofs exist for the functional fragment [20, 36], for extensions including references [6, 17, 35, 37], and for a monomorphic language including first-class continuations [8]. However, there are several drawbacks to the existing proofs. Proofs of type soundness are sensitive to the precise formulation of the semantics of the language—different techniques are required for denotational versus operational formulations of the semantics, and even for different languages within the same semantic framework. The proofs for two different languages are difficult to reconcile to prove soundness for a language including features of both. The existing techniques are complicated, and the resulting proofs are lengthy and error-prone. For an illustration of the difficulties with such proofs, we refer the reader to Tofte’s [37] discussion of Damas’s [6] faulty proof of a type soundness theorem for a polymorphic language with references.

We present a simple approach to the proof of soundness for Hindley/Milner-style polymorphic type systems. Our approach is based on *subject reduction*, a classical result from combinatory logic [4], and on *rewriting* as a means to specifying operational semantics [11, 12, 13, 14]. To demonstrate the approach, we develop a proof of soundness for the core¹ of STANDARD ML, which extends a functional polymorphic language with references and exceptions. We also show soundness for an extension to STANDARD ML with first-class continuations. In principle, our approach is uniformly applicable to any language;

¹Mitchell and Harper refer to the functional polymorphic sublanguage as the *essence* of ML [24]. However, the difficulties of typing references and exceptions [6, 17, 35, 37] and the fact that they cannot be *expressed* by facilities of the functional core [10] indicate that they are equally important.

in practice, the resulting proofs are lengthy but simple, requiring only ordinary inductive techniques.

In the next section, we describe a prototypical functional language with a polymorphic type system, and discuss the various approaches that have previously been used to prove type soundness. Section 3 presents the essence of our approach to proving type soundness; the remaining sections develop illustrative proofs of type soundness for a specific functional language and various extensions.

2 Previous Approaches to Proving Type Soundness

Any proof of type soundness is intimately tied to the formulation of the semantics of the language. The earlier proofs relied on denotational semantics; later proofs used structural operational semantics. The following two subsections briefly present these approaches from a historical perspective. The third subsection discusses the problems with them. We begin with a brief introduction to the formulation of Hindley/Milner-style type systems, in order to introduce our notation.

Our prototypical functional language has the following (abstract) syntax:

$$e ::= c \mid x \mid \lambda x.e \mid e_1 e_2 \mid \text{let } x \text{ be } e_1 \text{ in } e_2$$

where $x \in \text{Var}$ and $c \in \text{Const}$. Constants (*Const*) include both data (integer, real, boolean, etc.) and operations ($+$, \div , \wedge , \vee , etc.). Variables (*Var*) are lexically scoped. Juxtaposition denotes application and is left associative; λ -expressions construct call-by-value procedural abstractions. Semantically the **let**-expression behaves like $((\lambda x.e_2) e_1)$, binding the value of e_1 to x in e_2 ; however, the type system allows x to be *polymorphic*: different occurrences of x in e_2 may be assigned different types.

A polymorphic type system for this language has types of the form:

$$\tau ::= \iota_1 \mid \dots \mid \iota_n \mid \tau_1 \rightarrow \tau_2$$

where ι_1, \dots, ι_n are ground types like *int* and *bool*. The type system is formulated as a deductive proof system that assigns types to expressions. The proof system produces conclusions, or *type judgments*, of the form $\Gamma \triangleright e : \tau$, meaning that expression e has type τ in type environment Γ . Type environments are finite maps from variables to types, and are used to give types to open expressions. The typing rules are:

$$\begin{array}{ll}
 \text{(var)} & \Gamma \triangleright x : \tau \text{ if } \Gamma(x) = \tau \\
 \text{(const)} & \Gamma \triangleright c : \tau \text{ if } \text{TypeOf}(c) = \tau \\
 \text{(abs)} & \frac{\Gamma[x \mapsto \tau_1] \triangleright e : \tau_2}{\Gamma \triangleright \lambda x.e : \tau_1 \rightarrow \tau_2} \\
 \text{(app)} & \frac{\Gamma \triangleright e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \triangleright e_2 : \tau_1}{\Gamma \triangleright e_1 e_2 : \tau_2} \\
 \text{(let)} & \frac{\Gamma \triangleright e_1 : \tau_1 \quad \Gamma \triangleright e_2[x \mapsto e_1] : \tau_2}{\Gamma \triangleright \text{let } x \text{ be } e_1 \text{ in } e_2 : \tau_2}
 \end{array}$$

The function *TypeOf* assigns types to constants. $\Gamma[x \mapsto \tau]$ denotes the functional extension or update of map Γ at x to τ ; $e[x \mapsto e']$ means the capture-avoiding substitution of e' for x in e . *Programs* are closed expressions. A program e is *well-typed* if it has a type τ in the empty type environment; we then write $\triangleright e : \tau$.

2.1 Proofs based on Denotational Semantics

Milner [20]: Milner formulated and proved a type soundness theorem for a functional language like the above based on a denotational semantics for the language of untyped expressions. The semantic domain is the solution of the following reflexive domain equation [32]:

$$\begin{aligned} \mathbf{V} &= \mathbf{B}_0 \oplus \dots \oplus \mathbf{B}_n \oplus \mathbf{F} \oplus \mathbf{W} \\ \mathbf{F} &= \mathbf{V} \rightarrow \mathbf{V} \end{aligned}$$

$\mathbf{B}_0, \dots, \mathbf{B}_n$ are basic domains (with bottom), such as integers and booleans; \mathbf{F} is the function domain; \mathbf{W} is a domain consisting of the single element `WRONG`; \oplus is the separated sum; and \rightarrow builds continuous functions. The meaning function \mathcal{E} is constructed so as to assign the denotation `WRONG` to programs that result in type errors. To assign meanings to expressions with free variables, the meaning function takes an environment $\rho : \text{Var} \rightarrow \mathbf{V}$, a finite map from variables to denotational values.

To establish soundness, Milner introduced a semantic relation \models . This relation identifies each type τ with an *ideal* \mathbf{V}^τ of the domain \mathbf{V} [18]. An ideal is simply a subset of the domain that is closed under certain operations (subset and least upper bounds of finite consistent subsets); each of the basic domains forms an ideal. A denotational value $v \in \mathbf{V}$ *possesses* type τ , written $\models v : \tau$, if v is a member of the ideal corresponding to τ , *i.e.*,

$$v \in \mathbf{V}^\tau \text{ iff } \models v : \tau.$$

An environment ρ *respects* a type environment Γ , written $\models \rho : \Gamma$, if the domain of ρ is at least as large as the domain of Γ , and for every variable x in the domain of Γ , $\models \rho(x) : \Gamma(x)$. A *semantic soundness* theorem states that if an expression e has type τ in type environment Γ , and if environment ρ respects Γ , then the denotation of e in ρ possesses type τ .

Theorem (Semantic Soundness) *If $\Gamma \triangleright e : \tau$ and $\models \rho : \Gamma$ then $\models \mathcal{E}[e]\rho : \tau$.*

This theorem is proved by induction on the structure of e . Strong soundness follows by restricting the theorem to closed expressions.

Theorem (Strong Soundness) *If $\triangleright e : \tau$ then $\models \mathcal{E}[e]\emptyset : \tau$.*

Weak soundness is a consequence of the fact that `WRONG` does not possess any type.

Theorem (Weak Soundness) *If $\triangleright e : \tau$ then $\mathcal{E}[e]\emptyset \neq \text{WRONG}$.*

Damas [5, 6]: Damas extended Milner’s results to a language with reference cells and destructive assignment. The proof technique is derived from Milner’s technique; however, the proof is significantly more complicated. The denotational semantics uses the following domain construction:

$$\begin{aligned} \mathbb{V} &= \mathbb{B}_0 \oplus \dots \oplus \mathbb{B}_n \oplus \mathbb{F} \oplus \mathbb{L} \oplus \mathbb{W} \\ \mathbb{F} &= \mathbb{V} \rightarrow \mathbb{S} \rightarrow \mathbb{V} \otimes \mathbb{S} \\ \mathbb{S} &= \mathbb{V}^* \\ \mathbb{L} &= \mathbb{N}_\perp \end{aligned}$$

where the domain of locations \mathbb{L} is the flat domain of natural numbers, \mathbb{S} is the domain of stores (sequences of values), and \otimes is the smash product. Procedures take both an input value and a store, and produce a result and a new store. The meaning function is parameterized over both a store and an environment. Due to the presence of stores, types are no longer ideals in \mathbb{V} , but are finite maps from store typings to subsets of \mathbb{V} .² A store typing is a finite map from locations to the types of values stored there. Since the semantic relation (\models) must involve the types of values in the store, it can no longer be defined by induction on types; its existence must be established by a category theoretical argument that generalizes the technique of inclusive predicates [31]. The complexity of the semantic relation complicates the proof; indeed, Tofte found a mistake in Damas’s proof, although the theorem is not thought to be false [37: page 2].

Abadi, Cardelli, Pierce, and Plotkin [1]: Abadi *et al.* demonstrate type soundness for a functional language with a *dynamic* type and related operations. Their proof with respect to a denotational semantics is similar to Milner’s proof, although it is complicated by the presence of *dynamic* values. A *dynamic* value is a pair consisting of a value v and a tag, drawn from a set *TypeCode*, that encodes the type of v . The domain equation includes a domain of *dynamic* values \mathbb{D} :

$$\begin{aligned} \mathbb{V} &= \mathbb{B}_0 \oplus \dots \oplus \mathbb{B}_n \oplus \mathbb{F} \oplus \mathbb{D} \oplus \mathbb{W} \\ \mathbb{F} &= \mathbb{V} \rightarrow \mathbb{V} \\ \mathbb{D} &= \textit{TypeCode} \times \mathbb{V} \end{aligned}$$

As type codes correspond to types, the denotation of a type code is a type. Types, in turn, denote ideals over \mathbb{V} , but due to the unusual element \mathbb{D} of the domain equation, establishing that types denote ideals requires extending the ideal model for recursive types [18]. The proof that the required fixed points exist in the new model involves a metric space argument, but is a straightforward extension of the original. Although the type system Abadi *et al.* use is not polymorphic, it should be possible to extend the type system and the soundness proof to include polymorphism.

Duba, Harper, and MacQueen [8]: Duba *et al.* present several languages extending a monomorphic functional core with first-class continuations in the spirit of Scheme [28]. They describe several approaches to proving type soundness for the languages. One is with

²This is an over-simplification; the reader interested in the precise definition is referred to Damas’s thesis [6].

respect to a continuation-passing denotational semantics, based on the following domain equations:

$$\begin{aligned} \mathbf{V} &= \mathbf{B}_0 \oplus \dots \oplus \mathbf{B}_n \oplus \mathbf{F} \oplus \mathbf{K} \\ \mathbf{F} &= \mathbf{V} \rightarrow \mathbf{K} \rightarrow \mathbf{A} \\ \mathbf{K} &= \mathbf{V} \rightarrow \mathbf{A} \\ \mathbf{A} &= \mathbf{V} \oplus \mathbf{W} \end{aligned}$$

where \mathbf{K} is the domain of continuations, and \mathbf{A} is the domain of answers. The meaning function \mathcal{E} takes both an environment and a continuation. Two semantic relations state what it means for a value v to possess a type τ , written $\models v : \tau$, and for a continuation κ to accept a value of type τ , written $\models \kappa :: \tau$. These relations are defined simultaneously by induction on types. The semantic soundness theorem states that if an expression e has type τ in type environment Γ , and if ρ respects Γ , and if continuation κ accepts values of type τ , then the meaning of e in environment ρ and continuation κ is not **WRONG**.

Theorem (Semantic Soundness) *If $\Gamma \triangleright e : \tau$ and $\models \rho : \Gamma$ and $\models \kappa :: \tau$ then $\mathcal{E}[[e]]\rho\kappa \neq \mathbf{WRONG}$.*

Unlike in the purely functional setting, the semantic soundness theorem only states that well-typed programs do not go **WRONG** (weak soundness), not that they produce answers of the expected type. To approach strong soundness, Duba *et al.* give an argument based on a continuation-passing-style (CPS) translation into the simply typed λ -calculus. For programs of ground type, the denotation of an expression is the same as the denotation of its CPS transform with the identity function as an initial continuation, *i.e.*,

$$\mathcal{E}[[e]]\emptyset = \mathcal{E}[[CPS(e) (\lambda x.x)]]\emptyset.$$

Since strong soundness holds for the simply typed λ -calculus, and $(CPS(e) (\lambda x.x))$ has the same type as e for programs of ground type, strong soundness holds for programs of ground type. However, the argument does not extend to higher types. For further details, we refer the reader to their paper [8: page 169].

2.2 Proofs based on Structural Operational Semantics

Tofte [36, 37]: Tofte reformulated Milner's functional language with a structural operational semantics [27]. The semantics is specified as a deductive proof system; a conclusion $E \vdash e \Rightarrow v$ of a deduction states that expression e evaluates to v in value environment E . A value environment is a finite map from variables to operational values. Values are either basic constants or closures, which result from the evaluation of λ -expressions. The soundness theorem again requires the definition of a semantic relation (\models) between operational values and types. For basic constants, the relation is explicitly specified by the function *TypeOf*; for closures, the relation is defined by induction on types, and by evaluation at correctly typed argument values. The semantic soundness theorem states that if an expression e has type τ in type environment Γ , and if environment E respects Γ , and if e evaluates to a value v in environment E , then v possesses type τ .

Theorem (Semantic Soundness) *If $\Gamma \triangleright e : \tau$ and $\models E : \Gamma$ and $E \vdash e \Rightarrow v$ then $\models v : \tau$.*

The proof is straightforward, and proceeds by induction on the depth of the deduction of $E \vdash e \Rightarrow v$. Strong and weak soundness follow as before.

Tofte turned to structural operational semantics in order to consider references. To this end, he reformulated the semantics with conclusions of the form $s, E \vdash e \Rightarrow v, s'$, meaning that in environment E and store s , expression e evaluates to value v and the new store s' . The semantic relation now involves the contents of the store and the store typing (ST), a map from locations to the types of values stored there. The semantic relation $s : ST \models v : \tau$ may be read as “given the ST -typed store s , value v possesses type τ .” Because of the possibility of cycles amongst references in the store, this relation is not definable by induction, but must be defined as the maximal fixed-point of a monotonic operator on the appropriate space. The semantic soundness theorem states that if an expression e has type τ in type environment Γ , and given store s of type ST , environment E respecting Γ , and that e evaluates to v and store s' , then a typing ST' can be found for s' such that v possesses type τ .

Theorem (Semantic Soundness) *If $\Gamma \triangleright e : \tau$ and $s : ST \models E : \Gamma$ and $s, E \vdash e \Rightarrow v, s'$ then there exists ST' such that $s' : ST' \models v : \tau$.*

The theorem is proved by induction on the depth of the deduction of $s, E \vdash e \Rightarrow v, s'$. Lemmas involving the semantic relation are proved using the technique of co-induction that Milner and Tofte developed for this purpose [21]. Leroy and Weis have successfully applied Tofte’s technique to show soundness for a different type system for polymorphic references [17]. Talpin and Jouvelot have also applied the technique to demonstrate soundness for a system that infers types, effects, and regions for references [35].

Duba, Harper, and MacQueen [8]: Duba *et al.* also present a proof of type soundness with respect to a structural operational semantics for a monomorphic language with continuations. The proof is an adaptation of Tofte’s technique; however, the semantics is significantly restructured, as the technique of defunctionalization [30] is used to represent the flow of control explicitly. The semantics has two judgment forms. The first kind of conclusion, $E; K \vdash e \Rightarrow v$, indicates that expression e evaluates to answer v in environment E and continuation K . The second kind of conclusion, $v \vdash K \Rightarrow v'$, indicates that continuation K evaluated at value v produces answer v' . As in the denotational case, two semantic relations are defined, indicating what it means for a value to possess a type ($\models v : \tau$), and for a continuation to accept a type ($\models K :: \tau$). Like the denotational case, the theorem only establishes weak soundness.

Theorem (Semantic Soundness) *If $\Gamma \triangleright e : \tau$ and $\models E : \Gamma$ and $\models K :: \tau$ and $E; K \vdash e \Rightarrow v$ then $v \neq \text{WRONG}$.*

It is not clear that the proof goes through in the presence of fixed-point operators [8: page 171], nor how to obtain strong soundness in this framework.

2.3 Discussion

In each case above, the proof of soundness for a language or for a different formulation of a language’s semantics involves the use of a different proof technique. The techniques

are often unrelated, so they provide no guidance in proving soundness for new languages or language features. A seemingly minor extension to a language may require a complete restructuring of its denotational or structural operational semantics, and may therefore require a completely new approach to re-establish soundness. For example, in introducing references to Milner’s functional language, Damas changed the domain equations for the denotational semantics to accommodate a store component, necessitating a completely different strategy for induction. In introducing type *dynamic*, Abadi *et al.* again changed the domain equations, and extended the ideal model of types to match. To accommodate continuations, Duba *et al.* changed the domain equations yet again, also foregoing the tools developed by Milner and Damas. They used a separate argument to establish strong soundness. In introducing references to an operational formulation of the functional language, Tofte changed the form of judgments obtained by the semantics, and used the technique of co-induction to establish a proof. For continuations in a structural operational semantics, Duba *et al.* completely altered the semantics,³ again changing the structure of the proof dramatically.

Each of the above proofs considers a language with a single extension to a functional core. It is natural to ask whether two extensions can be merged, such that a natural union of the two type systems is sound. However, since each of the proofs involves a different technique, it is generally impossible to merge the proofs. Tofte’s proof for references and Duba *et al.*’s proof for continuations give no direct assurance that STANDARD ML’s type system is sound, even if the other features of STANDARD ML are ignored.

3 A Syntactic Approach to Proving Type Soundness

Our approach to type soundness is based upon an operational formulation of a language’s semantics by *rewriting*. Each intermediate state of an evaluation of a program is itself a program, and the evaluation of a program is performed by successive *reductions* into a new state: $e_1 \mapsto e_2 \mapsto \dots$. Reduction may either continue forever ($e \uparrow$), or may reach a *final* state where no further evaluation is possible: $e_1 \mapsto e_k$ and $e_k \not\mapsto e_{k+1}$. Such a final state represents either an answer or a type error. Thus, proving type soundness reduces to proving that well-typed programs yield only well-typed answers.

Programming language calculi, like the λ -calculus, are the natural choice to specify the semantics of a language such that each intermediate step of evaluation is a program. Plotkin [26] shows how the semantics of a prototypical functional language relates to the λ -calculus. The λ -calculi extensions for state [12, 14] extend this strategy to languages with references and similar constructs. Likewise, the control calculi [13, 14] can be adapted to address non-local control facilities such as exceptions and first-class continuations.

Since the intermediate states of evaluation are programs, we may apply the type system to deduce a type for each state. Our strategy for proving type soundness rests upon two fundamental observations about the types of states. The first observation is that each intermediate state may be assigned the same type as the original program, *i.e.*, reductions preserve type:

³Furthermore, they take this as “evidence that the addition of [continuations] to Standard ML would be a substantial change, rather than an incremental modification, to the language” [8: page 170]. However, we believe that this only indicates a problem with their semantic framework.

if $\triangleright e_1 : \tau$ and $e_1 \longmapsto e_2$ then $\triangleright e_2 : \tau$.

In combinatory logic, this property is known as *subject reduction* [4].

The second observation is that sound type systems do not assign a type to irreducible expressions that are the sources of type errors, such as $1 + \text{true}$. Furthermore, any expression containing such a subexpression is untypable; we call expressions containing such irreducible subexpressions *faulty*:

if e is faulty, there is no τ such that $\triangleright e : \tau$.

Together, these two observations imply that all final states that can be reached from a well-typed program are well-typed answers. With a definition of *eval* that takes faulty expressions to `WRONG`, the type system is weakly sound.

As all answers produced are well-typed programs, subject reduction also establishes strong soundness. To state that an answer v lies within the correct subset V^τ of values, we employ the type system itself:

$$v \in V^\tau \text{ iff } \triangleright v : \tau.$$

This yields a soundness theorem based on a *syntactic* connection between answers and types:

if $\triangleright e : \tau$ then either $e \uparrow$ or $e \longmapsto v$ and $\triangleright v : \tau$.

Again, strong soundness follows with a definition of *eval* that takes faulty expressions to `WRONG`.

In summary, proving a syntactic soundness theorem involves:

- demonstrating subject reduction,
- characterizing answers and faulty expressions, and
- showing that the faulty expressions are untypable.

The structure of the proof always remains the same, whether or not the language includes some combination of references, exceptions, continuations, or other features.

3.1 Related Work

Curry and Feys [4] introduced the notion of subject reduction for combinatory logic. In the language of combinatory logic, the CL-term e of a deduction concluding $\triangleright e : \tau$ is called the *subject*, and the type τ is called the *predicate*. Subject reduction states that reduction of the subject of a deduction preserves the predicate. Subject reduction holds for terms in CL and the λ -calculus [33].

Mitchell and Plotkin [25] present a type preservation theorem, *i.e.*, subject reduction, for a variant of the second order polymorphic λ -calculus. This language has explicitly typed declarations; the proof of type preservation is significantly simpler in the presence of such declarations. They do not develop a type soundness theorem.

Abadi *et al.* [1] give a proof of soundness for type *dynamic* with respect to a structural operational semantics, in addition to their denotationally based proof. In their structural

operational semantics, substitutions are carried out immediately: λ -expressions evaluate to themselves, rather than closures, and judgments do not contain an environment component. Since answers are syntactic values, the type system can be applied to answers to show that evaluation preserves typing. Although this proof is similar to our proof of subject reduction, its extension to references or continuations apparently involves the same difficulties as Tofte's and Duba *et al.*'s proofs.

3.2 Road Map

The next three sections illustrate our approach by proving soundness for several languages. Section 4 considers a functional language, similar to that considered by Milner [20]. Section 5 considers an extension to references, an extension to exceptions, and shows how the results may be merged to consider a language with both. To our knowledge, this is the first proof of soundness for exceptions. Section 6 illustrates soundness for an extension providing first-class continuations, and is the first proof of strong type soundness for continuations in a polymorphic language. We conclude with a discussion of our technique, and suggestions for its application to other languages.

4 Functional ML

Functional ML is an applicative language with constants, call-by-value higher-order functions, and a Hindley/Milner-style polymorphic type system. A natural basis for a calculus of Functional ML is Plotkin's untyped λ_v -calculus [26].

Let Var be a denumerable set of *variables* and $Const$ be a set of *constants*. The *expressions* and *values* of Functional ML are:

$$\begin{array}{ll} (\text{Expressions}) & e ::= v \mid e_1 e_2 \mid \text{let } x \text{ be } e_1 \text{ in } e_2 \\ (\text{Values}) & v ::= c \mid x \mid Y \mid \lambda x.e \end{array}$$

where $x \in Var$ and $c \in Const$. The free variables, $FV(e)$, and bound variables of an expression are defined as usual, with λ - and *let*-expressions binding their variables. The *let*-expression binds x in e_2 but not e_1 , *i.e.*, *let* bindings are not recursive. The fixed-point combinator Y provides recursion. Following Barendregt [2], we adopt the convention that bound variables are always distinct from free variables in distinct expressions, and we identify expressions that differ only by a consistent renaming of the bound variables. $ClosedVal$ is the set of values with no free variables.

4.1 Semantics

The calculus for Functional ML is based upon four relations, called *notions of reduction*:

$$\begin{array}{ll} (\delta) & c v \longrightarrow \delta(c, v) \text{ if } \delta(c, v) \text{ is defined} \\ (\beta_v) & (\lambda x.e) v \longrightarrow e[x \mapsto v] \\ (\text{let}) & \text{let } x \text{ be } v \text{ in } e \longrightarrow e[x \mapsto v] \\ (Y) & Y v \longrightarrow v (\lambda x.(Y v) x) \end{array}$$

To abstract from the precise set of constants, we only assume the existence of a partial function:

$$\delta : \text{Const} \times \text{ClosedVal} \rightarrow \text{ClosedVal}$$

that interprets the application of functional constants to closed values and yields closed values. The β_v and *let* reductions use substitution: the notation $e[x \mapsto v]$ means the substitution of v for free occurrences of x in e , renaming bound variables of v as necessary to avoid capture. The Y reduction introduces an abstraction around $(Y v)$ as v must be applied to a value (x does not appear free in v by the variable conventions).

We refer to the union of these relations as \mathbf{v} , or simply \longrightarrow when there is no danger of confusion. The notion of reduction \mathbf{v} gives rise to a system of reductions, \longrightarrow . The relation \longrightarrow is the reflexive, transitive, and compatible closure of \longrightarrow :

$$(\longrightarrow) \quad \frac{e_1 \longrightarrow e_2}{e_1 \longrightarrow e_2} \quad e_1 \longrightarrow e_1 \quad \frac{e_1 \longrightarrow e_2 \quad e_2 \longrightarrow e_3}{e_1 \longrightarrow e_3} \quad \frac{e_1 \longrightarrow e_2}{C[e_1] \longrightarrow C[e_2]}$$

where

$$C ::= [] \mid C e \mid e C \mid \text{let } x \text{ be } C \text{ in } e \mid \text{let } x \text{ be } e \text{ in } C \mid \lambda x.C.$$

A *context* C is an expression with one subexpression replaced by a hole, denoted $[]$. The expression $C[e]$ results from placing an expression e in the hole of C ; this operation may involve capture of the free variables of e by C . An equational system may be constructed as the congruence closure of \mathbf{v} , but plays no rôle for our work.

With an appropriate choice of δ , the calculus satisfies Church-Rosser and Standardization properties; the proofs are variants of Plotkin's proofs for the λ_v -calculus [26]. By Standardization, we know that an expression reduces to a value precisely if a reduction of leftmost-outermost redexes outside of abstractions leads to a value. Based on this idea, we can use the calculus to define an evaluation function [11]. Let \mapsto be the relation:

$$(\mapsto) \quad E[e] \mapsto E[e'] \quad \text{iff} \quad e \longrightarrow e'$$

where

$$E ::= [] \mid E e \mid v E \mid \text{let } x \text{ be } E \text{ in } e.$$

The special contexts E are *evaluation contexts*: they ensure that the leftmost-outermost reduction is the only applicable reduction in the entire program. Evaluating from left to right in an application is consistent with STANDARD ML. Unlike with C contexts, there is no possibility of capture when placing an expression in the hole of an evaluation context, because the hole cannot appear inside a binding construct. Let \mapsto be the reflexive and transitive closure of \mapsto . Then the partial function *eval* is defined for closed expressions e as:

$$(\text{eval}) \quad \text{eval}(e) = v \quad \text{iff} \quad e \mapsto v.$$

It is easy to see that the decomposition of a program into a \mathbf{v} -redex and an evaluation context is unique. It follows that the stepping relation (\mapsto) is a function, and therefore that *eval* is a function, too.

4.2 Typing

The type system is a deductive proof system that assigns types to expressions. The types of Functional ML are:

$$(Types) \quad \tau ::= \iota_1 \mid \dots \mid \iota_n \mid \alpha \mid \tau_1 \rightarrow \tau_2$$

where ι_1, \dots, ι_n are ground types for basic constants (*int*, *bool*, *real*, etc.), and α ranges over a denumerable set of *type variables*. Type variables stand for some fixed but unknown type. Function types, $\tau_1 \rightarrow \tau_2$, are right associative.

The polymorphic type system allows certain variables to be used with a set of different types. For example, if the identity function $\lambda x.x$ is **let**-bound to the variable f , then f may be used as a function of type $int \rightarrow int$ and applied to values of type int , or f may be used as a function of type $bool \rightarrow bool$ and applied to values of type $bool$. *Type schemes* represent such sets of types (α^* means zero or more distinct type variables):

$$(TypeSchemes) \quad \sigma ::= \forall \alpha^*. \tau$$

The free type variables $FTV(\sigma)$ of a type or type scheme are defined as usual, where $\forall \alpha_1 \dots \alpha_n. \tau$ binds α_1 through α_n in τ . We identify type schemes that differ only by a consistent renaming of bound type variables, or by a reordering of the binding occurrences. We also identify $\forall. \tau$ with τ , so types may be regarded as type schemes with no bound type variables.

A type scheme $\forall \alpha_1 \dots \alpha_n. \tau$ represents the set of types that may be obtained from τ by substituting for α_1 through α_n . This is formalized by the notion of *generalization*: a type scheme σ generalizes a type τ , written $\sigma \succ \tau$, if there is a substitution for the bound variables of σ yielding τ . A substitution S is a (partial) function from type variables to types. A substitution may be applied to a type in the obvious way; we write $S\tau$. Substitution on type schemes respects bound type variables and is capture-avoiding. Generalization of a type is defined as:

$$(\succ) \quad \forall \alpha_1 \dots \alpha_n. \tau' \succ \tau \text{ iff } S\tau' = \tau \text{ and } \text{Dom}(S) = \{\alpha_1, \dots, \alpha_n\} \text{ for some } S.$$

Generalization is extended to type schemes:

$$(\succ) \quad \sigma \succ \sigma' \text{ iff whenever } \sigma' \succ \tau \text{ then } \sigma \succ \tau.$$

Some examples are:

$$\begin{array}{lcl} \tau & \succ & \tau \\ \forall \alpha. \alpha \rightarrow \alpha & \succ & int \rightarrow int \\ \forall \alpha_1 \alpha_2. \alpha_1 \rightarrow \alpha_2 & \succ & int \rightarrow bool \\ \forall \alpha. \alpha \rightarrow \alpha & \succ & \alpha_2 \rightarrow \alpha_2 \\ \forall \alpha_1 \alpha_2. \alpha_1 \rightarrow \alpha_2 & \succ & \forall \alpha_1. \alpha_1 \rightarrow \alpha_2 \end{array}$$

Generalization is a partial order on type schemes.

In order to assign types to open expressions, assumptions about the types of free variables are provided by *type environments*, denoted Γ . Type environments are finite maps from variables to type schemes. The free type variables $FTV(\Gamma)$ of a type environment

(var)	$\Gamma \triangleright x : \tau$ if $\Gamma(x) \succ \tau$
(const)	$\Gamma \triangleright c : \tau$ if $TypeOf(c) \succ \tau$
(Y)	$\Gamma \triangleright Y : ((\tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2$
(abs)	$\frac{\Gamma[x \mapsto \tau_1] \triangleright e : \tau_2}{\Gamma \triangleright \lambda x. e : \tau_1 \rightarrow \tau_2}$
(app)	$\frac{\Gamma \triangleright e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \triangleright e_2 : \tau_1}{\Gamma \triangleright e_1 e_2 : \tau_2}$
(let)	$\frac{\Gamma \triangleright e_1 : \tau_1 \quad \Gamma[x \mapsto Close(\tau_1, \Gamma)] \triangleright e_2 : \tau_2}{\Gamma \triangleright \text{let } x \text{ be } e_1 \text{ in } e_2 : \tau_2}$

$Close(\tau, \Gamma) = \forall \alpha_1 \dots \alpha_n. \tau$ where $\{\alpha_1, \dots, \alpha_n\} = FTV(\tau) \setminus FTV(\Gamma)$

Figure 1: Typing rules for Functional ML

are the free type variables of the type schemes in its range. For stating and proving various lemmas, it is convenient to define substitution on type environments pointwise, *i.e.*, $S(\{\langle x_1, \sigma_1 \rangle, \dots, \langle x_n, \sigma_n \rangle\}) = \{\langle x_1, S\sigma_1 \rangle, \dots, \langle x_n, S\sigma_n \rangle\}$.

The function $TypeOf : Const \rightarrow TypeSchemes$ associates a type scheme with every constant. Since we have not precisely specified the set of constants, Functional ML is really a class of languages parameterized by the set $Const$ and the functions δ and $TypeOf$. Just as the Standardization and Church-Rosser properties restrict the admissible set of δ functions, a particular choice of $Const$, δ , and $TypeOf$ must satisfy a *typability* condition for type soundness to hold. For every c, τ, τ' , and v ,

(δ -typability) if $TypeOf(c) \succ \tau' \rightarrow \tau$ and $\triangleright v : \tau'$,
then $\delta(c, v)$ is defined and $\triangleright \delta(c, v) : \tau$.

This condition requires that δ be defined for all constants of functional type and arguments of matching type, and restricts the set of results that δ may produce. This rules out, for the moment, functional constants such as division that are not defined on all values of their input type.

A *type judgment* $\Gamma \triangleright e : \tau$ for e an expression of Functional ML is the conclusion of a deduction constructed according to the inference rules in Figure 1. The judgment $\Gamma \triangleright e : \tau$ is read “in type environment Γ , expression e has type τ ”. If the type environment in a judgment is empty, we write $\triangleright e : \tau$. An expression e is a *well-typed program* if it is closed and there is a judgment $\triangleright e : \tau$.

In this formulation of the typing rules, only one rule is applicable to an expression. Hence if there is a deduction assigning an expression a specific type, that deduction is unique. This permits proofs by induction on the structure of a deduction for a type judgment to proceed by case analysis on the structure of the expression of the type judgment.

This type system has the important property that there is an algorithm \mathcal{W} to determine whether an expression has a type [20]. Given a type environment and an expression, the algorithm computes a substitution and a type. The algorithm is sound [20] with respect to

the type system, meaning that it infers only valid typings.

Theorem (Soundness of \mathcal{W}) [Milner] *If $(S, \tau) = \mathcal{W}(\Gamma, e)$ succeeds then $S\Gamma \triangleright e : \tau$.*

The algorithm is also complete [5], meaning that if an expression has a valid typing, then the algorithm will find a typing. Furthermore, the algorithm always terminates.

Theorem (Completeness of \mathcal{W}) [Damas, Milner] *If $S_1\Gamma \triangleright e : \tau_1$ then $(S, \tau) = \mathcal{W}(\Gamma, e)$ succeeds and there exists S' such that $S'(S\Gamma) = S_1\Gamma$ and $S'Close(\tau, S\Gamma) \succ \tau_1$. If $\mathcal{W}(\Gamma, e)$ does not succeed, then it stops with **fail**.*

Proving completeness involves demonstrating that if an expression has a type, then it has a *principal* type, of which all others are substitution instances.

Theorem (Principal Typing) *If $\Gamma \triangleright e : \tau_1$ then there exists τ such that $\Gamma \triangleright e : \tau$ and whenever $\Gamma \triangleright e : \tau_2$ then $Close(\tau, \Gamma) \succ \tau_2$. The type τ is principal for e in Γ .*

The existence of a sound and complete algorithm establishes the decidability of the type system. However, this is solely a property of the type system, and says nothing about the relation of the type system to the semantics of the language.

4.3 Type Soundness

Our proof of type soundness rests upon the notion of subject reduction [4]. The subject reduction property states that reductions preserve the type of expressions. Lemma 4.3 below extends subject reduction to Functional ML.

Subject reduction by itself is not sufficient for type soundness. In addition, we must prove that programs containing type errors are not typable. Put differently, if an expression e is irreducible due to some type error, say $(/ \ 1 \ 0)$, then the type system should not be able to assign a type to the expression. Otherwise, a well-typed program could reduce to such an expression without violating subject reduction, and still cause a type error to occur. We call such expressions with type errors *faulty expressions* and prove that faulty expressions cannot be typed.

Some obvious facts about deductions that we use with no more ado are:

- (i) if $\Gamma \triangleright C[e] : \tau$ then there exist Γ' and τ' such that $\Gamma' \triangleright e : \tau'$;
- (ii) if there are no Γ', τ' such that $\Gamma' \triangleright e : \tau'$, then there are no Γ, τ such that $\Gamma \triangleright C[e] : \tau$.

These follow from the facts that (1) there is exactly one inference rule for each expression form, and (2) each inference rule requires a proof for each subexpression of the expression in its conclusion. Several facts about generalization follow from the definition of *Close*:

- (i) if $\sigma' \succ \sigma$ then $Close(\tau_1, \Gamma[x \mapsto \sigma']) \succ Close(\tau_1, \Gamma[x \mapsto \sigma])$;
- (ii) if $x \notin \text{Dom}(\Gamma)$ then $Close(\tau_1, \Gamma) \succ Close(\tau_1, \Gamma[x \mapsto \sigma])$.

The following lemma states that extra variables in the type environment Γ of a judgment $\Gamma \triangleright e : \tau$ that are not free in the expression e may be ignored.

Lemma 4.1 *If $\Gamma(x) = \Gamma'(x)$ for all $x \in FV(e)$, then $\Gamma \triangleright e : \tau$ iff $\Gamma' \triangleright e : \tau$.*

A key lemma that we use in the proof of Subject Reduction for Functional ML and its extensions is the Replacement Lemma, adapted from Hindley and Seldin [16: page 181]. This allows the replacement of one of the subexpressions of a typable expression with another subexpression of the same type, without disturbing the type of the overall expression.

Lemma 4.2 (Replacement) *If:*

- (i) \mathcal{D} is a deduction concluding $\Gamma \triangleright C[e_1] : \tau$,
- (ii) \mathcal{D}_1 is a subdeduction of \mathcal{D} concluding $\Gamma' \triangleright e_1 : \tau'$,
- (iii) \mathcal{D}_1 occurs in \mathcal{D} in the position corresponding to the hole (\square) in C , and
- (iv) $\Gamma' \triangleright e_2 : \tau'$,

then $\Gamma \triangleright C[e_2] : \tau$.

Proof. We think of deductions as trees with the conclusion at the root. Let \mathcal{D}_2 be the deduction concluding $\Gamma' \triangleright e_2 : \tau'$. Cut the subtree \mathcal{D}_1 out of \mathcal{D} and replace it with \mathcal{D}_2 . Replace all relevant occurrences of e_1 in the resulting tree with e_2 . Then the resulting tree is a valid deduction concluding $\Gamma \triangleright C[e_2] : \tau$, as may be shown by induction on the height of the tree [16: page 181]. \blacksquare

Subject reduction for Functional ML states that the reductions of \mathbf{v} preserve type.

Main Lemma 4.3 (Subject Reduction) *If $\Gamma \triangleright e_1 : \tau$ and $e_1 \longrightarrow e_2$ then $\Gamma \triangleright e_2 : \tau$.*

Proof. The proof proceeds by case analysis according to the reduction $e_1 \longrightarrow e_2$.

Case $c v \longrightarrow \delta(c, v)$. Then $\Gamma \triangleright c : \tau' \rightarrow \tau$ and $\Gamma \triangleright v : \tau'$ follow from $\Gamma \triangleright c v : \tau$ by **app**. Hence $\Gamma \triangleright \delta(c, v) : \tau$ by the δ -typability condition.

Case $(\lambda x.e) v \longrightarrow e[x \mapsto v]$. From $\Gamma \triangleright (\lambda x.e) v : \tau$ we have $\Gamma \triangleright v : \tau'$ and $\Gamma \triangleright \lambda x.e : \tau' \rightarrow \tau$ by **app**. From the latter, $\Gamma[x \mapsto \tau'] \triangleright e : \tau$ follows by **abs**. Hence $\Gamma \triangleright e[x \mapsto v] : \tau$ by Lemma 4.4.

Case $\text{let } x \text{ be } v \text{ in } e \longrightarrow e[x \mapsto v]$. From $\Gamma \triangleright \text{let } x \text{ be } v \text{ in } e : \tau$ we have $\Gamma \triangleright v : \tau'$ and $\Gamma[x \mapsto \text{Close}(\tau', \Gamma)] \triangleright e : \tau$ by **let**. As $\text{Close}(\tau', \Gamma) = \forall \alpha_1 \dots \alpha_n. \tau'$ where $\{\alpha_1, \dots, \alpha_n\} = FTV(\tau') \setminus FTV(\Gamma)$, we have $\Gamma \triangleright e[x \mapsto v] : \tau$ by Lemma 4.4.

Case $\mathbf{Y} v \longrightarrow v (\lambda x.(\mathbf{Y} v) x)$. From $\Gamma \triangleright \mathbf{Y} v : \tau$, where $\tau = \tau_1 \rightarrow \tau_2$, by **Y** and **app**

$$(1) \quad \Gamma \triangleright v : (\tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2.$$

Then from $\Gamma \triangleright \mathbf{Y} v : \tau$,

$$\begin{array}{ll} \Gamma[x \mapsto \tau_1] \triangleright \mathbf{Y} v : \tau_1 \rightarrow \tau_2 & \text{by Lemma 4.1,} \\ \Gamma[x \mapsto \tau_1] \triangleright (\mathbf{Y} v) x : \tau_2 & \text{by } \mathbf{app}, \\ \Gamma \triangleright \lambda x.(\mathbf{Y} v) x : \tau_1 \rightarrow \tau_2 & \text{by } \mathbf{abs}, \\ \Gamma \triangleright v (\lambda x.(\mathbf{Y} v) x) : \tau_1 \rightarrow \tau_2 & \text{by } \mathbf{app} \text{ with (1).} \end{array}$$

This completes the essence of the proof. It remains to establish several lemmas. \blacksquare

A Substitution Lemma is the key to showing type preservation for reductions involving substitution.

Lemma 4.4 (Substitution) *If $\Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e : \tau'$ and $x \notin \text{Dom}(\Gamma)$ and $\Gamma \triangleright v : \tau$ and $\{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\Gamma) = \emptyset$ then $\Gamma \triangleright e[x \mapsto v] : \tau'$.*

Proof. We proceed by induction on the length of the proof of $\Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e : \tau'$, and case analysis on the last step.

Case $e = c$. Then $\text{TypeOf}(c) \succ \tau'$ by **const**, and thus $\Gamma \triangleright c : \tau'$ by **const**. Then $\Gamma \triangleright c[x \mapsto v] : \tau'$ since $c[x \mapsto v] = c$.

Case $e = x'$. If $x' \neq x$, then $\Gamma x' \succ \tau'$ by **var**, and $\Gamma \triangleright x' : \tau'$ again by **var**. Hence $\Gamma \triangleright x'[x \mapsto v] : \tau'$ since $x'[x \mapsto v] = x'$.

If $x' = x$, then by **var**

$$\begin{aligned} & \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau](x) \succ \tau' \\ \text{i.e. } & \forall \alpha_1 \dots \alpha_n. \tau \succ \tau', \end{aligned}$$

i.e. we can find a substitution S with domain $\{\alpha_1, \dots, \alpha_n\}$ such that $S\tau = \tau'$. Next, we have $S\Gamma \triangleright v : S\tau$ by Lemma 4.5 since $\Gamma \triangleright v : \tau$, and therefore $S\Gamma \triangleright v : \tau'$. But $\text{Dom}(S) \cap \text{FTV}(\Gamma) = \emptyset$, so $S\Gamma = \Gamma$. Also $x[x \mapsto v] = v$, hence $\Gamma \triangleright x[x \mapsto v] : \tau'$.

Case $e = \lambda x'. e_1$. By **abs**, it follows from the assumption $\Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright \lambda x'. e_1 : \tau'$ that $\tau' = \tau_1 \rightarrow \tau_2$, and

$$\begin{aligned} & \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau][x' \mapsto \tau_1] \triangleright e_1 : \tau_2 \\ \text{i.e. } & \Gamma[x' \mapsto \tau_1][x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_1 : \tau_2. \end{aligned}$$

Choose a substitution $S : \{\alpha_1, \dots, \alpha_n\} \rightarrow \{\alpha'_1, \dots, \alpha'_n\}$ such that $\alpha'_1, \dots, \alpha'_n, \alpha_1, \dots, \alpha_n$, and $\text{FTV}(\Gamma)$ are all distinct. Then

- (2) $\Gamma[x' \mapsto S\tau_1][x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_1 : S\tau_2$ by Lemma 4.5,
- (3) $\Gamma[x' \mapsto S\tau_1] \triangleright v : \tau$ by Lemma 4.1, and
- (4) $\text{FTV}(\Gamma[x' \mapsto S\tau_1]) \cap \{\alpha_1, \dots, \alpha_n\} = \emptyset$ by the choice of S .

Thus, $\Gamma[x' \mapsto S\tau_1] \triangleright e_1[x \mapsto v] : S\tau_2$ by the inductive hypothesis with (2), (3), and (4). Since S is a bijection, S^{-1} exists, hence

$$\begin{aligned} & S^{-1}(\Gamma[x' \mapsto S\tau_1]) \triangleright e_1[x \mapsto v] : S^{-1}(S\tau_2) \quad \text{by Lemma 4.5,} \\ \text{i.e. } & \Gamma[x' \mapsto \tau_1] \triangleright e_1[x \mapsto v] : \tau_2. \end{aligned}$$

But then

$$\begin{aligned} & \Gamma \triangleright \lambda x'. e_1[x \mapsto v] : \tau_1 \rightarrow \tau_2 && \text{by } \mathbf{abs}, \\ \text{i.e. } & \Gamma \triangleright (\lambda x'. e_1)[x \mapsto v] : \tau_1 \rightarrow \tau_2. \end{aligned}$$

Case $e = (e_1 e_2)$. From $\Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright (e_1 e_2) : \tau'$ by the first premise of **app**

$$(5) \quad \begin{array}{l} \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_1 : \tau_1 \rightarrow \tau' \\ \Gamma \triangleright e_1[x \mapsto v] : \tau_1 \rightarrow \tau' \end{array} \quad \text{by ind. hyp.}$$

By the second premise of **app**

$$(6) \quad \begin{array}{l} \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_2 : \tau_1 \\ \Gamma \triangleright e_2[x \mapsto v] : \tau_1 \end{array} \quad \text{by ind. hyp.}$$

Then by **app** with (5) and (6)

$$\begin{array}{l} \Gamma \triangleright (e_1[x \mapsto v] e_2[x \mapsto v]) : \tau' \\ \text{i.e. } \Gamma \triangleright (e_1 e_2)[x \mapsto v] : \tau'. \end{array}$$

Case $e = \text{let } x' \text{ be } e_1 \text{ in } e_2$. By the first premise of **let**

$$(7) \quad \begin{array}{l} \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_1 : \tau_1 \\ \Gamma \triangleright e_1[x \mapsto v] : \tau_1 \end{array} \quad \text{by ind. hyp.}$$

By the second premise of **let**

$$(8) \quad \begin{array}{l} \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau][x' \mapsto \text{Close}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])] \triangleright e_2 : \tau' \\ \text{i.e. } \Gamma[x' \mapsto \text{Close}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])][x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_2 : \tau'. \end{array}$$

Since $\Gamma \triangleright v : \tau$

$$(9) \quad \Gamma[x' \mapsto \text{Close}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])] \triangleright v : \tau \quad \text{by Lemma 4.1.}$$

Now

$$\begin{aligned} & \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\Gamma[x' \mapsto \text{Close}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])]) \\ & \subseteq \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(\Gamma) \cup \text{FTV}(\text{Close}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau]))) \\ & = \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\text{Close}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])) \\ & = \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(\tau_1) \setminus (\text{FTV}(\tau_1) \setminus \text{FTV}(\Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau]))) \\ & = \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\tau_1) \cap \text{FTV}(\Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau]) \\ & \subseteq \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(\Gamma) \cup \text{FTV}(\forall \alpha_1 \dots \alpha_n. \tau)) \\ & = \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\forall \alpha_1 \dots \alpha_n. \tau) \\ & = \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(\tau) \setminus \{\alpha_1, \dots, \alpha_n\}) \\ & = \emptyset \end{aligned}$$

so by the inductive hypothesis with (8) and (9)

$$\begin{array}{l} \Gamma[x' \mapsto \text{Close}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])] \triangleright e_2[x \mapsto v] : \tau' \\ \Gamma[x' \mapsto \text{Close}(\tau_1, \Gamma)] \triangleright e_2[x \mapsto v] : \tau' \quad \text{by Lemma 4.6,} \\ \Gamma \triangleright \text{let } x' \text{ be } v_1[x \mapsto v] \text{ in } e_2[x \mapsto v] : \tau' \quad \text{by } \mathbf{let} \text{ with (7),} \\ \text{i.e. } \Gamma \triangleright (\text{let } x' \text{ be } v_1 \text{ in } e_2)[x \mapsto v] : \tau'. \quad \blacksquare \end{array}$$

Lemma 4.5 establishes that typing is stable under substitution.

Lemma 4.5 *If $\Gamma \triangleright e : \tau$ and S is a substitution then $S\Gamma \triangleright e : S\tau$.*

Proof. The proof proceeds by induction on the length of the proof of $\Gamma \triangleright e : \tau$ and case analysis on the last step. The proof is an adaptation of Tofte's proof of a similar lemma [37: lemma 4.2]. ■

Finally, Lemma 4.6 shows that generalizing the type of a variable in the type environment has no impact on the conclusion of a deduction.

Lemma 4.6 *If $\Gamma[x \mapsto \sigma] \triangleright e : \tau$ and $\sigma' \succ \sigma$ then $\Gamma[x \mapsto \sigma'] \triangleright e : \tau$.*

Proof. The proof proceeds by induction on the length of the proof of $\Gamma[x \mapsto \sigma] \triangleright e : \tau$ and case analysis on the last step. The proof is an adaptation of Damas and Milner's proof of a similar lemma [5: lemma 1]. The only interesting case is the case for **let**-expressions.

Case $e = \text{let } x' \text{ be } e_1 \text{ in } e_2$. By the first premise of **let**

$$(10) \quad \begin{array}{l} \Gamma[x \mapsto \sigma] \triangleright e_1 : \tau_1 \\ \Gamma[x \mapsto \sigma'] \triangleright e_1 : \tau_1 \end{array} \quad \text{by ind. hyp.}$$

By the second premise of **let**

$$(11) \quad \begin{array}{l} \Gamma[x \mapsto \sigma][x' \mapsto \text{Close}(\tau_1, \Gamma[x \mapsto \sigma])] \triangleright e_2 : \tau \\ \Gamma[x \mapsto \sigma'][x' \mapsto \text{Close}(\tau_1, \Gamma[x \mapsto \sigma])] \triangleright e_2 : \tau \quad \text{by ind. hyp.} \\ \Gamma[x \mapsto \sigma'][x' \mapsto \text{Close}(\tau_1, \Gamma[x \mapsto \sigma'])] \triangleright e_2 : \tau \quad \text{by ind. hyp.} \end{array}$$

Then $\Gamma[x \mapsto \sigma'] \triangleright \text{let } x' \text{ be } e_1 \text{ in } e_2 : \tau$ by **let** with (10) and (11). ■

A corollary of Subject Reduction is that standard reduction steps preserve type.

Corollary 4.7 *If $\Gamma \triangleright e_1 : \tau$ and $e_1 \mapsto e_2$ then $\Gamma \triangleright e_2 : \tau$.*

Proof. First, if $e_1 \mapsto e_2$ then $e_1 = E[e]$ and $e_2 = E[e']$ and $e \mapsto e'$, so $\Gamma \triangleright e_2 : \tau$ by the Replacement Lemma. Then the result follows by induction on the length of the reduction sequence $e_1 \mapsto e_2$. ■

Subject reduction ensures that if we start with a typable expression, we cannot reach an untypable expression through any sequence of reductions. This by itself, however, does not yield type soundness. Subject reduction simply ensures that any properties implied by typability are preserved by reduction. The critical property we seek is that evaluation of a typable expression cannot get *stuck*.

Definition 4.8 (Stuck Expressions) *The evaluation of an expression e is stuck if e is not a value and there is no e' such that $e \mapsto e'$.*

Of course, whether an expression eventually reduces to a stuck expression is not a decidable property. We approximate the set of expressions that become stuck with a set of *faulty* expressions.

Definition 4.9 (Faulty Expressions) *The faulty expressions of Functional ML are the expressions containing a subexpression $(c\ v)$ where $\delta(c, v)$ is undefined.*

The idea is that any faulty expression *may* become stuck, *i.e.*, the property “reduces to a faulty expression” is a conservative approximation to “reduces to an expression that is stuck”. Thus the faulty expressions are a superset of the stuck expressions. For example, the expression $((\lambda y.2)\ (\lambda x.+ 1\ \text{true}))$ is faulty because of the subexpression $(+ 1\ \text{true})$, but is not stuck because it reduces to 2.

The behavior of evaluation is summarized by the following lemma. Let $e \uparrow$ mean that e diverges, *i.e.*, $e \mapsto e'$ for some e' , and for all e' such that $e \mapsto e'$, there exists e'' such that $e' \mapsto e''$.

Lemma 4.10 (Uniform Evaluation) *For closed e , if there is no e' such that $e \mapsto e'$ and e' is faulty, then either $e \uparrow$ or $e \mapsto v$.*

Proof. By induction on the length of the reduction sequence, we need only show that either e is faulty, $e \mapsto e'$ and e' is closed, or e is a value. From the definition of \mapsto , $e \mapsto e'$ iff $e = E[e_1]$, $e' = E[e'_1]$, and $e_1 \mapsto e'_1$. (Recall that $E ::= [] \mid E\ e \mid v\ E \mid \text{let } x\ \text{be } E\ \text{in } e$.) We proceed by induction on the structure of e .

Case $e = c, \mathsf{Y}$, or $\lambda x.e$. Then e is a value.

Case $e = x$. Since e is closed, this case cannot occur.

Case $e = \text{let } x\ \text{be } e_1\ \text{in } e_2$. By the inductive hypothesis with e_1 , there are three cases to consider. If e_1 is faulty, then $\text{let } x\ \text{be } e_1\ \text{in } e_2$ is faulty. If $e_1 = E_1[e']$ and $e' \mapsto e''$ then $e = E[e']$ where $E = \text{let } x\ \text{be } E_1\ \text{in } e_2$; thus $e \mapsto E[e'']$. Otherwise, if e_1 is a value, then $\text{let } x\ \text{be } e_1\ \text{in } e_2 \mapsto e_2[x \mapsto e_1]$.

Case $e = (e_1\ e_2)$. By the inductive hypothesis with e_1 , there are three cases to consider. If e_1 is faulty, then $(e_1\ e_2)$ is faulty. If $e_1 = E_1[e']$ and $e' \mapsto e''$ then $e = E[e']$ where $E = (E_1\ e_2)$; thus $e \mapsto E[e'']$. Otherwise, e_1 is a value. By the inductive hypothesis with e_2 , there are three subcases to consider. If e_2 is faulty, then $(e_1\ e_2)$ is faulty. If $e_2 = E_2[e']$ and $e' \mapsto e''$ then $e = E[e']$ where $E = (e_1\ E_2)$; thus $e \mapsto E[e'']$. Otherwise, both e_1 and e_2 are values, and the possibilities can be summarized as follows:

$e_1\ e_2$	c	$\lambda x.e$	Y	x
c	\mapsto or faulty	\mapsto or faulty	\mapsto or faulty	\times
$\lambda x.e$	\mapsto	\mapsto	\mapsto	\times
Y	\mapsto	\mapsto	\mapsto	\times
x	\times	\times	\times	\times

Cases marked with \times cannot occur because the expression is not closed. Cases marked with \mapsto indicate that the combination reduces with $E = []$. ■

To relate faulty expressions and typability, we show that no faulty expression is typable.

Main Lemma 4.11 (Faulty Expressions are Untypable) *If e is faulty, there are no Γ, τ such that $\Gamma \triangleright e : \tau$.*

Proof. It suffices to show that the subexpressions of e that cause e to be faulty are untypable. In general, we proceed by case analysis according to the form of the subexpression, but for Functional ML there is only one case.

Suppose $(c v)$ is faulty and $\Gamma \triangleright c v : \tau$. By **app**, $\Gamma \triangleright c : \tau' \rightarrow \tau$ and $\Gamma \triangleright v : \tau'$. Then by the δ -typability condition, $\delta(c, v)$ is defined, but this contradicts the assumption that $(c v)$ is faulty. ■

Since the faulty expressions are untypable and we have type preservation, the property “does not reduce to a faulty expression” is implied by typability.⁴

Theorem 4.12 (Syntactic Soundness) *If $\triangleright e : \tau$ then either $e \uparrow$ or $e \mapsto v$ and $\triangleright v : \tau$.*

Proof. By Uniform Evaluation, either $e \mapsto e'$ and e' is faulty, or $e \uparrow$, or $e \mapsto v$. Since $\triangleright e : \tau$, Subject Reduction implies $\triangleright v : \tau$ and $\triangleright e' : \tau$. Suppose $e \mapsto e'$ and e' is faulty. Since faulty expressions are untypable, $\triangleright e' : \tau$ is a contradiction, therefore this case cannot occur. Hence either $e \uparrow$ or $e \mapsto v$ and $\triangleright v : \tau$. ■

To state strong and weak soundness theorems, we must have a definition of evaluation that differentiates between programs that diverge and those that cause type errors. Let *answers* (a) be values or the special result `WRONG`, and define *eval'* as:

$$(eval') \quad eval'(e) = \begin{cases} \text{WRONG} & \text{if } e \mapsto e' \text{ and } e' \text{ is faulty;} \\ v & \text{if } e \mapsto v. \end{cases}$$

Programs that cause a type error return the special answer `WRONG`. Strong and weak soundness are now corollaries of Syntactic Soundness.

Theorem 4.13 (Strong Soundness) *If $\triangleright e : \tau$ and $eval'(e) = a$ then $\triangleright a : \tau$.*

Theorem 4.14 (Weak soundness) *If $\triangleright e : \tau$ then $eval'(e) \neq \text{WRONG}$.*

5 References and Exceptions

Besides functional facilities, first-class references and exceptions are the most important core programming facilities of STANDARD ML and similar languages. In this section we consider several extensions to Functional ML. The first subsection deals with Reference ML, an extension of Functional ML with references. It uses our previous work on a calculus of state, λ_v -S [12, 14], and in particular its cell-oriented variant [3]. The second subsection addresses Exception ML, an extension of Functional ML with exceptions. It uses a modified version of our control calculus, λ_v -C [13, 14], especially its fragment with prompts [9]. The final subsection presents Core ML, which includes both references and exceptions.

⁴Type preservation by itself does not guarantee type soundness. Consider the (trivial) type system where every type is a subtype of every other type. Type preservation holds since *every* expression can be assigned every type. However, the system does not prevent type errors, because it rejects no expressions.

5.1 Reference ML

Extending the λ_v -calculus to a calculus of functions and references requires extending the syntax with a new kind of expression and several new values (underlined):

$$\begin{array}{ll}
 (\text{Expressions}) & e ::= v \mid e_1 \ e_2 \mid \text{let } x \text{ be } e_1 \text{ in } e_2 \mid \underline{\rho\theta.e} \\
 (\text{Values}) & v ::= c \mid x \mid \mathbf{Y} \mid \lambda x.e \mid \underline{\text{ref}} \mid \underline{!} \mid \underline{:=} \mid \underline{:=} \ v \\
 & \underline{\theta} ::= \underline{\{ \langle x, v \rangle \}^*}
 \end{array}$$

The expression $\rho\langle x_1, v_1 \rangle \dots \langle x_n, v_n \rangle.e$ binds x_1, \dots, x_n in v_1, \dots, v_n and e . Above, θ represents a finite function from variables to values, *i.e.*, we treat θ as a set of pairs whose first components are distinct. We also identify all ρ -expressions that differ only by a consistent renaming of bound variables.

The values **ref**, **!**, and **:=** are the familiar operations of ML.⁵ The application of **ref** to a value creates a reference cell containing that value. The application of **!** to a cell returns the value contained in that cell. The binary assignment operator **:=** evaluates both its operands, the first of which must evaluate to a cell, and assigns the value of the second operand to that cell. Since all operations are carried, the application of the assignment operator to a variable *is* a value. Specifically, the expression $(:= c)$ may be thought of as a *capability* to assign to the cell c . The ρ -expression is semantically an abbreviation of a let-expression:

$$\begin{aligned}
 & \rho\langle x_1, v_1 \rangle \dots \langle x_n, v_n \rangle.e \\
 \equiv & \text{let } x_1 \text{ be ref } u_1 \text{ in} \\
 & \dots \\
 & \text{let } x_n \text{ be ref } u_n \text{ in} \\
 & (\lambda y_1 \dots y_n y_{n+1} \cdot y_{n+1}) (:= x_1 \ v_1) \dots (:= x_n \ v_n) \ e
 \end{aligned}$$

where u_1, \dots, u_n are arbitrary values (of the right type). In order to simplify the semantics, we do not treat it as such [14]. Intuitively, a ρ -expression plays the rôle of a store fragment during the reduction of a program with imperative assignment statements. A reference cell is represented as a variable; the pair $\langle x, v \rangle$ in a ρ -expression indicates that reference cell x contains value v .

5.1.1 Semantics

In addition to the reductions for **v**, the reductions for references are:⁶

$$\begin{array}{ll}
 (\text{ref}) & \text{ref } v \longrightarrow \rho\langle x, v \rangle.x \\
 (\text{deref}) & \rho\theta\langle x, v \rangle.R[! \ x] \longrightarrow \rho\theta\langle x, v \rangle.R[v] \\
 (\text{assign}) & \rho\theta\langle x, v_1 \rangle.R[:= \ x \ v_2] \longrightarrow \rho\theta\langle x, v_2 \rangle.R[v_2] \\
 (\rho_{\text{merge}}) & \rho\theta_1.\rho\theta_2.e \longrightarrow \rho\theta_1\theta_2.e \\
 (\rho_{\text{lift}}) & R[\rho\theta.e] \longrightarrow \rho\theta.R[e] \quad \text{if } R \neq []
 \end{array}$$

⁵Introducing **ref**, **!**, and **:=** as values rather than special forms like $(\text{ref } \epsilon)$, $(! \ \epsilon)$, and $(:= \ \epsilon_1 \ \epsilon_2)$ simplifies the proofs of several lemmas that proceed by induction on the structure of expressions. The latter approach requires duplicating much of the work of the case for applications.

⁶This semantics is due to Crank and Felleisen [3], which was derived from Felleisen and Hieb's work [14]. A similar definition of evaluation for dealing with state using rewriting techniques also appears in the work of Mason and Talcott [19].

By the variable conventions, x is not free in v in the *ref* reduction; the domains of θ_1 and θ_2 are disjoint in ρ_{merge} ; and the free variables of R are disjoint from the domain of θ in ρ_{lift} . We refer to the union of these five reductions as \mathbf{r} . The notions of reduction \mathbf{r} and \mathbf{v} (adapted *mutatis mutandis* to the full syntax) form the basis for a calculus of functions and references. We refer to the union as \mathbf{vr} , and write $\xrightarrow{\mathbf{vr}}$ for a reduction in \mathbf{vr} . The extended notion of reduction gives rise to a system of reductions and equations with the extension of contexts to:

$$C ::= [] \mid C \ e \mid e \ C \mid \text{let } x \text{ be } C \text{ in } e \mid \text{let } x \text{ be } e \text{ in } C \mid \lambda x.C \mid \rho\theta.C \mid \rho\theta\langle x, C \rangle.e.$$

The calculus satisfies the same basic properties as the λ_v -calculus [3].

The definition of \mathbf{r} relies on a set of *R contexts*:

$$R ::= [] \mid R \ e \mid v \ R \mid \text{let } x \text{ be } R \text{ in } e.$$

The use of *R contexts* in the new reductions reflects the additional sequencing restrictions that the introduction of side-effects in a programming language requires for the semantics to be deterministic. In particular, the creation, the dereferencing, and the updating operations on a reference cell must be ordered in a linear fashion, which implies some further ordering among the operations on distinct cells. The *R contexts* build this minimal order of evaluation into those reductions that refer to reference cells and their operations; following ML, we choose to perform operations from left to right.

For \mathbf{r} reductions, a cell is represented by an ordinary variable appearing in the domain of a store fragment θ . When a program creates a new cell via **ref**, the *ref* reduction introduces a ρ -expression that records the cell's current value. The dereference operator $!$ selects a cell's value from the *closest* ρ -expression (relative to *R contexts*). If the appropriate cell is bound in a ρ -expression that is not the closest one, the intervening ρ -expressions must first be merged with the outer one. This is accomplished with ρ_{lift} and ρ_{merge} reductions, which lift a partial store out of *R contexts*, and merge it with outer stores (by the variable convention, cells are renamed appropriately to avoid collisions). An assignment replaces a cell's current value in the closest ρ -expression, after performing all necessary lift and merge steps, and returns the assigned value.

For example, consider the following expression:

$$\begin{aligned} & ((\text{let } x \text{ be ref } (\lambda x. + !x \ 2) \text{ in } := x \ (\lambda x. + !x \ 3)) \\ & \quad (\text{let } x \text{ be ref } 2 \text{ in } (\lambda z.x)(:= x \ 4))) \end{aligned}$$

In this expression, the two assignments can happen in an arbitrary order, and the system of reductions admits both possibilities. Both assignments, however, must happen before the outermost application is reduced. Figure 2 gives one possible reduction sequence.

The above calculus only constrains the relative ordering of *ref*, *deref*, and *assign* reductions. By defining extended evaluation contexts:

$$E ::= [] \mid E \ e \mid v \ E \mid \text{let } x \text{ be } E \text{ in } e \mid \rho\theta.E$$

we fix the order in which reductions in \mathbf{v} and *ref*, *deref*, and *assign* may be applied. The stepping function $\xrightarrow{\mathbf{vr}}$ is defined as before. To show that $\xrightarrow{\mathbf{vr}}$ is a function requires a simple

$$\begin{array}{l}
((\text{let } x \text{ be ref } (\lambda x. + !x 2) \text{ in } := x (\lambda x. + !x 3)) \\
(\text{let } x \text{ be ref } 2 \text{ in } (\lambda z. x)(:= x 4))) \\
\longrightarrow (\rho\langle x, \lambda x. + !x 2 \rangle := x (\lambda x. + !x 3)) (\rho\langle x, 2 \rangle. (\lambda z. x)(:= x 4)) \quad 2 \times \text{ref}, \rho_{\text{lift}}, \text{let} \\
\longrightarrow (\rho\langle x, \lambda x. + !x 3 \rangle. \lambda x. + !x 3) (\rho\langle x, 2 \rangle. (\lambda z. x)(:= x 4)) \quad \text{assign} \\
\longrightarrow (\rho\langle x, \lambda x. + !x 3 \rangle. \lambda x. + !x 3) (\rho\langle x, 4 \rangle. x) \quad \text{assign}, \beta_v \\
\longrightarrow \rho\langle x, \lambda x. + !x 3 \rangle\langle y, 4 \rangle. (\lambda x. + !x 3) y \quad 2 \times \rho_{\text{lift}}, \rho_{\text{merge}} \\
\longrightarrow \rho\langle x, \lambda x. + !x 3 \rangle\langle y, 4 \rangle. + !y 3 \quad \beta_v \\
\longrightarrow \rho\langle x, \lambda x. + !x 3 \rangle\langle y, 4 \rangle. + 4 3 \quad \text{deref} \\
\longrightarrow \rho\langle x, \lambda x. + !x 3 \rangle\langle y, 4 \rangle. 7 \quad \delta
\end{array}$$

Figure 2: Reducing a program with side-effects

diamond theorem, since the order in which ρ_{lift} and ρ_{merge} reductions happen is not fixed, and the context R in the ρ_{lift} rule is selected in a nondeterministic fashion. The answers returned by evaluation of a program are no longer simply closed values, since an answer may be a reference cell. Answers are values or ρ -expressions enclosing values:

$$(\text{answers}) \quad a ::= v \mid \rho\theta.v.$$

Based on these definitions, evaluation may be defined like before:

$$(\text{eval}_{\mathbf{vr}}) \quad \text{eval}_{\mathbf{vr}}(e) = a \quad \text{iff} \quad e \xrightarrow{\mathbf{vr}} a.$$

5.1.2 Typing

Typing reference cells in the presence of polymorphism is not straightforward, as the obvious solution is unsound [6, 36, 37]. To assign types to reference cells, we extend the set of types with an additional constructor:

$$\tau ::= \iota_1 \mid \dots \mid \iota_n \mid \alpha \mid \tau_1 \rightarrow \tau_2 \mid \underline{\tau \text{ ref}}$$

The type $\tau \text{ ref}$ is the type of cells containing a value of type τ .

The obvious types for the three operators are indicated by their semantics:

$$\begin{array}{l}
\Gamma \triangleright \text{ref} : \tau \rightarrow \tau \text{ ref} \\
\Gamma \triangleright ! : \tau \text{ ref} \rightarrow \tau \\
\Gamma \triangleright := : \tau \text{ ref} \rightarrow \tau \rightarrow \tau
\end{array}$$

For example, the first argument of the assignment operator ($:=$) must be a reference cell. The second argument must be a value matching the type of the cell. The result is of the same type as the second argument since our assignment operator returns the value assigned. The following expression illustrates why this typing is unsound:

$$\begin{array}{l}
\text{let } f \text{ be ref } (\lambda y. y) \text{ in} \\
(\lambda z. ! f \text{ true}) (:= f (\lambda n. + 1 n))
\end{array}$$

If the type of the reference bound to x is generalized to $\forall\alpha.(\alpha \rightarrow \alpha)$ *ref*, then the reference cell can be treated as having type $(int \rightarrow int)$ *ref* for the assignment, and as having type $(bool \rightarrow bool)$ *ref* for the dereference, hence this expression is typable as *bool*. But evaluating this expression leads to a type error:

$$\begin{aligned}
& \text{let } f \text{ be ref } (\lambda y.y) \text{ in} \\
& \quad (\lambda z. ! f \text{ true}) (:= f (\lambda n.+ 1 n)) \\
\longrightarrow & \text{let } f \text{ be } \rho\langle x, \lambda y.y \rangle. x \text{ in} \\
& \quad (\lambda z. ! f \text{ true}) (:= f (\lambda n.+ 1 n)) \\
\longrightarrow & \rho\langle x, \lambda y.y \rangle. \\
& \quad (\lambda z. ! x \text{ true}) (:= x (\lambda n.+ 1 n)) \\
\longrightarrow & \rho\langle x, \lambda n.+ 1 n \rangle. ! x \text{ true} \\
\longrightarrow & \rho\langle x, \lambda n.+ 1 n \rangle. + 1 \text{ true}
\end{aligned}$$

As Tofte [36, 37] points out, the problem with the obvious typing is the generalization of type variables that appear free in the type of a value in the store.

The central idea of Tofte's solution is to ensure that the only storable values are those that will not be used polymorphically at run-time. To this end, type variables are partitioned into *imperative* and *applicative* variables: the former set of type variables is named *ImpTypeVar*; the latter *AppTypeVar*. Types are classified accordingly: an *imperative type* cannot contain any applicative type variables. Only values of imperative type can be stored. When the type of a value is generalized in a *let*-expression, type variables that might appear in the types of values in the store are required to be imperative, and are not generalized.

The notion of generalization requires appropriate adaptation. We still have:

$$(\succ) \quad \forall\alpha_1 \dots \alpha_n. \tau' \succ \tau \text{ if } S\tau' = \tau \text{ and } \text{Dom}(S) = \{\alpha_1, \dots, \alpha_n\} \text{ for some } S,$$

but the substitution S is required to map imperative type variables to imperative types. Thus, if $\underline{\alpha}$ is an imperative type variable, we have:

$$\begin{aligned}
& \forall \underline{\alpha}. \underline{\alpha} \rightarrow \underline{\alpha} \succ int \rightarrow int \\
& \text{and } \forall \underline{\alpha}. \underline{\alpha} \rightarrow \underline{\alpha} \succ \underline{\alpha} \rightarrow \underline{\alpha} \\
& \text{but not } \forall \underline{\alpha}. \underline{\alpha} \rightarrow \underline{\alpha} \succ \alpha \rightarrow \alpha
\end{aligned}$$

since α is potentially applicative (the meta-variable α still represents *any* type variable—imperative or applicative).

Like all other techniques for typing references, Tofte's system requires modifying the typing rule for *let*-expressions. The *let* rule is split into two rules, according to whether or not the right-hand side of the declaration is a value. Figure 3 gives the new rules. The second rule does not generalize over imperative type variables, and thus will not generalize the type of a value in the store. However, if the expression bound by a *let*-expression is a value, as in the first rule, then its evaluation cannot create a new cell, so generalization of imperative type variables in its type cannot generalize the type of a value in the store. Tofte uses the terminology *expansive* and *non-expansive* to denote this syntactic classification of expressions into those that may create new references, and those that definitely do not create new references. Like Tofte, we classify only values as non-expansive, but a stronger type system is possible by classifying more expressions as non-expansive [36, 37].

$$\begin{array}{l}
(\mathbf{let}_v) \quad \frac{\Gamma \triangleright v : \tau_1 \quad \Gamma[x \mapsto \mathit{Close}(\tau_1, \Gamma)] \triangleright e : \tau_2}{\Gamma \triangleright \mathbf{let } x \mathbf{ be } v \mathbf{ in } e : \tau_2} \\
(\mathbf{let}_e) \quad \frac{\Gamma \triangleright e_1 : \tau_1 \quad \Gamma[x \mapsto \mathit{AppClose}(\tau_1, \Gamma)] \triangleright e_2 : \tau_2 \quad e_1 \notin \mathit{Values}}{\Gamma \triangleright \mathbf{let } x \mathbf{ be } e_1 \mathbf{ in } e_2 : \tau_2} \\
\mathit{Close}(\tau, \Gamma) = \forall \alpha_1 \dots \alpha_n. \tau \text{ where } \{\alpha_1, \dots, \alpha_n\} = \mathit{FTV}(\tau) \setminus \mathit{FTV}(\Gamma) \\
\mathit{AppClose}(\tau, \Gamma) = \forall \alpha_1 \dots \alpha_n. \tau \text{ where } \{\alpha_1, \dots, \alpha_n\} = (\mathit{FTV}(\tau) \setminus \mathit{FTV}(\Gamma)) \cap \mathit{AppTypeVar}.
\end{array}$$

Figure 3: Modified let-expression typing rules

$$\begin{array}{l}
(\mathbf{ref}) \quad \Gamma \triangleright \mathbf{ref} : \tau \rightarrow \tau \mathit{ref} \text{ if } \tau \text{ is imperative} \\
(\mathbf{deref}) \quad \Gamma \triangleright ! : \tau \mathit{ref} \rightarrow \tau \\
(\mathbf{assign}) \quad \Gamma \triangleright := : (\tau \mathit{ref} \rightarrow \tau \rightarrow \tau) \\
(\mathbf{rho}) \quad \frac{\Gamma[x_1 \mapsto \tau_1 \mathit{ref}] \dots [x_n \mapsto \tau_n \mathit{ref}] \triangleright e : \tau \quad \Gamma[x_1 \mapsto \tau_1 \mathit{ref}] \dots [x_n \mapsto \tau_n \mathit{ref}] \triangleright v_i : \tau_i \quad \tau_i \text{ is imperative } \quad 1 \leq i \leq n}{\Gamma \triangleright \rho \langle x_1, v_1 \rangle \dots \langle x_n, v_n \rangle . e : \tau}
\end{array}$$

Figure 4: Additional typing rules for references

Figure 4 gives four additional typing rules for reference cells. The typing rules for **ref** and ρ -expressions ensure that any value placed in the store has an imperative type. The typing rules for **!** and **:=** do not need to be explicitly constrained, since their use is implicitly constrained by the type of a value already in the store.

In addition to the restrictions placed on the typing of constants by Functional ML, we also require that there be no constants of reference type ($\tau \mathit{ref}$). This ensures that all values of reference type are in fact reference cells, and can be assigned or dereferenced. This restriction is used in the proof of Theorem 5.6.

The new system assigns the same types as the old one to Functional ML expressions. Let $\overset{f}{\triangleright}$ indicate provability in the old, functional system, and let $\overset{r}{\triangleright}$ indicate provability in the system with references.

Proposition 5.1 *If e is an expression in Functional ML, and Γ and τ are purely applicative, then $\Gamma \overset{f}{\triangleright} e : \tau$ iff $\Gamma \overset{r}{\triangleright} e : \tau$.*

Since Functional ML expressions do not contain **ref** or ρ -expressions, no imperative type variables need be chosen to type an expression, and Close and $\mathit{AppClose}$ yield the same type scheme for applicative types.

5.1.3 Type Soundness for References

To prove type preservation for the extended language, we must show that the notion of reduction **v** extended to the new syntax preserves types, and that the reductions of **r**

preserve types. Showing that the extended \mathbf{v} preserves types amounts to re-proving the various lemmas used in section 4.3 for the extended system. In general, we refer to the extension of Lemma 4. n as Lemma 4. n ⁺. Showing that \mathbf{r} preserves types involves adding a case for each reduction to the proof of Subject Reduction. The cases involving R contexts use the fact that:

$$\text{if } \Gamma \triangleright R[e] : \tau \text{ then } \Gamma \triangleright e : \tau'.$$

The type environment Γ is the same in both the antecedent and consequent since R contexts do not bind variables.

Main Lemma 5.2 (Subject Reduction for \mathbf{vr}) *If $\Gamma \triangleright^r e_1 : \tau$ and $e_1 \xrightarrow{\mathbf{vr}} e_2$ then $\Gamma \triangleright^r e_2 : \tau$.*

Proof. The cases for the reductions of \mathbf{v} are the same as before, using an appropriately extended Replacement Lemma and a similarly extended Lemma 4.1.

Case $\mathbf{ref} v \longrightarrow \rho\langle x, v \rangle.x$. From $\Gamma \triangleright \mathbf{ref} v : \tau$, where $\tau = \tau' \mathit{ref}$ and τ' is imperative, $\Gamma \triangleright v : \tau'$ by **ref** and **app**, and $\Gamma[x \mapsto \tau' \mathit{ref}] \triangleright v : \tau'$ by Lemma 4.1⁺. Since $\Gamma[x \mapsto \tau' \mathit{ref}] \triangleright x : \tau' \mathit{ref}$ by **var**, $\Gamma \triangleright \rho\langle x, v \rangle.x : \tau' \mathit{ref}$ by **rho**.

In order to satisfy the imperative restriction on the typing of ρ -expressions, **ref** applications must yield values of imperative type, as ensured by **ref**.

Case $\rho\theta\langle x, v \rangle.R[! x] \longrightarrow \rho\theta\langle x, v \rangle.R[v]$. From $\Gamma \triangleright \rho\theta\langle x, v \rangle.R[! x] : \tau$ by **rho**

$$(12) \quad \Gamma' \triangleright v : \tau'$$

where τ' is imperative and $\Gamma'x = \tau' \mathit{ref}$. Thus $\Gamma' \triangleright x : \tau' \mathit{ref}$, hence $\Gamma' \triangleright ! x : \tau'$ by **deref** and **app**. With (12) by the Replacement⁺ Lemma, we have $\Gamma \triangleright \rho\theta\langle x, v \rangle.R[v] : \tau$.

Case $\rho\theta\langle x, v_1 \rangle.R[:= x v_2] \longrightarrow \rho\theta\langle x, v_2 \rangle.R[v_2]$. From $\Gamma \triangleright \rho\theta\langle x, v_1 \rangle.R[:= x v_2] : \tau$ by **rho**

$$(13) \quad \Gamma' \triangleright v_1 : \tau'$$

where τ' is imperative and $\Gamma'x = \tau' \mathit{ref}$. Thus $\Gamma' \triangleright x : \tau' \mathit{ref}$, and by **assign** and **app**

$$\begin{aligned} \Gamma' \triangleright := x : \tau' &\rightarrow \tau' \\ \Gamma' \triangleright := x v_2 : \tau' &\text{ and } \Gamma' \triangleright v_2 : \tau' && \text{by } \mathbf{app}. \end{aligned}$$

With (13) and using the Replacement⁺ Lemma twice, we have $\Gamma \triangleright \rho\theta\langle x, v_2 \rangle.R[v_2] : \tau$.

Case $\rho\theta_1.\rho\theta_2.e \longrightarrow \rho\theta_1\theta_2.e$. From $\Gamma \triangleright \rho\theta_1.\rho\theta_2.e : \tau$ by **rho**

$$(14) \quad \Gamma' \triangleright \rho\theta_2.e : \tau \quad \text{and} \quad \Gamma' \triangleright v_i : \tau_i$$

where $\theta_1 = \langle x_1, v_1 \rangle \dots \langle x_n, v_n \rangle$, $\Gamma' = \Gamma[x_1 \mapsto \tau_1 \mathit{ref}] \dots [x_n \mapsto \tau_n \mathit{ref}]$, and $1 \leq i \leq n$. Again by **rho**

$$(15) \quad \Gamma'' \triangleright e : \tau \quad \text{and} \quad \Gamma'' \triangleright v'_j : \tau'_j$$

where $\theta_2 = \langle x'_1, v'_1 \rangle \dots \langle x'_m, v'_m \rangle$, $\Gamma'' = \Gamma'[x'_1 \mapsto \tau'_1 \mathit{ref}] \dots [x'_m \mapsto \tau'_m \mathit{ref}]$, and $1 \leq j \leq m$. Since $\text{Dom}(\theta_1) \cap \text{Dom}(\theta_2) = \emptyset$ (by the variable conventions),

$$\begin{aligned} \Gamma'' \triangleright v_i : \tau_i &\text{ for } 1 \leq i \leq n && \text{by (14),} \\ \Gamma \triangleright \rho\theta_1\theta_2.e : \tau &&& \text{by } \mathbf{rho} \text{ with (15).} \end{aligned}$$

Case $R[\rho\theta.e] \longrightarrow \rho\theta.R[e]$. We have $\Gamma \triangleright R[\rho\theta.e] : \tau$, and proceed by induction on the structure of R to show $\Gamma \triangleright \rho\theta.R[e] : \tau$.

Case $R = []$. Then $R[\rho\theta.e] = \rho\theta.R[e]$.

Case $R = (R' e')$. Then $\Gamma \triangleright (R'[\rho\theta.e] e') : \tau$, and by **app**

$$(16) \quad \Gamma \triangleright R'[\rho\theta.e] : \tau' \rightarrow \tau$$

$$(17) \quad \text{and } \Gamma \triangleright e' : \tau'.$$

Then $\Gamma \triangleright \rho\theta.R'[e] : \tau' \rightarrow \tau$ by ind. hyp. with (16), and by **rho**

$$(18) \quad \Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright v_i : \tau_i$$

$$(19) \quad \text{and } \Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright R'[e] : \tau' \rightarrow \tau$$

where $\theta = \langle x_1, v_1 \rangle \dots \langle x_n, v_n \rangle$ and $1 \leq i \leq n$. Since $x_1, \dots, x_n \notin FV(e')$,

$$\Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright e' : \tau' \quad \text{by Lemma 4.1}^+ \text{ with (17),}$$

$$\Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright R'[e] e' : \tau \quad \text{by } \mathbf{app} \text{ with (19),}$$

$$\Gamma \triangleright \rho\theta.R'[e] e' : \tau \quad \text{by } \mathbf{rho} \text{ with (18).}$$

Case $R = (v R')$. Similar to the previous case.

Case $R = \text{let } x \text{ be } R' \text{ in } e'$. Then $\Gamma \triangleright \text{let } x \text{ be } R'[\rho\theta.e] \text{ in } e' : \tau$, and by **let_e**

$$(20) \quad \Gamma \triangleright R'[\rho\theta.e] : \tau'$$

$$(21) \quad \text{and } \Gamma[x \mapsto \text{AppClose}(\tau', \Gamma)] \triangleright e' : \tau.$$

Then $\Gamma \triangleright \rho\theta.R'[e] : \tau'$ by ind. hyp. with (20), and by **rho**

$$(22) \quad \Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright v_i : \tau_i$$

$$(23) \quad \text{and } \Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright R'[e] : \tau'$$

where $\theta = \langle x_1, v_1 \rangle \dots \langle x_n, v_n \rangle$ and each τ_i is imperative for $1 \leq i \leq n$. Since $x_1, \dots, x_n \notin FV(e')$,

$$\Gamma[x \mapsto \text{AppClose}(\tau', \Gamma)][x_i \mapsto \tau_i \text{ ref}] \triangleright e' : \tau \quad \text{by Lemma 4.1}^+ \text{ with (21)}$$

$$\text{i.e. } \Gamma[x_i \mapsto \tau_i \text{ ref}][x \mapsto \text{AppClose}(\tau', \Gamma)] \triangleright e' : \tau.$$

But each τ_i is imperative, due to the restriction on the typing of ρ -expressions, hence

$$\text{AppClose}(\tau', \Gamma) = \text{AppClose}(\tau', \Gamma[x_i \mapsto \tau_i \text{ ref}]).$$

If $R'[e]$ is expansive, then

$$\Gamma[x_i \mapsto \tau_i \text{ ref}][x \mapsto \text{AppClose}(\tau', \Gamma[x_i \mapsto \tau_i \text{ ref}])] \triangleright e' : \tau$$

$$\Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright \text{let } x \text{ be } R'[e] \text{ in } e' : \tau \quad \text{by } \mathbf{let}_e \text{ with (23)}$$

$$\Gamma \triangleright \rho\theta.\text{let } x \text{ be } R'[e] \text{ in } e' : \tau \quad \text{by } \mathbf{rho} \text{ with (22).}$$

Otherwise, $R'[e]$ is non-expansive, and

$$\Gamma[x_i \mapsto \tau_i \text{ ref}][x \mapsto \text{Close}(\tau', \Gamma[x_i \mapsto \tau_i \text{ ref}])] \triangleright e' : \tau$$

$$\Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright \text{let } x \text{ be } R'[e] \text{ in } e' : \tau \quad \text{by } \mathbf{let}_v \text{ with (23)}$$

$$\Gamma \triangleright \rho\theta.\text{let } x \text{ be } R'[e] \text{ in } e' : \tau \quad \text{by } \mathbf{rho} \text{ with (22).}$$

Thus $\Gamma \triangleright \rho\theta.R[e] : \tau$ by induction. ■

It remains to extend the proofs of the various lemmas. The only one that is not straightforward is the Substitution Lemma. As in the definition of generalization, substitutions S in the following proof are required to map imperative type variables to imperative types.

Lemma 5.3 (Substitution) *If $\Gamma[x \mapsto \forall\alpha_1 \dots \alpha_n.\tau] \triangleright^r e : \tau'$ and $x \notin \text{Dom}(\Gamma)$ and $\Gamma \triangleright^r v : \tau$ and $\{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\Gamma) = \emptyset$ then $\Gamma \triangleright^r e[x \mapsto v] : \tau'$.*

Proof. The cases from the proof for Functional ML are unchanged, with the exception that the **let**-expression case applies only when the bound expression is a value. There is one new case for expansive **let**-expressions, and new cases for the additional syntax.

Case $e = \text{ref}, !$, or $:=$. In each case, $\Gamma' \triangleright e : \tau'$ for any Γ' , hence $\Gamma \triangleright e : \tau'$. Also in each case, $e[x \mapsto v] = e$, so $\Gamma \triangleright e[x \mapsto v] : \tau'$.

Case $e = \rho\theta.e_1$. From $\Gamma[x \mapsto \forall\alpha_1 \dots \alpha_n.\tau] \triangleright \rho\theta.e_1 : \tau'$ by **rho**

$$\begin{aligned} & \Gamma[x \mapsto \forall\alpha_1 \dots \alpha_n.\tau][x_i \mapsto \tau_i \text{ ref}] \triangleright v_i : \tau_i \\ & \text{and } \Gamma[x \mapsto \forall\alpha_1 \dots \alpha_n.\tau][x_i \mapsto \tau_i \text{ ref}] \triangleright e_1 : \tau' \end{aligned}$$

where $\theta = \langle x_1, v_1 \rangle \dots \langle x_m, v_m \rangle$, x_1, \dots, x_m, x are distinct, and each τ_i is imperative for $1 \leq i \leq m$. Choose a substitution $S : \{\alpha_1, \dots, \alpha_n\} \rightarrow \{\alpha'_1, \dots, \alpha'_n\}$ such that $\alpha'_1, \dots, \alpha'_n$ are distinct type variables, $S\alpha_i = \alpha'_i$, and $\{\alpha'_1, \dots, \alpha'_n\} \cap (\{\alpha_1, \dots, \alpha_n\} \cup \text{FTV}(\Gamma) \cup \text{FTV}(\tau') \cup \bigcup_i \text{FTV}(\tau_i)) = \emptyset$. Then by Lemma 4.5⁺

$$\begin{aligned} & S(\Gamma[x \mapsto \forall\alpha_1 \dots \alpha_n.\tau][x_i \mapsto \tau_i \text{ ref}]) \triangleright v_i : S\tau_i \\ \text{i.e. } & \Gamma[x \mapsto \forall\alpha_1 \dots \alpha_n.\tau][x_i \mapsto S\tau_i \text{ ref}] \triangleright v_i : S\tau_i \\ \text{i.e. } & \Gamma[x_i \mapsto S\tau_i \text{ ref}][x \mapsto \forall\alpha_1 \dots \alpha_n.\tau] \triangleright v_i : S\tau_i. \end{aligned}$$

Also by Lemma 4.5⁺,

$$\begin{aligned} & S(\Gamma[x \mapsto \forall\alpha_1 \dots \alpha_n.\tau][x_i \mapsto \tau_i \text{ ref}]) \triangleright e_1 : S\tau' \\ \text{i.e. } & \Gamma[x \mapsto \forall\alpha_1 \dots \alpha_n.\tau][x_i \mapsto S\tau_i \text{ ref}] \triangleright e_1 : S\tau' \\ \text{i.e. } & \Gamma[x_i \mapsto S\tau_i \text{ ref}][x \mapsto \forall\alpha_1 \dots \alpha_n.\tau] \triangleright e_1 : S\tau'. \end{aligned}$$

By applying the inductive hypothesis to each,

$$\begin{aligned} & \Gamma[x_i \mapsto S\tau_i \text{ ref}] \triangleright v_i[x \mapsto v] : S\tau_i \\ & \text{and } \Gamma[x_i \mapsto S\tau_i \text{ ref}] \triangleright e_1[x \mapsto v] : S\tau'. \end{aligned}$$

Since S is a bijection, S^{-1} exists; hence

$$\begin{aligned} & \Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright v_i[x \mapsto v] : \tau_i \\ & \text{and } \Gamma[x_i \mapsto \tau_i \text{ ref}] \triangleright e_1[x \mapsto v] : \tau'. \end{aligned}$$

Then $\Gamma \triangleright \rho\theta[x \mapsto v].e_1[x \mapsto v] : \tau'$ by **rho**, thus $\Gamma \triangleright (\rho\theta.e_1)[x \mapsto v] : \tau'$.

Case $e = \text{let } x' \text{ be } e_1 \text{ in } e_2$ where $e_1 \notin \text{Values}$. By the first premise of \mathbf{let}_e

$$(24) \quad \begin{array}{l} \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_1 : \tau_1 \\ \Gamma \triangleright e_1[x \mapsto v] : \tau_1 \end{array} \quad \text{by ind. hyp.}$$

By the second premise of \mathbf{let}_e

$$\begin{array}{l} \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau][x' \mapsto \text{AppClose}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])] \triangleright e_2 : \tau' \\ \text{i.e. } \Gamma[x' \mapsto \text{AppClose}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])][x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_2 : \tau'. \end{array}$$

Choose a substitution $S : \{\alpha_1, \dots, \alpha_n\} \cap \text{ImpTypeVar} \rightarrow \{\alpha'_1, \dots, \alpha'_m\}$ such that $\alpha'_1, \dots, \alpha'_m$ are distinct imperative type variables, S is a bijection, and $\{\alpha'_1, \dots, \alpha'_m\} \cap (\text{FTV}(\Gamma) \cup \text{FTV}(\tau) \cup \{\alpha_1, \dots, \alpha_n\}) = \emptyset$. Then by Lemma 4.5⁺

$$(25) \quad \begin{array}{l} S(\Gamma[x' \mapsto \text{AppClose}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])][x \mapsto \forall \alpha_1 \dots \alpha_n. \tau]) \triangleright e_2 : S\tau' \\ \text{i.e. } \Gamma[x' \mapsto S\text{AppClose}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])][x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_2 : S\tau' \\ \text{i.e. } \Gamma[x' \mapsto \text{AppClose}(S\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])][x \mapsto \forall \alpha_1 \dots \alpha_n. \tau] \triangleright e_2 : S\tau' \end{array}$$

since the domain of S contains only imperative type variables. From $\Gamma \triangleright v : \tau$ since x' is not free in v

$$(26) \quad \Gamma[x' \mapsto \text{AppClose}(S\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])] \triangleright v : \tau.$$

Now,

$$\begin{aligned} & \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\Gamma[x' \mapsto \text{AppClose}(S\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])]) \\ & \subseteq \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(\Gamma) \cup \text{FTV}(\text{AppClose}(S\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau]))) \\ & = \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\text{AppClose}(S\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])) \\ & = \{\alpha_1, \dots, \alpha_n\} \cap \text{FTV}(\forall \alpha'_1 \dots \alpha'_m. S\tau_1) \\ & = \{\alpha_1, \dots, \alpha_n\} \cap (\text{FTV}(S\tau_1) \setminus \{\alpha''_1, \dots, \alpha''_m\}) \\ & = \emptyset \end{aligned}$$

where $\{\alpha''_1, \dots, \alpha''_m\} = (\text{FTV}(S\tau_1) \setminus \text{FTV}(\Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])) \cap \text{AppTypeVar}$. By the inductive hypothesis with (25) and (26), it follows that

$$\Gamma[x' \mapsto \text{AppClose}(S\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])] \triangleright e_2[x \mapsto v] : S\tau'.$$

Since S is a bijection, S^{-1} exists, so by Lemma 4.5⁺

$$\begin{array}{l} S^{-1}(\Gamma[x' \mapsto \text{AppClose}(S\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])]) \triangleright e_2[x \mapsto v] : S^{-1}(S\tau') \\ \text{i.e. } \Gamma[x' \mapsto \text{AppClose}(\tau_1, \Gamma[x \mapsto \forall \alpha_1 \dots \alpha_n. \tau])] \triangleright e_2[x \mapsto v] : \tau'. \end{array}$$

Then by Lemma 4.6⁺

$$\Gamma[x' \mapsto \text{AppClose}(\tau_1, \Gamma)] \triangleright e_2[x \mapsto v] : \tau'$$

and by \mathbf{let}_e with (24)

$$\begin{array}{l} \Gamma \triangleright \text{let } x \text{ be } e_1[x \mapsto v] \text{ in } e_2[x \mapsto v] : \tau' \\ \text{i.e. } \Gamma \triangleright (\text{let } x \text{ be } e_1 \text{ in } e_2)[x \mapsto v] : \tau'. \end{array} \quad \blacksquare$$

The evaluation of an expression in Reference ML can become stuck for several new reasons. The expanded definition of faulty expressions reflects this fact.

Definition 5.4 (Faulty Expressions) *The faulty expressions of Reference ML are those expressions containing a subexpression of the form:*

$$\begin{aligned} & (c \ v) \text{ where } \delta(c, v) \text{ is undefined;} \\ & (! \ v) \text{ where } v \notin \text{Var;} \\ & (:= \ v) \text{ where } v \notin \text{Var;} \text{ or} \\ & \rho\theta\langle x, v_2 \rangle.C[x \ v_1]. \end{aligned}$$

As before, we have a Uniform Evaluation Lemma, stating that programs either yield an answer, diverge, or reduce to a faulty expression.

Lemma 5.5 (Uniform Evaluation) *For closed e , if there is no e' such that $e \xrightarrow{\text{vr}} e'$ and e' is faulty, then either $e \uparrow$ or $e \xrightarrow{\text{vr}} a$.*

Proof. The proof proceeds by induction on the length of the reduction sequence, and is similar to the proof for Functional ML. \blacksquare

All of the expressions in the expanded definition are untypable.

Main Lemma 5.6 (Faulty Expressions are Untypable) *If e is faulty, there are no Γ and τ such that $\Gamma \overset{\tau}{\triangleright} e : \tau$.*

Proof. Again, it suffices to show that the subexpressions of e that cause e to be faulty are untypable. We proceed by case analysis according to the form of the subexpression, assuming for each case that the expression can be typed, and deriving a contradiction.

Case $(c \ v)$ where $\delta(c, v)$ is undefined. As for Functional ML.

Case $(! \ v)$ where $v \notin \text{Var}$. Assume that $\Gamma \triangleright ! \ v : \tau$. Then $\Gamma \triangleright v : \tau$ *ref* by **app** and **deref**, which implies v is a variable (since no other value can have type τ *ref*), contrary to the assumption.

Case $(:= \ v)$ where $v \notin \text{Var}$. Assume that $\Gamma \triangleright := \ v : \tau$. Then $\Gamma \triangleright v : \tau'$ *ref* where $\tau = \tau' \rightarrow \tau'$ by **app** and **assign**, which implies v is a variable, contrary to the assumption.

Case $\rho\theta\langle x, v' \rangle.C[x \ v]$. Assume that $\Gamma \triangleright \rho\theta\langle x, v' \rangle.C[x \ v] : \tau$. Then $\Gamma' \triangleright C[x \ v] : \tau$ by **rho**, where $\Gamma'(x) = \tau'$ *ref*. Since $\Gamma'' \triangleright x \ v : \tau_2$ where $\Gamma'(x) = \Gamma''(x)$, $\Gamma'' \triangleright x : \tau_1 \rightarrow \tau_2$ by **app**, hence $\Gamma''(x) \succ \tau_1 \rightarrow \tau_2$. This contradicts $\Gamma''(x) = \tau'$ *ref*. \blacksquare

Type soundness now follows as before.

Theorem 5.7 (Syntactic Soundness for Reference ML) *If $\overset{\tau}{\triangleright} e : \tau$ then either $e \uparrow$ or $e \xrightarrow{\text{vr}} a$ and $\overset{\tau}{\triangleright} a : \tau$.*

Proof. Exactly as for Theorem 4.12, using the appropriately extended lemmas. \blacksquare

5.2 Exception ML

Due to δ -typability, constant functions must be defined for *all* values of their input type. This precludes constants, such as integer or real division, that are defined on all but a few recognizable input values of correct type. STANDARD ML solves this problem by introducing exceptions.

Our exception extension provides named exceptions with parameters, a means of raising exceptions, and a means of handling exceptions. Our extension thus includes three new phrases:

$$\text{exception } x_1 \dots x_n \text{ in } e, \quad \text{raise } e_1 e_2, \quad e_3 \text{ handle } e_1 e_2.$$

The first phrase declares x_1, \dots, x_n to be exception names lexically bound in e . The second phrase requires e_1 to evaluate to an exception name, and raises that exception with a parameter, the value of e_2 . The parameter propagates along with the exception, and is used at the site where an exception is handled. The third phrase evaluates its subexpressions e_1 and e_2 *first*, and installs the value of e_2 as an exception handler for e_1 exceptions that are raised during the evaluation of e_3 . The subexpression e_1 must evaluate to an exception name, and e_2 must evaluate to a function that accepts parameters of e_1 exceptions. When an exception is raised, control transfers to the *dynamically* closest handler for that kind of exception. The handler function is applied to the exception parameter, and the result of the application is returned as the result of the **handle**-phrase. If there is no active handler for a raised exception, evaluation terminates with an “unhandled exception” answer.

To extend Functional ML with exceptions, we add several new phrases to the syntax:

$$\begin{aligned} (\text{Expressions}) \quad e &::= v \mid e_1 e_2 \mid \text{let } x \text{ be } e_1 \text{ in } e_2 \mid \underline{\text{exception } \chi \text{ in } e} \\ (\text{Values}) \quad v &::= c \mid x \mid Y \mid \lambda x.e \mid \underline{\text{raise}} \mid \underline{e \text{ handle}} \mid \underline{\text{raise } v} \mid \underline{e \text{ handle } v} \\ \chi &::= \underline{x^*} \end{aligned}$$

In the expression **exception** χ **in** e , the variables in χ are bound in e . We treat χ as a set of variables, and call these *exception names*. As **raise** is a curried binary operator, the application of **raise** to a value is a value, and is included in the syntactic class of values (like $:=$ in Reference ML). Likewise, the expression e **handle** is a value, and receives an exception name and handler function by ordinary application. Hence the application of e **handle** to a value is also a value.

5.2.1 Semantics

Raising and handling exceptions requires several new reductions:

$$\begin{aligned} (\text{raise}) \quad & U[\text{raise } x v] \longrightarrow \text{raise } x v \quad \text{if } U \neq [] \\ (\text{handle}) \quad & (\text{raise } x v_1) \text{ handle } x v_2 \longrightarrow v_2 v_1 \\ (\text{reraise}) \quad & \text{exception } \chi_{x_1 x_2} \text{ in } X[(\text{raise } x_1 v_1) \text{ handle } x_2 v_2] \longrightarrow \text{exception } \chi_{x_1 x_2} \text{ in } X[\text{raise } x_1 v_1] \quad x_1 \neq x_2 \\ (\text{unhandle}) \quad & v_1 \text{ handle } x v_2 \longrightarrow v_1 \end{aligned}$$

These new notions of reduction again rely on (two different kinds of) contexts for the enforcement of a specific order of evaluation for applications in connection with the raising

and handling of exceptions:

$$\begin{array}{ll}
 (\textit{raising}) & U ::= [] \mid U \ e \mid v \ U \mid \textit{let } x \ \textit{be } U \ \textit{in } e \\
 (\textit{handling}) & X ::= [] \mid X \ e \mid v \ X \mid \textit{let } x \ \textit{be } X \ \textit{in } e \mid X \ \textit{handle } x \ v.
 \end{array}$$

When an exception is raised, *raise* reductions propagate the exception outwards, eliminating pending computations. As U contexts do not include the phrase $U \ \textit{handle } x \ v$, *raise* reductions propagate the exception only to the closest exception handler. If the closest handler matches the raised exception, a *handle* reduction applies the handler function to the exception parameter. If the handler does not match the exception, a *reraise* reduction discards the handler, and the exception may continue to propagate by *raise* reductions. Should an expression protected by a handler evaluate to a value without raising an exception, the *unhandle* reduction discards the handler.

The **exception**-expression requires two additional reductions, similar to those for ρ -expressions:⁷

$$\begin{array}{ll}
 (\textit{exn}_{merge}) & \textit{exception } \chi_1 \ \textit{in } \textit{exception } \chi_2 \ \textit{in } e \longrightarrow \textit{exception } \chi_1 \chi_2 \ \textit{in } e \\
 (\textit{exn}_{ift}) & X[\textit{exception } \chi \ \textit{in } e] \longrightarrow \textit{exception } \chi \ \textit{in } X[e] \quad \textit{if } X \neq []
 \end{array}$$

Indeed, the only difference is that the **exception**-expression does not bind any values to its variables. The purpose of an **exception**-expression is to distinguish different exceptions in the *reraise* reduction, in order that evaluation be deterministic. If x_1 and x_2 are different but only λ - or **let**-bound, subsequent reductions may substitute them to the same **exception**-bound variable, making the *raise* reduction applicable. For example:

$$\begin{array}{l}
 \textit{exception } x \ \textit{in } (\lambda y. \lambda z. (\textit{raise } y \ v_1) \ \textit{handle } z \ v_2) \ x \ x \\
 \longrightarrow \textit{exception } x \ \textit{in } (\textit{raise } x \ v_1) \ \textit{handle } x \ v_2 \\
 \longrightarrow \textit{exception } x \ \textit{in } v_2 \ v_1.
 \end{array}$$

Insisting that x_1 and x_2 be **exception**-bound in *reraise* ensures that they refer to distinct exceptions. The reduction *handle* does not require its variables to be **exception**-bound, because substitution necessarily replaces both instances of x in the *handle* redex with the same exception name.

We use \mathbf{x} to refer to the six reductions for exceptions introduced above. Taking the union of \mathbf{v} (extended to the full syntax) and \mathbf{x} yields \mathbf{vx} , the notion of reduction for Exception ML. With evaluation contexts extended to:

$$E ::= [] \mid E \ e \mid v \ E \mid \textit{let } x \ \textit{be } E \ \textit{in } e \mid \textit{exception } \chi \ \textit{in } E \mid E \ \textit{handle } x \ v$$

the stepping function $\xrightarrow{\mathbf{vx}}$ may be defined along the same lines as before. Answers for $\textit{eval}_{\mathbf{vx}}$ are:

$$(\textit{answers}) \quad a ::= v \mid \textit{exception } \chi \ \textit{in } v \mid \textit{exception } \chi x \ \textit{in } \textit{raise } x \ v.$$

⁷Due to our use of the variable conventions for *exn_{merge}*, our exceptions are *generative* as in STANDARD ML [22, 23].

Answers of the third form are *unhandled exceptions*: since handlers are dynamically bound, it is possible for an exception to be raised when no handler for it is installed. Evaluation is defined as:

$$(eval_{\mathbf{v}\mathbf{x}}) \quad eval_{\mathbf{v}\mathbf{x}}(e) = a \quad \text{iff} \quad e \xrightarrow{\mathbf{v}\mathbf{x}} a.$$

We can now extend the domain of δ to admit functions such as division, by allowing δ to return an expression that raises an exception:⁸

$$\delta : Const \times ClosedVal \rightarrow ClosedVal \cup \{\text{exception } \mathbf{x} \text{ in raise } \mathbf{x} \ v \mid v \text{ is closed}\}.$$

Functions such as division can now be defined on every element of their input type, by returning an exception when their application does not make sense.

5.2.2 Typing

We extend the set of types with an additional type constructor:

$$\tau ::= \iota_1 \mid \dots \mid \iota_n \mid \alpha \mid \tau_1 \rightarrow \tau_2 \mid \underline{\tau \text{ } \mathit{exn}}$$

The type $\tau \text{ } \mathit{exn}$ is the type of exceptions with a parameter of type τ .

In typing exceptions in the presence of polymorphism, one encounters similar difficulties as in typing references. The phrase **raise** $e_1 \ e_2$ requires e_1 to be an exception of type $\tau \text{ } \mathit{exn}$, and e_2 to be a matching parameter of type τ . If e_3 has type τ' , the phrase **e_3 handle** $e_1 \ e_2$ requires e_1 to be an exception of type $\tau \text{ } \mathit{exn}$, and e_2 to be a handler for such exceptions. The handler must be a function of type $\tau \rightarrow \tau'$, since it takes the exception parameter as input, and returns a value to be returned in place of the value of e_3 . The obvious typing for exceptions allows the parameter types of exceptions to be fully polymorphic, permitting the following expression to be typed:⁹

$$\begin{aligned} & \text{let } \mathbf{rh} \text{ be exception } \mathbf{x} \text{ in pair (raise } \mathbf{x}) (\lambda \mathbf{a} . \lambda \mathbf{b} . (\mathbf{a} \ 3) \text{ handle } \mathbf{x} \ \mathbf{b}) \\ & \text{in (snd } \mathbf{rh}) (\lambda \mathbf{d} . (\mathbf{fst} \ \mathbf{rh}) \ \mathbf{true}) (\lambda \mathbf{y} . + \ 1 \ \mathbf{y}) \\ \longrightarrow & \text{exception } \mathbf{x} \text{ in } + \ 1 \ \mathbf{true} \end{aligned}$$

Here, functions to raise and handle the same \mathbf{x} -exception are returned as a pair, bound by a **let**-expression to **rh**, and then used. In the body of the **let**-expression, the expression **(raise** \mathbf{x}) has type $\alpha_1 \rightarrow \alpha_2$, and **($\lambda \mathbf{a} . \lambda \mathbf{b} . (\mathbf{a} \ 3)$ handle** $\mathbf{x} \ \mathbf{b}$) has type $(int \rightarrow \alpha_3) \rightarrow (\alpha_1 \rightarrow \alpha_3) \rightarrow \alpha_3$. If the exception parameter type α_1 is generalized, then in the body of the **let**-expression an \mathbf{x} -exception may be raised with an argument of type *bool* by **(fst** \mathbf{rh}), and caught by **(snd** \mathbf{rh}), which is expecting an argument of type *int*.

Since the problems of typing exceptions and references are similar, it is not surprising that the same modifications to the typing rules for functional expressions apply: types and type variables are classified as imperative or applicative, and the **let** rule is split according to whether the bound expression is expansive (see Figure 3). Figure 5 presents the additional typing rules for exceptions. Again, as with Reference ML, we require the additional restriction that there be no constants of exception type ($\tau \text{ } \mathit{exn}$).

⁸Usually a set of exceptions that may be raised by constant functions is defined in an initial environment; we take an alternative approach to simplify the presentation.

⁹The constants **pair**, **fst**, and **snd** provide pairing and projection operations.

(raise)	$\Gamma \triangleright \text{raise} : \tau_1 \text{ } \text{exn} \rightarrow \tau_1 \rightarrow \tau_2$
(handle)	$\frac{\Gamma \triangleright e : \tau_2}{\Gamma \triangleright e \text{ handle} : \tau_1 \text{ } \text{exn} \rightarrow (\tau_1 \rightarrow \tau_2) \rightarrow \tau_2}$
(exn)	$\frac{\Gamma[x_1 \mapsto \tau_1 \text{ } \text{exn}] \dots [x_n \mapsto \tau_n \text{ } \text{exn}] \triangleright e : \tau \quad \tau_i \text{ is imperative} \quad 1 \leq i \leq n}{\Gamma \triangleright \text{exception } x_1 \dots x_n \text{ in } e : \tau}$

Figure 5: Additional typing rules for exceptions

5.2.3 Type Soundness for Exceptions

To establish soundness for the exception typing, we proceed as for references: we establish type preservation for \mathbf{vx} , and show that the faulty expressions of Exception ML are untypable.

Main Lemma 5.8 (Subject Reduction for \mathbf{vx}) *If $\Gamma \overset{x}{\triangleright} e_1 : \tau$ and $e_1 \xrightarrow{\mathbf{vx}} e_2$ then $\Gamma \overset{x}{\triangleright} e_2 : \tau$.*

Proof. The cases for the reductions of \mathbf{v} are the same as before.

Case $U[\text{raise } x \ v] \longrightarrow \text{raise } x \ v$. From $\Gamma \triangleright U[\text{raise } x \ v] : \tau$

$$\Gamma \triangleright \text{raise } x \ v : \tau'$$

by **app**. Hence

$$(27) \quad \Gamma \triangleright \text{raise } x : \tau_1 \rightarrow \tau_2$$

$$(28) \quad \text{and } \Gamma \triangleright v : \tau_1$$

by **app**. From (27)

$$\begin{array}{ll} \Gamma \triangleright x : \tau_1 \text{ } \text{exn} & \text{by } \mathbf{app} \text{ and } \mathbf{raise}, \\ \Gamma \triangleright \text{raise } x : \tau_1 \rightarrow \tau & \text{by } \mathbf{app} \text{ and } \mathbf{raise}, \\ \Gamma \triangleright \text{raise } x \ v : \tau & \text{by } \mathbf{app} \text{ with (28)}. \end{array}$$

Case exception χ_1 in exception χ_2 in $e \longrightarrow$ exception $\chi_1 \chi_2$ in e . Similar to ρ_{merge} in Lemma 5.2.

Case $X[\text{exception } \chi \text{ in } e] \longrightarrow$ exception χ in $X[e]$. Similar to ρ_{lift} in Lemma 5.2.

Case $v_1 \text{ handle } x \ v_2 \longrightarrow v_1$. $\Gamma \triangleright v_1 : \tau$ from $\Gamma \triangleright v_1 \text{ handle } x \ v_2 : \tau$ by **handle** and **app**.

Case $(\text{raise } x \ v_1) \text{ handle } x \ v_2 \longrightarrow v_2 \ v_1$. From $\Gamma \triangleright (\text{raise } x \ v_1) \text{ handle } x \ v_2 : \tau$ by **handle**, **app**, and **raise**

$$\Gamma \triangleright v_1 : \tau' \quad \text{and} \quad \Gamma \triangleright v_2 : \tau' \rightarrow \tau$$

where $\Gamma x = \tau' \text{ } \text{exn}$. Then $\Gamma \triangleright v_2 \ v_1 : \tau$ by **app**.

Case $\text{exception } \chi x_1 x_2 \text{ in } X[(\text{raise } x_1 v_1) \text{ handle } x_2 v_2] \longrightarrow \text{exception } \chi x_1 x_2 \text{ in } X[\text{raise } x_1 v_1]$. Then

$$\begin{array}{ll} \Gamma' \triangleright (\text{raise } x_1 v_1) \text{ handle } x_2 v_2 : \tau' & \\ \Gamma' \triangleright (\text{raise } x_1 v_1) \text{ handle} : \tau_2 \text{ exn} \rightarrow (\tau_2 \rightarrow \tau') \rightarrow \tau' & \text{by app,} \\ \Gamma' \triangleright \text{raise } x_1 v_1 : \tau' & \text{by handle,} \\ \Gamma \triangleright \text{exception } \chi x_1 x_2 \text{ in } X[\text{raise } x_1 v_1] : \tau & \text{by Replacement}^+. \end{array}$$

With the extension of the various lemmas, the proof of type preservation for Exception ML is complete. The lemmas go through with simple adaptations to the new syntax. \blacksquare

There are several new kinds of faulty expressions for Exception ML.

Definition 5.9 (Faulty expressions) *The faulty expressions of Exception ML are those expressions containing a subexpression of the form:*

$$\begin{array}{l} (c \ v) \text{ where } \delta(c, v) \text{ is undefined,} \\ (\text{raise } v) \text{ where } v \notin \text{Var,} \\ (e \ \text{handle } v) \text{ where } v \notin \text{Var, or} \\ \text{exception } \chi x \text{ in } C[x \ v], \end{array}$$

where

$$\begin{array}{l} C ::= [] \mid C \ e \mid e \ C \mid \text{let } x \ \text{be } C \ \text{in } e \mid \text{let } x \ \text{be } e \ \text{in } C \mid \lambda x. C \\ \mid \text{exception } \chi \ \text{in } C \mid C \ \text{handle.} \end{array}$$

Again, a uniform evaluation lemma can be shown, and the faulty expressions can be proven untypable. Type soundness for Exception ML follows as before.

Theorem 5.10 (Syntactic Soundness for Exception ML) *If $\overset{x}{\triangleright} e : \tau$ then either $e \uparrow$ or $e \xrightarrow{\text{vx}} a$ and $\overset{x}{\triangleright} a : \tau$.*

5.3 Core ML

We can combine references and exceptions in one language, Core ML, that has all the essential features of STANDARD ML. In combining the reference and exception extensions, we must ensure that they interact appropriately.

5.3.1 Semantics

To combine the calculi for the two extensions, the R contexts of the reference fragment must be extended to include phrases from the exception fragment; likewise, X contexts must be extended to include phrases from the reference fragment:

$$\begin{array}{l} R ::= \dots \mid R \ \text{handle } v_1 \ v_2 \\ X ::= \dots \mid \rho\theta.X \end{array}$$

The phrase `exception χ in R` is not included in R in order that exceptions in the store do not escape their bindings, *i.e.*, `exception`-expressions can move through ρ -expressions by the exn_{ift} reduction, with appropriate variable renaming:

$$\rho\theta.\text{exception } \chi x \text{ in } e \longrightarrow \text{exception } \chi x \text{ in } \rho\theta.e$$

but not vice versa:

$$\text{exception } \chi x \text{ in } \rho\theta\langle r, x \rangle.e \not\rightarrow \rho\theta\langle r, x \rangle.\text{exception } \chi x \text{ in } e.$$

If the second were a permissible reduction, the occurrence of x in the store would escape its binding.

After this extension of the R and X contexts, the resulting notion of reduction for the full syntax is \mathbf{vrX} . The E evaluation contexts used in defining $\xrightarrow{\mathbf{vrX}}$ and $eval_{\mathbf{vrX}}$ must also be extended:

$$E ::= [] \mid E e \mid v E \mid \text{let } x \text{ be } E \text{ in } e \mid E \text{ handle } v_1 v_2 \mid \rho\theta.E \mid \text{exception } \chi \text{ in } E.$$

Answers are:

$$\begin{aligned} (\text{answers}) \quad a ::= & \{ \text{exception } \chi \text{ in } \} \{ \rho\theta. \} v \\ & \mid \{ \text{exception } \chi x \text{ in } \} \{ \rho\theta. \} \text{raise } x v \end{aligned}$$

(where $\{ \textit{phrase} \}$ means *phrase* may be omitted). The complete calculus for Core ML may be found in the appendix.

5.3.2 Typing

To combine the typing rules for the two fragments is easy, since they both rely on the same classification of types as imperative or applicative, and do not interfere. The complete typing rules for Core ML may be found in the appendix.

5.3.3 Type Soundness for Core ML

The type soundness of the resulting calculus requires re-establishing the various lemmas and theorems for the extended syntax and extended evaluation contexts. These proofs are all straightforward. There is a combinatorial increase in the kinds of faulty expressions because of the interaction between the two fragments: references cannot be raised or handled; exceptions cannot be dereferenced nor assigned. The structure of the proof, however, stays the same.

Definition 5.11 (Faulty Expressions) *The faulty expressions of Core ML are those expressions containing a subexpression of the form:*

$$\begin{aligned} & (c \ v) \text{ where } \delta(c, v) \text{ is undefined;} \\ & (! \ v) \text{ where } v \notin \text{Var}, \quad (\text{raise } v) \text{ where } v \notin \text{Var}, \\ & (:= \ v) \text{ where } v \notin \text{Var}, \quad (e \ \text{handle } v) \text{ where } v \notin \text{Var}, \\ & \rho\theta\langle x, v_2 \rangle.C[x \ v_1], \quad \text{exception } \chi x \text{ in } C[x \ v], \\ & \text{exception } \chi x \text{ in } C[! \ x], \quad \rho\theta\langle x, v \rangle.C[\text{raise } x], \\ & \text{exception } \chi x \text{ in } C[:= \ x], \quad \rho\theta\langle x, v \rangle.C[e \ \text{handle } x], \end{aligned}$$

where

$$C ::= [] \mid C \ e \mid e \ C \mid \text{let } x \text{ be } C \text{ in } e \mid \text{let } x \text{ be } e \text{ in } C \mid \lambda x. C \\ \mid \rho\theta.C \mid \rho\theta\langle x, C \rangle.e \mid \text{exception } \chi \text{ in } C \mid C \text{ handle.}$$

Type soundness follows from subject reduction, uniform evaluation, and faulty expressions being untypable, as before.

Theorem 5.12 (Syntactic Soundness for Core ML) *If $\triangleright^{\tau x} e : \tau$ then either $e \uparrow$ or $e \xrightarrow{\text{vrx}} a$ and $\triangleright^{\tau x} a : \tau$.*

6 Control ML

In this section we present an extension to Functional ML providing first-class continuations. The typing of our extension is similar to that described by Duba *et al.* [8], and implemented in STANDARD ML OF NEW JERSEY [34]. It is a simple matter to merge this extension with Core ML.

Control ML extends Functional ML's syntax with two new constructs:

$$\begin{array}{ll} (\text{Expressions}) & e ::= v \mid e_1 \ e_2 \mid \text{let } x \text{ be } e_1 \text{ in } e_2 \mid \underline{\text{abort}} \ e \\ (\text{Values}) & v ::= c \mid x \mid Y \mid \lambda x. e \mid \underline{\text{callcc}} \end{array}$$

An **abort**-expression evaluates its subexpression to a value, and returns this value as the result of the program. When the **callcc** operator (*call-with-current-continuation*) is applied to a function f , it captures a representation of the current continuation (or program control stack), encapsulates this continuation in an abstraction, and applies f to this abstraction. If this abstraction is later applied to a value, control transfers to the captured continuation.

6.1 Semantics

While defining calculus reductions for first-class continuations is possible [13, 14], the resulting calculus is inappropriate for our problem. Since we are only interested in evaluation, we take a simpler approach [11]. We give top-level evaluation rules, *program reductions*, for the control operators by extending the stepping relation (\mapsto) *directly*:

$$\begin{array}{ll} (\delta, \beta_v, \text{let}, Y) & E[e] \mapsto E[e'] \text{ iff } e \xrightarrow{v} e' \\ (\text{callcc}) & E[\text{callcc } v] \mapsto E[v \ (\lambda x. \text{abort } E[x])] \\ (\text{abort}) & E[\text{abort } e] \mapsto e \end{array}$$

where

$$E ::= [] \mid E \ e \mid v \ E \mid \text{let } x \text{ be } E \text{ in } e.$$

In this style of semantics, the evaluation context E represents the current continuation. An *abort* reduction simply discards the current continuation, continuing evaluation with its subexpression. The *callcc* reduction captures the current continuation, encapsulates it in a function, and applies **callcc**'s argument to this function. The continuations created by **callcc** are *abortive*: when invoked, they discard the continuation surrounding their application

```

let product be  $\lambda x$ . callcc
  ( $\lambda k$ . letrec p be  $\lambda z$ . if nil? z then 1
    else if = 0 (car z) then k 0
    else * (car z) (p (cdr z))
    in p x)
in product (cons 2 (cons 3 (cons 0 (cons 5 nil))))

```

Figure 6: A program illustrating a simple use of `callcc`

(via an *abort* reduction), installing instead the captured continuation. Answers are simply values, and *eval* is defined as for Functional ML.

Figure 6 illustrates a simple use of `callcc`. This program computes the product of a list of numbers, performing no multiplications if the answer is zero. The constants `nil`, `cons`, `nil?`, `car`, and `cdr` are the usual constants and operations for lists. The expressions `if...then...else...` and `letrec...be...in...` are syntactic sugar for appropriate constant applications. This program computes the product of a list of numbers, performing no multiplications if the answer is zero. The program escapes from pending multiplications by jumping out of the recursion to the continuation of the application of `product`.

6.2 Typing

The obvious approach to typing `callcc` leads to the following rule:

(**naïve-callcc**) $\Gamma \triangleright \text{callcc} : ((\tau_1 \rightarrow \tau_2) \rightarrow \tau_1) \rightarrow \tau_1$

As indicated by the *callcc* reduction, `callcc` takes a function whose argument is a procedural abstraction of the evaluation context, the *continuation*. If the evaluation context expects a value of type τ_1 , then the continuation has type $\tau_1 \rightarrow \tau_2$ for any type τ_2 . If `callcc`'s argument ignores this continuation, it must produce a value of type τ_1 , which implies that `callcc`'s argument must have type $(\tau_1 \rightarrow \tau_2) \rightarrow \tau_1$. The result type of the continuation, τ_2 , is arbitrary, because the application of *k* built into the continuation ensures that it never returns.

However, this obvious typing is unsound, as illustrated by the following example:¹⁰

```

let x be callcc ( $\lambda k$ . pair ( $\lambda x$ .x) ( $\lambda f$ . k (pair f ( $\lambda d$ .5)))))
in ( $\lambda z$ . snd x ( $\lambda x$ .+ x 1)) (fst x true)

```

The `callcc` application returns a pair of type $(\alpha \rightarrow \alpha) \times ((\alpha \rightarrow \alpha) \rightarrow \text{int})$, consisting of the identity function and a function that applies a continuation. Since α is not free in the type environment, it is closed over, and `x` has type scheme $\forall \alpha. (\alpha \rightarrow \alpha) \times ((\alpha \rightarrow \alpha) \rightarrow \text{int})$ in the body of the `let`-expression. In evaluating the body, `fst x` is the polymorphic identity function, so it may validly be applied to `true`. Then the second function of the pair is applied

¹⁰A variation of this example was discovered by Robert Harper and Mark Lillibridge [sm1 electronic mailing list, July 8, 1991]. Despite widespread use of a naïvely typed `callcc` extension in STANDARD ML OF NEW JERSEY, it took two years until this problem was discovered.

to $(\lambda x. + x \ 1)$, and the continuation so invoked rebinds x to the pair consisting of $(\lambda x. + x \ 1)$ and $(\lambda d.5)$. The second evaluation of $\text{fst } x \ \text{true}$ results in an attempt to add 1 to true .

As with references and exceptions, a correct typing for continuations builds upon the classification of types as imperative or applicative, and requires the **let** typing rules of Figure 3. The correct typing rules are:

(**callcc**) $\Gamma \triangleright \text{callcc} : ((\tau_1 \rightarrow \tau_2) \rightarrow \tau_1) \rightarrow \tau_1$ if τ_1 is imperative.

(**abort**)
$$\frac{\Gamma \triangleright e : \tau_1}{\Gamma \triangleright \text{abort } e : \tau_2}$$

An **abort**-expression may have any type, regardless of the type of its subexpression, because when evaluated it never returns.

6.3 Weak Type Soundness

Because **abort**-expressions can return a value of any type, proving type soundness for the continuation extension is not as simple as for the previous extensions. For example, if e is an expression of type *bool*, the program:

if e then 1 else (**abort true**)

is typable according to the above typing rules, but returns either *int* or *bool* according to whether e is true or false. Thus subject reduction in the usual sense does not hold; however, a weaker lemma does hold, stating that *typability* is preserved.

Main Lemma 6.1 (Typability Preservation) *If $\triangleright e_1 : \tau$ and $e_1 \mapsto e_2$ then $\triangleright e_2 : \tau'$.*

Proof. We need only consider the additional cases for the *callcc* and *abort* reductions. The others are simply adapted from Lemma 4.3 (Subject Reduction for Functional ML).

Case $E[\text{abort } e] \mapsto e$. Then $\Gamma \triangleright e : \tau'$ for some τ' . But since $E[\text{abort } e]$ is closed and E does not bind variables, e is closed, hence $\triangleright e : \tau'$.

Case $E[\text{callcc } v] \mapsto E[v (\lambda x. \text{abort } E[x])]$. From $\triangleright E[\text{callcc } v] : \tau$ we have $\triangleright \text{callcc } v : \tau_1$, and by **callcc** and **app**

$$(29) \quad \triangleright v : (\tau_1 \rightarrow \tau_2) \rightarrow \tau_1$$

where τ_1 is imperative.

Since E is closed, x does not appear free in E , and we claim that $[x \mapsto \tau_1] \triangleright E[x] : \tau$. The proof of this claim proceeds by induction on the structure of E .

Case $E = []$. Then $\tau = \tau_1$, and $[x \mapsto \tau_1] \triangleright x : \tau$ by **var**.

Case $E = (E' e')$. Then $\triangleright (E'[\text{callcc } v] e') : \tau$, and by **app**

$$(30) \quad \triangleright E'[\text{callcc } v] : \tau_3 \rightarrow \tau$$

$$(31) \quad \text{and } \triangleright e' : \tau_3$$

for some τ_3 . Then $[x \mapsto \tau_1] \triangleright E'[x] : \tau_3 \rightarrow \tau$ by ind. hyp. with (30). Also $[x \mapsto \tau_1] \triangleright e' : \tau_3$ by Lemma 4.1⁺ with (31) since $x \notin FV(E)$. Hence

$$[x \mapsto \tau_1] \triangleright E'[x] e' : \tau \quad \text{by } \mathbf{app}.$$

Case $E = (v E')$. Similar to the previous case.

Case $E = \text{let } x' \text{ be } E' \text{ in } e'$. Then $\triangleright \text{let } x' \text{ be } E'[\text{callcc } v] \text{ in } e' : \tau$. Since $(\text{callcc } v)$ is expansive, $E'[\text{callcc } v]$ is expansive. Thus by let_e

$$(32) \quad \triangleright E'[\text{callcc } v] : \tau_3$$

$$(33) \quad \text{and } [x' \mapsto \text{AppClose}(\tau_3, \emptyset)] \triangleright e' : \tau$$

for some τ_3 . Then

$$(34) \quad [x \mapsto \tau_1] \triangleright E'[x] : \tau_3 \quad \text{by ind. hyp. with (32),}$$

and

$$\begin{aligned} & [x' \mapsto \text{AppClose}(\tau_3, \emptyset)][x \mapsto \tau_1] \triangleright e' : \tau \quad \text{by Lemma 4.1}^+ \text{ with (33)} \\ \text{i.e. } & [x \mapsto \tau_1][x' \mapsto \text{AppClose}(\tau_3, \emptyset)] \triangleright e' : \tau. \end{aligned}$$

By the critical restriction that τ_1 is imperative,

$$\text{AppClose}(\tau_3, \emptyset) = \text{AppClose}(\tau_3, [x \mapsto \tau_1]),$$

hence

$$\begin{aligned} & [x \mapsto \tau_1][x' \mapsto \text{AppClose}(\tau_3, [x \mapsto \tau_1])] \triangleright e' : \tau \\ & [x \mapsto \tau_1] \triangleright \text{let } x' \text{ be } E'[x] \text{ in } e' : \tau \quad \text{by } \text{let}_e \text{ with (34).} \end{aligned}$$

This completes the proof of the claim. Hence

$$\begin{aligned} & [x \mapsto \tau_1] \triangleright \text{abort } E[x] : \tau_2 && \text{by } \text{abort}, \\ & \triangleright (\lambda x. \text{abort } E[x]) : \tau_1 \rightarrow \tau_2 && \text{by } \text{abs}, \\ & \triangleright v (\lambda x. \text{abort } E[x]) : \tau_1 && \text{by } \text{app} \text{ with (29)} \\ & \triangleright E[v (\lambda x. \text{abort } E[x])] : \tau && \text{by Replacement}^+. \blacksquare \end{aligned}$$

The faulty expressions for Control ML are the same as those of Functional ML, extended to the new syntax.

Definition 6.2 (Faulty Expressions) *The faulty expressions of Control ML are the expressions containing a subexpression $(c v)$ where $\delta(c, v)$ is undefined.*

As with Functional ML, the faulty expressions are untypable and Uniform Evaluation holds. However, since we only have typability preservation, we obtain a weaker version of syntactic soundness that does not indicate the type of answers.

Theorem 6.3 (Weak Syntactic Soundness) *If $\triangleright e : \tau$ then either $e \uparrow$ or $e \longrightarrow v$.*

Proof. By Uniform Evaluation, either $e \longrightarrow e'$ and e' is faulty, or $e \uparrow$, or $e \longrightarrow v$. Since $\triangleright e : \tau$, Typability Preservation implies $\triangleright v : \tau'$ and $\triangleright e' : \tau'$. Suppose $e \longrightarrow e'$ and e' is faulty. Since faulty expressions are untypable, $\triangleright e' : \tau'$ is a contradiction, therefore this case cannot occur. Hence either $e \uparrow$ or $e \longrightarrow v$. \blacksquare

This weaker theorem establishes weak soundness, but does not imply strong soundness. In fact, if the use of **abort** is unrestricted, it is not possible to predict the type of an answer, as the example at the beginning of this section illustrates.

6.4 Strong Type Soundness

Examining the proof of Typability Preservation reveals that every reduction with the exception of *abort* preserves type. In particular, the *calcc* reduction preserves type, which suggests that **calcc** by itself is strongly sound. To obtain a proof of strong soundness for **calcc**, we must eliminate **abort** from the surface language available to programmers, but retain it in the underlying language of evaluation for **calcc** reductions. Because the **abort** expressions introduced by **calcc** applications have a restricted shape, they never return values of other than the top-level type of the program, and do not compromise strong soundness.

To establish a subject reduction lemma, we define an augmented type system that infers both the ordinary type of an expression and a set of *abort types*. The abort types of an expression e are the types of the immediate subexpressions of **abort**-expressions in e . The augmented system does not permit abort types to be generalized, thereby ensuring that each syntactic occurrence of **abort** can produce only one type of answer. This restriction eliminates expressions such as the following:

```

let a be  $\lambda x.$  abort x
in if  $e$  then a 1
   else a true

```

Hence if a well-typed program aborts, the answer produced is one of the program's abort types. Type judgments for this augmented system have the form $\Gamma \stackrel{a}{\triangleright} e : \tau, T$, meaning that in type environment Γ , expression e has ordinary type τ and abort types T , where T is a set of types. Figure 7 presents the typing rules of the augmented system.

The augmented type system corresponds closely to our original type system for Control ML. The rule for typing **abort**-expressions is the only rule that connects the ordinary type and the abort types of an expression. In each axiom (**var**, **const**, **Y**, **calcc**) the abort set is completely unconstrained; the inference rules simply propagate the abort set. The **let**-expression rules do not generalize type variables in the set of abort types, hence the augmented system does not permit polymorphic uses of **abort**, and accepts a subset of the expressions accepted by the ordinary system. A Correspondence Lemma establishes this connection between the ordinary system (\triangleright) and the augmented system ($\stackrel{a}{\triangleright}$).

Lemma 6.4 (Correspondence)

- (i) If $\Gamma \triangleright e : \tau$ and e contains no **abort**-expressions then $\Gamma \stackrel{a}{\triangleright} e : \tau, T$ for any T ;
- (ii) If $\Gamma \stackrel{a}{\triangleright} e : \tau, T$ then $\Gamma \triangleright e : \tau$.

Proof. The proof of each direction shows how to construct a deduction for the consequent from a deduction of the antecedent. Both proofs are straightforward. ■

Subject Reduction for the augmented system establishes that reduction preserves the set of abort types. As the **calcc** reduction introduces an **abort**-expression with the top-level type of the program, the set of abort types preserved by Subject reduction includes the top-level type.

(var)	$\Gamma \triangleright^a x : \tau, T$ if $\Gamma(x) \succ \tau$
(const)	$\Gamma \triangleright^a c : \tau, T$ if $TypeOf(c) \succ \tau$
(Y)	$\Gamma \triangleright^a Y : ((\tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2, T$
(abs)	$\frac{\Gamma[x \mapsto \tau_1] \triangleright^a e : \tau_2, T}{\Gamma \triangleright^a \lambda x. e : \tau_1 \rightarrow \tau_2, T}$
(app)	$\frac{\Gamma \triangleright^a e_1 : \tau_1 \rightarrow \tau_2, T \quad \Gamma \triangleright^a e_2 : \tau_1, T}{\Gamma \triangleright^a e_1 e_2 : \tau_2, T}$
(let _v)	$\frac{\Gamma \triangleright^a v : \tau_1, T \quad \Gamma[x \mapsto Close(\tau_1, \Gamma, T)] \triangleright^a e : \tau_2, T}{\Gamma \triangleright^a \text{let } x \text{ be } v \text{ in } e : \tau_2, T}$
(let _e)	$\frac{\Gamma \triangleright^a e_1 : \tau_1, T \quad \Gamma[x \mapsto AppClose(\tau_1, \Gamma, T)] \triangleright^a e_2 : \tau_2, T \quad e_1 \notin Values}{\Gamma \triangleright^a \text{let } x \text{ be } e_1 \text{ in } e_2 : \tau_2, T}$
(callcc)	$\Gamma \triangleright^a \text{callcc} : ((\tau_1 \rightarrow \tau_2) \rightarrow \tau_1) \rightarrow \tau_1, T$ if τ_1 is imperative
(abort)	$\frac{\Gamma \triangleright^a e : \tau_1, T \quad \tau_1 \in T}{\Gamma \triangleright^a \text{abort } e : \tau_2, T}$

	$Close(\tau, \Gamma, T) = \forall \alpha_1 \dots \alpha_n. \tau$
where $\{\alpha_1, \dots, \alpha_n\}$	$= FTV(\tau) \setminus (FTV(\Gamma) \cup FTV(T))$
	$AppClose(\tau, \Gamma, T) = \forall \alpha_1 \dots \alpha_n. \tau$
where $\{\alpha_1, \dots, \alpha_n\}$	$= (FTV(\tau) \setminus (FTV(\Gamma) \cup FTV(T))) \cap AppTypeVar$

Figure 7: The augmented type system for Control ML

Main Lemma 6.5 (Subject Reduction) *If $\triangleright^a e_1 : \tau, T$ and $\tau \in T$ and $e_1 \mapsto e_2$ then $\triangleright^a e_2 : \tau', T$ and $\tau' \in T$.*

Proof. The proof proceeds by case analysis according to the reduction $e_1 \mapsto e_2$. Each case is easily adapted from the corresponding case in the proof of Typability Preservation with straightforward adaptations of the necessary lemmas. ■

With Uniform Evaluation and the fact that faulty expressions are untypable (in the ordinary system), we obtain a syntactic soundness theorem for **abort**-free expressions.

Theorem 6.6 (Syntactic Soundness) *If $\triangleright e : \tau$ and e contains no **abort**-expressions then either $e \uparrow$ or $e \mapsto v$ and $\triangleright v : \tau$.*

Proof. By Uniform Evaluation, either $e \mapsto e'$ and e' is faulty, or $e \uparrow$, or $e \mapsto v$. Since $\triangleright e : \tau$ and e contains no **abort**-expressions, by Correspondence $\triangleright^a e : \tau, \{\tau\}$. Subject Reduction then implies $\triangleright^a e' : \tau, \{\tau\}$. Similarly $\triangleright^a v : \tau, \{\tau\}$, and by Correspondence $\triangleright e' : \tau$ and $\triangleright v : \tau$. Suppose $e \mapsto e'$ and e' is faulty. Since faulty expressions are untypable, $\triangleright e' : \tau$ is a contradiction, therefore this case cannot occur. Hence either $e \uparrow$ or $e \mapsto v$ and $\triangleright v : \tau$. ■

7 Discussion

Subject reduction is the key lemma in our approach to proving type soundness. In order to demonstrate subject reduction, it must be possible to assign a type to each intermediate evaluation state of a program. This is most easily accomplished by specifying evaluation as rewriting. Rewriting may be specified as local reductions, as our calculus for Core ML, or as program (or top-level) rewriting, as in Control ML.

Through most of this paper, we present the semantics of the various languages with calculi that permit local reductions. It is also possible to use our proof technique with a semantics that specifies only program reductions (\longrightarrow), as in Control ML. Such a semantics for Reference ML is slightly simpler, as the *ref*, ρ_{merge} , and ρ_{lift} reductions coalesce into one program reduction, but the structure of the proof is essentially the same. Our work on an alternative type system for references uses this approach [38]. We chose to use calculi in this paper as the resulting proofs are more regular in structure.

In specifying the semantics of references as a calculus, we use an additional expression form, the ρ -expression, that is not present in ML. As noted earlier, ρ -expressions may be regarded as abbreviations, both from a semantic perspective and a typing perspective (the typing rule for ρ -expressions can be derived from the abbreviation). In principle, it is possible to eliminate ρ -expressions from the syntax, however the specification of redexes in the reduction rules becomes complicated. Alternatively, ρ -expressions may be considered as belonging to the state space of evaluation, and not to the language of programs that may be written by a programmer. In a presentation of the typing rules for consumption by programmers, the typing rule for ρ -expressions may be deleted. However, unlike ρ -expressions, the **abort**-expressions of the continuation fragment cannot be considered as abbreviations [10]. To obtain a strong soundness theorem, **abort**-expressions must be considered as belonging only to the evaluation space, and not to the syntax of programs. The exception fragment has no such additional expressions.

Other operational or denotational techniques for specifying semantics introduce additional semantic objects, such as closures and stores, rather than additional expression forms. Since the type system does not apply to such objects, stating and proving strong type soundness requires introducing a semantic relation (\models) between semantic objects and types (see Section 2). We believe that introducing additional expression forms and typing rules is the simpler choice, and offer the simpler proofs produced by our method as evidence.

While in this paper we have concentrated on the essential aspects of ML, there are other features of static type systems that we believe our technique can address. STANDARD ML contains a datatype specification mechanism that allows the definition of new types and associated data constructors. This mechanism is indispensable when writing non-trivial ML programs. Subtyping, inheritance, and type inference for records are a strong focus of recent research efforts to explain object-oriented languages, as many popular object-oriented languages have unsound static type systems. Reppy [29] has successfully addressed concurrency with our technique; it should also be possible to treat nondeterminism and distributed computing. We have used our technique to prove an alternative type system for references sound [38]. Finally, it may be possible to adapt our technique to demonstrate the consistency of module systems like that of STANDARD ML.

Acknowledgements

We would like to thank Bob Harper for pointing out that type preservation is known as subject reduction in combinatory logic, and Hans Boehm and the anonymous referees for comments on earlier drafts of this paper.

Appendix: The Core of STANDARD ML

Syntax

$$\begin{aligned}
e &::= v \mid e_1 e_2 \mid \text{let } x \text{ be } e_1 \text{ in } e_2 \mid \rho\theta.e \mid \text{exception } \chi \text{ in } e \\
v &::= c \mid x \mid Y \mid \lambda x.e \mid \text{ref } \mid ! \mid := \mid \text{raise } \mid e \text{ handle } := v \mid \text{raise } v \mid e \text{ handle } v \\
\theta &::= \{ \langle x, v \rangle \}^* \\
\chi &::= x^* \\
a &::= \{ \text{exception } \chi \text{ in } \} \{ \rho\theta. \} v \mid \{ \text{exception } \chi x \text{ in } \} \{ \rho\theta. \} \text{raise } x v
\end{aligned}$$

Semantics

$$\begin{aligned}
\text{eval}(e) = a &\text{ iff } e \longrightarrow a \\
E[e] \longmapsto E[e'] &\text{ iff } e \longrightarrow e' \\
E &::= [] \mid E e \mid v E \mid \text{let } x \text{ be } E \text{ in } e \mid E \text{ handle } v_1 v_2 \mid \rho\theta.E \mid \text{exception } \chi \text{ in } E
\end{aligned}$$

Functional ML

$$\begin{aligned}
(\delta) & \quad c v \longrightarrow \delta(c, v) \quad \text{if } \delta(c, v) \text{ is defined} \\
(\beta_v) & \quad (\lambda x.e) v \longrightarrow e[x \mapsto v] \\
(\text{let}) & \quad \text{let } x \text{ be } v \text{ in } e \longrightarrow e[x \mapsto v] \\
(Y) & \quad Y v \longrightarrow v (\lambda x.(Y v) x)
\end{aligned}$$

References

$$\begin{aligned}
(\text{ref}) & \quad \text{ref } v \longrightarrow \rho \langle x, v \rangle . x \\
(\text{deref}) & \quad \rho\theta \langle x, v \rangle . R[! x] \longrightarrow \rho\theta \langle x, v \rangle . R[v] \\
(\text{assign}) & \quad \rho\theta \langle x, v_1 \rangle . R[:= x v_2] \longrightarrow \rho\theta \langle x, v_2 \rangle . R[v_2] \\
(\rho_{\text{merge}}) & \quad \rho\theta_1 . \rho\theta_2 . e \longrightarrow \rho\theta_1 \theta_2 . e \\
(\rho_{\text{lift}}) & \quad R[\rho\theta.e] \longrightarrow \rho\theta . R[e] \quad \text{if } R \neq []
\end{aligned}$$

$$R ::= [] \mid R e \mid v R \mid \text{let } x \text{ be } R \text{ in } e \mid R \text{ handle } v_1 v_2$$

Exceptions

$$\begin{aligned}
(\text{raise}) & \quad U[\text{raise } x v] \longrightarrow \text{raise } x v \quad \text{if } U \neq [] \\
(\text{handle}) & \quad (\text{raise } x v_1) \text{ handle } x v_2 \longrightarrow v_2 v_1 \\
(\text{reraise}) & \quad \text{exception } \chi_{x_1 x_2} \text{ in } X[(\text{raise } x_1 v_1) \text{ handle } x_2 v_2] \longrightarrow \text{exception } \chi_{x_1 x_2} \text{ in } X[\text{raise } x_1 v_1] \quad x_1 \neq x_2 \\
(\text{unhandle}) & \quad v_1 \text{ handle } x v_2 \longrightarrow v_1 \\
(\text{exn}_{\text{merge}}) & \quad \text{exception } \chi_1 \text{ in exception } \chi_2 \text{ in } e \longrightarrow \text{exception } \chi_1 \chi_2 \text{ in } e \\
(\text{exn}_{\text{lift}}) & \quad X[\text{exception } \chi \text{ in } e] \longrightarrow \text{exception } \chi \text{ in } X[e] \quad \text{if } X \neq []
\end{aligned}$$

$$\begin{aligned}
U &::= [] \mid U e \mid v U \mid \text{let } x \text{ be } U \text{ in } e \\
X &::= [] \mid X e \mid v X \mid \text{let } x \text{ be } X \text{ in } e \mid X \text{ handle } v_1 v_2 \mid \rho\theta.X
\end{aligned}$$

Typing*Functional ML*

(var)	$\Gamma \triangleright x : \tau$ if $\Gamma(x) \succ \tau$
(const)	$\Gamma \triangleright c : \tau$ if $TypeOf(c) \succ \tau$
(Y)	$\Gamma \triangleright Y : ((\tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2) \rightarrow \tau_1 \rightarrow \tau_2$
(abs)	$\frac{\Gamma[x \mapsto \tau_1] \triangleright e : \tau_2}{\Gamma \triangleright \lambda x. e : \tau_1 \rightarrow \tau_2}$
(app)	$\frac{\Gamma \triangleright e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \triangleright e_2 : \tau_1}{\Gamma \triangleright e_1 e_2 : \tau_2}$
(let_v)	$\frac{\Gamma \triangleright v : \tau_1 \quad \Gamma[x \mapsto Close(\tau_1, \Gamma)] \triangleright e : \tau_2}{\Gamma \triangleright \text{let } x \text{ be } v \text{ in } e : \tau_2}$
(let_e)	$\frac{\Gamma \triangleright e_1 : \tau_1 \quad \Gamma[x \mapsto AppClose(\tau_1, \Gamma)] \triangleright e_2 : \tau_2 \quad e_1 \notin Values}{\Gamma \triangleright \text{let } x \text{ be } e_1 \text{ in } e_2 : \tau_2}$

$Close(\tau, \Gamma) = \forall \alpha_1 \dots \alpha_n. \tau$ where $\{\alpha_1, \dots, \alpha_n\} = FTV(\tau) \setminus FTV(\Gamma)$

$AppClose(\tau, \Gamma) = \forall \alpha_1 \dots \alpha_n. \tau$ where $\{\alpha_1, \dots, \alpha_n\} = (FTV(\tau) \setminus FTV(\Gamma)) \cap AppTypeVar$

References

(ref)	$\Gamma \triangleright \text{ref} : \tau \rightarrow \tau \text{ ref}$ if τ is imperative
(deref)	$\Gamma \triangleright ! : \tau \text{ ref} \rightarrow \tau$
(assign)	$\Gamma \triangleright := : (\tau \text{ ref} \rightarrow \tau \rightarrow \tau)$
(rho)	$\frac{\Gamma[x_1 \mapsto \tau_1 \text{ ref}] \dots [x_n \mapsto \tau_n \text{ ref}] \triangleright e : \tau \quad \Gamma[x_1 \mapsto \tau_1 \text{ ref}] \dots [x_n \mapsto \tau_n \text{ ref}] \triangleright v_i : \tau_i \quad \tau_i \text{ is imperative} \quad 1 \leq i \leq n}{\Gamma \triangleright \rho \langle x_1, v_1 \rangle \dots \langle x_n, v_n \rangle. e : \tau}$

Exceptions

(raise)	$\Gamma \triangleright \text{raise} : \tau_1 \text{ exn} \rightarrow \tau_1 \rightarrow \tau_2$
(handle)	$\frac{\Gamma \triangleright e : \tau_2}{\Gamma \triangleright \text{handle} : \tau_1 \text{ exn} \rightarrow (\tau_1 \rightarrow \tau_2) \rightarrow \tau_2}$
(exn)	$\frac{\Gamma[x_1 \mapsto \tau_1 \text{ exn}] \dots [x_n \mapsto \tau_n \text{ exn}] \triangleright e : \tau \quad \tau_i \text{ is imperative} \quad 1 \leq i \leq n}{\Gamma \triangleright \text{exception } x_1 \dots x_n \text{ in } e : \tau}$

References

- [1] ABADI, M., CARDELLI, L., PIERCE, B., AND PLOTKIN, G. Dynamic typing in a statically-typed language. *ACM Transactions on Programming Languages and Systems* 13, 2 (April 1991), 237–268. Previously appeared in: *Proceedings of the 16th Annual Symposium on Principles of Programming Languages* (January 1989), 213–227.
- [2] BARENDREGT, H. P. *The Lambda Calculus: Its Syntax and Semantics*, revised ed., vol. 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, 1984.
- [3] CRANK, E., AND FELLEISEN, M. Parameter-passing and the lambda calculus. *Proceedings of the 18th Annual Symposium on Principles of Programming Languages* (January 1991), 233–244.
- [4] CURRY, H. B., AND FEYS, R. *Combinatory Logic, Volume I*. North-Holland, Amsterdam, 1958.
- [5] DAMAS, L., AND MILNER, R. Principal type schemes for functional programs. *Proceedings of the 9th Annual Symposium on Principles of Programming Languages* (January 1982), 207–212.
- [6] DAMAS, L. M. M. *Type Assignment in Programming Languages*. PhD thesis, University of Edinburgh, 1985.
- [7] DONAHUE, J., AND DEMERS, A. Data types are values. *ACM Transactions on Programming Languages and Systems* 7, 3 (July 1985), 426–445.
- [8] DUBA, B. F., HARPER, R., AND MACQUEEN, D. Typing first-class continuations in ML. *Proceedings of the 18th Annual Symposium on Principles of Programming Languages* (January 1991), 163–173.
- [9] FELLEISEN, M. The theory and practice of first-class prompts. *Proceedings of the 15th Annual Symposium on Principles of Programming Languages* (1988), 180–190.
- [10] FELLEISEN, M. On the expressive power of programming languages. *Science of Computer Programming* 17 (1991), 35–75. Preliminary version in: *Proceedings of the European Symposium on Programming, LNCS 432* (1990), 134–151.
- [11] FELLEISEN, M., AND FRIEDMAN, D. P. Control operators, the SECD-machine, and the λ -calculus. In *Formal Description of Programming Concepts III*, M. Wirsing, Ed. Elsevier Science Publishers B.V. (North-Holland), Amsterdam, 1986, pp. 193–217.
- [12] FELLEISEN, M., AND FRIEDMAN, D. P. A syntactic theory of sequential state. *Theoretical Computer Science* 69, 3 (1989), 243–287. Preliminary version in: *Proceedings of the 14th Annual Symposium on Principles of Programming Languages*, 1987, 314–325.
- [13] FELLEISEN, M., FRIEDMAN, D. P., KOHLBECKER, E. E., AND DUBA, B. A syntactic theory of sequential control. *Theoretical Computer Science* 52, 3 (1987), 205–237. Preliminary version in: *Proceedings of the Symposium on Logic in Computer Science*, 1986, 131–141.

- [14] FELLEISEN, M., AND HIEB, R. The revised report on the syntactic theories of sequential control and state. Tech. Rep. TR-100, Rice University, June 1989. To appear in: *Theoretical Computer Science*, 1992.
- [15] HINDLEY, R. The principal type-scheme of an object in combinatory logic. *Transactions of the American Mathematical Society* 146 (December 1969), 29–60.
- [16] HINDLEY, R. J., AND SELDIN, J. P. *Introduction to Combinators and λ -Calculus*. Cambridge University Press, 1986.
- [17] LEROY, X., AND WEIS, P. Polymorphic type inference and assignment. *Proceedings of the 18th Annual Symposium on Principles of Programming Languages* (January 1991), 291–302.
- [18] MACQUEEN, D. B., PLOTKIN, G., AND SETHI, R. An ideal model for recursive polymorphic types. *Proceedings of the 11th Annual Symposium on Principles of Programming Languages* (January 1984), 165–174.
- [19] MASON, I., AND TALCOTT, C. Programming, transforming, and proving with function abstractions and memories. In *Proceedings of the International Conference on Automata, Languages, and Programming, LNCS 372* (1989), Springer-Verlag, pp. 574–588.
- [20] MILNER, R. A theory of type polymorphism in programming. *Journal of Computer and System Sciences* 17 (1978), 348–375.
- [21] MILNER, R., AND TOFTE, M. Co-induction in relational semantics. *Theoretical Computer Science* 87 (1991), 209–220.
- [22] MILNER, R., AND TOFTE, M. *Commentary on Standard ML*. MIT Press, Cambridge, Massachusetts, 1991.
- [23] MILNER, R., TOFTE, M., AND HARPER, R. *The Definition of Standard ML*. MIT Press, Cambridge, Massachusetts, 1990.
- [24] MITCHELL, J. C., AND HARPER, R. The essence of ML. *Proceedings of the 15th Annual Symposium on Principles of Programming Languages* (January 1988), 28–46.
- [25] MITCHELL, J. C., AND PLOTKIN, G. D. Abstract types have existential type. *ACM Transactions on Programming Languages and Systems* 10, 3 (July 1988), 470–502. Also appeared in: *Proceedings of the 11th Annual Symposium on Principles of Programming Languages*, 1984, 37–51.
- [26] PLOTKIN, G. D. Call-by-name, call-by-value and the lambda-calculus. *Theoretical Computer Science* 1 (1975), 125–159.
- [27] PLOTKIN, G. D. A structural approach to operational semantics. Tech. Rep. DAIMI FN-19, Århus University, September 1981.
- [28] REES, J., AND CLINGER, W. Revised³ report on the algorithmic language Scheme. *SIGPLAN Notices* 21, 12 (December 1986), 37–79.

- [29] REPPY, J. H. *Higher-order Concurrency*. PhD thesis, Cornell University, 1991.
- [30] REYNOLDS, J. Definitional interpreters for higher order programming languages. *ACM Conference Proceedings* (1972), 717–740.
- [31] REYNOLDS, J. C. On the relation between direct and continuation semantics. *Proceedings of the International Conference on Automata, Languages, and Programming* (1974), 141–156.
- [32] SCOTT, D. Data types as lattices. *SIAM Journal of Computing* 3, 5 (1976), 522–586.
- [33] SELDIN, J. P. A sequent calculus for type assignment. *Journal of Symbolic Logic* 42 (1977), 11–28.
- [34] STANDARD ML OF NEW JERSEY release notes (version 0.75). AT&T Bell Laboratories, November 1991.
- [35] TALPIN, J.-P., AND JOUVELOT, P. The type and effect discipline. Tech. Rep. EMP-CRI A/206, Ecole des Mines de Paris, July 1991.
- [36] TOFTE, M. *Operational Semantics and Polymorphic Type Inference*. PhD thesis, University of Edinburgh, 1987.
- [37] TOFTE, M. Type inference for polymorphic references. *Information and Computation* 89, 1 (November 1990), 1–34.
- [38] WRIGHT, A. K. Typing references by effect inference. In *Proceedings of the European Symposium on Programming, LNCS 582* (1992), Springer-Verlag, pp. 473–491.