A Type System for Expressive Security Policies*

David Walker
Cornell University

Abstract

Certified code is a general mechanism for enforcing security properties. In this paradigm, untrusted agent code carries annotations that allow a host to verify its trustworthiness. Before running the agent, the host checks the annotations and proves that they imply the host's security policy. Despite the flexibility of this scheme, so far, compilers that generate certified code have focused on simple memory and control-flow safety rather than more general security properties.

Security automata can enforce an expressive collection of security policies including access control policies and resource bounds policies. In this paper, we show how to take specifications in the form of security automata and automatically transform them into signatures for a typed lambda calculus that will enforce the corresponding safety property. Moreover, we describe how to instrument typed source language programs with security checks and typing annotations so that the resulting programs are provably secure and can be mechanically checked. This work provides a foundation for the process of automatically generating certified code for expressive security policies in a type-theoretic framework.

1 Introduction

Strong type systems such as those supported by Java or ML provide provable guarantees about the run-time behaviour of programs. If we type check programs before executing them, we know they "won't go wrong." Usually, the notion "won't go wrong" implies *memory safety* (programs only access memory that has been allocated for them), *control flow safety* (programs only jump to and execute valid code), and *abstraction preservation* (programs use abstract data types only as their interfaces allow). These properties are essential building blocks for any secure system such as a web browser, extensible operating system, or server that may download, check and execute untrusted programs. However, in order to build such systems, we must restrict program behaviour much more drastically; standard type safety properties are not sufficient to enforce realistic access control policies or to restrict the dissemination of secret information.

Certified code is a general framework for verifying security properties in untrusted code. To use this security architecture, programmers and compilers must attach a collection of annotations to the code they produce. These annotations can be proofs, types, or annotations from some other kind of formal system. Regardless, there must be some way of reconstructing a proof that the code obeys a certain security policy, for upon receiving annotated code, an untrusting web browser or operating system will use a mechanical checker to verify that the program is safe before executing it.

In theory, certified code is very general, but in practice, compilers that emit certified code have focused on a relatively limited set of properties. For example, Necula and Lee's proof-carrying code (PCC) implementation [17, 16] uses a first-order logic and they have shown that they can check many interesting properties of hand-coded assembly language programs including access control and resource bound policies [19]. However, the main focus of their proof-generating compiler Touchstone [18] is the generation of efficient code; the security policy they enforce is the standard type and memory safety. Other frameworks for producing certified code, including Kozen's efficient code certification (ECC) [7] and Morrisett *et al.*'s [15, 12, 11] Typed Assembly Language (TAL), concentrate exclusively on standard type safety properties.

The main reason that certified code has been used in this restricted fashion is that automated theorem provers are not powerful enough to infer properties of arbitrary programs and constructing proofs by hand is prohibitively expensive. Researchers have been able to produce proofs of type safety for their code because they place heavy restrictions on the programming languages that they compile. Furthermore, when they cannot prove type safety statically, they insert dynamic checks that ensure code cannot violate certain invariants. For example, in order to handle arrays safely, a compiler must prove each array access is in bounds. The Touchstone compiler uses a theorem prover to attempt to complete this proof statically, but when it cannot do so, it inserts a run-time check.

In order to construct a tractable system for verifying expressive security policies, we will follow the model developed for standard type safety proofs: Instrument the program with run-time security checks and then eliminate those checks that a theorem prover can verify statically. For an interesting class of security properties, this strategy will guar-

*This material is based on work supported in part by the AFOSR grant F49620-97-1-0013. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not reflect the views of this agency.

antee that any program can be automatically rewritten so that it is provably safe; the programmer need not be burdened by extensive proof obligations.

1.1 Security Automata

Unlike memory safety and control-flow safety, properties that must be enforced across all applications, other security policies may vary from one application to the next. Schneider [22] has proposed *security automata* as a flexible mechanism for specifying a much larger class of security policies than are possible in traditional type systems. Security automata are defined similarly to other automata [6]. In the model we will use in this paper, these machines possess either a finite or a countably infinite set of states, and rules for making transitions from one state to the next. The author of the security policy determines the set of states and the transition relation. One of these states is designated as the *bad* state and entry into this state is defined to be a violation of the security policy.

Security automata enforce safety properties by monitoring programs as they execute. Before an untrusted program is allowed to execute a security-sensitive or *protected* operation, the security automaton checks to see if the operation will cause a transition to the *bad* state. If so, the automaton terminates the program. If not, the program is allowed to execute the operation and the security automaton makes a transition to a new state. For example, a web browser might allow applets to open and read files, and send bits on the network. However, assuming the web browser wants to ensure some level of privacy, the open, read, and send operations will be designated protected operations. Before an untrusted applet can execute one of these functions, the browser will check the security automaton definition to ensure the operation is allowed in the current state.

By monitoring programs in this way, security automata are sufficiently powerful that they can restrict access to private files or bound the use of resources. They can also enforce the safety properties typically implied by type systems such as memory safety and control-flow safety. However, security automata can only enforce safety policies. Hence, some interesting security policies including information flow, resource availability, and general liveness properties cannot be enforced by this mechanism. However, Schneider [22] points out that some of these applications can still use a security automaton: The automaton must simply enforce a stronger property than is required. For example, the policy that admits any program that allocates a finite amount of memory cannot be enforced by a security automaton. However, security automata can enforce the policy that prevents a program from allocating more than an *a priori* fixed amount (such as 1 MB) of memory. The latter policy reduces the number of legal programs that can be written (a program that allocates 1.1 MB will be prematurely terminated), but it may be effective in practice.

Figure 1 depicts a security automaton that enforces the simple policy that programs must not perform a send on the network after reading a file. The security automaton actually has three states: the *start* state, the *has_read* state, and the *bad* state, which is not shown in the diagram. For the purpose of this example, there are only two protected operations: the *send* operation and the *read* operation. Each of the arcs in the graph is labeled with one of these operations and indicates the state transition that occurs when the operation is invoked. If there is no outgoing arc from a certain state labeled with the appropriate operation, the automaton

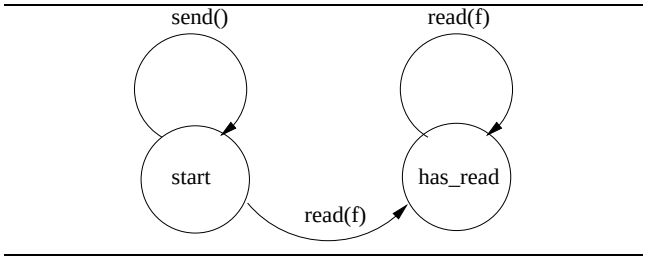

Figure 1: Security Automaton for a Simple File System Policy

makes a transition to the *bad* state. For example, there is no *send* arc emanating from the *has_read* state. Consequently, if a program tries to use the network in the *has_read* state, it will be terminated.

1.2 Implementing Security Automata

We can enforce the policies specified by security automata by instrumenting programs with run-time security checks. For example, consider a program that executes the *send* operation at some point during computation. According to the security policy, we must be in the *start* state in order for the *send* to be safe. A program instrumentation tool could enforce this policy by wrapping the code invoking *send* with security checks:

```
let new_state = check_send(current_state) in
if new_state = bad then
    halt
else
    send()
```

The first statement invokes the security automaton to determine the next state given that a *send* operation is invoked in the current state. The second statement tests the next state to make sure it is not the *bad* one. If it is *bad*, then program execution is terminated.

After performing the initial transformation that wraps all protected operations with run-time security checks, a program optimizer might attempt to eliminate redundant checks by performing standard program optimizations such as loop-invariant removal and common subexpression elimination [1]. An optimizer might also use its knowledge of the special structure of a security automaton to eliminate more checks than would otherwise be possible.

An implementation developed by Erlingsson and Schneider [24] attests to the fact that security automata can enforce a broad, practical set of security policies. Their tool, SASI, automatically instruments untrusted code with checks dictated by a security automaton specification and optimizes the output code to eliminate checks that can be proven unnecessary. SASI is both flexible and efficient and they have implemented a variety of security policies from the literature. For example, using SASI for the Intel Pentium architecture, they have specified the memory and control-flow safety policy enforced by Software Fault Isolation (SFI) [25]. The SASI-instrumented code is only slightly slower than the code produced by the special-purpose, hand-coded MiSFIT tool for SFI [23]. As another example, using SASI for the Java Virtual Machine, they have been able to reimplement the security manager for Sun's Java 1.1. The SASI-instrumented code is equally as efficient as the Java security manager in some cases and more efficient in others. SASI

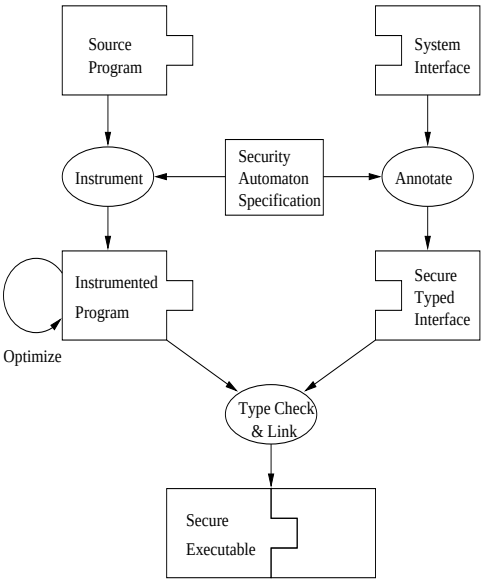

Figure 2: Architecture of a Certifying Compiler

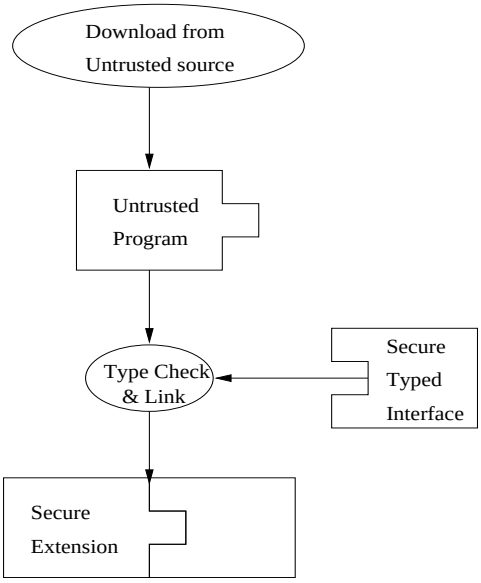

Figure 3: Architecture of a Secure Extensible System

is more flexible than the Java security model because individual systems or users can customize the set of protected operations rather than having that set dictated by the Java language specification.

1.3 Overview

This paper presents a type-based framework for automatically compiling programs into certified code that satisfies security automaton specifications. Figure 2 presents the architecture of the system that we propose. The left-hand side of the diagram shows the process of instrumenting programs that do not necessarily follow the security policy with run-time security checks and typing annotations. The instrumentation procedure requires a security automaton specification as its input. After the initial instrumentation pass has been completed, the program can be optimized to remove the redundant run-time checks. The right-hand side of the diagram depicts the process of transforming the host system interface, which contains declarations of the security-sensitive functions, into a secure typed interface. The new interface specifies additional preconditions that must be satisfied before these functions can be invoked.

Both transformed application programs and system interface are written in a secure typed language. The soundness of its type system guarantees that if the transformed program type checks against the system interface, then it obeys the security policy implied by that interface. Hence, the system administrator must trust that the security automaton specification implies the desired semantic properties and that the translation of the interface and the implementation of the type checker are correct. However, because the type system is defined independently of the instrumentation algorithm, the system administrator does not have to trust the program instrumentor or optimizer. In fact, untrusted parties can write their own instrumentation and optimization transformations. Figure 3 depicts an extensible system that imports code from an unknown source. The system does not know or care how the code it receives is generated; it may have been produced by the compiler sketched in Figure 2 or by some entirely different means. Regardless of its source, if the code type checks against the secure interface, it will not violate the security policy. Thus, the type system provides a new way to certify that untrusted code obeys expressive security policies.

The central technical contribution of this paper is the formulation of this secure typed target language. This language must be flexible enough to allow optimization and yet ensure that malicious programs cannot circumvent security checks. In order to accomplish these goals, the type system of the target language contains a language of predicates. Each protected operation is given a type that uses the predicates to specify a precondition restricting the application of the function. More specifically, the precondition states that the function must not cause the program to enter the *bad* state.

Continuing our example from the previous section, we will attach the following preconditions to the *send* operation:

$$P_1 : transition_{send}(current_state, new_state)$$
$$P_2 : new_state \neq bad$$

These predicates state that executing the *send* operation causes a transition from the *current_state* to the *new_state*, and moreover, that the *new_state* is not *bad* so the security policy will not be violated.

Each time *send* or another protected function is invoked, the type checker must be able to prove that the calling context satisfies the precondition. If it can do so, the program satisfies the security policy. However, in general, given an arbitrary program, a theorem prover will not be able to prove all preconditions are satisfied in all cases. Consequently, programs will have to contain run-time security checks. These security checks will be given types that express post-conditions containing information about the automaton transition function. The post-conditions can be used to help prove the preconditions on the protected operations.

The following example demonstrates how we can verify that the wrapper code for the *send* operation is safe:

base types	b	$\in$	BaseType
types	τ	$::=$	$b \mid (\tau_1, \ldots, \tau_n) \to \mathtt{0}$
constants	a	$\in$	A
protected ops	f	$\in$	F
value vars	g, x	$\in$	ValueVar
values	v	$::=$	$a \mid f \mid x \mid$ $\mathtt{fix}\, g(x_1{:}\tau_1, \ldots, x_n{:}\tau_n).e$
expressions	e	$::=$	$v_0(v_1, \ldots, v_n) \mid \mathtt{halt}$

Figure 4: Source Language Syntax

```
let new_state = check_send(current_state) in
P1:  transition_send(current_state, new_state)
if new_state = bad then
    halt
else
    P2: new_state ≠ bad
    send()
```

Predicate P_1 is the post-condition for the $check_{send}$ function and predicate P_2 can be inferred given the test in the conditional statement. The predicates P_1 and P_2, together with the information that the automaton is in the state designated *current_state*, are sufficient to prove that the precondition on the *send* operation has been satisfied. Therefore, the code will type check and the security policy cannot be violated.

The remaining sections of this paper describe the approach in more detail. We begin by describing a source language for program instrumentation (Section 2). Next, we present a formal model for security automata based on work by Alpern and Schneider [2, 22] (Section 3). At this point, we define a typed language, $\lambda_{\mathcal{A}}$, for encoding security automata and specify how to construct a typed interface that will specialize the language so that it enforces the policy specified by a particular automaton (Section 4). Finally, we show how to instrument insecure programs and discuss possible optimizations to the procedure (Section 5). The last section discusses related and future work.

2 The Insecure Source Language

Before we can describe security automata formally, we must describe the language that we wish to make secure. For the purposes of this paper, we use a simply-typed lambda calculus augmented with a finite set of base types and some constants. The syntax of the language appears in Figure 4.

There are two classes of constants, a countably infinite set of objects a of base type and finite set of function constants f. The latter denotes the security-sensitive operations that are part of the host system interface. For example, these operations may send bits on the network or read files. We will continue refer to these functions as the *protected* functions or operations, and in the next section we will show how to use security automata to restrict access to them. Ordinary (unrestricted) functions may be constructed using the notation:

$$\mathtt{fix}\, g(x_1{:}\tau_1, \ldots, x_n{:}\tau_n).e$$

where the arguments x_1 to x_n have types τ_1 to τ_n and g, the name of the function itself, may appear recursively in e. When an argument is unused or a function is not recursive, we will often use an underscore in place of the name of the argument or function.

The only distinctive element of our lambda calculus is that it is presented in *continuation-passing style* (CPS) [21]. Functions written in a CPS style never "return" to their caller; instead, they are passed an auxiliary function, or *continuation*. When the function has completed its computation, it calls this continuation. We signify the fact that our functions never return using the notation "$\to\mathtt{0}$" in their type. The symbol `halt` terminates computation. Although CPS is not strictly necessary, this decision will simplify the presentation of the target language. In particular, pre- and post-conditions can be handled uniformly, the latter being the pre-condition for a continuation. For the remainder of the paper, we will assume familiarity with CPS.

We will assume a static semantics for the language given by a typing judgements $\Gamma \vdash v : \tau$ and $\Gamma \vdash e$ where Γ is a finite map from value variables to types. The former judgement concludes value v is well-formed and has type τ while the latter judgement says simply that e is well-formed. CPS expressions do not return values and consequently, their judgements need not specify a type.

The types for constants are given by a signature $\mathcal{C}_{source}$ where $\mathcal{C}_{source}(a)$ has the form b and $\mathcal{C}_{source}(f)$ has the form $(b_1, \ldots, b_n, (b) \to \mathtt{0}) \to \mathtt{0}$. For the purposes of this paper, the protected operations f operate on objects of base type and accept a single continuation. This restriction is not a fundamental limitation of the work, but allowing higher-order protected operations does cause some complications.[1] Otherwise, the typing rules for the language are completely standard and have been omitted.

We also assume a standard small-step operational semantics for our language denoted by $e_1 \longmapsto e_2$. Function constants must obey their signature and are assumed to be total. Hence:

$$\begin{array}{l}\text{if } \mathcal{C}_{source}(f) = (b_1, \ldots, b_n, (b) \to \mathtt{0}) \to \mathtt{0} \\ \text{then } f(a_1, \ldots, a_n, v_{cont}) \longmapsto v_{cont}(a) \\ \qquad\qquad (\text{for some } a \text{ with type } b)\end{array}$$

when a_i has type b_i. As we mentioned above, `halt` terminates computation. In other words, there does not exist an expression e such that $\mathtt{halt} \longmapsto e$. We will use this instruction to terminate programs that misbehave. Once again, the remaining operational rules are standard and have been omitted.

2.1 An Example: The Taxation Applet

Using our simple source language, we can write a "taxation applet." When invoked, this applet sends a request out for tax forms. After sending the request, the applet reads a private file containing the customer salary before computing the taxes owed. We will assume files and integers (*int*) are available as base types; *send* (across the network for the tax forms) and *read* (file) are the two protected operations.[2] The applet follows:

[1] The needed extensions complicate the semantics of the target language, but offer little additional insight into the solution of the problem. See Section 4.4 for more explanation.

[2] In this example and others, we use the notation $\mathtt{let}\, x_1, \ldots, x_n = v(v_1, \ldots, v_n)\, \mathtt{in}\, e$ as an abbreviation for the function application:

$$v(v_1, \ldots, v_n, \mathtt{fix}_(x_1{:}\tau_1, \ldots, x_n{:}\tau_n).e)$$

Q	a finite or countably infinite set of states
q_0	a distinguished initial state
bad	the single "bad" state
F	a finite set of function symbols (f)
A	a countable set of constants (a)
δ	a computable (deterministic) total function: $\mathrm{F} \to (Q \times \vec{\mathrm{A}}) \to Q$

Figure 5: Elements of a Security Automaton

```
fix_(secret:file, x_cont:(int) → 0).
  let_      = send() in        % send for forms
  let salary = read(secret) in % read salary
  let taxes  = salary in       % compute taxes!!
  x_cont(taxes)
```

The customer would like to ensure information about his or her salary is not leaked on the network. In the following sections, we will satisfy the customer by showing how to specify and enforce the "no send after read" policy that was discussed informally in the introduction.

3 Security Automata

Our definition of security automata is derived from the work of Alpern and Schneider [2] who use similar automata to define and reason about safety and liveness properties. Schneider [22] extended this work by defining the class of security enforcement mechanisms that monitor program execution. He proved that these enforcement mechanisms are specified by security automata and that they can enforce any safety property.

For our purposes, a security automaton ($\mathcal{A}$) is a 6-tuple containing the fields summarized in Figure 5. Like an ordinary finite automaton, a security automaton has a finite or a countably infinite set of states Q, and a distinguished initial state q_0. The automaton also has a single bad state (bad). Entrance into the bad state indicates that the security policy has been violated. All other states are considered "good" or "accepting" states. The automaton's inputs correspond to the application of a function symbol f to arguments $a_1, \ldots, a_n$ where f is taken from the set F (the set of protected operations) and $a_1, \ldots, a_n$ are taken from the set A.

A security automaton defines allowable program behaviours by specifying a transition function δ. Formally, δ is a deterministic, total function with a signature $\mathrm{F} \to (Q \times \vec{\mathrm{A}}) \to Q$ where $\vec{\mathrm{A}}$ denotes the set of lists $a_1, \ldots, a_n$. Upon receiving an input $f(a_1, \ldots, a_n)$, an automaton makes a transition from its current state to the next state as dictated by this transition function. If the security policy permits the operation $f(a_1, \ldots, a_n, v_{cont}) \longmapsto v_{cont}(a)$ in the current state then the next state will be one of the "good" ones. On the other hand, if the security policy disallows the action $f(a_1, \ldots, a_n, v_{cont}) \longmapsto v_{cont}(a)$ then $\delta(f)(q, a_1, \ldots, a_n)$ will equal bad. All programs must respect the basic typing rules so at minimum, for all f, q, a_1, ..., a_n there exists a state q' such that q' is not bad and $\delta(f)(q, a_1, \ldots, a_n) = q'$ only if $\mathcal{C}(f) = (b_1, \ldots, b_n, (b) \to 0) \to 0$ and $\mathcal{C}(a_i) = b_i$. Furthermore, once the automaton enters the bad state, it stays there. Formally, for all f and $a_1, \ldots, a_n$, $\delta(f)(bad, a_1, \ldots, a_n) = bad$. Finally, the transition function δ must be computable. Moreover, the implementor of the security policy must supply code for a family of functions δ_f such that $\delta_f(q_1, a_1, \ldots, a_n)$ equals $\delta(f)(q_1, a_1, \ldots, a_n)$. In the following sections, we will use these functions to instrument untrusted code with security checks.

The language accepted by the automaton $\mathcal{A}$, written $\mathcal{L}(\mathcal{A})$, is a set of strings where a string is a finite sequence of symbols $s_1, \ldots, s_n$ and each symbol s_i is of the form $f(a_1, \ldots, a_m)$. The string $s_1, \ldots, s_n$ belongs to $\mathcal{L}(\mathcal{A})$ if the predicate $Accept(q_0, s_1, \ldots, s_n)$ holds where $Accept$ is the least predicate such that:

Definition 1 (String Acceptance) *For all states q and (possibly empty) sequences of symbols $s_1, \ldots, s_n$, $Accept(q, s_1, \ldots, s_n)$ if $q \neq bad$ and:*

1. *$s_1, \ldots, s_n$ is the empty sequence or*
2. *$s_1 = f(a_1, \ldots, a_m)$ and $\delta(f)(q, a_1, \ldots, a_m) = q'$ and $Accept(q', s_2, \ldots, s_n)$*

We are now in a position to define the security policy for our taxation applet. Informally, the policy we desire is "no network send after any file has been read." Furthermore, for the sake of future examples, access to some files will be restricted; in order to determine whether or not the applet has been granted access to the file f, we will have to invoke the function $read?(f)$, which has been provided by the implementor of the file system.

The corresponding security automaton has three states: $start$, the initial state; has_read, the state we enter after any file read; and, of course, the bad state. The protected operations, F, are $send$ and $read$ and the constants, A, include all files and the integers. The transition function δ, written in pseudo-code, is the following:

```
δ_send(q) =
    if q = start then
        start
    else
        bad

δ_read(q, f) =
    if q = start ∧ read?(f) then
        has_read
    else if q = has_read ∧ read?(f) then
        has_read
    else
        bad
```

4 The Secure Target Language

So far, we have defined two languages, a lambda calculus for writing applications and a specification language for describing security properties. This section presents a third language, $\lambda_{\mathcal{A}}$, that serves as a target for the compilation of the other two. Application programs, such as the taxation applet, are compiled into $\lambda_{\mathcal{A}}$ expressions by inserting typing annotations and run-time checks. At the same time, the security automaton specification is used to construct an interface that gives types to the protected functions. The types on protected functions specify sufficient preconditions that the type system for the new language can enforce the security policy. Because the instrumented application programs type check against this interface, they are secure.

In the following section, we will describe the main constructs in the language $\lambda_{\mathcal{A}}$ independently of any security

$$
\begin{array}{lrcl}
\textit{kinds} & \kappa & ::= & \texttt{Val} \mid \texttt{State} \mid \mathcal{B} \mid (\kappa, \ldots, \kappa) \to \mathcal{B} \\
\textit{indices} & \iota, \hat{a}, \hat{q} & & \\
\textit{index sig} & \mathcal{I} & : & \textit{indices} \to \textit{kinds} \\
\textit{pred. vars} & \rho, \varrho & & \\
\textit{predicates} & P & ::= & \rho \mid \iota \mid \iota(P_1, \ldots, P_n) \\
\textit{predicate ctxt} & \Delta & ::= & \cdot \mid \Delta, \rho{:}\kappa \mid \Delta, P \\
\textit{base types} & b & \in & \text{BaseType} + \texttt{S} \\
\textit{types} & \tau & ::= & b(P) \mid \forall[\Delta].(P, \tau_1, \ldots, \tau_n) \to \texttt{0} \mid \exists \rho{:}\kappa.\tau \\
 & & & \\
\textit{constants} & a, q, f & & \\
\textit{constant sig} & \mathcal{C} & : & \textit{constants} \to \textit{types} \\
\textit{value vars} & g, x & & \\
\textit{values} & v & ::= & x \mid a \mid \texttt{fix}\, g[\Delta].(P, x_1{:}\tau_1, \ldots, x_n{:}\tau_n).e \mid v[P] \mid v[\cdot] \mid \texttt{pack}[P, v] \textit{ as } \tau \\
\textit{expressions} & e & ::= & v_0(v_1, \ldots, v_n) \mid \texttt{halt} \mid \texttt{let}\, \rho, x = \texttt{unpack}\, v\, \texttt{in}\, e \mid \texttt{if}\, v\, (q \to e_1 \mid _ \to e_2)
\end{array}
$$

Figure 6: Syntax of $\lambda_{\mathcal{A}}$

policy. In section 4.2, we will show how to specialize the language by constructing the typed interface for a particular security automaton. Section 5 describes an algorithm for instrumenting application programs.

4.1 The Syntax and Semantics of $\lambda_{\mathcal{A}}$

The secure language contains three main parts: the predicates P, the types τ, and the term level constructs. We will explain each of these parts in succession. Figure 6 presents the syntax of the entire language.

Predicates Predicates, P, may be variables ρ or ϱ, indices (constant predicates) ι, or functions of a number of arguments $\iota(P_1, \ldots, P_n)$. There are three distinct kinds of predicates:

- Predicates that correspond to values of base type (kind `Val`).
- Predicates that correspond to security automaton states (kind `State`).
- Predicates that describe relations between values and/or states (kind $(\kappa_1, \ldots, \kappa_n) \to \mathcal{B}$ where $\mathcal{B}$ is the boolean kind).

Most security policies depend upon the properties of particular values, and consequently, we must track these values in the type system. This is the role of the predicates with kind `Val`. For each value a of base type, there is a corresponding index that we will write using the notation $\hat{a}$. For example, the file foo has an associated index $\widehat{foo}$. Using the index $\widehat{foo}$, we can specify precise properties of foo including "foo is readable" or "foo is writeable." In our examples, we will use the meta-variable ρ for predicate variables that range over indices of kind `Val`.

We are careful to distinguish between indices and values in order to separate computation performed at compile time (type checking) and computation performed at run time. Indices, and more generally, predicates are used exclusively at compile time; we erase these annotations before running the program. If we want to use the information contained in a predicate at run time, we must use the corresponding value instead. This design has the advantage that we need only pass run-time data around when we actually need it rather than because we required it to type check the program.

The second kind of predicate allows us to specify information about automaton states. Again, for every automaton state q, there is an associated index $\hat{q}$. We will usually use the predicate $\hat{q}$ to indicate the program is currently executing in automaton state q. Sometimes in our examples (and in particular for the states *start* and *has_read*), we will omit the hat notation; context is enough to discriminate between states and their associated indices. The meta-variable ϱ will range over variables of kind `State`.

Finally, the language of predicates contains a set of relations. These relations will serve two purposes. The first purpose is to describe the transition function of the security automaton. For example, the transition function for the automaton of Section 3 can be described by two predicates, $\delta_{send}(P_{q_1}, P_{q_2})$ and $\delta_{read}(P_{q_1}, P_{q_2}, P_{file})$. The former predicate may be read "in state P_{q_1}, executing the *send* operation causes a transition to state P_{q_2}." The latter predicate is similar. The predicate $\delta_{send}(start, start)$ states that performing a *send* operation in the *start* state causes a transition to the *start* state. These predicates will be discussed in more detail in the next section when we describe how to instantiate the language signature for a particular automaton.

The second relation we will need is the special predicate $\neq (\cdot, \cdot)$ of kind $(\texttt{State}, \texttt{State}) \to \mathcal{B}$. We will use this predicate to denote the fact that some state q is not equal to the *bad* state and consequently that it is safe to execute an operation that causes a transition into q.

We specify the well-formedness of predicates using the judgement $\Phi \vdash P : \kappa$ where Φ is a type-checking context containing three components: a predicate context Δ, a finite map Γ from value variables to types, and another predicate P' indicating the current state of the automaton. The latter two components are not used for specifying the well-formedness of predicates (they will be used for type-checking terms). The signature $\mathcal{I}$ assigns kinds to the indices. Kinds for variables are determined by the predicate context Δ. The kind of a function symbol must agree with the kinds of the predicates to which it is applied. The formal rules are uninteresting so we have removed them to Appendix A.

Figure 7 gives the rules for provability of predicates. The judgement $\Phi \vdash P_{bool}$ indicates that the boolean-valued predicate P_{bool} is true, and the judgement $\Phi \vdash P_{state}$ in_state indicates that the program is currently executing in the automaton state P_{state}. We will elaborate on how these judgements are used when we describe the static semantics of functions.

$\boxed{\Phi \vdash P \qquad \Phi \vdash P \text{ in_state}}$

$$\frac{}{\Delta_1, P, \Delta_2; \Gamma; P' \vdash P} \quad (1)$$

$$\frac{}{\Phi \vdash \neq (\hat{q}, \hat{q'})} \ (\hat{q} \neq \hat{q'}) \quad (2)$$

$$\frac{}{\Delta; \Gamma; P \vdash P \text{ in_state}} \quad (3)$$

Figure 7: Static Semantics: Provability

Aside from the special predicate $\neq (\cdot, \cdot)$, our predicates are completely uninterpreted. This decision makes it trivial to show the decidability of the type system. However, some optimizations may not be possible without a stronger logic. To remedy this situation, implementers are free to add axioms to the type system provided they also supply a decision procedure for the richer logic. In Section 5.2, we show how to add security policy-specific axioms that allow many unnecessary security checks to be eliminated.

Types As mentioned above, security policies often depend upon the properties of particular values. In order to reflect values into the type structure, we use singleton types $b(P)$ where P is either a index or a variable of kind `Val` and b is one of the base types. For example, our file foo will be assigned the type $file(\widehat{foo})$.

In many situations, we will not be able to infer the state of the security automaton statically. As a result, we will have to pass representations of automaton states around at run time and check them dynamically to determine their values. These state representations will have the singleton type $\mathtt{S}(P)$ where P has kind `State` and `S` is the new base type for automaton states.

In many circumstances, we may not know or even care which state or value we are manipulating. For instance, the math library may not contain any security-sensitive operations. We would simply like these functions to manipulate integers without being specific about which ones. In this case, we will use the existential type $\exists\rho{:}\mathtt{Val}.int(\rho)$ to indicate we have an integer, but we do not know which one. As we will see in Section 5, this existential encodes the generic (non-singleton) base types from the source language.

The final type constructor is a modified function type $\forall[\Delta].(P, \tau_1, \ldots, \tau_n) \to \mathtt{0}$. The predicate context Δ abstracts a series of predicate variables for unknown values or states and requires that a sequence of boolean-valued predicates be satisfied before the function can be invoked. The predicate P in the first argument position is not an argument to the function. Rather, it is another precondition requiring that the function be called in the state associated with P. The actual arguments to the function must have types τ_1 through τ_n.

We specify the well-formedness of types using the judgement $\Phi \vdash \tau$. A type is well-formed if Δ contains the free predicate variables of the type. Function types have the form $\forall[\Delta].(P, \tau_1, \ldots, \tau_n) \to \mathtt{0}$. They require that P has kind `State` and that the predicates occuring in Δ have kind $\mathcal{B}$. Once again, the formal rules may be found in Appendix A. Because predicates are uninterpreted, we can use standard syntactic equality of types up to alpha conversion of bound variables.

Values and Expressions The typing rules for values have the form $\Phi \vdash v : \tau$ and state that the value v has type τ in the given context. The judgement $\Phi \vdash e$ states that the expression e is well-formed. Recall that CPS expressions do not return values, and hence the latter judgement is not annotated with a return type. Figure 8 presents the formal rules. In these judgements, we use the notation $\Phi, \rho{:}\kappa$ to denote a new context in which the binding $\rho{:}\kappa$ has been appended to the list of assumptions in Φ. The operation is undefined if ρ appears in Φ. The notations Φ, P and $\Phi, x{:}\tau$ and the extension to Φ, Δ are similar, although P may already appear in Φ.

The values include variables and constants. The treatment of variables is standard, and as in the insecure language, a signature $\mathcal{C}$ gives types to the constants.

The value $v[P]$ is the instantiation of the polymorphic value v with the predicate P. We consider this instantiation a value because predicates are used only for type-checking purposes; they have no run-time significance. The value $v[\cdot]$ is somewhat similar: if v has type $\forall[P, \Delta].(\cdots) \to \mathtt{0}$ and we can prove the predicate P is valid in the current context, we give $v[\cdot]$ the type $\forall[\Delta].(\cdots) \to \mathtt{0}$. Again, the notation $[\cdot]$ is used only to specify that the type-checker should attempt to prove the precondition; it will not influence the execution of programs. In a system with a more sophisticated logic than the one presented in this paper, we might not want to trust the correctness of complex decision procedures for the logic. In this case, we would replace $[\cdot]$ with a proof of the precondition and replace the type checker's decision procedure with a much simpler proof-checker.

Target language function values differ from source language functions in that they specify a list of preconditions using the predicate context Δ. Every function also expresses a state precondition P. The function can only be called in the state denoted by P. Static semantics rule (10) contains the the judgement $\Phi \vdash P$ in_state, which ensures this invariant is maintained. This rule also enforces the standard constraints that argument types must match the types of the formal parameters. Finally, because the predicate context is empty, any preconditions the function might have specified must have already been proven valid.

Rule (6) states that we type check the body of a function assuming its preconditions hold. In this rule, we use the notation $\Phi \leftarrow P$ to denote a context Φ' in which the state component of Φ has been *replaced* by P. For example, suppose a function g is defined in the context $\Delta'; \Gamma'; P'$. The type checker can use any of the predicates in Δ' to help prove g is well-formed but it cannot assume that g will be called in the state P'. The function g is defined here, but may not be used until much later in the computation when the state is different (P'' perhaps).

It is tempting to define a predicate "$in_state(P)$" and to include this predicate in the list of function preconditions Δ. Using this mechanism, it may appear as though we could eliminate the special-purpose state component of the type-checking context. Unfortunately, this simplification is unsound. Consider the following informal example:

```
% Assume current state = start
let g:∀[in_state(start)].(τ1, ..., τn) → 0 = ··· in
% Prove precondition:
let g':∀[].(τ1, ..., τn) → 0 = g[·] in
% Change the state to q' where q' ≠ start:
let _ = op() in
% g is not called in the start state!
g'(v1, ..., vn)
```

$\boxed{\Phi \vdash v : \tau}$

$$\frac{}{\Phi \vdash x : \tau} \ (\Phi(x) = \tau) \tag{4}$$

$$\frac{}{\Phi \vdash a : \tau} \ (\mathcal{C}(a) = \tau) \tag{5}$$

$$\frac{\begin{array}{c}\Phi \vdash \tau_g \\ (\Phi, \Delta, g{:}\tau_g, x_1{:}\tau_1, \ldots, x_n{:}\tau_n) \leftarrow P \vdash e\end{array}}{\Phi \vdash \texttt{fix}\, g[\Delta].(P, x_1{:}\tau_1, \ldots, x_n{:}\tau_n).e : \tau_g} \tag{6}$$
$$(\text{where } \tau_g = \forall[\Delta](P, \tau_1, \ldots, \tau_n) \to \mathtt{0})$$

$$\frac{\Phi \vdash v : \forall[\rho{:}\kappa, \Delta].(P', \tau_1, \ldots, \tau_n) \to \mathtt{0} \quad \Phi \vdash P : \kappa}{\Phi \vdash v[P] : (\forall[\Delta].(P', \tau_1, \ldots, \tau_n) \to \mathtt{0})[P/\rho]} \tag{7}$$

$$\frac{\Phi \vdash v : \forall[P, \Delta].(P', \tau_1, \ldots, \tau_n) \to \mathtt{0} \quad \Phi \vdash P}{\Phi \vdash v[\cdot] : \forall[\Delta].(P', \tau_1, \ldots, \tau_n) \to \mathtt{0}} \tag{8}$$

$$\frac{\Phi \vdash P : \kappa \quad \Phi \vdash v : \tau[P/\rho]}{\Phi \vdash \texttt{pack}[P, v] \ as \ \exists\rho{:}\kappa.\tau : \exists\rho{:}\kappa.\tau} \tag{9}$$

$\boxed{\Phi \vdash e}$

$$\frac{\begin{array}{c}\Phi \vdash v_0 : \forall[].(P, \tau_1, \ldots, \tau_n) \to \mathtt{0} \\ \Phi \vdash v_1 : \tau_1 \quad \cdots \quad \Phi \vdash v_n : \tau_n \\ \Phi \vdash P \text{ in_state}\end{array}}{\Phi \vdash v_0(v_1, \ldots, v_n)} \tag{10}$$

$$\frac{}{\Phi \vdash \texttt{halt}} \tag{11}$$

$$\frac{\Phi \vdash v : \exists\rho{:}\kappa.\tau \quad \Phi, \rho{:}\kappa, x{:}\tau \vdash e}{\Phi \vdash \texttt{let}\, \rho, x = \texttt{unpack}\, v\, \texttt{in}\, e} \tag{12}$$

$$\frac{\begin{array}{cc}\Phi \vdash v : \mathtt{S}(\rho) & \Phi \vdash q : \mathtt{S}(\hat{q}) \\ \Phi' \vdash e_1[\hat{q}/\rho] & \Phi, \neq (\rho, \hat{q}) \vdash e_2\end{array}}{\Phi \vdash \texttt{if}\, v\, (q \to e_1 \mid _ \to e_2)} \tag{13}$$
$$(\text{where } \Phi = \Delta, \rho{:}\texttt{State}, \Delta'; \Gamma; P \text{ and } \Phi' = \Delta, (\Delta'[\hat{q}/\rho]); \Gamma[\hat{q}/\rho]; P[\hat{q}/\rho])$$

$$\frac{\begin{array}{cc}\Phi \vdash v : \mathtt{S}(\hat{q}) & \Phi \vdash q : \mathtt{S}(\hat{q}) \\ \Phi \vdash e_1 & \end{array}}{\Phi \vdash \texttt{if}\, v\, (q \to e_1 \mid _ \to e_2)} \tag{14}$$

$$\frac{\begin{array}{cc}\Phi \vdash v : \mathtt{S}(P) & \Phi \vdash q : \mathtt{S}(\hat{q}) \\ \Phi, \neq (P, \hat{q}) \vdash e_2 & \end{array}}{\Phi \vdash \texttt{if}\, v\, (q \to e_1 \mid _ \to e_2)} \ \begin{pmatrix} P \neq \rho \\ P \neq \hat{q} \end{pmatrix} \tag{15}$$

Figure 8: Static Semantics: Values and Expressions

In the last line, the function g is invoked in a state q' when the function definition assumed it would be invoked in the *start* state. The example highlights the main difference between the state predicates and the others: The validity of the predicates in Δ is invariant throughout the execution of the program whereas the validity of a state predicate varies during execution because it depends implicitly on the current state of the machine.

Existential values are handled in standard fashion [9]. The value $\texttt{pack}[P, v]$ *as* $\exists\rho{:}\kappa.\tau$ creates an existential package that hides P in τ using ρ. The corresponding elimination form, $\texttt{let}\, \rho, x = \texttt{unpack}\, v'\, \texttt{in}\, e$ unpacks the existential v', substituting v for x and P for ρ into the remaining expression e. As with polymorphic types, we assume a type-erasure interpretation of existentials.

Finally, the conditional $\texttt{if}\, v\, (q \to e_1 \mid _ \to e_2)$ tests an automaton state v to determine whether v is the state q or not. If so, the program executes e_1 (see rule (13)) and if not, the program executes e_2. A variant of Harper and Morrisett's typecase [5] operator, `if` also performs type refinement. If v has type $\mathtt{S}(\rho)$ then it refines the type-checking context with the information that $\rho = \hat{q}$ by substituting $\hat{q}$ for ρ. On the other hand, if v is not q, the second branch is taken and the context is refined with the information $\neq (\rho, \hat{q})$. Programs can use this mechanism to dynamically check whether or not they are about to enter the bad state and prevent it.

There is no need to use the $\texttt{if}\, v$ construct if we know which state a value v represents. For example, we know the expression $\texttt{if}\, q\, (q \to e_1 \mid _ \to e_2)$ will reduce to e_1 and therefore e_2 is dead code and the test is wasted computation. However, during the proof of soundness of the type system, such configurations arise and cause difficulties. To avoid these difficulties, we follow the strategy of Crary *et al.* [3] and add the *trivialization* rules (14) and (15) which deal with these redundant cases. Each rule type checks only the branch of the if statement that will be taken.

Operational Semantics The operational semantics for the language is given by the relation $e \longmapsto_s e'$ (see Figure 9). The symbol s is either empty $(\cdot)$ or it is a protected function symbol applied to some number of arguments $(f(a_1, \ldots, a_n))$. Most operations, and, in fact, all of the operations shown in Figure 9, emit the empty symbol. However, this figure does not show the operation of the protected functions. In the next section, we will explain the operational semantics of the protected functions f in the context of a signature for a particular security automaton.

We have given a typed operational semantics to facilitate the proof of soundness of the system. However, inspection of the rules will reveal that evaluation does not depend upon types or predicates, provided the expressions are well-formed. Therefore we can type-check a program and then erase the types before executing it.

4.2 The Security Automaton Signature

In order to specialize the generic language, we construct a typed interface or *signature* for the constants in the language. The signature for a security automaton $\mathcal{A}$, shown in Figure 10, has three parts: the type signature $\mathcal{I}$, the value signature $\mathcal{C}$, and the operational signature. These signatures are derived from the insecure language signature $\mathcal{C}_{insecure}$ for constants and from the security automaton specification.

The type signature gives each constant $\hat{a}$ from the insecure language and state $\hat{q}$ from the security automaton kinds `Val` and `State` respectively. Furthermore, for each

$$v(v_1,\ldots,v_n) \longmapsto_{\cdot} e_m[v, v_1, \ldots, v_n/g, x_1, \ldots, x_n] \qquad \begin{array}{l}\text{if } v = v'\,\phi_1 \cdots \phi_m \\ \text{and } v' = \mathtt{fix}\, g[\theta_1,\ldots,\theta_m].(x_1{:}\tau_1,\ldots,x_n{:}\tau_n).e_0 \\ \text{and for } 1 \le i \le m, \\ \quad \phi_i = [\cdot] \text{ and } \theta_i = P_i \text{ and } e_i = e_{i-1}, \text{ or} \\ \quad \phi_i = [P_i'] \text{ and } \theta_i = \rho_i{:}\kappa_i \text{ and } e_i = e_{i-1}[P_i'/\rho_i]\end{array}$$

$$\mathtt{let}\,\rho, x = \mathtt{unpack}\,(\mathtt{pack}[P, v]\ as\ \tau)\,\mathtt{in}\,e \longmapsto_{\cdot} e[P, v/\rho, x]$$

$$\mathtt{if}\,q'\,(q \to e_1 \mid _ \to e_2) \longmapsto_{\cdot} e_1 \qquad \text{if } q' = q$$

$$\mathtt{if}\,q'\,(q \to e_1 \mid _ \to e_2) \longmapsto_{\cdot} e_2 \qquad \text{if } q' \neq q$$

Figure 9: Operational Semantics for $\lambda_{\mathcal{A}}$

Type and Value Signature

$$\begin{array}{lcll}
\mathcal{I}(\hat{a}) & = & \mathtt{Val} & \text{for } a \in \mathrm{A} \\
\mathcal{I}(\hat{q}) & = & \mathtt{State} & \text{for } q \in Q \\
\mathcal{I}(\delta_f) & = & (\mathtt{State}, \mathtt{State}, \overbrace{\mathtt{Val}, \ldots, \mathtt{Val}}^{n}) \to \mathcal{B} & \\
 & & \quad \text{if } \mathcal{C}_{insecure}(f) = (b_1, \ldots, b_n, (b) \to \mathtt{0}) \to \mathtt{0} & \\
 & & & \\
\mathcal{C}_{target}(a) & = & b(\hat{a}) \quad \text{if } \mathcal{C}_{insecure}(a) = b & \\
\mathcal{C}_{target}(q) & = & \mathtt{S}(\hat{q}) \quad \text{if } q \in Q & \\
\mathcal{C}_{target}(\delta_f) & = & \forall[\varrho_1{:}\mathtt{State}, \rho_1{:}\mathtt{Val}, \ldots, \rho_n{:}\mathtt{Val}, \neq (\varrho_1, bad)].(\varrho_1, S(\varrho_1), b_1(\rho_1), \ldots, b_n(\rho_n), & \\
 & & \quad \forall[\varrho_2{:}\mathtt{State}, \delta_f(\varrho_1, \varrho_2, \rho_1, \ldots, \rho_n)](\varrho_1, S(\varrho_2)) \to \mathtt{0}) \to \mathtt{0} & \\
 & & \quad \text{if } \mathcal{C}_{insecure}(f) = (b_1, \ldots, b_n, (b) \to \mathtt{0}) \to \mathtt{0} & \\
\mathcal{C}_{target}(f) & = & \forall[\varrho_1{:}\mathtt{State}, \varrho_2{:}\mathtt{State}, \rho_1{:}\mathtt{Val}, \ldots, \rho_n{:}\mathtt{Val}, \neq (\varrho_2, bad), \delta_f(\varrho_1, \varrho_2, \rho_1, \ldots, \rho_n)]. & \\
 & & \quad (\varrho_1, b_1(\rho_1), \ldots, b_n(\rho_n), \forall[].(\varrho_2, \exists\rho{:}\mathtt{Val}.b(\rho)) \to \mathtt{0}) \to \mathtt{0} & \\
 & & \quad \text{if } \mathcal{C}_{insecure}(f) = (b_1, \ldots, b_n, (b) \to \mathtt{0}) \to \mathtt{0} &
\end{array}$$

Operational Signature

$$\begin{array}{ll}
\delta_f[\hat{q_1}][\hat{a_1}] \cdots [\hat{a_n}][\cdot](q_1, a_1, \ldots, a_n, v_{cont}) & \text{if } \delta(f)(q_1, a_1, \ldots, a_n) = q_2 \\
\quad \longmapsto_{\cdot} v_{cont}[\hat{q_2}][\cdot](q_2) & \text{and for } 1 \le i \le n, \mathcal{C}_{insecure}(a_i) = b_i \\
 & \text{and } \mathcal{C}_{insecure}(f) = (b_1, \ldots, b_n, (b) \to \mathtt{0}) \to \mathtt{0} \\
 & \\
f[\hat{q_1}][\hat{q_2}][\hat{a_1}] \cdots [\hat{a_n}][\cdot][\cdot](a_1, \ldots, a_n, v_{cont}) & \text{if for } 1 \le i \le n, \mathcal{C}_{insecure}(a_i) = b_i \\
\quad \longmapsto_{f(a_1,\ldots,a_n)} v_{cont}(\mathtt{pack}[\hat{a}, a]\ as\ \exists\rho{:}\mathtt{Val}.b(\rho)) & \text{and } \mathcal{C}_{insecure}(f) = (b_1, \ldots, b_n, (b) \to \mathtt{0}) \to \mathtt{0} \\
\quad \text{(for some } a \text{ such that } \mathcal{C}_{insecure}(a) = b) &
\end{array}$$

Figure 10: Signature for $\lambda_{\mathcal{A}}$

protected function f, the type signature specifies a predicate δ_f. The invariant the type system ensures is that for all $\hat{q_1}, \hat{q_2}, \hat{a}_1, \ldots, \hat{a}_n$, we will be able to prove the predicate $\delta_f(\hat{q_1}, \hat{q_2}, \hat{a}_1, \ldots, \hat{a}_n)$ only if the corresponding automaton transition holds. In other words, only if:

$$\delta(f)(q_1, a_1, \ldots, a_n) = q_2$$

The value signature specifies that objects a and states q are given the correct singleton types. For each protected function symbol f in the insecure language, the signature gives types to a pair of function values. The function δ_f is supplied by the implementer of the security policy; it is used to dynamically determine the automaton state transition function given the program executes the function f in the state ϱ_1 with arguments identified by $\rho_1, \ldots, \rho_n$. When δ_f has computed the transition, it calls its continuation, passing it the next state ϱ_2 so the continuation can test this state to determine whether it is the *bad* state. The continuation assumes the predicate $\delta_f(\varrho_1, \varrho_2, \rho_1, \ldots, \rho_n)$. Before calling the function f itself, we require that the type-checker be able to prove that f will make a transition to some state other than the bad state. Hence the precondition on f states that we must know the automaton transition we are making $(\delta_f(\varrho_1, \varrho_2, \rho_1, \ldots, \rho_n))$ and moreover that that the new state ϱ_2 is not equal to *bad*.

4.3 Properties of $\lambda_{\mathcal{A}}$

A predicate P is valid with respect to an automaton $\mathcal{A}$, written $\mathcal{A} \models P$, if:

- P is $\neq (\hat{q}, \hat{q'})$ and $q \neq q'$, or
- P is $\delta_f(\hat{q_1}, \hat{q_2}, \hat{a}_1, \ldots, \hat{a}_n)$ and $\mathcal{A}.\delta(f)(q_1, a_1, \ldots, a_n) = q_2$

We say that an expression e is *secure* with respect to a security automaton $\mathcal{A}$ in state q, written $\mathcal{A}; q \vdash e$, if

1. $q \neq bad$

and there exist predicates $P_1, \ldots, P_n$ such that:

2. $\mathcal{A} \models P_i$, for $1 \leq i \leq n$
3. $P_1, \ldots, P_n; \cdot; \hat{q} \vdash e$

If we can prove that an expression e is secure in our deductive system then the expression should not violate the security policy when it executes. The soundness theorem below formalizes this notion. The first part of the theorem, *Type Soundness*, ensures that programs obey a basic level of control-flow safety. More specifically, it ensures that expressions do not get *stuck* during the course of evaluation. An expression e is *stuck* if e is not `halt` and there does not exist an e' such that $e \longmapsto_s e'$. Hence, Type Soundness implies a program will only halt when it executes the halt instruction and not because we have applied a function to too few or the wrong types of arguments. The second part of the theorem, *Security*, ensures programs obey the policy specified by the security automaton $\mathcal{A}$. In other words, the sequence of protected operations executed by the program must form a string in the language $\mathcal{L}(\mathcal{A})$. In this second statement, we use the notation $|s_1, \ldots, s_n|$ to denote the subsequence of the symbols $s_1, \ldots, s_n$ with all occurences of $\cdot$ removed.

Theorem 1 (Soundness)
If $\mathcal{A}; q_0 \vdash e_1$ *then*

1. (Type Soundness) *For all evaluation sequences* $e_1 \longmapsto_{s_1} e_2 \longmapsto_{s_2} \cdots \longmapsto_{s_n} e_{n+1}$, *the expression* e_{n+1} *is not* stuck.
2. (Security) *If* $e_1 \longmapsto_{s_1} e_2 \longmapsto_{s_2} \cdots \longmapsto_{s_n} e_{n+1}$ *then* $|s_1, s_2, \ldots, s_n| \in \mathcal{L}(\mathcal{A})$

Soundness can be proven syntactically in the style of Wright and Felleisen [27] using the following two lemmas. The proof appears in a companion technical report [26].

Lemma 2 (Progress) *If* $\mathcal{A}; q \vdash e$ *then either:*

1. $e \longmapsto_s e'$ *or*
2. $e =$ `halt`

Lemma 3 (Subject Reduction) *If* $\mathcal{A}; q \vdash e$ *and* $e \longmapsto_s e'$ *then*

1. *if* $s = \cdot$ *then* $\mathcal{A}; q \vdash e'$
2. *and if* $s = f(a_1, \ldots, a_n)$ *then* $\mathcal{A}; q' \vdash e'$ *where* $\delta(f)(q, a_1, \ldots, a_n) = q'$

Finally, inspection of the typing rules will reveal that for any expression or value, there is exactly one typing rule that applies and that the preconditions for the rules only depend upon subcomponents of the terms or values (with possibly a predicate substitution). Judgements for the well-formedness of types and predicates are also well-founded so the type system is decidable:

Theorem 4 *It is decidable whether or not* $\Phi \vdash e$.

4.4 Language Extensions

If security policies depend upon higher-order functions or immutable data structures such as tuples and records, we will have to track the values of these data structures in the type system using singleton types as we did with values of base type. The simplest way to handle this extension is to use an allocation semantics [13, 14]. In this setting, when a function closure $\texttt{fix}\, g[\Delta](\cdots).e$ is allocated, it is bound to a new address (ℓ). Instead of substituting the closure through the rest of the code as we do now, we would substitute the address (ℓ) through the code and give it the singleton type $\tau(\hat{\ell})$ where τ is $\forall[\Delta](\cdots) \to \texttt{0}$. All the other mechanisms remain unchanged. We decided not to present this style of semantics in this paper because it adds extra mechanism but gives little additional insight into the solution of the problem. An allocation semantics is common in low-level typed languages such as Morrisett's Typed Assembly Language [15] that must reason about memory structure.

There are a number of possibilities for handling mutable data structures. The main principle is that if security-sensitive operations depend upon mutable data then the state of that data must be encoded in the state of the automaton. The assignment operator must be designated as a protected operation that changes the state.

5 Program Instrumentation

It is straightforward to design a translation that instruments our insecure source language with security checks (see Figure 11) now that we have set up the appropriate type-theoretic machinery in the secure target language.

$$
\begin{array}{lcl}
|b| & = & \exists\rho{:}\mathtt{Val}.b(\rho) \\
|(\tau_1,\ldots,\tau_n)\to \mathtt{0}| & = & \forall[\varrho{:}\mathtt{State},\neq(\varrho,bad)].(\varrho,\mathtt{S}(\varrho),|\tau_1|,\ldots,|\tau_n|)\to\mathtt{0} \\
 & & \\
|x| & = & x \\
|a| & = & \mathtt{pack}[\hat{a},a]\ as\ \exists\rho{:}\mathtt{Val}.b(\rho) \quad \text{if } \mathcal{C}_{insecure}(a)=b \\
|\mathtt{fix}\, g(x_1{:}\tau_1,\ldots,x_n{:}\tau_n).e| & = & \mathtt{fix}\, g[\varrho{:}\mathtt{State},\neq(\varrho,bad)].(\varrho,x{:}S(\varrho),x_1{:}|\tau_1|,\ldots,x_n{:}|\tau_n|).|e|_{\varrho,x} \\
|f| & = & \mathtt{fix}_[\varrho_1{:}\mathtt{State},\neq(\varrho_1,bad)](\varrho_1,x_0{:}S(\varrho_1),x_1{:}|b_1|,\ldots,x_n{:}|b_n|,x_{n+1}{:}|(b)\to\mathtt{0}|). \\
 & & \quad \mathtt{let}\,\rho_1,x_1' = \mathtt{unpack}\,x_1\,\mathtt{in} \\
 & & \quad \ldots \\
 & & \quad \mathtt{let}\,\rho_n,x_n' = \mathtt{unpack}\,x_n\,\mathtt{in} \\
 & & \quad \mathtt{let}\,\varrho_2,\delta_f(\varrho_1,\varrho_2,\rho_1,\ldots,\rho_n),x_{\varrho_2} = \delta_f[\varrho_1][\rho_1]\cdots[\rho_n][\cdot](x_0,x_1',\ldots,x_n')\,\mathtt{in} \\
 & & \quad \mathtt{if}\, x_{\varrho_2}(\\
 & & \qquad bad\to\mathtt{halt} \\
 & & \qquad |\ _\to \mathtt{let}\, x = f[\varrho_1][\varrho_2][\rho_1]\cdots[\rho_n][\cdot][\cdot](x_1',\ldots,x_n')\,\mathtt{in} \\
 & & \qquad\quad x_{n+1}[\varrho_2][\cdot](x_{\varrho_2},x)) \\
 & & \text{if } \mathcal{C}_{insecure}(f)=(b_1,\ldots,b_n,(b)\to\mathtt{0})\to\mathtt{0} \\
 & & \\
|v_0(v_1,\ldots,v_n)|_{P,v} & = & |v_0|[P][\cdot](v,|v_1|,\ldots,|v_n|) \\
|\mathtt{halt}\,|_{P,v} & = & \mathtt{halt}
\end{array}
$$

Figure 11: Program Instrumentation

The interesting portion of the type translation involves the translation of function types. The static semantics maintains the invariant that programs never enter the bad state and we naturally express this fact as a precondition to every function call. Hence the type translation of the function type $(\tau_1,\ldots,\tau_n)\to\mathtt{0}$ is:

$$\forall[\varrho{:}\mathtt{State},\neq(\varrho,bad)].(\varrho,\mathtt{S}(\varrho),|\tau_1|,\ldots,|\tau_n|)\to\mathtt{0}$$

In general, we may not know the current state statically so we quantify over all states ϱ, provided $\varrho\neq bad$. In order to determine state transfers do not go wrong, we will also have to thread a representation of the state ($\mathtt{S}(\varrho)$) through the computation.

Most of the work in the value translation is accomplished during the translation of the protected operations. We unpack the arguments so the individual values can be tracked through the type system and then use δ_f to determine the state transition that will occur if we execute f on these arguments. After checking to ensure we do not enter the bad state, we execute f itself passing it a continuation that executes in state ϱ_2. In this translation, we use the abbreviation:

$$\begin{array}{l}\mathtt{let}\,\Delta,x_1,\ldots,x_n = v(v_1,\ldots,v_n)\,\mathtt{in}\,e \equiv \\ v(v_1,\ldots,v_n,\mathtt{fix}\,_[\Delta].(x_1{:}\tau_1,\ldots,x_n{:}\tau_n).e)\end{array}$$

and we assume predicate and value variables bound by `let` are fresh.

Instrumented programs type check and thus they are *secure* in the sense made precise in the last section:

Theorem 5 *If* $\vdash_{insecure} e$ *then* $\mathcal{A};q_0\vdash|e|_{\hat{q}_0,q_0}$.

This property can be proven using a straightforward induction on the typing derivation of the source term.

5.1 Instrumenting The Taxation Applet

The code below shows the results of instrumentating the taxation applet from Section 2 with checks from the security automaton of Section 3. We have simplified the output of the formal translation slightly to make it more readable. In particular, we have inlined the functions that the translation wraps around each protected function symbol.

```
fix _[ϱ1:State, ≠ (ϱ1, bad)]
     (ϱ1, xϱ1:S(ϱ1), secret:∃ρ:Val.file(ρ),
     xcont:τcont).
  let ϱ2, δsend(ϱ1, ϱ2), xϱ2 = δsend[ϱ1][·](x) in
  if xϱ2 (
    bad → halt
    | _→
      let _ = send[ϱ1][ϱ2][·][·]() in
      let ρ, secret′ = unpack secret in
      let ϱ3, δread(ϱ2, ϱ3, ρ), xϱ3 =
        δread[ϱ2][·](xϱ2, secret′) in
      if xϱ3 (
        bad → halt
        | _→
          let salary = read[ϱ2][ϱ3][·][·](secret′) in
          let taxes = salary in
          xcont[ϱ3][·](xϱ3, taxes)))
```

where $\tau_{cont}=$

```
  ∀[ϱcont, ≠ (ϱcont, bad)].
    (ϱcont, S(ϱcont), ∃ρ′:Val.int(ρ′)) → 0
```

The translation does not assume that the taxation applet is invoked in the initial automaton state and consequently the resulting function abstracts the input state ϱ_1. Also, as specified by the translation, objects of base type, like the file *secret* become existentials. The main point of interest in this example is that before each of the protected operations *send* and *read*, the corresponding automaton function determines the next state. Then the `if` construct checks that these states are not *bad*. If the dynamic check succeeds, the type checker introduces information into the context that allows it to infer that executing the *read* and *send* operations is safe. When reading the code calling *send* or *read*, notice that the instantiation of predicate variables indicates the state transition that occurs. For example, execution of the expression $send[\varrho_1][\varrho_2][\cdot][\cdot]()$ causes the automaton to make a transition from ϱ_1 to ϱ_2.

5.2 Optimization

Many security automata exhibit special structure that allows us to optimize secure programs by eliminating checks that are inserted by the naive program instrumentation procedure [24]. One common case is an operation f that always succeeds in a given state q and transfers control to a new state q' regardless of its arguments. In this situation, we can make the following axiom available to the type checker:

$$\frac{}{\Phi \vdash \delta_f(q, q', P_1, \ldots, P_n)} \text{ (for all } P_1, \ldots, P_n)$$

If we know we are in state q, we can use the axiom above and the fact that $q \neq bad$ to satisfy the precondition on f; there is no need to perform a run-time check.

The *send* operation in the security automaton in Section 3 has this property. When invoked in the initial state, *send* always succeeds and execution continues in the initial state. Therefore, we can safely add the axiom:

$$\frac{}{\Phi \vdash \delta_{send}(start, start)}$$

Now, if we know our taxation function is only invoked in the *start* state, we can rewrite it, eliminating one of the run-time checks:

```
% Optimization 1:
fix _[](start, x_start:S(start), secret:∃ρ:Val.file(ρ),
        x_cont:τ_cont).
  let _ = send[start][start][·][·]() in
  ...
```

The type checker can prove *send* is executed in the *start* state and that the predicate $\delta_{send}(start, start)$ and the predicate $\neq (start, bad)$ are valid. Therefore, the optimized applet continues to type-check.

A second important way to optimize $\lambda_{\mathcal{A}}$ programs is to perform a control-flow analysis that propagates provable predicates statically through the program text. Using this technique, we can further optimize the taxation applet. Assume the calling context can prove the predicate $\delta_{read}(start, has_read, \rho)$ (perhaps a run-time check was performed at some earlier time) where ρ is the value predicate corresponding to the file *secret*. In this case, the caller can invoke a tax applet with a stronger precondition that includes the predicate $\delta_{read}(start, has_read, \rho)$. Moreover, with this additional information, an optimizer can eliminate the redundant check surrounding the file read operation:

```
% Optimization 2:
fix _[ρ:Val, δ_read(start, has_read, ρ)](start,
        secret:file(ρ),
        x_cont:∀[].(has_read, ∃ρ:Val.int(ρ)) → 0).
  let _ = send[start][start][·][·]() in
  let salary = read[start][has_read][·][·](secret) in
  let taxes = salary in
  x_cont(taxes)
```

In the code above, the optimizer rewrites the applet precondition with the necessary information. The caller is now obligated to prove the additional precondition before the applet can be invoked. The caller also unpacks the secret file before making the call so that the type checker can make the connection between the arguments to the δ_{read} predicate and this particular file. Finally, because the automaton state transitions are statically known throughout this program, we do not need to thread the state representation through the program. We assumed an optimizer was able to detect this unused argument and eliminate it. After performing all these optimizations, the resulting code is operationally equivalent to the original taxation applet from section 2, but provably secure.

The flexibility in the type system is particularly useful when a program repeatedly performs the same restricted operations. A more sophisticated tax applet might need to make a series of reads from the secret file (for charitable donations, number of dependents, etc.). If we assume the recursive function *read_a_lot* performs these additional reads, we need no additional security checks:

```
fix read_a_lot
    [ϱ:State, ρ:Val, δ_read(ϱ, has_read, ρ), ≠ (ϱ, bad)]
    (ϱ, secret:file(ρ),
    x_cont:∀[](has_read, ∃ρ:Val.int(ρ)) → 0).
  % In unknown state ϱ
  let info = read[ϱ][has_read][·][·]() in
  % In known state has_read
  ...
  % Must prove δ(has_read, has_read, ρ)
  read_a_lot[has_read][ρ][·][·](secret, x_cont)
```

The *read_a_lot* function can be invoked in a good state ϱ (*i.e.* either *start* or *has_read*) when we can prove the predicate $\delta_{read}(\varrho, has_read, \rho)$. Using the δ_{read} predicate in the function precondition, the type checker infers that the read operation transfers control from the ϱ state to the *has_read* state. Before the recursive call, the type checker has the obligation to prove $\delta_{read}(has_read, has_read, \rho)$ but it cannot do so because it only knows that $\delta_{read}(\varrho, has_read, \rho)$! Fortunately, we can remedy this problem by adding another policy-specific axiom to the type-checker:

$$\frac{\Phi \vdash \neq (P, bad) \quad \Phi \vdash \delta_{read}(P, has_read, P_f)}{\Phi \vdash \delta_{read}(P', has_read, P_f)} \text{ (for all } P, P', P_f)$$

This axiom states that if we can read a file P_f in one state (P), then we can read it in any state (except the *bad* one) and we always move to the *has_read* state. This condition is easily decidable.

In practice, Erlingsson and Schneider's untyped optimizer analyzes security automaton structure and performs optimizations similar to the ones discussed above [24]. Once the optimizer has obtained the information necessary for a particular transformation, this information can also be used to automatically generate the policy-specific axioms that we have discussed.

One significant optimization that is used in practice that we cannot encode in this typed framework is inlining. The run-time security checks, δ_f, are abstract constants in our framework because we have made a decision to trust the implementor of the security policy. When the policy writer implements the functions δ_f, he does not have to supply a formal proof that the functions imply some semantic property; the implementer merely asserts that some abstract predicate $\delta_f(q_1, q_2, \rho_1, \ldots, \rho_n)$ has been satisfied. If we inline δ_f into untrusted code, the corresponding assertion will also be inlined *into untrusted code*. In the current framework, there is no way to verify that these assertions are placed correctly by untrusted code. If performance is critical, the user will have to rely on a trusted just-in-time compiler to inline checks when they cannot be proven redundant statically.

6 Related and Future Work

The design of $\lambda_{\mathcal{A}}$ was inspired, in part, by Xi and Pfenning's Dependent ML (DML) [29, 28]. As in DML, we track the identity of values using dependent refinement types and singleton types. However, rather than applying the technology to array bounds check elimination and dead-code elimination, we have applied it to the problem of expressing security policies.

The secure language $\lambda_{\mathcal{A}}$ also bears resemblance to monadic systems [10, 20]. Both paradigms can be used to thread state through a computation. Normally, monadic types do not express the details of the state transformation. In contrast, the $\lambda_{\mathcal{A}}$ types describe the effect of the translation precisely: A function precondition specifies the input state and its continuation specifies the output state.

Leroy and Rouaix [8] also consider security in the context of strongly-typed languages. Their main concern is proving that standard strongly-typed languages provide certain security properties. For example, they show that a program written in a typed lambda calculus augmented with references cannot modify unreachable (in the sense of tracing garbage collection) locations. They also show that they can wrap dynamic checks around writes to sensitive locations and ensure "bad" values are not written to these locations. They did not investigate mechanisms for mechanically checking that their instrumented programs are safe, nor did they study the broader range of security properties that can be enforced using security automata.

Evans and Twyman [4] have developed the Naccio system for specifying security policies. Like SASI, Naccio allows users to add security state to untrusted programs and to define operations that perform security checks. Naccio does not produce certified code and, like SASI, requires a control-flow safety enforcement mechanism (such as SFI or a strong type system) to ensure that malicious programs cannot bypass security checks.

The inability to inline security checks suggests an interesting direction for future research. The central problem is that the inlined code is trusted and has more priviledges than code written by an outsider. In particular, the inlined code has the priviledge to assert a state change in the security automaton. In order to inline we need a way to distinguish trusted and untrusted segments of code and to ensure that the trusted segments have not been tampered with. Recent work by Zdancewic *et al.* [30] describes how to distinguish between code segments with varying privileges. It may be possible to augment their system with some syntactic rules for reasoning about code equivalence. Using these techniques, inlining and then further optimization such as constant folding may be possible.

In order to test the practical consequences of our language design, we are interested in porting Erlingsson and Schneider's untyped SASI compiler [24] to the type-preserving Popcorn compiler [11] and using Typed Assembly Language as the secure target language. The current TAL implementation contains many of the typing constructs that we will need including polymorphic function types, existentials, and singleton types. We are also currently augmenting the system with a first-order predicate logic that interprets integer inequalities. We believe we will be able extend this framework with the additional constructs necessary to generate a secure, yet flexible and efficient Typed Assembly Language.

Acknowledgements

I greatly appreciate the time Úlfar Erlingsson and Fred Schneider have spent explaining their security automaton model and SASI system to me. Karl Crary, Úlfar Erlingsson, Neal Glew, Dan Grossman, and Steve Zdancewic have made extremely helpful suggestions on previous drafts of this work. I thank my advisor Greg Morrisett for the energy and enthusiasm he has dedicated to my education.

References

[1] Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools.* Addison-Wesley, 1986.

[2] Bowen Alpern and Fred Schneider. Recognizing safety and liveness. *Distributed Computing*, 2:117–126, 1987.

[3] Karl Crary, Stephanie Weirich, and Greg Morrisett. Intensional polymorphism in type-erasure semantics. In *ACM SIGPLAN International Conference on Functional Programming*, pages 301–312, Baltimore, September 1998.

[4] David Evans and Andrew Twyman. Flexible policy-directed code safety. In *IEEE Security and Privacy*, Oakland, CA, May 1999.

[5] Robert Harper and Greg Morrisett. Compiling polymorphism using intensional type analysis. In *Twenty-Second ACM Symposium on Principles of Programming Languages*, pages 130–141, San Francisco, January 1995.

[6] John E. Hopcroft and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation.* Addison-Wesley, 1979.

[7] Dexter Kozen. Efficient code certification. Technical Report TR98-1661, Cornell University, January 1998.

[8] Xavier Leroy and Francois Rouaix. Security properties of typed applets. In *Twenty-Fifth ACM Symposium on Principles of Programming Languages*, pages 391–403, San Diego, January 1998.

[9] John C. Mitchell and Gordon D. Plotkin. Abstract types have existential type. *ACM Transactions on Progamming Languages and Systems*, 10(3):470–502, July 1988.

[10] Eugenio Moggi. Notions of computation and monads. *Information and Computation*, 93:55–92, 1991.

[11] Greg Morrisett, Karl Crary, Neal Glew, Dan Grossman, Richard Samuels, Frederick Smith, David Walker, Stephanie Weirich, and Steve Zdancewic. TALx86: A realistic typed assembly language. In *ACM SIGPLAN Workshop on Compiler Support for System Software*, pages 25–35, Atlanta, GA, May 1999.

[12] Greg Morrisett, Karl Crary, Neal Glew, and David Walker. Stack-based typed assembly language. In *Second International Workshop on Types in Compilation*, pages 95–117, Kyoto, March 1998. Published in Xavier Leroy and Atsushi Ohori, editors, *Lecture Notes in Computer Science*, volume 1473, pages 28-52. Springer-Verlag, 1998.

[13] Greg Morrisett, Matthias Felleisen, and Robert Harper. Abstract models of memory management. In *ACM Conference on Functional Programming and Computer Architecture*, pages 66–77, La Jolla, June 1995.

[14] Greg Morrisett and Robert Harper. Semantics of memory management for polymorphic languages. In A.D. Gordon and A.M. Pitts, editors, *Higher Order Operational Techniques in Semantics*, Publications of the Newton Institute. Cambridge University Press, 1997.

[15] Greg Morrisett, David Walker, Karl Crary, and Neal Glew. From System F to Typed Assembly Language. In *Twenty-Fifth ACM Symposium on Principles of Programming Languages*, pages 85–97, San Diego, January 1998.

[16] George Necula. Proof-carrying code. In *Twenty-Fourth ACM Symposium on Principles of Programming Languages*, pages 106–119, Paris, 1997.

[17] George Necula and Peter Lee. Safe kernel extensions without run-time checking. In *Proceedings of Operating System Design and Implementation*, pages 229–243, Seattle, October 1996.

[18] George Necula and Peter Lee. The design and implementation of a certifying compiler. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 333 – 344, Montreal, June 1998.

[19] George C. Necula and Peter Lee. Safe, untrusted agents using proof-carrying code. *LNCS 1419: Special Issue on Mobile Agent Security*, October 1997.

[20] Simon L. Peyton Jones and Philip Wadler. Imperative functional programming. In *Twentieth ACM Symposium on Principles of Programming Languages*, Charleston, South Carolina, January 1993.

[21] John C. Reynolds. Definitional interpreters for higher-order programming languages. In *Conference Record of the 25th National ACM Conference*, pages 717–740, Boston, August 1972.

[22] Fred Schneider. Enforceable security policies. Technical Report TR98-1664, Cornell University, January 1998.

[23] Chris Small. MiSFIT: A tool for constructing safe extensible C++ systems. In *Proceedings of the Third USENIX Conference on Object-Oriented Technologies*, Portland, OR, June 1997.

[24] Úlfar Erlingsson and Fred B. Schneider. SASI enforcement of security policies: A retrospective. Submitted for publication, April 2, 1999.

[25] Robert Wahbe, Steven Lucco, Thomas Anderson, and Susan Graham. Efficient software-based fault isolation. In *Fourteenth ACM Symposium on Operating Systems Principles*, pages 203–216, Asheville, December 1993.

[26] David Walker. A type system for expressive security policies. Technical Report TR99-1740, Cornell University, April 1999.

[27] Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.

[28] Hongwei Xi. *Dependent Types in Practical Programming*. PhD thesis, Carnegie Mellon University, 1999.

[29] Hongwei Xi and Frank Pfenning. Dependent types in practical programming. In *Twenty-Sixth ACM Symposium on Principles of Programming Languages*, pages 214–227, San Antonio, TX, January 1999.

[30] Steve Zdancewic, Dan Grossman, and Greg Morrisett. Principals in programming languages. Submitted to the Fourth ACM SIGPLAN International Conference on Functional Programming, 1999.

A Type Formation for $\lambda_{\mathcal{A}}$

$Dom(\Delta)$	=	The set of variables ρ such that $\rho{:}\kappa$ appears in Δ
Γ		A finite map of value vars to types
$\Gamma, x{:}\tau$		Update Γ so that $\Gamma(x) = \tau$
Φ	::=	$\Delta; \Gamma; P$
$(\Delta; \Gamma; P), \rho{:}\kappa$	=	$\Delta, \rho{:}\kappa; \Gamma; P$ if $\rho \notin \Delta$
$(\Delta; \Gamma; P), x{:}\tau$	=	$\Delta; \Gamma, x{:}\tau; P$ if $x \notin \Gamma$
$(\Delta; \Gamma; P), P'$	=	$\Delta, P'; \Gamma; P$
$(\Delta; \Gamma; P) \leftarrow P'$	=	$\Delta; \Gamma; P'$

Figure 12: Type Checking Contexts

$\boxed{\Phi \vdash P : \kappa}$

$$\frac{}{\Delta; \Gamma; P \vdash \rho : \kappa} \ (\Delta(\rho) = \kappa) \tag{16}$$

$$\frac{}{\Delta; \Gamma; P \vdash \iota : \kappa} \ (\mathcal{I}(\iota) = \kappa) \tag{17}$$

$$\frac{\Phi \vdash \iota : (\kappa_1, \ldots, \kappa_n) \to \kappa \quad \Phi \vdash P_1 : \kappa_1 \quad \cdots \quad \Phi \vdash P_n : \kappa_n}{\Phi \vdash \iota(P_1, \ldots, P_n) : \kappa} \tag{18}$$

$\boxed{\Phi \vdash \Delta}$

$$\frac{}{\Phi \vdash \cdot} \tag{19}$$

$$\frac{\Delta; \Gamma; P \vdash \Delta'}{\Delta; \Gamma; P \vdash \Delta', \rho{:}\kappa} \ (\rho \notin Dom(\Delta) \cup Dom(\Delta')) \tag{20}$$

$$\frac{\Phi \vdash \Delta \quad \Phi, \Delta \vdash P : \mathcal{B}}{\Phi \vdash \Delta, P} \tag{21}$$

$\boxed{\Phi \vdash \tau}$

$$\frac{\Phi \vdash P : \mathtt{Val}}{\Phi \vdash b(P)} \ (b \neq \mathtt{S}) \tag{22}$$

$$\frac{\Phi \vdash P : \mathtt{State}}{\Phi \vdash \mathtt{S}(P)} \tag{23}$$

$$\frac{\Phi \vdash \Delta \quad \Phi, \Delta \vdash P : \mathtt{State} \quad \Phi, \Delta \vdash \tau_1 \quad \cdots \quad \Phi, \Delta \vdash \tau_n}{\Phi \vdash \forall[\Delta].(P, \tau_1, \ldots, \tau_n) \to \mathtt{0}} \tag{24}$$

$$\frac{\Phi, \rho{:}\kappa \vdash \tau}{\Phi \vdash \exists \rho{:}\kappa.\tau} \tag{25}$$

Figure 13: Contexts, Types, and Predicates