

An Effective Theory of Type Refinements

Yitzhak Mandelbaum
Princeton University

David Walker *
Princeton University

Robert Harper †
Carnegie Mellon University

ABSTRACT

We develop an explicit two level system that allows programmers to reason about the behavior of effectful programs. The first level is an ordinary ML-style type system, which confers standard properties on program behavior. The second level is a conservative extension of the first that uses a *logic of type refinements* to check more precise properties of program behavior. Our logic is a fragment of intuitionistic linear logic, which gives programmers the ability to reason *locally* about changes of program state. We provide a generic resource semantics for our logic as well as a sound, decidable, syntactic refinement-checking system. We also prove that refinements give rise to an optimization principle for programs. Finally, we illustrate the power of our system through a number of examples.

Categories and Subject Descriptors

D.3.1 [Programming Languages] : Formal Definitions and Theory—semantics; F.3.1 [Logic and Meanings of Programs] : Specifying and Verifying and Reasoning about Programs—pre- and post-conditions, invariants; F.3.2 [Logic and Meanings of Programs] : Semantics of Programming Languages—operational semantics

General Terms

Languages, Design, Theory

Keywords

Effectful Computation, Type Refinement, Linear Logic, Local Reasoning, Type Theory

*Supported in part by NSF Trusted Computing grant CCR-0208601, NSF Career grant CCR-0238328 (Programming Languages for Secure and Reliable Component Software) and a generous gift from Microsoft Research.

†Supported by the National Science Foundation under grant CCR-0204248: Type Refinements from September 1, 2002 until August 31, 2005.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICFP'03, August 25–29, 2003, Uppsala, Sweden.

Copyright 2003 ACM 1-58113-756-7/03/0008 ...\$5.00.

1. INTRODUCTION

One of the major goals of programming language design is to allow programmers to express and enforce properties of the execution behavior of programs. Conventional type systems, especially those with polymorphism and abstract types, provide a simple yet remarkably effective means of specifying program properties. However, there remain many properties which, while apparent at compile-time, cannot be checked using conventional type systems.

For this reason, there has been substantial interest in the formulation of *refinements* of conventional types that allow programmers to specify such properties. For example, Davies and Pfenning [6] show how to extend ML with intersection types and domain-specific predicates, and Xi and Pfenning [39] popularized the use of singleton types. They also present compelling applications including static array-bounds checking [38].

A separate research thread has shown how to use “type-and-effect” to check properties of programs involving state. These properties include safe region-based memory management [35], safe locking [11], and the correctness of correspondence assertions for communication protocols [15].

However, neither refinement types nor the many type-and-effect systems attempt to provide a general-purpose logical framework for reasoning about effectful computations. Xi and Pfenning’s dependent type system and related work [8, 2, 5, 6] only seek to capture properties of values and pure computations, rather than properties of *effectful* computations. For example, they are unable to describe *protocols* that require effectful functions to be used in a specified order. Therefore, these systems cannot be used to enforce important invariants such as the fact that a lock be held before a data structure is accessed or that a file is opened before being read or closed. On the other hand, with few exceptions, type-and-effect systems, which clearly take state into account, have focused on applying the type-and-effect methodology to solve specific problems, rather than on supporting a parameterized theory and general-purpose logic for encoding domain-specific program invariants.

We propose a new system of type refinements that allows the programmer to reason about program values and state in a unified manner, without being tied to a particular problem area. A type refinement combines a detailed description of the value produced by a computation with a formula describing the state of program resources after the computation’s execution. The formula is written in a general logic for reasoning about state, whose atomic predicates are determined by the programmer. We extend our

system to higher-order functions by including state pre- and post-conditions in refinements of function types.

Our work can be seen as a continuation of earlier work on refinement types, as well as a general language in which to design and use type-and-effect systems. It also serves to provide a semantic framework for understanding practical work in this area such as the Vault programming language [7]. Overall, the goal of our research is to provide a general, robust and extensible theory of type refinements that captures sound techniques for local reasoning about program state. We describe the main contributions of our system below.

A Two-Level System Based on Conservative Extension. We formalize the notion of a *type refinement* and construct a two level system for checking properties of programs. The first level involves simple type checking and the second level introduces our *logic of refinements* for reasoning about program properties that cannot be captured by conventional types. We establish a formal correspondence between types, which define the structure of a language, and refinements, which define domain-specific properties of programs written in a language. Only Denney [8] has explicitly considered such a two level system in the past, but he restricted his attention to pure computations.

We are careful to ensure that type refinements are a *conservative extension* of types. In other words, type refinements refine the information provided by the underlying, conventional type system rather than replace it with something different. The principle of conservative extension makes it possible for programmers to add type refinements gradually to legacy programs or to developing programs, to make these programs more robust.

A Parameterized Theory. The computational lambda calculus [20] serves as our basic linguistic framework. We parameterize this base language with a set of abstract base types, effectful operators over these types, and possible worlds. Consequently, our theorems hold for a very rich set of possible effects and effectful computations. In addition, we have separated our central type-checking rules from the specifics of the logic of refinements. Our theorems will hold for a variety of fragments of linear logic and we conjecture that similar substructural logics can be used in its place with little or no modification to the core system.

Support for Local Reasoning. It is essential that it be possible to focus attention on a single facet of program state without explicit reference to any others. For example, we may wish to reason about an I/O protocol separately from a locking protocol or a state protocol. Moreover, we may wish to reason separately about two different instances of the same state protocol. Formalisms based on classical logic are only sound for *global* reasoning about the entire state of a program. Not only is this unworkable as a practical matter, it is incompatible with modularity, which is essential for all but the simplest programs.

To support modularity it is necessary to employ a logic for *local* reasoning about independent facets of a program’s state [17, 26, 31]. Non-linear logic does not permit local reasoning, essentially because it validates the structural principles of weakening and contraction. Substructural logics, such as Linear Logic [14], or Bunched Implications (BI) [25, 17], do not validate these principles, and so are good candidates for a logic supporting local reasoning. We employ a fragment of linear logic that is adequate for many practical purposes.

For local reasoning to be sound, the underlying effectful operators in the language must in fact “act locally.” More precisely, we have identified a crucial *locality condition* on effectful operators that is necessary for soundness in the presence of local reasoning. We have proven the soundness of refinement checking in the presence of this locality condition. The soundness of refinement checking not only provides a means for checking certain correctness criteria, it also entails an optimization principle for effectful operators.

A Decidable System of Type Refinements. While we focus in this paper on a declarative presentation of our system of type refinements, we have developed a decidable, algorithmic refinement-checking system and proven it both sound and complete. A key aspect of the algorithmic system is the introduction of refinement annotations to the syntax of the language. These annotations allow the programmer to guide the checker so that it may search for type-refinement derivations in a deterministic fashion.

A Semantics for an Important Fragment of Vault. We provide a number of examples to demonstrate the expressiveness of our system. Based on these examples, our refinements appear to subsume the state-logic used in the Vault programming language [7] (although our idealized language does not contain the array of data structures present in Vault, nor the specialized type-inference techniques). Hence, our system suggests a semantics for an important fragment of Vault.

A Simple Example: File Access

Before digging in to the technical details of our framework, we present a simple example that introduces a number of important concepts. The example revolves around enforcing a simple resource usage protocol for file access and the following interface defines the types of each operation over files.

```
fnew : String → File
fopen : File → unit
fclose: File → unit
fwrite: (File, String) → unit
```

Informally, the protocol for using these operators requires that a file must be open before being written or closed, and should be closed before the program terminates. However, the interface, as presented, has no way to enforce this protocol. We demonstrate this point with a short function and use of that function that type check but do not meet the requirements of the file interface. The key problem is that the function makes an informal assumption about its arguments and the assumption is not satisfied before the function is used. The example is written in a syntax similar to that of Java. Note that `header` and `footer` are intended as strings defined elsewhere in the program.

```
// f must be open
saveEntry(f:File, entry:String) : unit {
    fwrite(f, header);
    fwrite(f, entry);
    fwrite(f, footer);
}

:
:

File myFile = fnew('data.txt');
saveEntry(myFile, myString); // Error: f not open.
```

To address this type of problem, we begin by specializing the file interface with type refinements that express the access protocol mentioned above.

```

fnew  : (String; 1)  $\rightarrow$   $\exists[f : \mathbf{File}]$ (Its(f); closed(f))
fopen : (Its(f); closed(f))  $\rightarrow$  (unit; open(f))
fclose: (Its(f); open(f))  $\rightarrow$  (unit; closed(f))
fwrite: (Its(f), String; open(f))  $\rightarrow$  (unit; open(f))

```

In the refined interface, the function arguments now have two components. The first component is either a refinement that corresponds to a conventional type or a more specific *singleton type* that specifies the exact value of the argument. The second component is a logical formula that describes the state of the system. Function results also have two components, with the added detail that results can be existentially quantified, as is the case with **fnew**. Note that all refinements are implicitly universally quantified over the free variables in the refinement. For example, **fopen** is implicitly quantified with $\forall[f : \mathbf{File}]$.

Two kinds of formulas appear in the example: predicates and the formula **1**. As a precondition, **1** states that the function requires, and therefore affects, none of the existing state. As a postcondition it states that the function produces no state.

We can now explain the refined interface. The operator **fnew** takes a **String** (the file name) and requires no existing state. It returns some **File** object *f* and assures that *f* is initially closed. **fopen** takes a particular *f* and changes its state from closed to open. **fclose** does the reverse. Finally, **fwrite** requires that *f* be open, but does not change any state relating to *f*.

Next, we add type refinements to the **saveEntry** function and its use, and see how a refinement checker would catch the protocol violation. Notice how we can now formally check what was previously an informal assumption.

```

saveEntry(f:File, entry:String; open(f)): (unit; open(f)) {
  fwrite(f, header); // State: open(f)
  fwrite(f, entry); // State: open(f)
  fwrite(f, footer); // State: open(f)
}

```

```

:
:

```

```

File myFile = fnew(‘‘data.txt’’); // State: closed(f)
saveEntry(myFile, myString); // Error: closed(f) != open(f)

```

Finally, we fix the use of **saveEntry** by inserting a call to **fopen** in the appropriate place, as shown below:

```

File myFile = fnew(‘‘data.txt’’); // State: closed(f)
fopen(myFile); // State: open(f)
saveEntry(myFile, myString); // State: open(f)

```

Our logic of refinements is much richer than the simple fragment we have used in this example as it includes all of the multiplicative and additive connectives of linear logic and some restricted uses of the modality “!”. For some more sophisticated examples, the reader may wish to skip ahead to section 4.

In the next section, we introduce our parameterized base language and its conventional type system (Section 2). In Section 3, we provide the syntax for general first-order refinements and provide a semantics for world (state) refinements. Next, we give a declarative account of refinement

checking and discuss how the declarative account is converted to an algorithmic one. Finally, we show that our refinements are a conservative extension of the underlying type system, state the main soundness theorems relating to our system and discuss how refined operators may be optimized. In the last section, we indicate our current research directions and comment further on related work.

Finally, due to space considerations, we have omitted from this paper the proofs of our theorems as well as a number of details that were not critical to the presentation of our work. Readers are encouraged to see our companion technical report [19] for complete details.

2. BASE LANGUAGE

We use Moggi’s computational λ -calculus [20] as a basic linguistic framework, as reformulated by Pfenning and Davies [28]. The framework is enriched with recursive functions and a base type of booleans. In order to consider a variety of different sorts of effects, we parameterize the language by a collection of abstract types **a**, constants *c* with type **a** and a set of multi-ary operators **o** over these abstract types.

2.1 Abstract Syntax

The abstract syntax of the language is defined by the following grammar:

<i>Types</i>	<i>A</i>	$::=$	a Bool $A_1 \rightarrow A_2$ $A_1 \multimap A_2$
<i>Var’s</i>	<i>X</i>	$::=$	<i>x</i> <i>y</i> ...
<i>Values</i>	<i>V</i>	$::=$	<i>X</i> <i>c</i> true false $\lambda(X).M$ fun <i>X</i> ($X_1:A_1$) : A_2 is <i>E</i>
<i>Terms</i>	<i>M</i>	$::=$	<i>V</i> if <i>M</i> then M_1 else M_2 $M(M_1)$
<i>Exp’s</i>	<i>E</i>	$::=$	M o (M_1, \dots, M_k) let <i>X</i> be E_1 in E_2 end app (<i>M</i> , M_1) if <i>M</i> then E_1 else E_2

The binding conventions are as expected; we identify expressions up to consistent renaming of bound variables. The type $A_1 \rightarrow A_2$ is the type of “pure” functions, which always terminate without effect, and the type $A_1 \multimap A_2$ is the type of “impure” functions, which may not terminate and may have an effect when applied.

2.2 Abstract Resources

Our language is parameterized by a set of abstract, effectful operators, which manipulate some abstract resource or set of resources. We may reason about an instance of the language by specifying an interface for and implementation of these operators and resources. In the future, we intend to extend our language with a full-fledged module system and an internal means of defining new resources.

We summarize the language parameters in Figure 1. An interface Σ defines a set of abstract types \mathcal{B} , a set of constants \mathcal{C} , and a set of operators \mathcal{O} . The interface also provides a signature Σ_A that gives types to the constants and operators. When we come to checking refinements, we will do so with respect to a set of predicates \mathcal{P} , an interface Σ_p to specify the types of predicate arguments and finally, a signature Σ_ϕ to define the refinements of each constant or operator.

An implementation $\mathcal{M} = (\mathcal{W}, \mathcal{T})$ defines a set \mathcal{W} of worlds *w*, and a transition function \mathcal{T} between these worlds that specifies the behavior of the operators over constants of the appropriate types.

A world *w* is a pair (**Per**(*w*), **Eph**(*w*)) where **Per**(*w*) is a set of *persistent* facts and **Eph**(*w*) is a multiset of *ephemeral*

Interface	Contents
\mathcal{B} \mathcal{C} \mathcal{O} Σ_A \mathcal{P} Σ_p Σ_ϕ	Base Types Constant Names Operator Names Constant and Operator Types Predicates Predicate Types Constant and Operator Refinements
Implementation	Meaning
\mathcal{W} $\text{Per}(w)$ $\text{Eph}(w)$ $\mathcal{T}(o)$	Worlds w 's Persistent Facts w 's Ephemeral Facts o 's Behavior

Figure 1: Language Parameters

facts. The persistent facts of a world w will remain true in all worlds that can be reached through a computation starting with w . The ephemeral facts of a world may or may not hold in its future worlds.

The notation $w_1 + w_2$ denotes a world containing the union of the persistent facts from w_1 and w_2 , and the multi-set union of ephemeral facts from w_1 and w_2 .

If an operator is given type $\mathbf{a}_1, \dots, \mathbf{a}_n \rightarrow \mathbf{a}$ by an interface, then the transition function $\mathcal{T}(o)$ is a total function from a sequence of constants with types $\mathbf{a}_1, \dots, \mathbf{a}_n$ and world w to a constant with type \mathbf{a} and world w' . We use the symbol \rightarrow to note that while these operators always terminate, they may have effects on the world. We require that these functions act monotonically on the persistent facts in the world. In other words, if $\mathcal{T}(o)(c_1, \dots, c_n; w) = (c; w')$ then $\text{Per}(w) \subseteq \text{Per}(w')$.

The transition function $\mathcal{T}(o)$ must also obey a *locality* condition. In general, it may only have an effect on a part of the world, rather than the entire world. Most operators that one would like to define obey this locality condition. However, some useful operators do not. For example, in our system, programmers may not reason statically about a function such as `gc(roots)`, which deletes all resources except the resources referenced from the variable `roots`. We defer a formal explanation of this condition to Section 3.5 where we prove the soundness of refinement checking.

Example: File Access (continued). We now consider parameterizing the language with the file access primitives shown earlier in Section 1. We require three base types: a type for files **File**, for data **String** and the unit type **unit**. Our constants include a countable set of file handles and data (we use metavariables f and d to range over each of these sets, respectively) and a unit value $()$. The signature Σ_A provides the types for the operations and constants.

$$\begin{aligned}
 \mathcal{T}(\mathbf{fnew})(); w &= (f; w') \\
 &\text{where } \text{Eph}(w') = \text{Eph}(w) + \{\text{closed}(f)\} \\
 &\text{and } f \notin \text{Eph}(w) \\
 \mathcal{T}(\mathbf{fopen})(f; w) &= ((); w') \\
 &\text{if } w = w'' + \{\text{closed}(f)\} \\
 &\text{and } \text{Eph}(w') = \text{Eph}(w'') + \{\text{open}(f)\} \\
 \mathcal{T}(\mathbf{fopen})(f; w) &= ((); w) \\
 &\text{if } w \neq w'' + \{\text{closed}(f)\} \\
 \mathcal{T}(\mathbf{fclose})(f; w) &= ((); w') \\
 &\text{if } w = w'' + \{\text{open}(f)\} \\
 &\text{and } \text{Eph}(w') = \text{Eph}(w'') + \{\text{closed}(f)\} \\
 \mathcal{T}(\mathbf{fclose})(f; w) &= ((); w) \\
 &\text{if } w \neq w'' + \{\text{open}(f)\}
 \end{aligned}$$

Figure 2: Transition Function of File Operators

$$\begin{aligned}
 \Sigma_A(); &= \mathbf{unit} \\
 \Sigma_A(f) &= \mathbf{File} \\
 \Sigma_A(d) &= \mathbf{String} \\
 \Sigma_A(\mathbf{fnew}) &= \mathbf{unit} \rightarrow \mathbf{File} \\
 \Sigma_A(\mathbf{fopen}) &= \mathbf{File} \rightarrow \mathbf{unit} \\
 \Sigma_A(\mathbf{fclose}) &= \mathbf{File} \rightarrow \mathbf{unit} \\
 \Sigma_A(\mathbf{fwrite}) &= (\mathbf{File}, \mathbf{String}) \rightarrow \mathbf{unit}
 \end{aligned}$$

The set of predicates \mathcal{P} is exactly the predicates `closed` and `open` from our earlier example.

In the implementation component, we must specify the set of worlds and the behavior of the operators. Notice that there are no persistent predicates, and so the persistent set of any world is always empty. However, once a file f is initialized with a call to `fnew`, either `open(f)` or `closed(f)` will be in the ephemeral facts any future world. Therefore, the \mathcal{W} of the implementation is the countably infinite set of worlds w , where $\text{Per}(w) = \emptyset$ and $\text{Eph}(w)$ describes a set of files that are all either open or closed, but not both. The transition function \mathcal{T} specifies the dynamic semantics for each operator and is shown in Figure 2. Note that the side condition $f \notin \text{Eph}(w)$ indicates that f does not appear in any of the predicates in $\text{Eph}(w)$ (i.e. f is “fresh”).

A key aspect of our definition of \mathcal{T} is that each of the operators are defined to be *total* functions on the entire domain of worlds. If they were not total functions we would be unable to prove a generic soundness theorem for our language. Later (see Section 3.7), we will prove an optimization principle that allows programmers to replace these total functions with the appropriate partial functions when their program has the necessary refinement.

2.3 Semantics

Since our base language semantics is almost entirely standard, we merely state the forms of the various judgments. Complete details can be found in our companion technical report [19].

The static semantics is given by the following two judgment forms.

$$\begin{aligned}
 \Gamma \vdash_M M : A &\quad \text{Term } M \text{ has type } A \text{ in } \Gamma \\
 \Gamma \vdash_E E : A &\quad \text{Expression } E \text{ has type } A \text{ in } \Gamma
 \end{aligned}$$

The meta-variable Γ ranges over finite functions from variables x to types A .

The dynamic semantics is given by the following two evaluation judgements.

$$\begin{array}{l} M \Downarrow V \quad \text{the term } M \text{ evaluates to value } V \\ E @ w \Downarrow V @ w' \quad \text{in } w \text{ the expression } E \text{ evaluates to } V \\ \text{and changes to } w' \end{array}$$

The following rule defines evaluation for effectful operators using the language parameter \mathcal{T} .

$$\frac{M_i \Downarrow c_i \quad (\text{for } 1 \leq i \leq n) \quad \mathcal{T}(\mathbf{o})(c_1, \dots, c_n; w) = c'; w'}{\mathbf{o}(M_1, \dots, M_n) @ w \Downarrow c @ w'} \quad (\text{D-E-OP})$$

Aside from this rule, the dynamic semantics is entirely standard. The language as a whole satisfies the standard type safety theorem [19].

3. REFINEMENTS

In order to define and check further, more specific, properties of values and computations than supported by the type system alone, we introduce a *logic of refinements* that may be layered on top of the computational lambda calculus described in the previous section.

3.1 Syntax

Our logic contains three classes of refinements. A *term refinement* is a predicate over a type, describing a more specific property of a term than a type alone. A *world refinement* is a formula describing the (implicit) state of a world. An *expression refinement* is a predicate over both the type of an expression and the state of the world after the expression is executed. The table below describes the syntax of term, world and expression refinements.

<i>Binding</i>	$b ::= c:\mathbf{a}$
<i>Term Refs</i>	$\phi ::= \mathbf{a} \mid \mathbf{Bool} \mid \mathbf{Its}(c) \mid \pi$
<i>Function Refs</i>	$\pi ::= \phi_1 \rightarrow \phi_2 \mid (\phi; \psi) \rightarrow \eta \mid \forall b. \pi$
<i>World Refs</i>	$\psi ::= p(c_1, \dots, c_n) \mid !p(c_1, \dots, c_n) \mid$ $\mathbf{1} \mid \psi_1 \otimes \psi_2 \mid \psi_1 \multimap \psi_2 \mid$ $\top \mid \psi_1 \& \psi_2 \mid \mathbf{0} \mid \psi_1 \oplus \psi_2$
<i>Expr. Refs</i>	$\eta ::= \exists[\vec{b}](\phi; \psi)$

Since we are concentrating on properties of effectful computations, we have chosen a minimalist logic of term refinements. There is a refinement that corresponds to each type in the base language as well as *singleton types* denoted $\mathbf{Its}(c)$. Partial functions are refined in order to specify a precondition for the state of the world on input and a postcondition consisting of an expression refinement. The precondition for a partial function could also have been an (existentially quantified) expression refinement, but this extension provides no gain in expressive power. We allow function refinements to be prefixed with first-order universal quantification.

The world refinements consist of the multiplicative-additive fragment of linear logic augmented with intuitionistic predicates $!p(c_1, \dots, c_n)$. The connectives $\mathbf{1}$, \otimes and \multimap form the multiplicative fragment of the logic whereas the connectives \top , $\&$, $\mathbf{0}$, and \oplus are known as the additives. Both \otimes and $\&$ are forms of conjunction. Intuitively, a world can be described by the formula $\psi_1 \otimes \psi_2$ if it can be split into two disjoint parts such that one part can be described by ψ_1 and the other part can be described by ψ_2 . On the other hand, a world satisfies $\psi_1 \& \psi_2$ if it can be described by both

$$\begin{array}{c} \frac{}{\vec{b} \vdash \mathbf{Bool} \sqsubseteq \mathbf{Bool}} \quad \frac{}{\vec{b} \vdash \mathbf{a} \sqsubseteq \mathbf{a}} \quad \frac{\Sigma_A(c) = \mathbf{a} \text{ or } c:\mathbf{a} \in \vec{b}}{\vec{b} \vdash \mathbf{Its}(c) \sqsubseteq \mathbf{a}} \\ \frac{\vec{b} \vdash \phi_1 \sqsubseteq A_1 \quad \vec{b} \vdash \phi_2 \sqsubseteq A_2}{\vec{b} \vdash \phi_1 \rightarrow \phi_2 \sqsubseteq A_1 \rightarrow A_2} \quad \frac{\vec{b} \vdash \phi_1 \sqsubseteq A_1 \quad \vec{b} \vdash \eta_2 \sqsubseteq_E A_2}{\vec{b} \vdash (\phi_1; \psi_1) \rightarrow \eta_2 \sqsubseteq A_1 \rightarrow A_2} \\ \frac{\vec{b}, c:\mathbf{a} \vdash \pi \sqsubseteq A}{\vec{b} \vdash \forall c:\mathbf{a}. \pi \sqsubseteq A} \quad \frac{\vec{b} \vdash \phi \sqsubseteq A}{\vec{b} \vdash (\phi; \psi) \sqsubseteq_E A} \quad \frac{\vec{b}, c:\mathbf{a} \vdash \exists[\vec{b}_1](\phi; \psi) \sqsubseteq_E A}{\vec{b} \vdash \exists[c:\mathbf{a}, \vec{b}_1](\phi; \psi) \sqsubseteq_E A} \end{array}$$

Figure 3: A Refinement of a Type

ψ_1 and ψ_2 simultaneously. The formulas $\mathbf{1}$ and \top are the identities for \otimes and $\&$ respectively. The formula \oplus is a disjunction and $\mathbf{0}$ is its identity.

When \vec{b} is the empty sequence in some expression refinement $\exists[\vec{b}](\phi; \psi)$, we often use the abbreviation $(\phi; \psi)$. We use the notation $\text{FV}(\phi)$ to denote the set of free variables appearing in the term refinement ϕ . We use a corresponding notation for world and expression refinements. We use the notation $[c'/b]X$ to denote capture-avoiding substitution of c' for c in term or world refinement X when $b = (c:\mathbf{a})$ and $\Sigma_A(c') = \mathbf{a}$. We extend this notation to substitution to a sequence of bindings as in $[c'_1, \dots, c'_n/\vec{b}]X$ or $[\vec{b}'/\vec{b}]X$. In either case, constants substituted for variables must have the correct type and the sequences must have the same length or else the substitution is undefined. We also extend substitution to persistent and ephemeral contexts in the ordinary way.

Every refinement refines a particular type. To formalize this relationship, we define two new judgments $\vec{b} \vdash \phi \sqsubseteq A$ and $\vec{b} \vdash \eta \sqsubseteq_E A$, which indicate that a term or expression refinement refines the type A given the set of bindings \vec{b} . Figure 3 defines these relations. Below, we present a lemma stating that any refinement refines a unique type.

Lemma 1

- If $\vec{b} \vdash \phi \sqsubseteq A_1$ and $\vec{b} \vdash \phi \sqsubseteq A_2$ then $A_1 = A_2$.
- If $\vec{b} \vdash \eta \sqsubseteq_E A_1$ and $\vec{b} \vdash \eta \sqsubseteq_E A_2$ then $A_1 = A_2$.

Finally, for every type A , there is a trivial refinement $\text{triv}(A)$ that refines it.

$$\begin{array}{l} \text{triv}(\mathbf{Bool}) = \mathbf{Bool} \\ \text{triv}(\mathbf{a}) = \mathbf{a} \\ \text{triv}(A_1 \rightarrow A_2) = \text{triv}(A_1) \rightarrow \text{triv}(A_2) \\ \text{triv}(A_1 \multimap A_2) = (\text{triv}(A_1); \top) \multimap (\text{triv}(A_2); \top) \end{array}$$

3.2 Semantics of Refinements

We were inspired to define a semantics for our world refinements by the work of Ishtiaq and O'Hearn [17]. The semantics appears in Figure 4. The model (world) w used in the semantics is that described earlier in section 2.2. The fragment of the logic without the modality $!$ is an instance of Simon Ambler's resource semantics [1, p. 30-32]. It relies upon an abstract relation \lesssim which defines the relationship between primitive facts. For example, in a system containing arithmetic predicates such as $\text{less}(x, y)$, the relation would include $\text{less}(x, 3) \lesssim \text{less}(x, 5)$. In most of our examples, the

$w \models \psi$ if and only if

- $\psi = p(c_1, \dots, c_n)$, $\text{Eph}(w) = \{X\}$ and $X \lesssim p(c_1, \dots, c_n)$
- $\psi = !p(c_1, \dots, c_n)$, $X \in \text{Per}(w)$ and $X \lesssim p(c_1, \dots, c_n)$ and $\text{Eph}(w) = \emptyset$
- $\psi = \mathbf{1}$ and $\text{Eph}(w) = \emptyset$
- $\psi = \psi_1 \otimes \psi_2$ and there exist w_1, w_2 , such that $w = w_1 + w_2$ and $w_1 \models \psi_1$ and $w_2 \models \psi_2$
- $\psi = \psi_1 \multimap \psi_2$ and for all worlds w_1 such that $w_1 \models \psi_1$, $w_1 + w \models \psi_2$
- $\psi = \top$ (and no other conditions need be satisfied)
- $\psi = \psi_1 \& \psi_2$ and $w \models \psi_1$ and $w \models \psi_2$
- $\psi = \mathbf{0}$ and false (this refinement can never be satisfied).
- $\psi = \psi_1 \oplus \psi_2$ and either $w \models \psi_1$ or $w \models \psi_2$.

$w \models \Omega$ iff $\text{Per}(w) \supseteq \Omega$.

$w \models \cdot$ iff $\text{Eph}(w) = \emptyset$.

$w \models \psi_1, \dots, \psi_n$ iff there exist w_1, \dots, w_n such that

- $w = w_1 + \dots + w_n$ and
- $w_1 \models \psi_1, \dots, w_n \models \psi_n$.

$w \models \Omega; \Delta$ iff $w \models \Omega$ and $w \models \Delta$.

Figure 4: Semantics of World Refinements and Contexts

relation \lesssim will simply be the identity relation. In other words, our predicates are usually left uninterpreted.

The semantics of world refinements is extended to closed persistent contexts Ω (lists of predicates $p(\vec{c})$) and ephemeral contexts Δ (lists of world refinements), and appears in Figure 4. We treat both kinds of contexts as equivalent up to reordering of their elements.¹

We will show later that linear logical entailment is sound with respect to our semantics. However, as noted by Amblar [1, p. 32], there is no sense in which linear logical reasoning is complete with respect to this semantics. Despite this deficiency, linear logic has proven to be very useful for many applications. We leave definition of a sound and complete logic for our resource semantics to future work.

3.3 Declarative Refinement-Checking

In this section, we give a declarative account of how to check that a (possibly open) term or expression has a given refinement. Refinement checking of open terms will occur within a context of the following form. Whenever we consider the semantics of refinements or refinement checking, we presuppose that the values, terms and expressions in question are well-formed with an appropriate type.

$$\begin{array}{l} \text{Persistent Ctxt } \Omega \quad ::= \cdot \mid \Omega, c:\mathbf{a} \mid \Omega, x:\phi \mid \Omega, p(\vec{c}) \\ \text{Ephemeral Ctxt } \Delta \quad ::= \cdot \mid \Delta, \psi \end{array}$$

Furthermore, we define a derivative form of context, Ω_b to be a vector, \vec{b} , consisting of all elements in Ω of the form $c:\mathbf{a}$.

Persistent contexts are constrained so that the variables c and x appear at most once to the left of any $:$ in the context and no variables c may shadow any of the constants in \mathcal{C} . When necessary, we will implicitly alpha-vary bound

¹When we extend Ω to open contexts which include constant declarations, reordering must respect the dependencies introduced by such declarations (see Section 3.3).

$\vdash \Sigma \text{ ok}$	<i>Signature Σ is well-formed</i>
$\vdash \Omega \text{ ok}$	<i>Context Ω is well-formed</i>
$\Omega \vdash \Delta \text{ ok}$	<i>Context Δ is well-formed in Ω</i>
$\Omega \vdash \phi \text{ ok}$	<i>Refinement ϕ is well-formed in Ω</i>
$\Omega \vdash \psi \text{ ok}$	<i>World ref. ψ is well-formed in Ω</i>
$\Omega \vdash \eta \text{ ok}$	<i>Expression ref. η is well-formed in Ω</i>
$\Omega \gg_M M : \phi$	<i>Term M has refinement ϕ in Ω</i>
$\Omega; \Delta \gg_E E : \eta$	<i>Expression E has ref. η in $\Omega; \Delta$</i>
$\Omega; \phi \implies_M \phi'$	<i>Term refinement ϕ entails ϕ' in Ω</i>
$\Omega; \Delta \implies_W \psi$	<i>Context Δ entails ψ in Ω</i>
$\Omega; \Delta; \eta \implies_E \eta'$	<i>Expression ref. η entails η' in $\Omega; \Delta$</i>
$\Omega; \Delta \rightsquigarrow (\Omega_i; \Delta_i)_n$	<i>Context $\Omega; \Delta$ reduces to the context list $(\Omega_i; \Delta_i)_n$ in one step</i>
$\Omega; \Delta \rightsquigarrow^* (\Omega_i; \Delta_i)_n$	<i>Context $\Omega; \Delta$ reduces to the context list $(\Omega_i; \Delta_i)_n$ in 0 or more steps</i>

Figure 5: Refinement-Checking Judgments

variables to maintain this invariant. We treat contexts that differ only in the order of the elements as equivalent and do not distinguish them (provided both contexts in question are well formed; in other words, reordering must respect dependencies.). Occasionally, we call the persistent context *unrestricted* and the ephemeral context *linear*. Both contraction and weakening hold for the unrestricted context while neither of these structural properties holds for the linear context.

Declarative refinement checking is formulated using the judgment forms in Figure 5. All but the first judgment are implicitly parameterized by a fixed, well-formed interface Σ .

The first six judgments in the list are relatively standard. They simply check that each sort of type or context is well-formed in the context Ω . This check amounts to the fact that constants and variables that appear in a type or context appear bound previously in the context or in the signature. The formal rules appear in the technical report.

The next two judgments form the heart of the system. They check terms and expressions to ensure that they have appropriate refinements. The term refinement-checking rules may be found in Figure 6. We point out a few important details. In rule (R-T-CONST) constants c are given very precise *singleton types*, following work by Xi and Pfenning [39]. Also, rule (R-T-TAPP) does not consider the case that the function in an application has a polymorphic refinement. This possibility is taken care of by the (R-T-SUB) rule, which instantiates universal quantifiers implicitly. Such instantiations can be resolved by standard first-order unification. Furthermore, rule (R-T-IF) does not check that the first term M has a boolean refinement, because we assume that refinement checking is preceded by ordinary type checking.

The expression refinement-checking rules appear in Figure 7. Rule (R-E-TERM) defines the interface between pure and effectful computations. As pure terms neither require nor produce state, they are checked in an empty ephemeral context and given the world refinement $\mathbf{1}$. We use the (R-E-SUB) rule (discussed in more detail below) to properly check terms within a non-empty ephemeral context. The rule for checking operators requires that we guess a sequence of con-

$$\begin{array}{c}
\frac{}{\Omega, x : \phi \gg_M x : \phi} \quad (\text{R-T-VAR}) \\
\\
\frac{c \in \text{Dom}(\Sigma_\phi)}{\Omega \gg_M c : \text{Its}(c)} \quad (\text{R-T-CONST}) \\
\\
\frac{}{\Omega \gg_M \text{true} : \text{Bool}} \quad (\text{R-T-TRUE}) \\
\\
\frac{}{\Omega \gg_M \text{false} : \text{Bool}} \quad (\text{R-T-FALSE}) \\
\\
\frac{\Omega_b, \vec{b} \vdash \phi_1 \sqsubseteq A \quad \Omega, \vec{b} \vdash \phi_1 \text{ ok} \quad \Omega, \vec{b}, x : \phi_1 \gg_M M : \phi_2 \quad (\phi = \forall \vec{b} \cdot \phi_1 \rightarrow \phi_2)}{\Omega \gg_M \lambda(x:A).M : \phi} \quad (\text{R-T-LAM}) \\
\\
\frac{\Omega_b \vdash \phi \sqsubseteq A_1 \rightarrow A \quad \Omega \vdash \phi \text{ ok} \quad \Omega, x : \phi, \vec{b}, x_1 : \phi_1; \psi_1 \gg_E E : \eta \quad (\phi = \forall \vec{b} \cdot (\phi_1; \psi_1) \rightarrow \eta)}{\Omega \gg_M \text{fun } x (x_1:A_1) : A \text{ is } E : \phi} \quad (\text{R-T-FUN}) \\
\\
\frac{\Omega \gg_M M_1 : \phi \quad \Omega \gg_M M_2 : \phi}{\Omega \gg_M \text{if } M \text{ then } M_1 \text{ else } M_2 : \phi} \quad (\text{R-T-IF}) \\
\\
\frac{\Omega \gg_M M : \phi_1 \rightarrow \phi_2 \quad \Omega \gg_M M_1 : \phi_1}{\Omega \gg_M M (M_1) : \phi_2} \quad (\text{R-T-TAPP}) \\
\\
\frac{\Omega \gg_M M : \phi \quad \Omega; \phi \Longrightarrow_M \phi'}{\Omega \gg_M M : \phi'} \quad (\text{R-T-SUB})
\end{array}$$

Figure 6: Refinement-Checking for Terms

starts to substitute for the polymorphic parameters in the operator refinement. Given this substitution, we check that operator arguments may be given refinements equal to their corresponding formal parameter.

There are three expression checking rules that are not syntax-directed. (R-E-SUB) merits special attention as it is the key to local reasoning. The rule splits the context into two disjoint parts, Δ_1 and Δ_2 , where Δ_1 is used to check the expression E , and Δ_2 passes through unused. As a result, the computation may be written in ignorance of the total global state. It need only know how to process the local state in Δ_1 . In fact, in the case that Δ_1 is empty, the computation may be completely pure. Additionally, (R-E-SUB) serves as a conventional subsumption rule in which we check that one expression refinement entails the other. (R-E-CUT) is the logical cut rule: If we can prove some intermediary result (ψ) which in turn makes it possible to demonstrate our final goal ($E : \eta$) then we should be able to prove our final goal from our original premises. Since Δ contains linear hypotheses that must not be duplicated, we split the context into two parts Δ_1 and Δ_2 , one part for each premise in the rule.

Finally, since proofs in substructural logics require careful manipulation of the context, we introduce a new rule (R-E-CONTEXT) to control context evolution during type checking. This rule depends upon the judgment $\Omega; \Delta \rightsquigarrow^* (\Omega_i; \Delta_i)_n$ which encodes the action of all natural left rules from the sequent calculus for linear logic. The notation $(\Omega_i; \Delta_i)_n$ stands for a list of (possibly zero) contexts $(\Omega_1; \Delta_1), \dots,$

$$\begin{array}{c}
\frac{\Omega \gg_M M : \phi}{\Omega; \cdot \gg_E M : (\phi, \mathbf{1})} \quad (\text{R-E-TERM}) \\
\\
\frac{\Omega \gg_M M_i : [\vec{c}/\vec{b}] \phi_i \quad (\text{for } 1 \leq i \leq n) \quad (\Sigma_\phi(\mathbf{o}) = \forall \vec{b} \cdot (\phi_1, \dots, \phi_n; \psi_1) \rightarrow \eta)}{\Omega; [\vec{c}/\vec{b}] \psi_1 \gg_E \mathbf{o}(M_1, \dots, M_n) : [\vec{c}/\vec{b}] \eta} \quad (\text{R-E-OP}) \\
\\
\frac{\Omega; \Delta \gg_E E_1 : \eta_1 \quad \Omega, \vec{b}_1, x : \phi_1; \psi_1 \gg_E E_2 : \eta_2 \quad (\vec{b}_1 \notin \text{FV}(\eta_2))}{\Omega; \Delta \gg_E \text{let } x \text{ be } E_1 \text{ in } E_2 \text{ end} : \eta_2} \quad (\eta_1 = \exists[\vec{b}_1](\phi_1; \psi_1)) \quad (\text{R-E-LET}) \\
\\
\frac{\Omega \gg_M M : (\phi_1; \psi_1) \rightarrow \eta \quad \Omega \gg_M M_1 : \phi_1}{\Omega; \psi_1 \gg_E \text{app}(M, M_1) : \eta} \quad (\text{R-E-PAPP}) \\
\\
\frac{\Omega; \Delta \gg_E E_1 : \eta \quad \Omega; \Delta \gg_E E_2 : \eta}{\Omega; \Delta \gg_E \text{if } M \text{ then } E_1 \text{ else } E_2 : \eta} \quad (\text{R-E-IF}) \\
\\
\frac{\Omega; \Delta \rightsquigarrow^* (\Omega_i; \Delta_i)_n \quad \Omega_i; \Delta_i \gg_E E : \eta \quad (\text{for } 1 \leq i \leq n)}{\Omega; \Delta \gg_E E : \eta} \quad (\text{R-E-CONTEXT}) \\
\\
\frac{\Omega; \Delta_2 \Longrightarrow_W \psi \quad \Omega; \Delta_1, \psi \gg_E E : \eta}{\Omega; \Delta_1, \Delta_2 \gg_E E : \eta} \quad (\text{R-E-CUT}) \\
\\
\frac{\Omega; \Delta_1 \gg_E E : \eta \quad \Omega; \Delta_2; \eta \Longrightarrow_E \eta'}{\Omega; \Delta_1, \Delta_2 \gg_E E : \eta'} \quad (\text{R-E-SUB})
\end{array}$$

Figure 7: Refinement-Checking for Expressions

$(\Omega_n; \Delta_n)$. We specifically use the word *reduces* since every valid judgment of this form reduces the number of connectives in the context when read from left to right. Most of the rules produce one context. However, the rule for disjunction produces two contexts (and E must have the same refinement in both of them) and the rule for falsehood produces no context (and we can choose any well-formed expression refinement for E without further checking). We extend the one-step context reduction judgment to its reflexive and transitive closure, which we denote $\Omega; \Delta \rightsquigarrow^* (\Omega_i; \Delta_i)_n$.

The last five judgments involved in refinement checking specify the logical component of the system. We have already discussed the context reduction judgments. This judgment is combined with the right rules from the sequent calculus and the cut rule in the judgment $\Omega; \Delta \Longrightarrow_W \psi$ to provide a full proof system for our fragment of linear logic. The judgment $\Omega; \phi \Longrightarrow_M \phi'$ is the corresponding proof system for term refinements. Notice that these rules do not depend upon the linear context Δ . Since terms are pure, their refinements should not depend upon ephemeral state. Finally, the judgment for expression refinement entailment $\Omega; \Delta; \eta \Longrightarrow_E \eta'$ combines the world and term proof systems with rules for existentials. These judgments are formally defined in Figures 8, 9, 10 and 11.

3.4 Algorithmic Refinement-Checking

We have developed an algorithmic refinement-checking system that is both sound and complete with respect to the

$$\begin{array}{c}
\frac{}{\Omega; \mathbf{a} \Longrightarrow_M \mathbf{a}} \quad (\text{L-T-BASE}) \\
\frac{}{\Omega; \mathbf{Bool} \Longrightarrow_M \mathbf{Bool}} \quad (\text{L-T-BOOL}) \\
\frac{}{\Omega; \mathbf{Its}(c) \Longrightarrow_M \mathbf{Its}(c)} \quad (\text{L-T-ITS}) \\
\frac{\Sigma_A(c) = \mathbf{a} \text{ or } \Omega(c) = \mathbf{a}}{\Omega; \mathbf{Its}(c) \Longrightarrow_M \mathbf{a}} \quad (\text{L-T-ITSBASE}) \\
\frac{\Omega; \phi'_1 \Longrightarrow_M \phi_1 \quad \Omega; \phi_2 \Longrightarrow_M \phi'_2 \quad \Omega \vdash \phi'_1 \text{ ok}}{\Omega; \phi_1 \rightarrow \phi_2 \Longrightarrow_M \phi'_1 \rightarrow \phi'_2} \quad (\text{L-T-TARR}) \\
\frac{\Omega; \phi'_1 \Longrightarrow_M \phi_1 \quad \Omega; \psi'_1 \Longrightarrow_W \psi_1 \quad \Omega; ; \eta \Longrightarrow_E \eta' \quad \Omega \vdash \phi'_1 \text{ ok} \quad \Omega \vdash \psi'_1 \text{ ok}}{\Omega; (\phi_1; \psi_1) \rightarrow \eta \Longrightarrow_M (\phi'_1; \psi'_1) \rightarrow \eta'} \quad (\text{L-T-PAARR}) \\
\frac{\Omega; [c'/c:\mathbf{a}]\pi \Longrightarrow_M \pi'}{\Omega; \forall c:\mathbf{a} \cdot \pi \Longrightarrow_M \pi'} \quad (\text{L-T-ALLL}) \\
\frac{\Omega; c:\mathbf{a}; \pi \Longrightarrow_M \pi'}{\Omega; \pi \Longrightarrow_M \forall c:\mathbf{a} \cdot \pi'} \quad (\text{L-T-ALLR})
\end{array}$$

Figure 8: Entailment for Term Refinements

$$\frac{}{\Omega; \Delta; !p(c_1, \dots, c_n) \rightsquigarrow \Omega; p(c_1, \dots, c_n); \Delta} \quad (\text{CR-!})$$

$$\frac{}{\Omega; \Delta; \mathbf{1} \rightsquigarrow \Omega; \Delta} \quad (\text{CR-1})$$

$$\frac{}{\Omega; \Delta; \psi_1 \otimes \psi_2 \rightsquigarrow \Omega; \Delta; \psi_1, \psi_2} \quad (\text{CR-MAND})$$

$$\frac{\Omega; \Delta_1 \Longrightarrow_W \psi_1}{\Omega; \Delta_1, \Delta_2, \psi_1 \multimap \psi_2 \rightsquigarrow \Omega; \Delta_2, \psi_2} \quad (\text{CR-IMP})$$

$$\frac{}{\Omega; \Delta; \psi_1 \& \psi_2 \rightsquigarrow \Omega; \Delta; \psi_1} \quad (\text{CR-AND1})$$

$$\frac{}{\Omega; \Delta; \psi_1 \& \psi_2 \rightsquigarrow \Omega; \Delta; \psi_2} \quad (\text{CR-AND2})$$

$$\frac{}{\Omega; \Delta; \mathbf{0} \rightsquigarrow} \quad (\text{CR-ZERO})$$

$$\frac{}{\Omega; \Delta; \psi_1 \oplus \psi_2 \rightsquigarrow (\Omega; \Delta; \psi_1), (\Omega; \Delta; \psi_2)} \quad (\text{CR-OR})$$

$$\frac{}{\Omega; \Delta \rightsquigarrow^* \Omega; \Delta} \quad (\text{CR*-REFLEX})$$

$$\frac{\Omega; \Delta \rightsquigarrow (\Omega_j; \Delta_j)_m \quad \Omega_j; \Delta_j \rightsquigarrow^* (\Omega_{j_k}; \Delta_{j_k})_{n_j} \quad (\text{for } 1 \leq j \leq m)}{\Omega; \Delta \rightsquigarrow^* (\Omega_{1_k}; \Delta_{1_k})_{n_1}, \dots, (\Omega_{m_k}; \Delta_{m_k})_{n_m}} \quad (\text{CR*-TRANS})$$

Figure 9: Context Reduction and Its Closure

$$\begin{array}{c}
\frac{}{\Omega; \psi \Longrightarrow_W \psi} \quad (\text{L-E-HYP}) \\
\frac{}{\Omega; p(c_1, \dots, c_n); \cdot \Longrightarrow_W !p(c_1, \dots, c_n)} \quad (\text{L-E-!R}) \\
\frac{}{\Omega; \cdot \Longrightarrow_W \mathbf{1}} \quad (\text{L-E-1R}) \\
\frac{\Omega; \Delta_1 \Longrightarrow_W \psi_1 \quad \Omega; \Delta_2 \Longrightarrow_W \psi_2}{\Omega; \Delta_1, \Delta_2 \Longrightarrow_W \psi_1 \otimes \psi_2} \quad (\text{L-E-MANDR}) \\
\frac{\Omega; \Delta; \psi_1 \Longrightarrow_W \psi_2 \quad \Omega \vdash \psi_1 \text{ ok}}{\Omega; \Delta \Longrightarrow_W \psi_1 \multimap \psi_2} \quad (\text{L-E-IMPR}) \\
\frac{}{\Omega; \Delta \Longrightarrow_W \top} \quad (\text{L-E-TR}) \\
\frac{\Omega; \Delta \Longrightarrow_W \psi_1 \quad \Omega; \Delta \Longrightarrow_W \psi_2}{\Omega; \Delta \Longrightarrow_W \psi_1 \& \psi_2} \quad (\text{L-E-ANDR}) \\
\frac{\Omega; \Delta \Longrightarrow_W \psi_1}{\Omega; \Delta \Longrightarrow_W \psi_1 \oplus \psi_2} \quad (\text{L-E-ORR1}) \\
\frac{\Omega; \Delta \Longrightarrow_W \psi_2}{\Omega; \Delta \Longrightarrow_W \psi_1 \oplus \psi_2} \quad (\text{L-E-ORR2}) \\
\frac{\Omega; \Delta \rightsquigarrow (\Omega_i; \Delta_i)_n \quad \Omega_i; \Delta_i \Longrightarrow_W \psi \quad (\text{for } 1 \leq i \leq n)}{\Omega; \Delta \Longrightarrow_W \psi} \quad (\text{L-E-LEFT}) \\
\frac{\Omega; \Delta_2 \Longrightarrow_W \psi_1 \quad \Omega; \Delta_1, \psi_1 \Longrightarrow_W \psi}{\Omega; \Delta_1, \Delta_2 \Longrightarrow_W \psi} \quad (\text{L-E-CUT})
\end{array}$$

Figure 10: Entailment for World Refinements

$$\frac{\Omega; \phi \Longrightarrow_M \phi' \quad \Omega; \Delta; \psi \Longrightarrow_W \psi'}{\Omega; \Delta; (\phi; \psi) \Longrightarrow_E (\phi'; \psi')} \quad (\text{L-ER-BASE})$$

$$\frac{\Omega; \Delta; \eta \Longrightarrow_E [c'/c:\mathbf{a}]\exists[\vec{b}](\phi; \psi)}{\Omega; \Delta; \eta \Longrightarrow_E \exists[c:\mathbf{a}, \vec{b}](\phi; \psi)} \quad (\text{L-ER-EXISTS R})$$

$$\frac{\Omega; c:\mathbf{a}; \Delta; \exists[\vec{b}](\phi; \psi) \Longrightarrow_E \eta}{\Omega; \Delta; \exists[c:\mathbf{a}, \vec{b}](\phi; \psi) \Longrightarrow_E \eta} \quad (\text{L-ER-EXISTS L})$$

Figure 11: Entailment for Expression Refinements

system presented above. There is one typing rule for each expression or term construct and all premises in the rules are fully determined, except those of the context-reduction judgment. We developed the system in two steps, outlined below.

The first step is cut elimination. We eliminate the two cut rules (the cut rule for expression refinement-checking and the cut rule for linear logic entailment) and show that the resulting system is sound and complete with respect to the original refinement-checking specification. We carry out the proof by modifying and extending the logical cut elimination proof in earlier work by Pfenning [27].

In the second step we eliminate the subsumption rule and introduce annotations in order to eliminate two critical sources of non-determinism present in the previous system. The first source is the non-syntax-directedness of the subsumption rule. We therefore incorporate the subsumption rule into the language in a syntax-directed manner, and modify the expression rules so that the context-splitting of the subsumption rule is deterministic.

The second source of non-determinism is the need for the refinement-checker to “guess” the refinement of a given term or expression when the refinement cannot be deduced. We therefore introduce type refinement annotations into the language, allowing the programmer to supply the checker with the missing refinement. In order to reduce the annotation burden, we have defined a bi-directional refinement checking algorithm. The essence of this system is the introduction of two new refinement-checking judgments: one for refinement inference and one for refinement checking. The former judgment infers a refinement for the given term or expression and produces it as an output. The latter judgment takes a refinement as input and checks the given term or expression against the refinement.

Our system also requires a second form of annotation to guide the use of the context rule in expressions. Essentially, these annotations specify when the refinement-checker must use the disjunctive left rule, which, when applied, causes the annotated expression to be rechecked in two different logical contexts. Rechecking the program text has the potential to be very expensive, so we place this facility under the control of the programmer. For a more detailed explanation of our type checking algorithm and full proofs of soundness and completeness with respect to the declarative system, we refer the reader to our technical report [19].

The algorithmic refinement-checking system is decidable modulo the three following aspects of the system:

1. Resolution of first-order existential variables.
2. Resource management.
3. Theorem proving in first-order MALL.

These sources of nondeterminism do not cause the system to be undecidable, as each can be solved independently (and has in the past). First, resolution of first-order existential variables can be done via either explicit instantiation or unification. Second, we must solve the resource, or context, management problem. This problem includes the issue of deciding how to split a linear context into parts in multiplicative rules such as the \otimes -right rule, (L-E-MANDR), and (R-E-SUB). There are several known approaches to efficient resource management in linear logic [3]. Third, theorem proving in the multiplicative-additive fragment of lin-

ear logic (MALL) has been proven decidable [18]. However, while our system is decidable, finding an efficient decision procedure for all three of the above problems will be challenging. We believe that further investigation should be done in the setting of a practical implementation.

3.5 Soundness

The proof of soundness of refinement checking requires the following soundness condition on the primitive operators.

Condition 2 (Soundness of Primitives)

Suppose

$$\Sigma_{\phi}(\mathfrak{o}) = \forall \vec{b}_1 \cdot ((\phi_{1,1}, \dots, \phi_{1,n}; \psi_1) \rightarrow \exists [\vec{b}_2](\phi_2; \psi_2))$$

If $w \vDash \Omega$; $[\vec{c}_1/\vec{b}_1]\psi_1$, and for $1 \leq i \leq n$, $\Omega \gg_M c'_i : [\vec{c}_1/\vec{b}_1]\phi_{1,i}$, and $\mathcal{T}(\mathfrak{o})(c'_1, \dots, c'_n; w + u) = c'; w'$ then there exist Ω' and \vec{c}_2 such that

1. $w' = w'' + u$;
2. $\Omega' \gg_M c' : [\vec{c}_2/\vec{b}_2][\vec{c}_1/\vec{b}_1]\phi_2$;
3. $w'' \vDash \Omega'$; $[\vec{c}_2/\vec{b}_2][\vec{c}_1/\vec{b}_1]\psi_2$.
4. $\Omega \subseteq \Omega'$

Informally, this condition states that the operator must behave as predicted by its type refinement, and, importantly, can have no effect on a part of the world that is not specified in the precondition of its refinement. Above, w satisfies the precondition of \mathfrak{o} 's refinement. Consequently, no extension, u , of the world may be modified during the operation of \mathfrak{o} at world $w + u$.

The following lemma expresses the relationship between our world semantics and logical judgments, stating that logical deduction respects the semantics of formulas.

Lemma 3 (Soundness of Logical Judgments)

If $w \vDash \Omega$; Δ and $\Omega; \Delta \Longrightarrow_W \psi$ then $w \vDash \psi$.

Finally, we may state and prove our refinement preservation theorem.

Theorem 4 (Refinement Preservation)

1. If $\Omega \gg_M M : \phi$ and $M \Downarrow V$ then $\Omega \gg_M V : \phi$.
2. If $\Omega; \Delta \gg_E E : \exists [\vec{b}](\phi; \psi)$, $w \vDash \Omega; \Delta$, and $E @ w + u \Downarrow V @ w'$, then there exist Ω' and \vec{c} such that $\Omega' \gg_M V : [\vec{c}/\vec{b}]\phi$, $w' = w'' + u$, $\Omega \subseteq \Omega'$ and $w'' \vDash \Omega'$; $[\vec{c}/\vec{b}]\psi$.

The following canonical forms theorem expresses the properties of values that refinement checking provides.

Theorem 5 (Refinement Canonical Forms)

If $\cdot \vdash V : A$ and $\Omega \gg_M V : \phi$ (with Ω containing only bindings and predicates) then one of the following holds:

1. $\phi = \mathbf{Bool}$ and $V = \mathbf{true}$ or $V = \mathbf{false}$;
2. $\phi = \mathbf{Its}(c)$ and $V = c$;
3. $\phi = \mathbf{a}$, $V = c$ and $\Sigma_A(c) = \mathbf{a}$;
4. $\phi = \forall \vec{b} \cdot \phi_1 \rightarrow \phi_2$ and $V = \lambda(x_1:A_1).M$;
5. $\phi = \forall \vec{b} \cdot (\phi_1; \psi_1) \rightarrow \eta$ and $V = \mathbf{fun } x (x_1:A_1) : A_2 \mathbf{ is } E$.

3.6 Conservative Extension

To capture the notion that refinements are a conservative extension of the type system, we present the theorems below. The first theorem states that any refinement given to a term (or expression) in our refinement checking system will always refine the type given to the term (or expression) in the type checking system. In this theorem, we define $\text{type}(\Omega)$ as the typing context Γ mapping all variables $x \in \text{Dom}(\Omega)$ to the type refined by their refinement in Ω . That is, if $x:\phi \in \Omega$ and $\Omega_b \vdash \phi \sqsubseteq A$ then $x:A \in \Gamma$.

Theorem 6

If $\Omega \gg_M M : \phi$ and $\text{type}(\Omega) \vdash_M M : A$ then $\Omega_b \vdash \phi \sqsubseteq A$. Similarly, if $\Omega; \Delta \gg_E E : \eta$ and $\text{type}(\Omega) \vdash_E E : A$ then $\Omega_b \vdash \eta \sqsubseteq A$.

The next theorem states that for any well-typed term, M (or expression, E), with type A , there exists a refinement-checking derivation for which M (E) has the trivial type associated with A . That is, any well-typed term (expression) can also be shown to be well-refined. In this theorem, $\text{triv}_\Gamma(\Gamma)$ is defined as the persistent context mapping elements $x \in \Gamma$ to the trivial refinement of their type in Γ . Also, $\text{triv}_\Sigma(\Sigma_A)$ is defined as the refinement interface containing the trivial refinements of the elements of Σ_A .

Theorem 7

If $\Gamma \vdash_M M : A$ and $\Sigma_\phi = \text{triv}_\Sigma(\Sigma_A)$ then $\text{triv}_\Gamma(\Gamma) \gg_M M : \text{triv}(A)$. Similarly, if $\Gamma \vdash_E E : A$ and $\Sigma_\phi = \text{triv}_\Sigma(\Sigma_A)$ then $\text{triv}_\Gamma(\Gamma); \top \gg_E E : (\text{triv}(A); \top)$.

3.7 Optimization

As well as helping programmers document and check their programs for additional correctness criteria, refinements provide language or library implementors with a sound optimization principle. When programs are checked to determine their refined type, implementors may replace the total function, $\mathcal{T}(o)$, implementing operator o , with a partial function, $\hat{\mathcal{T}}(o)$, that is only defined on the refined domain given by the refinement signature Σ_ϕ .

To be precise, we define the optimized function $\hat{\mathcal{T}}(o)$ as follows.

$$\begin{aligned} \hat{\mathcal{T}}(o)(c_1, \dots, c_n; w) &= \mathcal{T}(o)(c_1, \dots, c_n; w) \\ &\text{if } \Sigma_\phi(o) = \forall \vec{b}. (\phi_1, \dots, \phi_n; \psi) \rightarrow \eta, \\ &w \models \Omega; [\vec{c}/\vec{b}] \psi \\ &\text{and } \Omega \gg_M c_i : [\vec{c}/\vec{b}] \phi_i \text{ for } 1 \leq i \leq n \end{aligned}$$

$$\hat{\mathcal{T}}(o)(c_1, \dots, c_n; w) = \text{undefined otherwise}$$

We use the notation $\hat{\Downarrow}$ to denote the optimized evaluation of expressions with the transition function \mathcal{T} replaced by $\hat{\mathcal{T}}$ (that is, all operator implementations replaced by their optimized versions). We are able to prove that optimized and unoptimized evaluation are equivalent and therefore that it is safe for implementors to replace operator implementations with their optimized version.

Theorem 8 (Optimization)

If $\Omega; \Delta \gg_E E : \eta$ and $w \models \Omega; \Delta$ then $E @ w \Downarrow V @ w'$ if and only if $E @ w \hat{\Downarrow} V @ w'$.

4. EXAMPLES

In this section, we provide a number of examples that demonstrate the expressive power of our language. Our technical report presents several more examples including ML-style integer references, alias types and recursion counts. We omit leading universal quantifiers in our examples as they may easily be inferred in a manner similar to the way the Twelf system [29] infers leading quantifiers.

4.1 File Access Modes

We can extend our file I/O example from the introduction by tracking the access modes of open files and checking them before calling the read or write operators. We add the following access modes to our signature and change `open` to a binary predicate over file names and access modes. If `open(f, a)` is true at a particular program point, then file f is open in access mode a .

<code>r</code>	: Mode	Read Access
<code>w</code>	: Mode	Write Access
<code>rw</code>	: Mode	Read and Write Access
<code>open</code>	: (File, Mode) \rightarrow prop	Open Predicate

Below are the modified `fopen` and `fwrite` operators as well as a new `fread` operator. The refinement of `fopen` now captures the access mode a passed to `fopen` in the `open` predicate. In the refinements of `fread` and `fwrite`, we let `writable(f)` and `readable(f)` abbreviate `open(f, w) \oplus open(f, rw)` and `open(f, r) \oplus open(f, rw)`, respectively. We use the `&` connective in the refinement of `fread` in order to simultaneously ensure that f is readable before a call to `fread` and that whatever the access mode, it is preserved across the call. The `&` in the refinement of `fwrite` serves a parallel purpose.

$$\begin{aligned} \text{fopen} &: (\mathbf{Its}(f), \mathbf{Its}(a); \text{closed}(f)) \rightarrow (\mathbf{unit}; \text{open}(f, a)) \\ \text{fread} &: (\mathbf{Its}(f); \text{open}(f, a) \ \& \ \text{readable}(f)) \rightarrow \\ &\quad (\mathbf{String}; \text{open}(f, a)) \\ \text{fwrite} &: (\mathbf{Its}(f), \mathbf{String}; \text{open}(f, a) \ \& \ \text{writable}(f)) \rightarrow \\ &\quad (\mathbf{unit}; \text{open}(f, a)) \end{aligned}$$

4.2 File Buffering

To demonstrate the utility of persistent predicates, we add to our original file-I/O example support for layering a buffer on top of a file. We introduce three new predicates for this example: `bfs(b, f)` to indicate that a buffer b buffers a file f ; `fl` and `ufl` to conservatively estimate whether a buffer is flushed or unflushed (that is, contains unwritten data).

<code>bfs</code>	: (Buffer, File) \rightarrow prop	Buffers Relation
<code>fl</code>	: (Buffer) \rightarrow prop	Flushed Predicate
<code>ufl</code>	: (Buffer) \rightarrow prop	Unflushed Predicate

Correct use of a buffer requires that the file beneath the buffer be open when a program attempts to write to that buffer and that the buffer is flushed before being freed. We enforce the first requirement in two steps. First, upon creating a new buffer b over a file f , we assert a persistent fact `bfs(b, f)`, indicating that b buffers f . Next, when attempting to write to a buffer b , we check that b buffers some file f and that f is open. As `bfs` is invariant over the lifetime of a given buffer, it is natural to represent it as a persistent

fact. In addition, doing so allows us to reuse the fact without explicitly preserving it and forget about it once it is no longer needed.

Below is the buffer interface. Calling `bnew` on a file f creates a new, flushed buffer over that file. Calling `bwrite` on a buffer b acts as expected, but requires that the file buffered by b be open, although b can be either flushed or unflushed. In either case, we conservatively estimate that b is not flushed upon return from `bwrite`. The operator `bflush` flushes a buffer, changing its state from `ufl` to `fl`. Finally, `bfree` takes a flushed buffer and frees it.

$$\begin{aligned} \mathbf{bnew} &: (\mathbf{Its}(f); \mathbf{1}) \rightarrow \exists[b : \mathbf{Buffer}](\mathbf{Its}(b); !\mathbf{bfs}(b, f) \otimes \mathbf{fl}(b)) \\ \mathbf{bwrite} &: (\mathbf{Its}(b); !\mathbf{bfs}(b, f) \otimes \mathbf{open}(f) \otimes (\mathbf{fl}(b) \oplus \mathbf{ufl}(b))) \rightarrow \\ & \quad ((); \mathbf{open}(f) \otimes \mathbf{ufl}(b)) \\ \mathbf{bflush} &: (\mathbf{Its}(b); \mathbf{ufl}(b)) \rightarrow ((); \mathbf{fl}(b)) \\ \mathbf{bfree} &: (\mathbf{Its}(b); \mathbf{fl}(b)) \rightarrow ((); \mathbf{1}) \end{aligned}$$

4.3 Interrupt Levels

For their study of Windows device drivers, DeLine and Fahndrich extend Vault with a special mechanism for specifying “capability states” which are arranged in a partial order [7]. They use the partial order and bounded quantification to specify preconditions on kernel functions. Here we give an alternate encoding and reason logically about the same kernel functions and their preconditions.

First, we assume a signature with abstract constants that correspond to each interrupt level and also a predicate L over these levels. If $L(c)$ is true at a particular program point then the program executes at interrupt level c at that point.

<code>pass</code>	: level	Passive Level
<code>apc</code>	: level	APC Level
<code>dis</code>	: level	Dispatch Level
<code>dirql</code>	: level	DIRQL Level
<code>L</code>	: level \rightarrow prop	Level Predicate

Next we consider a variety of kernel functions and their type refinements. First, the `KeSetPriorityThread` function requires that the program be at Passive Level when it is called and also returns in Passive Level. The function takes arguments with type `thread` and `pr`, which we assume are defined in the current signature.

`KeSetPriorityThread`:
 $(\mathbf{thread}, \mathbf{pr}; L(\mathbf{pass})) \rightarrow (\mathbf{pr}; L(\mathbf{pass}))$

Function `KeReleaseSemaphore` is somewhat more complex since it may be called in Passive, APC or Dispatch level and it preserves its level across the call. We let `less(dis)` abbreviate the formula $L(\mathbf{pass}) \oplus L(\mathbf{apc}) \oplus L(\mathbf{dis})$.

`KeReleaseSemaphore`:
 $(\mathbf{sem}, \mathbf{pr}, \mathbf{long}; L(l) \otimes (L(l) \multimap \mathbf{less}(\mathbf{dis}))) \rightarrow (\mathbf{pr}; L(l))$

Finally, `KeAcquireSpinLock` also must be called in one of three states. However, it returns in the Dispatch state and also returns an object representing the initial state (l) that the function was called in.

`KeAcquireSpinLock`:
 $(\mathbf{sem}, \mathbf{pr}, \mathbf{long}; L(l) \otimes (L(l) \multimap \mathbf{less}(\mathbf{dis}))) \rightarrow (\mathbf{Its}(l); L(\mathbf{dis}))$

5. DISCUSSION

5.1 Related Work

A number of researchers have recently proposed strategies for checking that programs satisfy sophisticated safety properties. Each system brings some strengths and some weaknesses when compared with our own. Here are some of the most closely related systems.

Refinement Types. Our initial inspiration for this project was derived from work on refinement types by Davies and Pfenning[6] and Denney [8] and the practical dependent types proposed by Xi and Pfenning [38, 39]. Each of these authors proposed to sophisticated type systems that are able to specify many program properties well beyond the range of conventional type systems such as those for Java or ML. However, none of these groups considered the ephemeral properties that we are able to specify and check.

Safe Languages. CCured [23], CQual [13], Cyclone [16], ESC [9, 12], and Vault [7, 10] are all languages designed to verify particular safety properties. CCured concentrates on showing the safety of mostly unannotated C programs; Cyclone allows programmers to specify safe stack and region memory allocation; ESC facilitates program debugging by allowing programmers to state invariants of various sorts and uses theorem proving technology to check them; and Vault and CQual make it possible to check resource usage protocols. Vault has been applied to verification of safety conditions in device drivers and CQual has been applied to find locking bugs in the Linux kernel. One significant difference between our work and the others is that we have chosen to use a general substructural logic to encode program properties. Vault is the most similar since its type system is derived from the capability calculus [36] and alias types [33, 37], which is also an inspiration for this work. However, the capability logic is somewhat ad hoc whereas we base our type system directly on linear logic. As far as we are aware, the semantics of Vault has not been fully formalized. We hope this work is an effective starting point in that endeavour.

Igarashi and Kobayashi’s resource usage analysis [?] is another piece of work in this vein. They define a complex type system that keeps track of the “uses” of resources. They have a general trace-based semantics as opposed to our possible worlds-style resource semantics. It is difficult to compare the expressive power of our two proposals precisely as the set of formulas involved is quite different. They have some interesting modal operators and a recursion operator, but their logic is propositional whereas ours is first-order.

Proof-Carrying Type Systems. Shao et al. [32] and Crary and Vanderwaart [4] have both developed powerful type languages that include a fully general logical framework within the type structure. Both languages were inspired by Necula and Lee’s work on proof carrying code [22, 21] and are designed as a very general framework for coupling low-level programs with their proofs of safety. In contrast, our language is intended to be a high-level language for programmers. Hence, the design space is quite different. Our specification language is less general than either of these, but it does not require programmers to write explicit proofs that their programs satisfy the safety properties in question. Moreover, neither of these logics contain the full complement of linear logic’s left-asynchronous connectives ($\mathbf{1}$, \otimes , $\mathbf{0}$, \oplus , \exists), which we find the most useful in our applications.

Hoare Logic. Recent efforts on the reasoning about pointers in Hoare logic [17, 30] by Ishtiaq, O’Hearn and Reynolds, provided guidance in construction of our semantic model of refinements. However, they use bunched logic in their work whereas we use a subset of linear logic. One important difference between the logics is that linear logic contains the modality $!$, which we use to reason about *persistent* facts. A notion of persistence seems essential to allow one to reason about values, which, by their nature, remain unchanged throughout the computation. Also, since our work is based on type theory, it naturally applies to higher-order programs, which is not the case for Hoare logic. Moreover, programmers who use Hoare logic have no automated support whereas our system has a decidable type checking algorithm.

5.2 Future Work

There are many directions for future work. We have begun to investigate the following issues. Our most immediate concern is the development of an implementation of the ideas presented in this paper. One of the authors (Mandelbaum) has developed a preliminary implementation for small core subset of Java. The current implementation is built using Polyglot [24], an extensible compiler infrastructure for Java, and allows programmers to reason with a minimalist subset of the logic that includes $\mathbf{1}$, \otimes and \top . We are currently developing a surface language based on ML that incorporates all of the features presented in this paper.

One of the limitations of the current proposal is the inability to check the membership of an ephemeral fact in the ephemeral context while preserving the other facts in that context. For example, consider a function `writeBoth` that takes two file arguments and writes to both of them. Below is a possible refinement of such a function:

```
writeBoth :
  (Its(f), Its(g); open(f)  $\otimes$  open(g))  $\multimap$ 
  ((); open(f)  $\otimes$  open(g))
```

However, this refinement precludes the possibility that f and g are equal, as $\text{open}(f) \otimes \text{open}(g)$ describes two distinct facts. Therefore, `writeBoth` cannot be passed the same file for both of its arguments. An alternative precondition, $(\text{open}(f) \otimes \top) \& (\text{open}(g) \otimes \top)$, does not depend on the distinctness of f and g but fails to preserve information about the remaining members of the ephemeral context, due to the imprecision of \top .

One solution to this problem would be to provide a way to view an ephemeral fact as temporarily persistent, as we can check the membership of a fact in the persistent context while preserving the remaining persistent facts. Then, we could successfully check calls to `writeBoth` with aliased arguments based on the following refinement:

```
writeBoth :
  (Its(f), Its(g); !open(f)  $\otimes$  !open(g))  $\multimap$  (();  $\mathbf{1}$ )
```

We are investigating a sharing principle, similar to the world-splitting allowed by the subsumption rule, that would provide a safe method of viewing ephemeral facts as temporarily persistent.

An alternative solution would be to add a limited form of second-order quantification to the logic. We could then refine `writeBoth` as follows:

```
writeBoth :
  (Its(f), Its(g); (open(f)  $\otimes$   $\top$ )  $\&$  (open(g)  $\otimes$   $\top$ )  $\&$  w)  $\multimap$ 
  ((); w)
```

In essence, the refinement checks that $\text{open}(f)$ and $\text{open}(g)$ are members of the world w , while ensuring that w is preserved.

The next item of interest to us is the encoding of type-and-effect systems in our language. We believe that our language provides a general framework in which to encode many type-and-effect systems. Based on the encoding shown above, we have devised a translation from a variant of a well-known type-and-effect system concerning lock types for static enforcement of mutual exclusion [11], into our language (extended with second-order quantification). We thereby show that our refinements are at least as powerful. Our translation also helps us understand the connection between types and effects and recent research on sophisticated substructural type systems such as the one implemented in Vault [7].

Finally, our language is parameterized by a single interface and implementation that enables us to consider reasoning about a variety of different sorts of effects. The next step in the development of this project is to extend the language with an advanced module system that allows programmers to define their own logical safety policies and to reason compositionally about their programs.

Acknowledgments

David Walker would like to thank Frank Pfenning for teaching him about the structure of linear logic—without Frank’s instruction, it would have been impossible to write this paper. We would like to thank Nate Nystrom, Michael Clarkson and Andrew Myers, the implementers of Polyglot [24], for giving us access to their software. Polyglot dramatically increased the speed at which we were able to develop a preliminary implementation of our ideas for Java. We are also very grateful to Dan Grossman, who gave us some important feedback on an earlier draft of this paper. Finally, we would like to thank the anonymous reviewers for their helpful comments. In particular, we thank them for bringing to our attention the `writeBoth` example, discussed in the previous section.

6. REFERENCES

- [1] S. Ambler. *First-order Linear Logic in Symmetric Monoidal Closed Categories*. PhD thesis, University of Edinburgh, 1991.
- [2] L. Augustsson. Cayenne — a language with dependent types. In *ACM International Conference on Functional Programming*, pages 239–250, Baltimore, Sept. 1999. ACM Press.
- [3] I. Cervesato, J. S. Hodas, and F. Pfenning. Efficient resource management for linear logic proof search. In *Proceedings of the Fifth International Workshop on Extensions of Logic Programming — ELP’96*, pages 67–81, Leipzig, Germany, March 1996.
- [4] K. Crary and J. Vanderwaart. An expressive, scalable type theory for certified code. In *ACM International Conference on Functional Programming*, Pittsburgh, Oct. 2002. ACM Press.
- [5] K. Crary and S. Weirich. Flexible type analysis. In *ACM International Conference on Functional Programming*, pages 233–248, Paris, Sept. 1999.
- [6] R. Davies and F. Pfenning. Intersection types and computational effects. In *ACM International Conference*

- on *Functional Programming*, pages 198–208, Montreal, Sept. 2000. ACM Press.
- [7] R. Deline and M. Fähndrich. Enforcing high-level protocols in low-level software. In *ACM Conference on Programming Language Design and Implementation*, pages 59–69, Snowbird, Utah, June 2001. ACM Press.
- [8] E. Denney. *A Theory of Program Refinement*. PhD thesis, University of Edinburgh, Edinburgh, 1998.
- [9] D. L. Detlefs. An overview of the extended static checking system. In *The First Workshop on Formal Methods in Software Practice*, pages 1–9. ACM(SIGSOFT), Jan. 1996.
- [10] M. Fähndrich and R. Deline. Adoption and focus: Practical linear types for imperative programming. In *ACM Conference on Programming Language Design and Implementation*, pages 13–24, Berlin, June 2002. ACM Press.
- [11] C. Flanagan and M. Abadi. Types for safe locking. In S. Swierstra, editor, *Lecture Notes in Computer Science*, volume 1576, pages 91–108, Amsterdam, Mar. 1999. Springer-Verlag. Appeared in the Eighth European Symposium on Programming.
- [12] C. Flanagan, R. Leino, M. Lillibridge, G. N. J. Saxes, and R. Stata. Extended static checking for java. In *ACM Conference on Programming Language Design and Implementation*, Berlin, June 2002. ACM Press.
- [13] J. Foster, T. Terauchi, and A. Aiken. Flow-sensitive type qualifiers. In *ACM Conference on Programming Language Design and Implementation*, Berlin, June 2002. ACM Press.
- [14] J.-Y. Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- [15] A. Gordon and A. Jeffrey. Typing correspondence assertions for communication protocols. In *Mathematical Foundations of Programming Semantics 17*. Elsevier, 2001.
- [16] D. Grossman, G. Morrisett, T. Jim, M. Hicks, Y. Wang, and J. Cheney. Region-based memory management in cyclone. In *ACM Conference on Programming Language Design and Implementation*, Berlin, June 2002. ACM Press.
- [17] S. Ishtiaq and P. O’Hearn. BI as an assertion language for mutable data structures. In *Twenty-Eighth ACM Symposium on Principles of Programming Languages*, pages 14–26, London, UK, Jan. 2001.
- [18] P. Lincoln and A. Scedrov. First-order linear logic without modalities is NEXPTIME-hard. *Theoretical Computer Science*, 135:139–154, 1994.
- [19] Y. Mandelbaum, D. Walker, and R. Harper. An effective theory of type refinements. Technical Report TR-656-02, Princeton University, December 2002.
- [20] E. Moggi. Notions of computation and monads. *Information and Computation*, 93:55–92, 1991.
- [21] G. Necula. Proof-carrying code. In *Twenty-Fourth ACM Symposium on Principles of Programming Languages*, pages 106–119, Paris, 1997.
- [22] G. Necula and P. Lee. Safe kernel extensions without run-time checking. In *Proceedings of Operating System Design and Implementation*, pages 229–243, Seattle, Oct. 1996.
- [23] G. C. Necula, S. McPeak, and W. Weimer. Ccured: Type-safe retrofitting of legacy code. In *ACM Symposium on Principles of Programming Languages*, London, Jan. 2002. ACM Press.
- [24] N. Nystrom, M. Clarkson, and A. C. Myers. Polyglot: An extensible compiler framework for java. In *12th International Conference on Compiler Construction*, Apr. 2003. to appear.
- [25] P. O’Hearn and D. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, 1999.
- [26] P. O’Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *CSL ’01*, pages 1–19, Paris, 2001.
- [27] F. Pfenning. Structural cut elimination in linear logic. Technical Report CMU-CS-94-222, Department of Computer Science, Carnegie Mellon University, December 1994.
- [28] F. Pfenning and R. Davies. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science*, 11(4):511–540, 2001.
- [29] F. Pfenning and C. Schürmann. system description: Twelf — a metalogical framework for deductive systems. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction*, number 1632 in LNAI, pages 202–206, Trento, Italy, July 1999. Springer-Verlag.
- [30] J. C. Reynolds. Intuitionistic reasoning about shared mutable data structure. In *Millennial perspectives in computer science*, Palgrave, 2000.
- [31] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Symposium on Logic in Computer Science’02*, pages 55–74, 2002.
- [32] Z. Shao, B. Saha, V. Trifonov, and N. Papaspyrou. A type system for certified binaries. In *ACM Symposium on Principles of Programming Languages*, London, Jan. 2002. ACM Press.
- [33] F. Smith, D. Walker, and G. Morrisett. Alias types. In *European Symposium on Programming*, pages 366–381, Berlin, Mar. 2000.
- [34] M. Tofte and J.-P. Talpin. Implementation of the typed call-by-value λ -calculus using a stack of regions. In *Twenty-First ACM Symposium on Principles of Programming Languages*, pages 188–201, Portland, Oregon, Jan. 1994.
- [35] M. Tofte and J.-P. Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- [36] D. Walker, K. Crary, and G. Morrisett. Typed memory management in a calculus of capabilities. *ACM Transactions on Programming Languages and Systems*, 22(4):701–771, May 2000.
- [37] D. Walker and G. Morrisett. Alias types for recursive data structures. In *Workshop on Types in Compilation*, Montreal, Sept. 2000.
- [38] H. Xi and F. Pfenning. Eliminating array bound checking through dependent types. In *ACM Conference on Programming Language Design and Implementation*, pages 249–257, Montreal, June 1998.
- [39] H. Xi and F. Pfenning. Dependent types in practical programming. In *Twenty-Sixth ACM Symposium on Principles of Programming Languages*, pages 214–227, San Antonio, TX, Jan. 1999.