Verified Enforcement of Automaton-based Information Release Policies

Nikhil Swamy and Michael Hicks
University of Maryland, College Park; Department of Computer Science; Technical Report CS-TR-4903

Abstract

Many organizations specify information release *policies to describe the terms under which sensitive information may be released to other organizations. This paper presents a new approach for ensuring that security-critical software correctly enforces its information release policy. Our approach has two parts. First, an information release policy is specified as a security automaton written in a new language called* AIR. *Second, we enforce an* AIR *policy by translating it into an API for programs written in* λ_{AIR}*, a core formalism for a functional programming language.* λ_{AIR} *uses a novel combination of dependent, affine, and singleton types to ensure that the API is used correctly. As a consequence we can certify that programs written in* λ_{AIR} *meet the requirements of the original* AIR *policy specification.*

1 Introduction

Many organizations, including financial institutions, healthcare providers, the military, and even the organizers of this conference, wish to specify the terms under which sensitive information in their possession can be released to their partners, clients, or the public. Such a specification constitutes an *information release policy*.

These policies are often quite complex. For example, consider the policy that regulates the disclosure of military information to foreign governments as defined by the United States Department of Defense [1992]. This policy includes the following provisions: a release must be authorized by an official with disclosure authority who represents the "DoD Component that originated the information"; the system must "edit or rewrite data packages to exclude information that is beyond that which has been authorized for disclosure"; a disclosure shall not occur until the foreign government has submitted "a security assurance [...] on the individuals who are to receive the information"; and, that the release must take place in the Foreign Disclosure and Technical Information System in which both approvals and denials of a release request must be logged.

We would like to ensure that software systems that handle sensitive data—including military systems, but also programs like medical-record databases, online auction software, and network appliances—correctly enforce such a high-level policy. As a concrete example, consider a specific kind of application called a *cross-domain guard*. These are programs, like network firewalls, that mediate the transfer of information between organizations at different trust levels. Commercial guards, e.g., the Data Sync guard produced by BAE [Focke et al., 2006], do not enforce high-level policies but rather implement low-level "dirty keyword" filters.

The research community has only recently begun to consider the verified enforcement of release policies. For instance, FlowWall [Hicks et al., 2007] (arguably the research counterpart of a system like DataSync guard) is a firewall which, by virtue of being built with the Jif programming language [Chong et al., 2006], is sure to enforce a low-level filtering policy, but it does not appeal to high-level information release criteria. Augmenting information flow policies with high-level conditions that control information release has been proposed by Chong and Myers [2004] and, more recently, by Banerjee et al. [2008]. However, in both these cases, reasoning separately about high-level release decisions is difficult since the release policy is embedded within the program.

To fill this gap, this paper presents a methodology for building highly-assured software that acts in accordance with a high-level information release policy. Our approach has two parts. First, we define AIR, a formal language for defining information release policies separately from the program that is to be secured. AIR's design follows from the observation that an information release policy is a kind of stateful authorization policy naturally expressed as an automaton [Schneider, 2000] (AIR stands for *automata for information release*). Satisfaction of a release obligation advances the state of the automation, and once all obligations have been fulfilled, the automaton reaches the accepting state and the protected information can be released. AIR allows one to express such automata in a natural way.

Second, we define λAIR (pronounced "lair"), a core calculus in which type-correct programs can be shown to correctly enforce an AIR policy. We have mechanized the proof of soundness of λAIR using the Coq proof assistant. λAIR has three elements.

• First, λAIR provides *singleton types* to allow a programmer to associate sensitive data with an AIR automaton that protects that data. For example, an object x implementing a security automaton is given type ¡$Instance^N$, where N is a type-level name unique to x. Then, an integer i protected by x would be given type *Protected Int N*, which is essentially a kind of dynamic labeling [Zheng and Myers, 2004]. While the state of an automaton can change, its association to a protected value will not change until all policy obligations have been fulfilled and the data is released, thus ensuring a kind of *complete mediation*. Prior work on verified enforcement of security automata via type checking [Walker, 2000] or inlined reference monitors [Erlingsson, 2004] is less flexible and/or has a larger trusted computing base (Section 6).

• Second, a λAIR program can express a release obligation with a *dependent type*, where an object having that type serves as a proof that the obligation has been fulfilled. For example, data could be released to a principal p only if p acts for some principal q (where p and q are program variables that store public keys). A proof of this fact could be represented by an object with type *ActsFor p q*. Generally speaking, proof objects represent certificates which are used to produce a *certified evaluation* of stateful policy logic—every authorization decision is accompanied by a proof that all obligations mandated by the high-level policy have been met. Certificates can be produced locally or remotely. We combine *affine types* with dependent and singleton types to ensure that that stale evidence about old policy states are never used in authorization decisions.

• Finally, given these mechanisms, we provide a way to compile an AIR policy to an API in λAIR, where each API function corresponds to an automaton transition such that the type of that function precisely expresses the evidence necessary for a transition to succeed. Thus type-correct λAIR programs must use the compiled AIR API correctly and, as a consequence, meet the specifications of the high-level policy. More precisely, we prove that the sequence of events produced by a program's execution is a word in the language accepted by the AIR automaton.

Using our techniques, one could build a cross-domain guard that adheres to high-level policy prescriptions; e.g., it would release information only after confirming that appropriate security assurances have been received, that to-be-released data packages have been rewritten appropriately, and that audit logs have been updated.

Our use of AIR policies for information release departs from prior work on declassification policies in that we do not focus on establishing a noninterference-like property for programs. However, our work complements noninterference-oriented interpretations of information release. For example, additional support for robust declassification [Zdancewic and Myers, 2001] could ensure that an adversary never causes information to be released, and furthermore, when it is released, it always follows the prescription of the high-level AIR policy.

2 AIR: Automata for Information Release

This section presents AIR, our language for expressing information release policies as automata.

2.1 Syntax of AIR, by Example

An AIR policy consists of one or more *class declarations*. A program will contain *instances* of a class, where each instance protects some sensitive data via a labeling. Protected data can be accessed in two ways. First, each class `C` has an *owning principal P* such that P and all who *act for P* may access data protected by an instance of `C`. Second, each class defines a *release policy* by which its protected data can be released to an instance of a different class.

The release policy is expressed using rules that define a security automaton, which is a potentially infinite state machine in which states represent security-relevant configurations. In the case of AIR, the security automaton defines conditions that must hold before data can be released. Each class instance consists of its current state, and each condition that is satisfied transitions the automaton to the next state. These transitions ultimately end in a release rule that allows data to be released to a different class instance, potentially in a modified form. Because sensitive data is associated with instances rather than classes, multiple resources may be governed by the same policy template (i.e., the automaton defined by the class) but release decisions are made independently for each resource. Dually, related resources can be protected by the same instance, thereby allowing release decisions made with respect to one resource to affect the others.

The formal syntax of AIR policies is presented in Figure 1. We explain the syntax of AIR while stepping through a running example, shown in Figure 3. Throughout, we use the notation $\overrightarrow{a}$ to stand for the ordered sequence $a_1, \ldots, a_n$. Where the context is clear, we will also treat $\overrightarrow{a}$ as the set $\{a_1, \ldots, a_n\}$.

A class declaration consists of a class identifier, an identifier for the owning principal, a list of automaton states, and a

Metavariables

id	class and rule ids	P	principals		
$\mathscr{C}$	state constructors	n, i, j	integers	x, y, z	variables

Core language

Declarations	D	::=	class id = (principal:P; states: $\overrightarrow{S}$; $\overrightarrow{R}$)
States	S	::=	$\mathscr{C}$ | $\mathscr{C}$ of $\overrightarrow{t}$
Rules	R	::=	id : R | id : T
Release	R	::=	When G release e with next state A
Transition	T	::=	When G do e with next state A
Guards	G	::=	x requested for use at y and $\overrightarrow{\exists x{:}t.C}$
Conditions	C	::=	A_1 *IsClass* A_2 | A_1 *InState* A_2
		|	A_1 *ActsFor* A_2 | $A_1 \leq A_2$
Atoms	A	::=	n | x | id | P | $\mathscr{C}(\overrightarrow{A})$ | $A_1 + A_2$
		|	Self | Class(A) | Principal(A)

e is an expression and t is a type in λAIR. (cf. Figure 5)

Figure 1. Syntax of AIR

```
class Alice_Confidential =
   principal : Alice;    states : Count of Int;

   release_to_bob :
      When x requested for use at d and
   ...
```

Figure 2. A simple AIR **policy**

sequence of rules that define the automaton transitions. Our example declares a single class US_Army_Confidential, owned by the principal *US_Army*, that defines the policy for confidential data owned by the U.S. Army. For simplicity, our examples use a flat namespace for class identifiers, and abstract names for principals.

Automaton states are represented by terms constructed from an algebraic datatype. The example has two kinds of states. The nullary constructor Init represents the initial state of the automaton; all classes must have this state. The other kind of state is an application of the unary constructor Debt to an argument of type *Int*. Constructors of the form $\mathscr{C}$ of $\overrightarrow{t}$ may carry data as indicated by the types $\overrightarrow{t}$. Types t (such as *Int*) are drawn from the programming language λAIR in which programs using AIR policies are written; λAIR is discussed in the next section.

Each rule in an AIR class is given a name, and is either a *release rule* or a *transition rule*. Each rule begins with a clause "When x requested for use at d", which serves to bind variables x and d in the remainder of the rule. Here, x names the information protected by an instance of this class, requested for release to some other instance d (usually of another class). This clause is followed by a conjunction of conditions that restrict the applicability of a rule; we discuss these in more detail below. Following these conditions, the rule specifies a λAIR expression e that can either release information (perhaps after downgrading it by filtering or encryption) or do some other action (like logging), depending on whether the rule is a release rule or a transition rule. A rule concludes with the next state of the automaton.

The first rule in the US_Army_Confidential class is a release rule called Conf_secret. This rule is qualified by a condition expression Class(d) *IsClass* US_Army_Secret stating that the rule applies when releasing x to an instance d of a class named US_Army_Secret. If applicable, this rule allows x to be released without modification—the release expression is simply x, and not, some function that downgrades x. After the release, the automaton remains in its current state; i.e. the state Self.

We use a small ontology for conditions based on integers, principals, classes and their instances—*IsClass* mentioned above, is one such condition. Generally speaking, condition expressions C are typed binary predicates over atoms A. For example, A_1 *ActsFor* A_2 is defined for *Principal*-typed atoms A_1 and A_2, and asserts that A_1 acts for A_2 according to some acts-for hierarchy among principals (not explicitly modeled here). Atoms include integers n, variables x, identifiers id, principal constants P, state literals constructed from an application of a state constructor $\mathscr{C}$ to a list of atoms, addition of integers and the implicit variable Self. We also include two operators: Class(z) is the class of the argument z, a class instance; and, Principal(z), which is the principal that owns the class z. Finally, we permit a condition C to be prefixed by one or more existentially quantified variables—i.e., in $\exists x_1{:}t_1.C_1, \ldots, \exists x_n{:}t_n.C_n$, each x_i is a variable of type t_i and is in scope as far to the right as possible, until the end of the rule. We omit the quantifier prefix when no such variables exist.

```
class US_Army_Confidential =
  principal : US_Army;    states : Init, Debt of Int;

  Conf_secret :
    When x requested for use at d and
      Class(d) IsClass US_Army_Secret
    release x with next state Self

  Conf_init :
    When x requested for use at d and
      Self InState Init
    do _ with next state Debt(0)

  Conf_coalition :
    When x requested for use at d and
      Principal(Class(d)) ActsFor Coalition,
      ∃count:Int.Self InState Debt(count),
      count ≤ 10
    release
      (log(...x...d); encrypt (pubkey (principal (class d))) x)
    with next state Debt(count + 1)
```

Figure 3. A stateful information release policy in AIR

2.2 A Simple Stateful Policy in AIR

Taken as a whole, the class US_Army_Confidential can be thought of as implementing a simple kind of *risk-adaptive* access control [Cheng et al., 2007], in which information is released according to a *risk budget*, with the intention of quantifying the risks vs. the benefits of releasing sensitive information. This class maintains a current risk debt, as reflected in the state Debt of *Int*. Each time the class authorizes an information release we add an estimate of the risk associated with that release to the debt. When the accumulated risk debt exceeds a threshold then releases outside the U.S. Army are no longer permitted. The other two rules in the policy, Conf_init and Conf_coalition, implement this behavior.

The Conf_init transition rule applies when processing a release to an instance d and when the automaton is in the Init state. The "do" expression initializes the risk debt to 0 by transitioning the automaton to the Debt(0) state. The Conf_coalition rule allows information to be released to a coalition partner. In particular, if the release target class is owned by a principal that acts for the *Coalition* (expressed by Principal(Class(d)) *ActsFor Coalition*), then information can be released only if the current risk debt has not exceeded the budget, as expressed in the latter two conditions. The first of these requires the current state of the automaton to be Debt(*count*), where *count* is variable with type *Int* which holds the current risk debt. The last condition requires that *count* is not above the preallocated risk budget of 10. With these conditions satisfied, Conf_coalition logs the fact that a release has been authorized and permits release of the data after it has been downgraded using an encryption function. In this case, the downgrading expression encrypts x with the public key of the principal that owns the class of the instance d. Unlike releases to US_Army_Secret which do not alter the risk debt, Conf_coalition increments the risk debt by transitioning to the Debt($count + 1$) state, indicating that releases to the *Coalition* are more risky than upgrading to a higher classification level of the same organization (via rule Conf_secret).

AIR as presented here is particularly simple. We anticipate extending AIR with support for more expressive condition ontologies and release rules. For instance, instead of a fixed set of ontologies, we could embed a stateful authorization logic (say, in the style of SMP [Becker and Nanz, 2007]) to allow custom ontologies and release rules to be programmed within an AIR class. We could also introduce a set of downgrading and logging primitives to completely separate AIR from λ_{AIR}.

3 A Programming Model for AIR

Given a particular AIR policy, we would like to do two things. First, we must have a way of reflecting an AIR policy in a program by protecting sensitive resources with instances of an AIR class. Second, we must ensure that all uses of protected data adhere to the prescriptions of the AIR policy. Taken together, we can then claim that an AIR policy is correctly enforced by a program. To achieve these goals, we have defined a formal model for a language called λ_{AIR} in which one writes programs that use AIR policies. λ_{AIR}'s type system ensures that these policies are used correctly. The rest of this section defines the programming model for this language and the next two sections flesh out its syntax and semantics. Section 5.4 proves that type-correct programs act only in accordance with their AIR policies.

```
1   let x_a1, a1 = get_secret_file_and_policy () in
2   let a2, channel = get_request () in
3     (* generating evidence of policy compliance *)
4   let a2, a2_class = get_class a2 in
5   let ev1 = acts_for (principal a2_class) Coalition in
6   let a1, Debt(debt), ev2 = get_current_state a1 in
7   let ev3 = leq debt 10 in
8     (* supplying evidence to policy API and releasing data *)
9   let a1', a2, x_a2 = Conf_coalition a1 x_a1 a2 ev1 debt ev2 ev3 in
10    send channel x_a2
```

Figure 4. Programming with an AIR policy

The programming model for using AIR policies has two elements. First, programmers tie an AIR policy to data in the program by constructing instances of AIR classes and labeling one or more pieces of data with these instances. This association defines (1) the set of principals that may view the data (in particular, the principal P that owns the class, and any principals that may act for P), and (2) the rules that allow the data to be released. As in other security-typed languages, the labeling specification (expressed using type annotations) is part of the trusted computing base.

Second, programmers manipulate data protected by an AIR class instance through a class-specific API that is generated by compiling each AIR class definition to a series of program-level definitions. For example, each AIR class's release and transition rules are compiled to functions that can be used to release protected data. The types given to these functions ensure that a caller of the function must always provide evidence that the necessary conditions to release protected data have been met.

Figure 4 illustrates a program using the AIR policy of Figure 3, written using a ML-like notation. (Significantly, our examples omit type annotations where they do not help clarify the exposition. λ_{AIR} does not support type inference at all.) At a high level, this program processes requests to release information from a secret file. The files are stored on the file system together with a policy label that represents a particular AIR class instance. Before disclosing the information, the program must make sure that the automaton that protects the data is in a state that permits the release. The first two lines set up the scenario. At line 1, we read the contents of a secret file into the variable x_a1 and the automaton that protects this file into the variable a1. Initially, only the principals that act for the owner of the class of a1 can view these secrets. At line 2, the program blocks until a request is received. The request consists of an output channel and another automaton instance a2 that represents the policy under which the requested information will be protected after the release. In effect, the information, once released, will be under the protection of the principal that owns the class of a2.

Prior to responding to the request, on lines 4-7 we must establish that a1 is in a state that permits the release. At line 4, we extract the class of the instance a2. At line 5, we check that the owner of a2's class acts for the Coalition principal and, if this check succeeds, we obtain a certificate ev1 as evidence of this fact. At line 6, we extract the current state of the automaton a1, use pattern matching to check that it is of the form Debt(debt) (for some value of debt) and receive an evidence object ev2 that attests to the fact that a1 is currently in this state. At line 7, we check that the total debt associated with the current state of the automaton is not greater than 10 and obtain ev3 as evidence if the check succeeds.

At line 9 we call Conf_coalition, a function produced by compiling the AIR policy. We pass in the automaton a1 and the secret data x_a1; the automaton a2 to which x_a1 is to be released; and the certificates that serve as evidence for the release conditions. Conf_coalition returns a1' which represents the next state of the automaton (presumably in the Debt(debt+1) state); a2 the unchanged destination automaton; and finally, x_a2, which contains the suitably downgraded secret value. On the last line, we send the released information on the channel received with the request.

For programs like our example, we would like to verify that all releases of information are mediated by calls to the appropriate transition and release rules as defined by the AIR policy (functions like Conf_coalition). Additionally, we would like to verify that a program satisfies the mandates of an AIR policy rule by presenting evidence that justifies the appropriate release conditions. This evidence-passing style supports our goal of certifying the evaluation of all authorization decisions, while being flexible about the mechanism by which an obligation is fulfilled. To return to the DoD example from the introduction, this design gives us the flexibility to allow release authorizations to be obtained in one part of the system and security assurances from the recipient to be handled in another; the cross-domain guard must simply collect evidence from the other components rather than performing these operations itself. λ_{AIR}'s type system is designed so that type correctness ensures these goals are satisfied, i.e., a type-correct program uses its AIR policy correctly. The type system has three key elements:

Singleton types. First, in order to ensure complete mediation, we must be able to correctly associate data with the class instance that protects it. For example, Conf_coalition expects its first argument to be an automaton and the second to be data protected by that automaton. In an ML-like type system, this function's type might begin with $\forall\alpha.\mathit{Instance} \to \alpha \to \ldots$ But such a type is not sufficiently precise since it does not prescribe any relationship between the first and second argument, allowing the programmer to erroneously pass in a2 as the first argument, rather than a1, for example. To remedy this problem, we can give Conf_coalition a type like the following (as a first approximation):

$$\forall N, \alpha.\mathit{Instance}^N \to \mathit{Protected}\ \alpha\ N \to \ldots$$

Here, N is a unique type-level name for the class instance provided in the first argument. The second argument's type $\mathit{Protected}\ \alpha\ N$ indicates it is an α value protected by the instance N, making clear the association between policy and data. We can ensure that values of type $\mathit{Protected}\ \alpha\ N$ may only be accessed by principals P that act for the owner of the class instantiated by the instance named N. This approach is more flexible than implicitly pairing each protected object with its own (hidden) automaton. For example, with our approach one can encode policies like secret sharing, in which a set of related documents are all protected by the same automaton instance. Each document's type would refer to the same automaton, e.g., $\mathit{Protected\ Doc\ N}$. Information released about one document updates the state of the automaton named N and can limit releases of the other documents.

Dependent types. Arguments 4-7 of Conf_coalition represent evidence (proof certificates) that the owner of class instance a2 acts for Coalition, and that a1 is in a state authorized to release the given data. We give types to these arguments that reflect the propositions that they are supposed to witness. For example, we give the seventh argument (ev3) to Conf_coalition the type $\mathit{LEQ\ debt}\ 10$ where LEQ is a dependent type constructor applied to two *expressions*, debt and 10, where each has type Int. Data with type $\mathit{LEQ\ n\ m}$ represents a certificate that proves $n \leq m$. If we allow such certificate values to only be constructed by trusted functions that are known to correctly implement the semantics of integer inequality, then we can be sure that functions like Conf_coalition are only called with valid certificates—i.e., type correctness guarantees that all certificates are valid proofs of the propositions represented by their types and there is no need to inspect these certificates at run time. If we interface with other programs, we can check the validity of proof certificates at run time before allowing a call to proceed. Either way, the type system supports an architecture that enables certified evaluation of an AIR policy.

Affine types. The final piece of our type system is designed to cope with the stateful nature of an AIR policy. The main problem caused by a state change is illustrated by the value returned by the Conf_coalition function. In our example, a1' represents the state of the policy automaton that protects x_a1 *after a release* has been authorized. Thus, we need a way to break the association between x_a1 and the old, stale automaton state a1. We achieve this in two steps. First, even though our type system supports dependent types, as shown earlier, we use singleton types to give x_a1 the type $\mathit{Protected}\ \alpha\ N$, where N is a unique type name for a1 (rather than giving x_a1 a more-direct dependent type of the form $\mathit{Protected}\ \alpha$ a1). The second step is to use *affine types* (values with an affine type can never be used more than once) to consume stale automaton values, so that at any program point, there is only one usable automaton value that has the type-name N. Thus, we give both a1 and a1' the type ¡$\mathit{Instance}^N$, where ¡t denotes an affinely qualified type t. Once a1 is passed as an argument to Conf_coalition (which constitutes a use) it can no longer be used in the rest of the program; a1' is the only automaton that can be used in subsequent authorization checks for x_a1. Thus, a combination of singleton and affine types transparently takes care of relabeling data with new automaton instances. (One might also wonder how we deal with proof certificates that can become stale because of the changing automaton state; we discuss this issue in detail in Section 5.1.)

To illustrate how these singleton, dependent, and affine types interact we show the type of Conf_coalition, slightly simplified, below (the full type is discussed in Section 5.2).

$$\begin{aligned}\forall N, M, \alpha.\ \ & ¡\mathit{Instance}^N \to \mathit{Protected}\ \alpha\ N \to ¡\mathit{Instance}^M \to \\ & \ldots \to (\mathsf{debt} : \mathit{Int}) \to \ldots \to (\mathit{LEQ}\ \mathsf{debt}\ 10) \to \\ & (¡\mathit{Instance}^N \times ¡\mathit{Instance}^M \times \mathit{Protected}\ \alpha\ M)\end{aligned}$$

The first three arguments are the *affine* source automaton (a1), the data it protects (x_a1), and the *affine* destination automaton (a2). On the next line, we show the dependent type given to the evidence that the current debt of the automaton is not greater than 10. Finally, consider the return type of Conf_coalition. The first component of this three-tuple is a class instance with the same name N as the first argument. This returned value is the new state of the automaton named N—it protects all existing data of type $\mathit{Protected}\ \alpha\ N$ (such as x_a1). The second component of the three-tuple is the unchanged target automaton. The third component contains the data ready to be released—its type, $\mathit{Protected}\ \alpha\ M$, indicates that it is now protected by the target automaton instance M. In effect, λ_{AIR} models state modifications by requiring automata states to

Metavariables

B Base terms T Type constructors α, β, γ Type vars

Core language

Terms	e	::=	$x \mid \lambda x{:}t.e \mid e\ e \mid \Lambda\alpha{::}k.e \mid e\ [t] \mid B \mid e\ \{e\}$
		$\mid$	$\mathsf{case}\ e\ \mathsf{of}\ \overrightarrow{x{:}t}.e : e\ \mathsf{else}\ e \mid \bot \mid \mathsf{new}\ e$
Types	t	::=	$(x{:}t) \to t \mid \alpha \mid \forall\alpha{::}k \xrightarrow{\varepsilon} t \mid T$
		$\mid$	$t \Rightarrow t \mid q\ t \mid t\ t \mid t\ e \mid t^{\eta}$
Type names	η	::=	$\alpha \mid \circ$
Affinity	q	::=	¡ $\mid \cdot$
Simple kinds	k	::=	$\mathtt{U} \mid \mathtt{A} \mid \mathtt{N}$
Kinds	K	::=	$k \mid k \to K \mid t \to K$
Name constraints	ε	::=	$\cdot \mid \alpha \mid \varepsilon \uplus \varepsilon \mid \varepsilon \cup \varepsilon$

Signatures and typing environments

Phase index	φ	::=	$\mathsf{term} \mid \mathsf{type}$
Signatures	S	::=	$(B{:}t) \mid (T{::}K) \mid S,S$
Type env.	Γ	::=	$x{:}t \mid \alpha{::}k \mid \Gamma,\Gamma$
Affine env.	A	::=	$x \mid A,A$

Figure 5. Syntax of λ_{AIR}

be manipulated in a store-passing style, reminescent of a monadic treatment of side effects in a purely functional language. However, by imposing the additional discipline of affine types, we are able to ensure that the program always has a consistent view of an automaton's state, while still retaining the benefits of a well-understood and relatively simple functional semantics.

Adhering to the constraints of λ_{AIR}'s type system is surely more burdensome than when using a more typical programming language. Thus λ_{AIR} may be most appropriate for the security-critical kernel of an application, or even as the (certifiable) target language of a program transformation for inline reference monitoring. We leave to future work an exploration of support—e.g., type inference—for improving λ_{AIR}'s usability.

4 Syntax and Semantics of λ_{AIR}

λ_{AIR} extends a core System F^{ω} [Mitchell, 1996] with support for singleton, dependent, and affine types. λ_{AIR} is parameterized by a *signature* S that defines base terms B and type constructors T—each AIR class declaration D is compiled to a signature S_D that acts as the API for programs that use D. All AIR classes share some elements in common, like integers, which appear in a *prelude* signature S_0. We explain the core of λ_{AIR} using examples from the prelude. The next section describes the remainder of the prelude and shows how our example AIR policy is compiled.

4.1 Syntax

Figure 5 shows the syntax of λ_{AIR}. The core language expressions e are mostly standard, including variables x, lambda abstractions $\lambda x{:}t.e$, application $e\ e'$, type abstraction $\Lambda\alpha{::}k.e$ and type application $e\ [t]$. Functions have dependent type $(x{:}t) \to t'$ where x names the argument and may be bound in t'. Type variables are α. A type t universally quantified over all types α of kind k is denoted $\forall\alpha{::}k \xrightarrow{\varepsilon} t$. Here, ε is a name constraint that records the type names α given to automaton instances in the body of the abstraction; we discuss these in detail later. When the constraint is empty we write a universally quantified type as $\forall\alpha{::}k.t$. The signature S defines the legal base terms B and type constructors T, mapping them to their types t and kinds K, respectively. The prelude S_0 defines several standard terms and types which we use to illustrate some of λ_{AIR}'s main features.

The type constructor *Int* represents the type of integers, and is given `U` kind in the prelude (written *Int*::`U`). Kind `U` is one of three simple kinds k. A type t with simple kind `A` is affine in that the typing rules permit terms of type t to be used at most once. Affine types are written ¡t, to contrast with the "of course" modality in linear logic, which is typically denoted using "!". ¡t is an instance of the form $q\ t$ where $q =$ ¡. Terms whose types have kind `U` are unrestricted in their use. We explain kind `N`, the kind of type names, shortly.

The prelude also defines two base terms for constructing integers: $\mathsf{Zero} : \mathit{Int}$ represents the integer 0, while $\mathsf{Succ} : \mathit{Int} \Rightarrow \mathit{Int}$ is a unary *data constructor* that produces an *Int* given an *Int*. Data constructor application is written $e\ \{e\}$; thus the integer 1 is represented $\mathsf{Succ}\ \{\mathsf{Zero}\}$ (but we write $0, 1, 2$ etc. for brevity). Programs can pattern match data constructors applications using the expression form $\mathsf{case}\ e\ \mathsf{of}\ \overrightarrow{x{:}t}.e : e\ \mathsf{else}\ e$. This is mostly standard; details are in our technical report [Swamy and Hicks, 2008].

In addition to simple kinds k, kinds K more generally can classify functional type constructors, using the forms $k \to K$ and $t \to K$. A type constructor t_1 having the first form can be applied to another type using $t_1\ t_2$ to produce a (standard) type, while one of the second form can be applied to a term using $t\ e$ to produce a dependent type. As an example of the first case, the prelude defines a type constructor $\times$::U $\to$ U $\to$ U to model pairs; $\times$ *Int Int* is the type of a pair of integers (for clarity, from here on we will use infix notation and write a pair type as $t \times t'$). The prelude also defines a base term constructor Pair which has a polymorphic type $\forall\alpha,\beta$::U.$\alpha \Rightarrow \beta \Rightarrow \alpha \times \beta$ for constructing pair values.

Evidence for condition expressions in an AIR policy are given dependent types. For example, the prelude provides means to test inequalities $A_1 \leq A_2$ that appear in a policy and generate certificates that witness an inequality:

$$(\mathit{LEQ}{::}\mathit{Int} \to \mathit{Int} \to \mathtt{U}),$$
$$(\mathit{leq}{:}(x{:}\mathit{Int}) \to (y{:}\mathit{Int}) \to \mathit{LEQ}\ x\ y)$$

LEQ is a dependent-type constructor that takes two expressions of type *Int* as arguments and produces a type having kind U. This type is used to classify certificates that witness the inequality between the term arguments. These certificates are generated by the *leq* function, which has a dependent type: the labels x and y on the first two arguments appear in the returned type. Thus the call *leq* 3 4 would return a certificate of type *LEQ* 3 4 because 3 is indeed less than 4. An attempt to construct a certificate *LEQ* 4 3 by calling *leq* 4 3 would fail at run time, returning $\bot$ (an unrecoverable failure) in our semantics—we could use option types or add support for exceptions to handle failures more gracefully. The signature does not include a data constructor for the *LEQ* type, so its values cannot be constructed directly by programs—the only way is by calling the *leq* function.

We discuss the remaining constructs—including name constraints ε, named types t^η, and the new e construct—in conjunction with the type rules next.

$\Gamma; A \vdash_\varphi e : t$ — A φ-level expression e in environment Γ with affine assumptions A has type t and uses names ε

$$\frac{\Gamma \vdash \Gamma(x) :: U}{\Gamma;\cdot \vdash_\varphi x : \Gamma(x);\cdot}\ \text{(T-X)} \qquad \frac{}{\Gamma; x \vdash_\varphi x : \Gamma(x);\cdot}\ \text{(T-XA)} \qquad \frac{}{\Gamma;\cdot \vdash_{S,\mathsf{type}} x : \Gamma(x);\cdot}\ \text{(T-X-type)} \qquad \frac{\Gamma; A \vdash_{S,\mathsf{type}} e : t; \varepsilon_1 \uplus \varepsilon}{\Gamma; A \vdash_{S,\mathsf{type}} e : t; \varepsilon}\ \text{(T-NC-type)}$$

$$\frac{\Gamma; A \vdash_\varphi e : t; \varepsilon \qquad \Gamma \vdash t :: U \qquad \Gamma(\alpha) = \mathtt{N}}{\Gamma; A \vdash_\varphi \mathsf{new}\ e : ¡t^\alpha; \alpha \uplus \varepsilon}\ \text{(T-NEW)} \qquad \frac{\Gamma; A \vdash_\varphi e : t^\alpha; \varepsilon}{\Gamma; A \vdash_\varphi e : t^\circ; \varepsilon}\ \text{(T-DROP)} \qquad \frac{\Gamma; A \vdash_\varphi e : t; \varepsilon \qquad \varepsilon' \subseteq \mathrm{dom}(\Gamma)}{\Gamma; A, A' \vdash_\varphi e : t; \varepsilon \uplus \varepsilon'}\ \text{(T-WKN)}$$

$$\frac{\begin{array}{c}\Gamma, \alpha{::}k; A \vdash_\varphi e : t; \varepsilon \uplus \varepsilon' \\ \alpha \notin \varepsilon \qquad \varepsilon' \in \{\cdot, \alpha\} \qquad q = p(A, \varepsilon)\end{array}}{\Gamma; A \vdash_\varphi \Lambda\alpha{::}k.e : q(\forall\alpha{::}k \xrightarrow{\varepsilon'} t); \varepsilon}\ \text{(T-TAB)} \qquad \frac{\begin{array}{c}\Gamma \vdash t_x :: k \qquad q = p(A, \varepsilon) \\ \Gamma, x : t_x; A, a(x,k) \vdash_\varphi e : t_e; \varepsilon\end{array}}{\Gamma; A \vdash_\varphi \lambda x{:}t_x.e : q((x{:}t_x) \to t_e); \varepsilon}\ \text{(T-ABS)}$$

where
$a(x, \mathtt{A}) = x \quad a(x, \mathtt{U}) = \cdot$
$p(A, \varepsilon) = ¡ \quad p(\cdot,\cdot) = \cdot$

$$\frac{\Gamma; A \vdash_\varphi e : q(\forall\alpha{::}k \xrightarrow{\varepsilon'} t'); \varepsilon \qquad \Gamma \vdash t :: k}{\Gamma; A \vdash_\varphi e\ [t] : [\alpha \mapsto t]t'; \varepsilon \uplus ([\alpha \mapsto t]\varepsilon')}\ \text{(T-TAP)} \qquad \frac{\Gamma; A \vdash_\varphi e : q((x{:}t') \to t); \varepsilon_1 \qquad \Gamma; A' \vdash_\varphi e' : t'; \varepsilon_2}{\Gamma; A, A' \vdash_\varphi e\ e' : [x \mapsto e']t; \varepsilon_1 \uplus \varepsilon_2}\ \text{(T-APP)}$$

$\Gamma \vdash t :: K$ — A type t has kind K in environment Γ

$$\frac{\Gamma(\alpha) = k}{\Gamma \vdash \alpha :: k}\ \text{(K-A)} \qquad \frac{\Gamma \vdash t :: \mathtt{A} \qquad \Gamma(\eta) = \mathtt{N} \vee \eta = \circ}{\Gamma \vdash t^\eta :: \mathtt{A}}\ \text{(K-N)} \qquad \frac{\Gamma \vdash t :: \mathtt{U}}{\Gamma \vdash ¡t :: \mathtt{A}}\ \text{(K-AFN)} \qquad \frac{\Gamma \vdash t :: k \qquad \Gamma, x : t \vdash t' :: k'}{\Gamma \vdash (x{:}t) \to t' :: \mathtt{U}}\ \text{(K-FUN)}$$

$$\frac{\Gamma' = \Gamma, \alpha{::}k \qquad \Gamma' \vdash t :: k \qquad \alpha' \in \varepsilon \Rightarrow \Gamma'(\alpha') = \mathtt{N}}{\Gamma \vdash \forall\alpha{::}k \xrightarrow{\varepsilon} t :: \mathtt{U}}\ \text{(K-UNIV)} \qquad \frac{\Gamma \vdash t :: t' \to K \qquad \Gamma;\cdot \vdash_{S,\mathsf{type}} e : t';\cdot}{\Gamma \vdash t\ e :: K}\ \text{(K-DEP)}$$

Figure 6. Static semantics of λ_{AIR} (Selected rules)

4.2 Static Semantics

Figure 6 shows the main rules from the static semantics of λ_{AIR}, which consists of two judgments. The full semantics can be found in our technical report. Both judgments are parameterized by a signature S. The typing judgment is additionally parameterized by a *phase index* φ, which indicates whether the judgment applies to a term- or type-level expression. The

judgment giving an expression e a type t is written $\Gamma; A \vdash_{\varphi} e : t; \varepsilon$ where Γ is the standard typing environment, A is a list of affine assumptions, and ε is a name constraint that records the set of fresh type names assigned to automata instances in e. The second judgment, $\Gamma \vdash t :: K$ states that a type t has kind K in the environment Γ.

Recall that the type system must address three main concerns. First, we must correctly assign unique type names to automata instances and then associate these names with protected data. Next, for certified evaluation, we must be able to accurately type evidence using dependent types. Finally, to cope with automaton state changes, we must (via affine types) prevent stale automaton instances from being reused. We consider each of these aspects of the system in turn, first in the typing judgment and then in the kinding judgment.

Assigning unique names to automata. We construct new automata using new e. (T-NEW) assigns the name α to the type in the conclusion, ensuring (via $\alpha \uplus \varepsilon$) that α is distinct from all other names ε that have been assigned to other automata. We require α to be in the initial environment Γ, or to be introduced into the context by a type abstraction. Recall from Section 3 that protected values will refer to this name α in their types (e.g., *Protected Int* α). The resulting type ¡t^{α} is also affinely qualified; we discuss this shortly.

(T-DROP) allows the unique name associated with a type to be replaced with the distinguished constant name $\circ$. This is sound because although the name α of a type ¡t^{α} can be hidden, α cannot be reused as the type-level name of any other automaton (i.e., ε is unaffected). This form of subtyping is convenient for giving types to proof objects that witness properties of the state of an automaton, while keeping our language of kinds for type constructors relatively simple. Section 5.1 illustrates an example use of (T-DROP).

(T-TAB) is used to check type abstractions. The first premise checks the body of the abstraction e in a context that includes the abstracted type variable α. Since we treat type names and types uniformly, functions polymorphic in a type name can be written by quantifying over α::N—the interesting elements of this rule have to do with managing these names. If the body of the abstraction e constructs a new automaton assigned the name α in (T-NEW), then α will be recorded in $\varepsilon \uplus \varepsilon'$, the name constraints of e. In this case $\varepsilon' = \alpha$ and ε contains all the other names used in the typing derivation of e; otherwise ε' is empty. In the conclusion, we decorate the universally quantified type with ε' to signify that the abstracted name α is used in e. Type abstractions are destructed according to (T-TAP). In the premises we require the kind of the argument to match the kind of the formal type parameter. In the conclusion, we must instantiate all the abstracted names ε' used in the body e' and ensure that these are disjoint from all other names ε used in the body.

Two additional points about our formulation of type-level names are worth noting. First, universally quantified types can be decorated with arbitrary name constraints ε (rather than just singleton names α). We expect this to be useful when enforcing composite policies. The name instantiation constraint ε can ensure that a function always constructs automata that belong to a specific set of classes in a large policy. Second, we could support recursion by following an approach taken by Pratikakis et al. [2006]. This requires using existential quantification to abstract names in recursive data structures and including a means to forget names assigned to automata that go out of scope (e.g., in each iteration of a loop).

Dependently typed functions and evidence. (T-ABS) gives functions a dependent type, $(x{:}t) \to t'$. Here, x names the formal parameter and is bound in t'. When a function is applied, (T-APP) substitutes the actual argument e' for x in the return type. Thus, given a function f that has type $(\mathsf{debt} : \mathit{Int}) \to (\mathit{LEQ}\ \mathsf{debt}\ 10) \to t$, the application $(f\ 11)$ is given the type $(\mathit{LEQ}\ 11\ 10) \to t$. That is, the type of the second argument of f *depends on* the term passed as the first argument. Note that although λ_{AIR} permits arbitrary expressions to appear in types, type checking the enforcement of an AIR policy is decidable because we never have to reduce expressions that appear in types.

Affine types for consistent state updates. Finally, we consider how the type system enforces the "use at most once" property of affine types. First, (T-NEW) introduces affine types by giving new automaton instances the type ¡t^{α}. Values of affine type can be destructed in the same way as values of unrestricted type. For example, (T-APP) and (T-TAP) allow e to be applied irrespective of the affinity qualifier on e's type. However, we must make sure that variables that can be bound to affinely typed values are not used more than once. This is prevented by the type rules through the use of affine assumptions A, which lists the subset of variables with affine type in Γ which have not already been used. The use of an affine variable is expressed in the rule (T-XA), which types a variable x in the context of the single affine assumption x. To prevent variables from being used more than once, other rules, such as (T-APP), are forced to split the affine assumptions between their subexpressions. Affine assumptions are added to A by (T-ABS) using the function $a(x,k)$, where x is the argument to the function and k is the kind of its type. If the argument x's type has kind A then it is added to the assumptions, otherwise it is not. We include a weakening rule (T-WKN) that allows affine assumptions to be forgotten (and for additional names ε' to be consumed). Finally, the function $p(A,\varepsilon)$ is used to determine the affinity qualifier of an abstraction. If no affine assumptions from the environment are used in the body of the abstraction ($A = \cdot$) and if no new automata are constructed in the body ($\varepsilon = \cdot$), then

it is unrestricted. Otherwise, it has captured an assumption from the environment or encloses an affinely tracked automaton and should be called at most once.

Kinding judgment. In $\Gamma \vdash t :: K$, the rule (K-A) is standard. (K-N) allows a name to be associated with any affine type t. (K-AFN) checks an affinely-qualified type: types such as ¡¡t are not well-formed. (K-FUN) is standard for a dependent type system—it illustrates that x is bound in the return type t'. (K-UNIV) is mostly standard, except that we must also check that the constraint ε only contain names that are in scope. (K-DEP) checks the application of a dependent type constructor. Here, we have to ensure that the type of the argument e matches the type of the formal. However, since e is a type-level expression, we check it in a context with the phase index $\varphi = \mathsf{type}$. Since types are erased at run time, type-level expressions are permitted, via (T-X-type), to treat affine assumptions intuitionistically. Erasure of types also allows us to lift the name constraints for type-level expressions e—(T-NC-type) allows any subset ε_1 of the names used in e to be forgotten.

4.3 Dynamic Semantics

The dynamic semantics of λAIR defines a standard call-by-value, small-step reduction relation for a purely functional language, using a left-to-right evaluation order. The full definition can be found in our technical report. The form of the relation is :

$$M \vdash e \xrightarrow{l} e'$$

This judgment claims that a term e reduces in a single step to e' in the presence of a model M that interprets the base terms in a signature. The security-relevant reduction steps are annotated with a trace element l, which is useful for stating our security theorem. In this section, we briefly discuss the form of the model M and the trace elements l and state our type soundness result.

Following a standard approach for interpreting constants in a signature [Mitchell, 1996], we define a model M by axiomatizing the reductions of base term applications. In practice, we would implement the model in a real programming language. For example, we could do this in FABLE [Swamy et al., 2008], a language we specifically designed for programming *policy functions* that may coerce one protected type to another (like Conf_coalition) or may produce unforgeable certificates (like acts_for).

A model M contains equations $B : \mathscr{D} \rightsquigarrow e$, where $\mathscr{D}$ is a sequence of types and values. We require the types of the expressions in these equations to be consistent with the type given to B in the signature. An example of an equation is $leq : 4, 3 \rightsquigarrow \bot$ indicating that the expression $(leq\ 4\ 3)$ reduces to $\bot$. We also need a mechanism to construct a value that represents a proof certificate for a valid inequality; i.e., values that inhabit *LEQ* 3 4. In practice, one could either chose a concrete representation for these objects if proofs need to be checked at run time (for instance, when interfacing with type-unsafe code); or, if we are in a purely type-safe setting, we could chose an arbitrary value (like unit) to represent a proof certificate. In our technical report, we introduce a special value to stand for proof objects that facilitates our soundness proof.

The security-relevant actions in a program execution are the reduction steps that correspond to automaton state changes. As indicated earlier, each transition and release rule in a policy will be translated to a function-typed base term like Conf_coalition. Thus, every time we reduce an expression e using a base-term equation $B : \mathscr{D} \rightsquigarrow e'$, we record $l = B : \mathscr{D}$ in the trace: i.e., $M \vdash e \xrightarrow{B:\mathscr{D}} e'$.

The statement of our type soundness theorem is shown below.

Theorem (Type soundness)**.** *Given a set of type names* $\Gamma = \alpha_1 :: \mathtt{N}, \ldots, \alpha_n :: \mathtt{N}$ *such that* $\Gamma; \cdot \vdash_{S,\mathsf{term}} e : t; \varepsilon$*, and an interpretation* M *such that* M *and* S *are type-consistent, then* $\exists e'.M \vdash e \xrightarrow{l} e'$ *or* e *is a value. Moreover, if* $M \vdash e \xrightarrow{l} e'$ *then* $\Gamma; \cdot \vdash_{S,\mathsf{term}} e' : t; \varepsilon$.

A detailed proof sketch of this theorem is in our technical report. We have also mechanized the soundness proof using the Coq proof assistant [Bertot and Castéran, 2004]. Our formalization adapts a proof technique recently proposed by Aydemir et al. [2008]. In particular, we use a locally nameless approach for representing both term- and type-level bindings and rely on cofinite quantification to introduce fresh names. We rely on a set of libraries distributed by Aydemir et al. that provide basic support for working with environments and finite sets. As of this writing, our Coq proof is complete, modulo a collection of identities about finite sets and context splitting. The proofs of these identities are beyond the capabilities of the decision procedures in the finite set libraries that we use and, without automation, we have found proofs of these identities in Coq to be tedious and time consuming. To alleviate this difficulty, we are currently in the process of devising our own set of specialized decision procedures to discharge the proofs of these identities.

5 Translating AIR to λAIR

In this section, we show how we translate an AIR class to a λAIR API, describe how that API is to be used, and state our main security theorem.

5.1 Representing AIR Primitives

In order to enforce an AIR policy we must first provide a way to tie the policy to the program by protecting data with AIR automata. We must also provide a concrete representation for automata instances and a means to generate certificates that attest to the various release conditions that appear in the policy. These constructs are common to all λAIR programs and appear in the standard prelude, along with the integers and pairs discussed in Section 4.1.

Protecting data. As indicated in Section 3, we include the following type constructor to associate an automaton with some data: $(\mathit{Protected}{::}\mathtt{U} \to \mathtt{N} \to \mathtt{U})$. A term with type *Protected t* α is governed by the policy defined by an automaton instance with type-level name α. We would like to ensure that all operations on protected data are mediated by functions that correspond to AIR policy rules. For this reason, we do not provide an explicit data constructor for values of this type (ensuring that they cannot be destructed directly, say, via pattern matching). Values of this type are introduced only by assigning the appropriate types to functions that retrieve sensitive data—for instance, library functions that read secret files from the disk can be annotated so that they return values with a protected type.

In addition to functions corresponding to AIR class rules, we can provide functions that allow a program to perform secure computations over protected values. We have explored such functions in our work on FABLE and showed that computations that respect a variety of policies (ranging from access control to information flow) can be encoded [Swamy et al., 2008]; we do not consider these further here.

Next, we discuss our representation of an AIR automaton—these include representations of the class that the automaton instantiates and the principal that owns the class.

Principals. The nullary constructor *Prin* is used to type principal constants *P*; i.e., $(\mathit{Prin}{::}\mathtt{U}), (P{:}\mathit{Prin})$. As with integers, we need a way to test and generate evidence for acts-for relationships between principals. We include the dependent-type constructor and run-time check shown below.

$$(\mathit{ActsFor}{::}\mathit{Prin} \to \mathit{Prin} \to \mathtt{U})$$
$$(\mathit{acts\_for}{:}(x{:}\mathit{Prin}) \to (y{:}\mathit{Prin}) \to \mathit{ActsFor}\ x\ y)$$

AIR **classes.** A class consists of a class identifier *id* and a principal *P* that owns the class. The type constructors $(\mathit{Id}{::}\mathtt{U}), (\mathit{Class}{::}\mathtt{U})$ are used to type identifiers and classes. Classes are constructed using the data constructor $(\mathsf{Class}{:}\mathit{Id} \Rightarrow \mathit{Prin} \Rightarrow \mathit{Class})$. The translation of an AIR class introduces nullary data constructors like $\mathsf{US\_Army\_Confidential}{:}\mathit{Id}$ and $\mathsf{US\_Army}{:}\mathit{Prin}$, from which we can construct the class $\mathsf{USAC} = \mathsf{Class}\ \{\mathsf{US\_Army\_Confidential}\}\ \{\mathit{US\_Army}\}$. Finally, we use a dependent-type constructor and run-time check to generate evidence that two classes are equal.

$$(\mathit{IsClass}{::}\mathit{Class} \to \mathit{Class} \to \mathtt{U}),$$
$$(\mathit{is\_class}{:}(x{:}\mathit{Class}) \to (y{:}\mathit{Class}) \to \mathit{IsClass}\ x\ y)$$

Class instances. Instances are typed using the *Instance*::U type constructor. Each instance must identify the class it instantiates and the current state of its automaton. For each state in a class declaration, we generate a data constructor in the signature that constructs an *Instance* from a *Class* and any state-specific arguments. For example, we have:

$$\mathsf{Init}{:}\mathit{Class} \Rightarrow \mathit{Instance}, \mathsf{Debt}{:}\mathit{Class} \Rightarrow \mathit{Int} \Rightarrow \mathit{Instance}$$

Thus the expression new $\mathsf{Init}\ \{\mathsf{USAC}\}$ constructs a new instance of a class. According to (T-NEW), this expression has the affine type $¡\mathit{Instance}^{\alpha}$, where the unique type-level name α allows us to protect some data with this automaton. Since we wish to allow data to be protected by automata that instantiate arbitrary AIR classes, we give all instances, regardless of their class, a type like $¡\mathit{Instance}^{\alpha}$, for some α. This has the benefit of flexibility—we can easily give types to library functions that can return data (like file system objects) protected by automata of different classes. However, we must rely on a run-time check to examine the class of an instance since it is not evident from the type.

The prelude includes the the following two elements to construct and type evidence about the class of an automaton instance:

$$\mathit{ClassOf}{::}\mathtt{N} \to \mathit{Class} \to \mathtt{U}$$
$$\mathit{class\_of\_inst}{:}\forall\alpha{::}\mathtt{N}.(x{:}¡\mathit{Instance}^{\alpha}) \to (¡\mathit{Instance}^{\alpha} * c{:}\mathit{Class} * \mathit{ClassOf}\ \alpha\ c)$$

The function *class_of_inst* extracts a *Class* value c from an instance named α and produces evidence (of type *ClassOf* α c) that α is an instance of c. The return type of this function is interesting for two reasons. First, because the returned value

relates the class object in the second component of the tuple to the evidence object in the third component, we give the returned value the type of a *dependently typed tuple*, (designated by the symbol $*$). Although we do not directly support these tuples, they can be easily encoded using dependently typed functions [Swamy et al., 2008]. Second, notice that even though *class_of_inst* does not cause a state transition, the first component of the tuple it returns contains an automaton instance with the same type as the argument x. This is a common idiom when programming with affine types; since the automaton instance is affine and can only be used once, functions like *class_of_inst* simply return the affine argument x back to the caller for further use.

The following constructs in the prelude allow a program to inspect the current state of an automaton instance.

$$\begin{array}{l} \mathit{InState}{::}¡\mathit{Instance}^{\circ} \to \mathit{Instance} \to \mathtt{U} \\ \mathit{state\_of\_inst}{:}\forall\alpha{::}\mathtt{N}.(x{:}¡\mathit{Instance}^{\alpha}) \to \\ \qquad\qquad (z{:}¡\mathit{Instance}^{\alpha} * y{:}\mathit{Instance} * \mathit{InState}\ z\ y) \end{array}$$

These constructs are similar to the forms shown for examining the class of an instance, but with one important difference. Since the state of an automaton is transient (it can change as transition rules are applied), we must be careful when producing evidence about the current state. This is in contrast to the class of an automaton which never changes despite changes to the current state. Thus, we must ensure that stale evidence about an old state of the automaton can never be presented as valid evidence about the current state.

The distinction between evidence about the class of an automaton and evidence about its current state is highlighted by the first argument to the type constructor *InState*. Unlike the first argument of the *ClassOf* constructor (which can be some type-level name $\alpha{::}N$), the first argument of *InState* is an *expression* with an affine type $¡\mathit{Instance}^{\circ}$ (introduced via subsumption in (T-DROP)) that stands for an automaton instance that has been assigned some name. Using this form of subtyping allows us to use *InState* to type evidence about the current state of any automaton. An alternative would be to enhance the kind language by allowing type constructors to be have polymorphic kinds—we chose this form of subtyping to keep the presentation simpler.

As described further in the next subsection, functions that correspond to AIR rules take an automaton instance a_1 (say, in state Init) as an argument, and produce a new instance a_1' as a result (say, in state $\mathsf{Debt}(0)$). Importantly, both a_1 and a_1' are given the type $¡\mathit{Instance}^{\alpha}$—i.e., the association between the type-level name α and the automaton instance is fixed and is invariant with respect to state transitions. Since the class of an automaton never changes (both a_1 and a_1' are instances of USAC) it is safe to give evidence about the class of an instance the type *ClassOf* α USAC—i.e., evidence about the class of an automaton can never become stale. On the other hand, evidence about the current state of the automaton can become stale. If we were to type this evidence using types of the form *InStateBad* α Init, then this evidence may be true of a_1 but it is not true of a_1'. Therefore, we make *InState* a dependent-type constructor to be applied to an automaton instance rather than a type-level name.

5.2 Translating Rules in an AIR Class

Our technical report defines a translation procedure from an AIR class to a λ_{AIR} signature. Space constraints preclude a presentation of the translation judgment here. Instead, we discuss the signature corresponding to the policy of Figure 3.

Release rules. Each release rule r in a class declaration is translated to a function-typed constant f_r in the signature. At a high-level, the rules have the following form. In response to a request to release data x, protected by instance a_1, to an instance a_2, the programmer must provide evidence for each of the conditions in the rule r. If such evidence can be produced, then f_r returns a new automaton state a_1', downgrades x as specified in the policy and returns x under the protection of a_2. As an example, consider the full type of the Conf_coalition rule shown below.

Conf_coalition :

$$\begin{array}{l} 1\ \ \forall \mathit{src}{::}\mathtt{N}, \mathit{dst}{::}\mathtt{N}, \alpha{::}\mathtt{U}. \\ 2\ \ (a1{:}¡\mathit{Instance}^{\mathit{src}}) \to ¡(x{:}\mathit{Protected}\ \alpha\ \mathit{src}) \to ¡(a2{:}¡\mathit{Instance}^{\mathit{dst}}) \to \\ 3\ \ ¡(e1{:}\mathit{ClassOf}\ \mathit{src}\ \mathsf{USAC}) \to ¡(\mathit{cd}{:}\mathit{Class}) \to ¡(e2{:}\mathit{ClassOf}\ \mathit{dst}\ \mathit{cd}) \to \\ 4\ \ ¡(e3{:}\mathit{ActsFor}\ (\mathit{principal}\ \mathit{cd})\ \mathit{Coalition}) \to ¡(\mathit{debt}{:}\mathit{Int}) \to \\ 5\ \ ¡(e4{:}\mathit{InState}\ a1\ (\mathsf{Debt}\ \{\mathsf{USAC}\}\ \{\mathit{debt}\})) \to ¡(e5{:}\mathit{LEQ}\ \mathit{debt}\ 10) \to \\ 6\ \ \qquad (¡\mathit{Instance}^{\mathit{src}} \times ¡\mathit{Instance}^{\mathit{dst}} \times \mathit{Protected}\ \alpha\ \mathit{dst}) \end{array}$$

The first two lines of this type were shown previously— x is the data to be released from the protection of automaton $a1$ (with type-level name *src*) to the automaton $a2$ (with type-level name *dst*). Since the argument $a1$ is affine, we require every

function type to the right of $a1$ to also be affine, since they represent closures that capture the affine value $a1$. At line 3, the argument $e1$ is evidence that shows that the source automaton is an instance of the USAC class; cd is another class object and $e2$ is evidence that the class of the destination automaton is indeed cd. At line 4, $e3$ stands for evidence of the first condition expression, which requires that the owning principal of the destination automaton acts for the *Coalition* principal. Line 5 contains evidence $e4$ that $a1$ is in some state Debt($debt$), where, from $e5$, $debt \leq 10$. The return type, as discussed before, contains the new state of the source automaton, the destination automaton $a2$ threaded through from the argument, and the data value x, downgraded according to the policy and with a type showing that it is protected by the dst automaton.

Transition rules. Each transition rule r in a class declaration is also translated to a function-typed constant f_r in the signature. However, instead of downgrading and coercing the type of some datum x, a transition function only returns the new state of the source automaton and an unchanged destination automaton. That is, instead of returning a three-tuple like Conf_coalition, a transition rule like Conf_init returns a pair $(¡Instance^{src} \times ¡Instance^{dst})$, where the first component is the new state of the source automaton and the second component is the unchanged destination automaton threaded through from the argument. The full type of Conf_init is shown below.

Conf_init :
1 $\forall src{::}\mathtt{N}, dst{::}\mathtt{N}, \alpha{::}\mathtt{U}.$
2 $(a1{:}¡Instance^{src}) \rightarrow ¡(x{:}Protected\ \alpha\ src) \rightarrow ¡(a2{:}¡Instance^{dst}) \rightarrow$
3 $¡(e1{:}ClassOf\ src\ \mathsf{USAC}) \rightarrow ¡(cd{:}Class) \rightarrow ¡(e2{:}ClassOf\ dst\ cd) \rightarrow$
4 $¡(e4{:}InState\ a1\ \mathsf{Init}) \rightarrow (¡Instance^{src} \times ¡Instance^{dst})$

5.3 Programming with the AIR API

The following example program, a revision of the program in Figure 4, illustrates how a client program interacts with the API generated for an AIR policy.

```
1  let x_a1, a1:¡Instance^src = get_usac_file_and_policy () in
2  let a2:¡Instance^dst, channel = get_request () in
3  let a1,USAC,ca1_ev = class_of_inst [src] a1 in
4  let a2,ca2,ca2_ev = class_of_inst [dst] a2 in
5  let actsfor_ev = acts_for (principal ca2) Coalition in
6  let a1, Debt{USAC}{debt}, a1_state_ev = state_of_inst [src] a1 in
7  let debt_ev = leq debt 10 in
8  let a1',a2,x_a2 = Conf_coalition [src][dst][Int] a1 x_a1 a2
9                               ca1_ev ca2 ca2_ev actsfor_ev
10                              debt a1_state_ev debt_ev in
11   send [Int] [dst] channel x_a2
```

As previously, the first two lines represent boilerplate code, where we read a file and its automaton policy and then block waiting for a release request. At line 3, we generate evidence a1_class_ev that a1 is an instance of the USAC class and at line 4 we retrieve a2's class ca2 and evidence ca2_ev that witnesses the relationship between ca2 and a2. At line 5, we check that the destination automaton is owned by a principal acting for the Coalition. At lines 6 and 7 we check that a1 is in the state Debt{USAC}{debt}, for some value of debt ≤ 10. If all the run-time checks succeed (i.e., calls to functions like leq), then we call Conf_coalition, instantiating the type variables, passing in the automata, the data to be downgraded and evidence for all the release conditions. We get back the new state of the *src* automaton a1', a2 is unchanged, and x_a2 which has type *Protected Int dst*. We can give the channel a type such as *Channel Int dst*, indicating that it can be used to send integers to the principal that owns the automaton *dst*. The send function can be given the type shown below:

$$\text{send}{:}\forall\alpha{::}\mathtt{U}, \beta{::}\mathtt{N}.Channel\ \alpha\ \beta \rightarrow Protected\ \alpha\ \beta \rightarrow Unit$$

This ensures that x_a1 cannot be sent on the channel. But, if the call to Conf_coalition succeeds, then the downgraded x_a2 has type *Protected Int dst*, which allows it to be sent.

5.4 Correctness of Policy Enforcement

In this section, we present a condensed version of our main security theorem and discuss its implications. The full statement and proof can be found in our technical report.

Theorem (Security)**.** *Given all of the following: (1) an* AIR *declaration D of a class with identifier* C *owned by principal P, and its translation to a signature S_D; (2) a model M_D consistent with S_D; (3) $\Gamma = src{::}\mathtt{N}, dst{::}\mathtt{N}, s : ¡Instance^{src}$; (4) $\Gamma; s \vdash_{S_D,\mathsf{term}} e : t; \varepsilon$ where $src \notin \varepsilon$; and (5) $M \vdash ((s \mapsto v)e) \xrightarrow{l_1} e_1 \ldots \xrightarrow{l_n} e_n$ where $v = \mathsf{new}\ \mathsf{Init}\ \{\mathsf{Class}\ \{\mathsf{C}\}\ \{P\}\}$. Then the string $l_1, \ldots, l_n$ is accepted by the automaton defined by D.*

The first condition relies on our translation judgment that produces a signature S_D from a class declaration D. The second condition is necessary for type soundness. Conditions (3) and (4) state that e is a well-typed expression in a context with a single free automaton $s : ¡Instance^{src}$ and two type name constants *src* and *dst*. By requiring that $src \notin \varepsilon$ we ensure that e does not give the name *src* to any other automaton instance. This theorem asserts that when e is reduced in a context where s is bound to an instance of the C class in the Init state, then the trace $l_1, \ldots, l_n$ of the reduction sequence is a word in the language accepted by the automaton of D.

The trace acceptance judgment has the form $A; D \models l_1, \ldots, l_n; A'$, which informally states that an automaton defined by the class D, in initial state A, accepts the trace $l_1, \ldots, l_n$ and transitions to the state A'. Recall that the trace elements l_i record base terms B that stand for security-relevant actions and sets of values that certify that the action is permissible. The trace acceptance judgment allows a transition from A to A' only if each transition is justified by all the evidence required by the rules in the class. This condition is similar to the one used by Walker [2000].

The security theorem as presented here is a general purpose result that applies to the enforcement of all AIR policies in λAIR. One of our near-term goals is to integrate our paper proof of this security theorem with our Coq proof of type-soundness. For the longer term, we aim to investigate policy-specific security properties and to rely on our infrastructure of λAIR metatheory in Coq to partially automate the proofs of these properties as well.

6 Related Work

The specification and enforcement of policies that control information release has received much recent attention. Sabelfeld and Sands [2005] survey many of these efforts and provide a useful way of organizing the various approaches. AIR policies address, to varying degrees, the *what*, *who*, *where* and *when* of declassification, the four dimensions identified by Sabelfeld and Sands. Most of this work approaches information release from the perspective of information flow policies [Denning, 1976], and most of the proposed security properties can be thought of as bisimulations. By contrast, our security theorem states that the program's actions are in accord with a high-level policy, not that these actions enforce an extensional security property (like noninterference). We believe that the two approaches are complementary. In combination, we could show a noninterference-like security theorem (e.g., noninterference until conditions [Chong and Myers, 2004], or robust declassification [Zdancewic and Myers, 2001]) while being able to reason that a high-level protocol for releasing information is correctly followed.

AIR policies are defined separately from programs that use them, allowing them to be reasoned about in isolation. Most related work embeds declassification policies within programs that use them, obscuring high-level intent. One exception is work on *trusted declassifiers* [Hicks et al., 2006]. Here, all possible information flows are specified as part of a graph in which nodes consist of either downgrading functions or principals, and edges consists of trust relationships. Paths through the graph indicate how data may be released. AIR classes generalize this approach in restricting which paths may occur in the graph, and in specifying release conditions in addition to downgrading functions.

Chong and Myers [2004] propose declassification policies as labels consisting of sequences of atomic labels separated by conditions c. Initially, labeled data may be viewed with the privileges granted by the first atomic label, but when a condition c is satisfied, the data may be relabeled to the next label in the sequence, and viewed at its privileges. Declassification labels are thus similar to AIR classes, with the main difference that our approach is more geared toward run-time checking: we support dynamically-checked conditions (theirs must be provable statically) and run-time labels (theirs are static annotations).

Security automata were first proposed by Schneider [2000] as a means of specifying and enforcing safety properties. AIR policies are actually a more general form of security automata called edit automata [Ligatti et al., 2003] because they may modify data before releasing it. To our knowledge, no prior work has used automata to specify the protection level and release conditions of sensitive data. Walker [2000] defines a type-based approach for enforcing security automata policies in which the definition of a single automaton is embedded in the type-checking judgment. Our approach allows multiple automata policies to be easily defined separately. Automata policies have also been enforced using inlined reference monitors, as in SASI and PSLang/PoET [Erlingsson, 2004]. Our approach is in contrast with SASI in that we support local policy state—Erlingsson identifies SASI's global policy restriction as a main obstacle towards making it practical. PSLang/PoET does support local policy state, but unlike λAIR, PSLang/PoET augments the run-time representation of protected data to include the policy. Dynamic labels in λAIR are more expressive (as discussed in Section 3, we can easily enforce secret sharing

policies on related data) and provide a way to verify that automata and protected data are always correctly manipulated. As such, one could imagine putting λAIR to use to certify that IRMs correctly enforce their policies.

There has also been much work on tracking the state of objects in types, dating back to Strom and Yemini [1986]. The calculus of capabilities [Crary et al., 1999] provides a way of tracking typestate, using singleton and linear types (a variant of affine types) to account for aliasing. The Vault [DeLine and Fähndrich, 2001] and Cyclone [Jim et al., 2002] programming languages implement typestate checkers in a practical setting to enforce proper API usage and correct manual memory management, respectively. λAIR's use of singleton and affine types is quite close to these systems. However, in these systems the state of a resource is a static type annotation, while in λAIR a policy automaton is first-class, allowing its state to be unknown until run time. Additionally, λAIR's use of dependent types permits more precise specifications, which is useful for certifying authorization decisions.

Certified evaluation of authorization decisions has been explored in a number of contexts. For instance, certified evaluation is a feature of the SD3 trust-management system proposed by Jim [2001]. More recently, the Aura system of Vaughan et al. [2008] proposes maintaining audit logs to record evidence that justifies authorization decisions made during the system's execution. The architecture we propose for certified evaluation in λAIR is closely related to both these approaches. While more investigation is required, λAIR's ability to accurately track evidence in the presence of state modifications opens the possibility of certified evaluation of a wider class of stateful authorization policies, like those expressible in SMP, a stateful authorization logic recently proposed by Becker and Nanz [2007].

7 Conclusions

This paper has presented AIR, a simple policy language for expressing stateful information release policies. We have defined a core formalism for a programming language called λAIR, in which stateful authorization policies like AIR can be certifiably enforced. In future work, we plan to add support for λAIR-style policy enforcement to our secure web-programming language, SELINKS [Swamy et al., 2008].

Acknowledgements: We thank Jeff Foster, Elnatan Reisner, and the anonymous reviewers for helpful comments on a draft of this paper. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U. S. Government.

References

Brian Aydemir, Arthur Charguéraud, Benjamin C. Pierce, Randy Pollack, and Stephanie Weirich. Engineering formal metatheory. In *POPL '08: Proceedings of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM, 2008.

Anindya Banerjee, David A. Naumann, and Stan Rosenberg. Expressive declassification policies and modular static enforcement. *IEEE Symposium on Security and Privacy*, 2008.

Moritz Y. Becker and Sebastian Nanz. A logic for state-modifying authorization policies. In Joachim Biskup and Javier Lopez, editors, *ESORICS*, volume 4734 of *Lecture Notes in Computer Science*. Springer, 2007.

Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development. Coq'Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. Springer Verlag, 2004.

Pau-Chen Cheng, Pankaj Rohatgi, Claudia Keser, Paul A. Karger, Grant M. Wagner, and Angela Schuett Reninger. Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In *IEEE Symposium on Security and Privacy*, 2007.

Stephen Chong and Andrew C. Myers. Security policies for downgrading. In *CCS '04: Proceedings of the 11th ACM conference on Computer and communications security*, pages 198–209. ACM, 2004.

Stephen Chong, Andrew C. Myers, Nathaniel Nystrom, Lantian Zheng, and Steve Zdancewic. Jif: Java + information flow. Software release, July 2006.

Karl Crary, David Walker, and Greg Morrisett. Typed memory management in a calculus of capabilities. In *POPL '99: Proceedings of the 26th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 262–275. ACM, 1999.

Robert DeLine and Manuel Fähndrich. Enforcing high-level protocols in low-level software. In *PLDI '01: Proceedings of the ACM SIGPLAN 2001 conference on Programming language design and implementation*, pages 59–69. ACM, 2001.

Dorothy E. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5):236–243, May 1976.

Ulfar Erlingsson. *The inlined reference monitor approach to security policy enforcement*. PhD thesis, 2004. Cornell University.

Michael W. Focke, James E. Knoke, Paul A. Barbieri, Robert D. Wherley, John G. Ata, and Dwight B. Engen. Trusted computing system. United States Patent No. 7,103,914, 2006. issued to BAE Systems Information Technology LLC.

Boniface Hicks, Dave King, Patrick McDaniel, and Michael Hicks. Trusted declassification:: high-level policy for a security-typed language. In *PLAS '06: Proceedings of the 2006 workshop on Programming languages and analysis for security*, pages 65–74. ACM, 2006.

Boniface Hicks, Tim Misiak, and Patrick McDaniel. Channels: Runtime system infrastructure for security-typed languages. *Computer Security Applications Conference, 2007. ACSAC 2007. Twenty-Third Annual*, 2007.

Trevor Jim. SD3: A trust management system with certified evaluation. In *SP '01: Proceedings of the 2001 IEEE Symposium on Security and Privacy*, page 106. IEEE Computer Society, 2001.

Trevor Jim, J. Greg Morrisett, Dan Grossman, Michael W. Hicks, James Cheney, and Yanling Wang. Cyclone: A safe dialect of c. In *Proceedings of the General Track of the USENIX Annual Technical Conference*. USENIX Association, 2002.

Jay Ligatti, Lujo Bauer, and David Walker. Edit automata: Enforcement mechanisms for run-time security policies. *International Journal of Information Security*, 2003.

John C. Mitchell. *Foundations of Programming Languages*. MIT Press, 1996.

Polyvios Pratikakis, Jeffrey S. Foster, and Michael Hicks. Context-sensitive correlation analysis for detecting races. In *Proceedings of the ACM Conference on Programming Language Design and Implementation (PLDI)*, pages 320–331, 2006.

Andrei Sabelfeld and David Sands. Dimensions and principles of declassification. In *IEEE Computer Security Foundations Workshop (CSFW)*, 2005.

Fred B. Schneider. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.*, 3(1):30–50, 2000.

Robert E. Strom and Shaula Yemini. Typestate: A programming language concept for enhancing software reliability. *IEEE Trans. Softw. Eng.*, 12(1), 1986.

Nikhil Swamy and Michael Hicks. Fable: A language for enforcing user-defined security policies (extended version). University of Maryland, Technical Report, CS-TR-4876; `http://www.cs.umd.edu/projects/PL/fable/TR.pdf`, 2007.

Nikhil Swamy and Michael Hicks. Verified enforcement of automaton-based information release policies, 2008. CS-TR-4906, CS Dept., U. Maryland.

Nikhil Swamy, Brian J. Corcoran, and Michael Hicks. Fable: A language for enforcing user-defined security policies. In *IEEE Symposium on Security and Privacy*, 2008.

United States Department of Defense. Department of defense directive number 5230.11, 1992.

Jeffrey A. Vaughan, Limin Jia, Karl Mazurak, and Steve Zdancewic. Evidence-based audit. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium*, Pittsburgh, PA, USA, June 2008.

David Walker. A type system for expressive security policies. In *ACM Symposium on Principles of Programming Languages*, 2000.

Steve Zdancewic and Andrew C. Myers. Robust declassification. In *IEEE Computer Security Foundations Workshop (CSFW)*, 2001.

Lantian Zheng and Andrew C. Myers. Dynamic security labels and noninterference. In *Proceedings of the Workshop on Formal Aspects in Security and Trust (FAST)*, 2004.

A Soundness of λ_{AIR}

Note: This technical report does not formalize refining types based on the results of pattern matching. We have formalized this separately elsewhere [Swamy and Hicks, 2007].

Definition 1 (Well-formed environment)**.** $\Gamma; A$ *is well-formed with respect to a signature S if and only if*

$$\begin{array}{lll} (i.) & \textit{All names bound in } \Gamma \textit{ are distinct} & \\ (ii.) & \Gamma = \Gamma_1, x:t, \Gamma_2 & \Rightarrow \quad FV(t) \subseteq dom(\Gamma_1) \\ (iii.) & A = A_1, x, A_2 & \Rightarrow \quad \Gamma = \Gamma_1, x:t, \Gamma_2 \end{array}$$

Definition 2 (Type consistency of a signature)**.** *A signature S and its model M are type consistent if and only if for each* $\mathscr{D} \rightsquigarrow e \in \overrightarrow{E}$*, where* $B : \overrightarrow{E} \in M$*, we have for* $\Gamma = \alpha_1 :: N, \ldots, \alpha_n :: N$*:*

1. $\Gamma; \cdot \vdash [\![B]\!]^{\mathscr{D}} : t; \varepsilon \iff \Gamma; \cdot \vdash e : t; \varepsilon$
2. *For every* $B \in dom(S)$*, if B is not a data constructor* B*, then* $B : \overrightarrow{E} \in M$*.*
3. *For every* $(T :: K) \in S$*,* $\vdash K\ ok$ *where*

$$\frac{}{\vdash k\ ok} \qquad \frac{\vdash K\ ok}{\vdash k \to K\ ok} \qquad \frac{\cdot \vdash t :: k \qquad \vdash K\ ok}{\vdash t \to K\ ok}$$

4. *For every* $(B : t) \in S$*,* $\cdot \vdash t :: k$*.*
5. $\Gamma; \cdot \vdash_{\varphi} [\![B]\!]^{\mathscr{D}} : t; \varepsilon \ \wedge\ \varepsilon \neq \cdot \ \Rightarrow\ (\Gamma \vdash t :: \mathtt{A} \ \vee\ \Gamma; \cdot \vdash_{\varphi} [\![B]\!]^{\mathscr{D}} : t; \cdot)$

Theorem 3 (Progress)**.** *Given* $\Gamma = \alpha_1 :: N, \ldots, \alpha_n :: N$ *such that (A1)* $\Gamma; \cdot \vdash_{\mathsf{term}} e : t$*; and (A2) given an interpretation M such that M and S are type consistent, then either* $\exists e'.M \vdash e \xrightarrow{l} e'$ *or* $\exists v.e = v$*.*

Proof. By induction on the structure of (A1).

Case (T-B): If B is a data constructor B, then it is a value and we are done. Otherwise, from Definition 2, we can satisfy the premise of (E-DEL) and take a step to $[\![B]\!]$.

Case (T-X), (T-XA): By assumption, Γ only contains type names α_i; i.e., e is a closed term. So, these cases are impossible.

Case (T-NEW): If in $\mathsf{new}\ e$, e is an expression, then by the induction hypothesis on the first premise we have $M \vdash e \xrightarrow{l} e'$. Then, by the syntactic form of the evaluation contexts $\mathscr{E}$ and by the congruence rule (E-CTX) we have the result. On the other hand, if e is a value v then $\mathsf{new}\ v$ is also a value.

Case (T-TAB): $\Lambda\alpha :: k.e$ is a value.

Case (T-TAP):

$$\frac{\Gamma; A \vdash_{\varphi} e : q(\forall \alpha :: k \xrightarrow{\varepsilon'} t'); \varepsilon \qquad \Gamma \vdash t :: k}{\Gamma; A \vdash_{\varphi} e\ [t] : [\alpha \mapsto t]t'; \varepsilon \uplus ([\alpha \mapsto t]\varepsilon')}\ \text{(T-TAP)}$$

If e is not a value, then by the induction hypothesis on the first premise, we have $M \vdash e \xrightarrow{l} e'$, and so, by (E-CTX) we have $M \vdash e\ [t] \xrightarrow{l} e'\ [t]$.

If e is a value, then by canonical forms of values of universally quantified type (applied to the first premise), we have two sub-cases:

Sub-case $e = \Lambda\alpha :: k.e$: In this case, (E-TAP) is applicable and we step to $(\alpha \mapsto t)e$.

Sub-case $e = [\![B]\!]^{\mathscr{D}}$: In this case, either (E-B3) or (E-B4) is applicable. We first need to show that $B : \overrightarrow{E} \in M$. But, this follows from (A1) which requires $[\![B]\!]^V$ to be well-typed. Thus, we have that $B \in \text{dom}(S)$ and by the assumption of type-consistency of M and S, (A2), we have that $B : \overrightarrow{E} \in M$. The premises of (E-B3) and (E-B4) are mutually exclusive and total. So, a step using one or the other must be possible.

Case (T-BOT): Step using (E-INF).

Case (T-CAP):

Models, certificates, values, and evalutation contexts

equation	E	::=	$\mathscr{D} \rightsquigarrow e$	certs	e	::=	$\ldots \mid [\![B]\!]^{\mathscr{D}}$
eqn. domain	$\mathscr{D}$	::=	$v \mid t \mid \mathscr{D}, \mathscr{D} \mid \cdot$	values	v	::=	$\mathsf{B} \mid [\![B]\!]^{\mathscr{D}} \mid \lambda x{:}t.e \mid \Lambda\alpha{::}k.e \mid v\ \{v'\} \mid \mathsf{new}\ v$
model	M	::=	$B : \overrightarrow{E} \mid M, M$	eval ctxt	$\mathscr{E}$	::=	$\bullet \mid \bullet\ e \mid v\ \bullet \mid \bullet\ [t] \mid \bullet\ \{e\} \mid v\ \{\bullet\} \mid \mathsf{case}\ \bullet\ \mathsf{of} \ldots \mid \mathsf{new}\ \bullet$

$\Gamma; A \vdash_\varphi e : t; \varepsilon$ — A φ-level expression e in environment Γ with affine assumptions A has type t and uses names ε

$$\frac{S(B) = t}{\Gamma; \cdot \vdash_\varphi B : t; \cdot}\ \text{(T-B)} \qquad \frac{\Gamma \vdash \Gamma(x) :: U}{\Gamma; \cdot \vdash_\varphi x : \Gamma(x); \cdot}\ \text{(T-X)} \qquad \frac{}{\Gamma; x \vdash_\varphi x : \Gamma(x); \cdot}\ \text{(T-XA)} \qquad \frac{}{\Gamma; \cdot \vdash_{\mathsf{type}} x : \Gamma(x); \cdot}\ \text{(T-X-type)}$$

$$\frac{\Gamma; A \vdash_{\mathsf{type}} e : t; \varepsilon_1 \uplus \varepsilon}{\Gamma; A \vdash_{\mathsf{type}} e : t; \varepsilon}\ \text{(T-NC-type)} \qquad \frac{\Gamma \vdash t :: k \qquad k \neq N}{\Gamma; \cdot \vdash_\varphi \bot : t; \cdot}\ \text{(T-BOT)}$$

$$\frac{\Gamma; A \vdash_\varphi e : t; \varepsilon \qquad \Gamma \vdash t :: U \qquad \Gamma(\alpha) = \mathtt{N}}{\Gamma; A \vdash_\varphi \mathsf{new}\ e : ¡t^\alpha; \alpha \uplus \varepsilon}\ \text{(T-NEW)} \qquad \frac{\Gamma; A \vdash_\varphi e : t^\alpha; \varepsilon}{\Gamma; A \vdash_\varphi e : t^\circ; \varepsilon}\ \text{(T-DROP)} \qquad \frac{\Gamma; A \vdash_\varphi e : t; \varepsilon \qquad \varepsilon' \subseteq \mathrm{dom}(\Gamma)}{\Gamma; A, A' \vdash_\varphi e : t; \varepsilon \uplus \varepsilon'}\ \text{(T-WKN)}$$

$$\frac{\begin{array}{c}\Gamma, \alpha{::}k; A \vdash_\varphi e : t; \varepsilon \uplus \varepsilon' \\ \alpha \notin \varepsilon \qquad \varepsilon' \in \{\cdot, \alpha\} \qquad q = p(A, \varepsilon)\end{array}}{\Gamma; A \vdash_\varphi \Lambda\alpha{::}k.e : q(\forall \alpha{::}k \xrightarrow{\varepsilon'} t); \varepsilon}\ \text{(T-TAB)} \qquad \frac{\begin{array}{c}\Gamma \vdash t_x :: k \qquad q = p(A, \varepsilon) \\ \Gamma, x : t_x; A, a(x,k) \vdash_\varphi e : t_e; \varepsilon\end{array}}{\Gamma; A \vdash_\varphi \lambda x{:}t_x.e : q((x{:}t_x) \to t_e); \varepsilon}\ \text{(T-ABS)}$$

where
$a(x, \mathtt{A}) = x \qquad a(x, \mathtt{U}) = \cdot$
$p(A, \varepsilon) = ¡ \qquad p(\cdot, \cdot) = \cdot$

$$\frac{\Gamma; A \vdash_\varphi e : q(\forall \alpha{::}k \xrightarrow{\varepsilon'} t'); \varepsilon \qquad \Gamma \vdash t :: k}{\Gamma; A \vdash_\varphi e\ [t] : [\alpha \mapsto t]t'; \varepsilon \uplus ([\alpha \mapsto t]\varepsilon')}\ \text{(T-TAP)} \qquad \frac{\Gamma; A \vdash_\varphi e : q((x{:}t') \to t); \varepsilon_1 \qquad \Gamma; A' \vdash_\varphi e' : t'; \varepsilon_2}{\Gamma; A, A' \vdash_\varphi e\ e' : [x \mapsto e']t; \varepsilon_1 \uplus \varepsilon_2}\ \text{(T-APP)}$$

$$\frac{\begin{array}{c}\Gamma; A \vdash_\varphi e : q(t' \Rightarrow t); \varepsilon \\ \Gamma; A' \vdash_\varphi e' : t'; \varepsilon'\end{array}}{\Gamma; A, A' \vdash_\varphi e\ \{e'\} : t; \varepsilon \uplus \varepsilon'}\ \text{(T-CAP)} \qquad \frac{\begin{array}{ccc}\Gamma; A \vdash_\varphi e : t_e; \varepsilon & \cdot; \cdot \vdash t_i :: k_i & \forall i.a(x_i, k_i) \in A'' \\ \overrightarrow{x : t}; A'' \vdash e_{pat} : t_e; \cdot & \Gamma, \overrightarrow{x : t}; A', A'' \vdash_\varphi e' : t; \varepsilon' & \Gamma; A' \vdash_\varphi e'' : t; \varepsilon''\end{array}}{\Gamma; A, A' \vdash_\varphi \mathsf{case}\ e\ \mathsf{of}\ \overrightarrow{x{:}t}.e_{pat} : e'\ \mathsf{else}\ e'' : t; \varepsilon \uplus (\varepsilon' \cup \varepsilon'')}\ \text{(T-CASE)}$$

$$\frac{\Gamma; A \vdash_\varphi [\![B]\!]^{\mathscr{D}}\ v : t; \varepsilon}{\Gamma; A \vdash_\varphi [\![B]\!]^{\mathscr{D}, v} : t; \varepsilon}\ \text{(T-B1)} \qquad \frac{\Gamma; A \vdash_\varphi [\![B]\!]^{\mathscr{D}}\ [t] : t'; \varepsilon}{\Gamma; A \vdash_\varphi [\![B]\!]^{\mathscr{D}, t} : t'; \varepsilon}\ \text{(T-B2)} \qquad \frac{\Gamma; A \vdash_\varphi B : t; \varepsilon}{\Gamma; A \vdash_\varphi [\![B]\!] : t; \varepsilon}\ \text{(T-B3)} \qquad \frac{\Gamma; A \vdash_\varphi e : t; \varepsilon \qquad t \cong t'}{\Gamma; A \vdash_\varphi e : t'; \varepsilon}\ \text{(T-CONV)}$$

$t \cong t'$ — Congruence of types t and t' under reduction of type-level expressions.

Type contexts $T ::= \bullet \mid x{:}\bullet \to t \mid x{:}t \xrightarrow{\varepsilon} \bullet \mid \forall \alpha{::}k \to \bullet \mid \bullet \Rightarrow t \mid t \Rightarrow \bullet \mid q\ \bullet \mid \bullet\ t \mid t\ \bullet \mid \bullet^\eta$

$$t \cong t\ \text{(TE-ID)} \qquad \frac{t \cong t'}{t' \cong t}\ \text{(TE-SYM)} \qquad \frac{t \cong t'}{T \cdot t \cong T \cdot t'}\ \text{(TE-CTX)} \qquad \frac{\cdot; \cdot \vdash_{\mathsf{type}} e : t, \varepsilon \qquad M \vdash e \xrightarrow{l} e' \qquad \cdot; \cdot \vdash_{\mathsf{type}} e' : t, \varepsilon}{T \cdot e \cong T \cdot e'}\ \text{(TE-RED)}$$

$\Gamma \vdash t :: K$ — A type t has kind K in environment Γ

$$\frac{\Gamma(\alpha) = k}{\Gamma \vdash \alpha :: k}\ \text{(K-A)} \qquad \frac{\Gamma \vdash t :: \mathtt{A} \qquad \Gamma(\eta) = \mathtt{N} \vee \eta = \circ}{\Gamma \vdash t^\eta :: \mathtt{A}}\ \text{(K-N)} \qquad \frac{S(T) = K}{\Gamma \vdash T :: K}\ \text{(K-TC)} \qquad \frac{\Gamma \vdash t :: \mathtt{U}}{\Gamma \vdash ¡t :: \mathtt{A}}\ \text{(K-AFN)}$$

$$\frac{\Gamma \vdash t :: k \qquad \Gamma, x : t \vdash t' :: k'}{\Gamma \vdash (x{:}t) \to t' :: \mathtt{U}}\ \text{(K-FUN)} \qquad \frac{\Gamma' = \Gamma, \alpha{::}k \qquad \Gamma' \vdash t :: k \qquad \alpha' \in \varepsilon \Rightarrow \Gamma'(\alpha') = \mathtt{N}}{\Gamma \vdash \forall \alpha{::}k \xrightarrow{\varepsilon} t :: \mathtt{U}}\ \text{(K-UNIV)}$$

$$\frac{\begin{array}{c}\Gamma \vdash t :: t' \to K \\ \Gamma; \cdot \vdash_{\mathsf{type}} e : t'; \cdot\end{array}}{\Gamma \vdash t\ e :: K}\ \text{(K-DEP)} \qquad \frac{\Gamma \vdash t :: k \to K \qquad \Gamma \vdash t' :: k}{\Gamma \vdash t\ t' :: K}\ \text{(K-TAP)} \qquad \frac{\begin{array}{c}\Gamma \vdash t_1 :: U \\ \Gamma \vdash t_2 :: U \qquad t_2 \in \{t \Rightarrow t', T\}\end{array}}{\Gamma \vdash t_1 \Rightarrow t_2 :: U}\ \text{(K-CON)}$$

Figure 7. Static semantics of λ_{AIR}

$M \vdash e \xrightarrow{l} e'$ Given an interpretation M of a signature, an expression e reduces to e' recording l in the trace.

$$\frac{M \vdash e \xrightarrow{l} e' \quad e' \neq \bot}{M \vdash \mathscr{E} \cdot e \xrightarrow{l} \mathscr{E} \cdot e'} \text{(E-CTX)} \qquad \frac{M \vdash e \xrightarrow{l} \bot}{M \vdash \mathscr{E} \cdot e \xrightarrow{l} \bot} \text{(E-BOT)} \qquad \frac{}{M \vdash \bot \longrightarrow \bot} \text{(E-INF)}$$

$$\frac{e' = (x \mapsto v)\, e}{M \vdash \lambda x{:}t.e\ v \longrightarrow e'} \text{(E-APP)} \qquad \frac{e' = (\alpha \mapsto t)e}{M \vdash \Lambda\alpha{::}k.e\ [t] \longrightarrow e'} \text{(E-TAP)}$$

$$\frac{\textbf{if } (v \succ e_{pat} : \sigma) \textbf{ then } e = \sigma(e') \textbf{ else } e = e''}{M \vdash \mathsf{case}\ v\ \mathsf{of}\ \overrightarrow{x{:}t}.e_{pat} : e'\ \mathsf{else}\ e'' \longrightarrow e} \text{(E-CASE)} \qquad \frac{B : \overrightarrow{E} \in M}{M \vdash B \longrightarrow [\![B]\!]} \text{(E-DEL)} \qquad \frac{\mathscr{D}, v \rightsquigarrow e \in \overrightarrow{E} \quad l = B : \mathscr{D}, v}{M, B : \overrightarrow{E}, M' \vdash [\![B]\!]^{\mathscr{D}}\ v \xrightarrow{l} e} \text{(E-B1)}$$

$$\frac{\mathscr{D}, v \rightsquigarrow e \notin \overrightarrow{E}}{M, B : \overrightarrow{E}, M' \vdash [\![B]\!]^{\mathscr{D}}\ v \longrightarrow [\![B]\!]^{\mathscr{D}, v}} \text{(E-B2)} \qquad \frac{\mathscr{D}, t \rightsquigarrow e \in \overrightarrow{E} \quad l = B : \mathscr{D}, t}{M, B : \overrightarrow{E}, M' \vdash [\![B]\!]^{\mathscr{D}}\ [t] \xrightarrow{l} e} \text{(E-B3)} \qquad \frac{\mathscr{D}, t \rightsquigarrow e \notin \overrightarrow{E}}{M, B : \overrightarrow{E}, M' \vdash [\![B]\!]^{\mathscr{D}}\ [t] \longrightarrow [\![B]\!]^{\mathscr{D}, t}} \text{(E-B4)}$$

$v \succ e_p : \sigma$ Pattern matching data constructors.

$$v \succ v : \cdot \text{ (U-ID)} \qquad v \succ x : x \mapsto v \text{ (U-VAR)} \qquad \frac{v \succ e :: \sigma \quad v' \succ \sigma\ e' : \sigma'}{v\ \{v'\} \succ e\ \{e'\} : \sigma, \sigma'} \text{(U-CON)}$$

Figure 8. Dynamic semantics of λ_{AIR}

$$\frac{\Gamma; A \vdash_\varphi e : q(t' \Rightarrow t); \varepsilon \quad \Gamma; A' \vdash_\varphi e' : t'; \varepsilon'}{\Gamma; A, A' \vdash_\varphi e\ \{e'\} : t; \varepsilon \uplus \varepsilon'} \text{(T-CAP)}$$

If e is not a value, by the induction hypothesis in the first premise we have, $M \vdash e_1\ \{e_2\} \longrightarrow e_1'\ \{e_2\}$, using (E-CTX). Similarly, if e_2 is not a value. If both are values, then $v_1\ \{v_2\}$ is a value too.

Case (T-CASE): We have $\mathsf{case}\ e\ \mathsf{of}\ \overrightarrow{x{:}t}.e_{pat} : e'\ \mathsf{else}\ e''$. If e is not a value, then by the penultimate form of $\mathscr{E}$ and the congruence (E-CTX) we can take a step. Otherwise, if e is a value, then we can always take a step using (E-CASE), since its premise is a tautology.

Case (T-ABS): $\lambda x{:}t.e$ is a value.

Case (T-APP):

$$\frac{\Gamma; A \vdash_\varphi e : q((x{:}t') \to t); \varepsilon_1 \quad \Gamma; A' \vdash_\varphi e' : t'; \varepsilon_2}{\Gamma; A, A' \vdash_\varphi e\ e' : [x \mapsto e']t; \varepsilon_1 \uplus \varepsilon_2} \text{(T-APP)}$$

If either e_1 or e_2 are not values, then we can reduce using the (E-CTX) congruence rule. Otherwise, by canonical forms of function-typed values (applied to the first premise), we get that $e = \lambda x{:}t.e$ or $e = [\![B]\!]^{\mathscr{D}}$. In both cases the reasoning proceeds similarly to the (T-TAP) case, except using (E-APP) and (E-B1) or (E-B2).

Case (T-B1), (T-B2), (T-B3): $[\![B]\!]^{\mathscr{D}}$ is a value.

Case (T-DROP), (T-WKN), (T-CONV): Induction hypothesis applied to the hypothesis.

Case (T-X-type), (T-NC-type): Inapplicable.

□

Proposition 4 (Well-formedness of environments)**.** *If $\Gamma; A$ is well-formed, and (A1) $\Gamma; A \vdash e : t; \varepsilon$ contains a premise of the form $\Gamma'; A' \vdash e' : t'; \varepsilon'$ then $\Gamma'; A'$ is well-formed and $\exists k.\Gamma' \vdash t' :: k$. Similarly, if (A1) contains a premise of the form $\Gamma'' \vdash t'' :: K$, then Γ'' is well-formed and $\vdash K$ ok.*

Proposition 5 (Weakening)**.** *If $\Gamma; A$ is well-formed, and $\Gamma; A \vdash e : t; \varepsilon$. Then, for all Γ', A' such that $\Gamma, \Gamma'; A, A'$ is well-formed, $\Gamma, \Gamma'; A, A' \vdash e : t; \varepsilon$. Similarly, if $\Gamma \vdash t :: K$, then $\Gamma, \Gamma' \vdash t :: K$.*

Proposition 6 (Inversion of empty name constraints)**.** *If $\Gamma; \cdot$ is well-formed, and $\Gamma; \cdot \vdash_\varphi v : t; \cdot$. Then, $\Gamma \vdash t ::$ U.*

Proposition 7 (Inversion of non-empty name constraints)**.** *If $\Gamma; \cdot$ is well-formed, and $\Gamma; \cdot \vdash_\varphi v : t; \varepsilon$ and $\varepsilon \neq \cdot$. Then, either $\Gamma \vdash t ::$ A or $\Gamma; \cdot \vdash_\varphi v : t; \cdot$.*

Proposition 8 (Uniqueness of kinding)**.** *If* $\Gamma \vdash t :: k$ *and* $\Gamma \vdash t :: k'$*. Then* $k = k'$*.*

Lemma 9 (Substitution)**.** *Given* $\Gamma;\cdot$ *well-formed, and* $\Gamma';A$ *such that:*

(A1) $\Gamma,\Gamma';A$ *is well-formed.*

(A2) $\Gamma,x:t_x,\Gamma';A_x,A$*, well-formed, where* $\Gamma \vdash t_x :: k$*, and* $A_x = \cdot \vee A_x = a(x,k)$*.*

(A3) $\Gamma,x:t_x,\Gamma';A_x,A \vdash_\varphi e : t;\varepsilon$

(A4) $\Gamma;\cdot \vdash v : t_x;\varepsilon'$

(A5) $\varepsilon \uplus \varepsilon'$

Then, for $\sigma = x \mapsto v$*,*

$$\Gamma,\sigma(\Gamma');A \vdash \sigma(e) : \sigma(t);\varepsilon \uplus \varepsilon'$$

Proof. By mutual induction on the structure of (A3), together with Lemma 10.

Throughout, we will use $\sigma(\Gamma) = \Gamma$, since $x \notin \mathrm{dom}(\Gamma)$

Case (T-B): Base terms B are closed; so using (T-B) we can get

$$\Gamma,\sigma(\Gamma') \vdash \sigma(B) : \sigma(t);\cdot$$

To ensure that the name constraint is still $\varepsilon \uplus \varepsilon' = \varepsilon'$ we conclude with (T-AFN) that allows bound names ε' to be added at will.

Case (T-X):

$$\frac{\Gamma,x:t_x,\Gamma'(y) = t \qquad \Gamma,x:t_x,\Gamma' \vdash t :: U}{\Gamma,x:t_x,\Gamma';\cdot \vdash_\varphi y : t;\cdot}\ \text{(T-X)}$$

We consider two sub-cases, depending on whether or not $y = x$.

Sub-case $y \neq x$: Thus, $\sigma(y) = y$. We have two sub-cases depending on whether y appears in Γ or in Γ'.

Sub-case (i): $y : t \in \Gamma'$. In this case, $FV(t) \cap \mathrm{dom}(\sigma) \neq \emptyset$; thus our conclusion is of the form $\Gamma,\sigma(\Gamma');\cdot \vdash y : \sigma(t);\cdot$, since we have $y : \sigma(t) \in \Gamma'$, to satisfy the first premise, and from Lemma 10, we have $\Gamma,\sigma(\Gamma') \vdash \sigma(t) :: U$ for the second premise. Finally, if $\varepsilon' \neq \cdot$, we conclude with (T-AFN) and introduce the additional effects.

Sub-case (ii): $y : t \in \Gamma$. From our initial remark, we know that $\sigma\Gamma = \Gamma$; thus, $\sigma(t) = t$. Our conclusion is of the form $\Gamma,\sigma(\Gamma');\cdot \vdash y : t,\cdot$, with the first premise satisfied trivially. The second premise follows from the mutual induction hypothesis of Lemma 8 to establish that $\Gamma,\sigma\Gamma' \vdash \sigma t :: \mathtt{U}$. Finally, as previously, we conclude with (T-AFN) for the effects, if necessary.

Sub-case $y = x$: Thus, $\sigma(y) = v$. From (A4) we have $\Gamma;\cdot \vdash v : t_x;\varepsilon'$ and, from weakening, we can establish $\Gamma,\sigma(\Gamma');\cdot \vdash v : t_x;\varepsilon'$. Finally, since $x \notin \mathrm{dom}(\Gamma)$ we can conclude from Proposition 4 that $x \notin FV(t_x)$, and thus $\sigma(t_x) = t_x$.

Case (T-XA):

$$\frac{\Gamma,x:t_x,\Gamma'(y) = t}{\Gamma,x:t_x,\Gamma';y \vdash_\varphi y : t;\cdot}\ \text{(T-XA)}$$

Again, we proceed by cases on whether $x = y$.

Sub-case $y \neq x$: Identical to the same sub-case of (T-X).

Sub-case $y = x$: By weakening, we have $\Gamma,\sigma(\Gamma');\cdot \vdash v : t;\varepsilon'$, which is sufficient since $A_x = y$ and $A = \cdot$.

Case (T-NEW):

$$\frac{\begin{array}{c}\Gamma,x:t_x,\Gamma';A_x,A \vdash e : t;\varepsilon \\ \Gamma,x:t_x,\Gamma' \vdash t :: U \qquad \Gamma,x:t_x,\Gamma'(\alpha) = \mathtt{N}\end{array}}{\Gamma,x:t_x,\Gamma';A_x,A \vdash \mathsf{new}\ e : ¡t^\alpha;\alpha \uplus \varepsilon}\ \text{(T-NEW)}$$

By the induction hypothesis we have $\Gamma, \sigma(\Gamma'); A \vdash e : \sigma(t); \varepsilon \uplus \varepsilon'$, since, by assumption, ε' is disjoint from $\alpha \uplus \varepsilon$. From mutual induction with Lemma 10 we have $\Gamma, \sigma(\Gamma') \vdash \sigma(t) :: \mathtt{U}$, and since $\alpha \notin \mathrm{dom}(\sigma)$ the third premise is also satisified. This suffices to establish the conclusion.

$$\Gamma, \sigma(\Gamma'); A \vdash \sigma(\mathsf{new}\ e) : \sigma(!t^{\alpha}); \alpha \uplus \varepsilon \uplus \varepsilon'$$

Case (T-ABS):

$$\frac{\begin{array}{c}\Gamma, x{:}t_x, \Gamma' \vdash t_y :: k \qquad q = p((A_x, A), \varepsilon) \\ \Gamma, x{:}t_x, \Gamma', y{:}t_y; A_x, A, a(y,k) \vdash e : t_e; \varepsilon\end{array}}{\Gamma, x{:}t_x, \Gamma'; A_x, A \vdash_{\varphi} \lambda y{:}t_y.e : q((y{:}t_y) \to t_e); \varepsilon} \text{ (T-ABS)}$$

Using Lemma 10 on the first premise we get

$$(P1') \quad \Gamma, \sigma(\Gamma') \vdash \sigma(t_y) :: k$$

We proceed on depending on whether or not ε' is empty.

Sub-case $\varepsilon' = \cdot$:

Using the induction hypothesis on the third premise we can get

$$(P3') \quad \Gamma, \sigma(\Gamma', y{:}t_y); A, a(y,k) \vdash \sigma(e) : \sigma(t_e); \varepsilon$$

From Proposition 6 and $\varepsilon' = \cdot$ we can establish $\Gamma \vdash t_x{::}\mathtt{U}$. From the uniqueness of kinding, Proposition 8, we can establish that $k = \mathtt{U}$ and thus $(A_x, A) = A$. Thus $q((A_x, A), \varepsilon \uplus \varepsilon') = q(A, \varepsilon)$ and we can conclude with the appropriate affinity qualifier on the function type.

Sub-case $\varepsilon' \neq \cdot$: We further divide this into sub-cases.

Sub-sub-case $x \in A_x$: Using the induction hypothesis on the third premise we can get

$$(P3') \quad \Gamma, \sigma(\Gamma', y{:}t_y); A, a(y,k) \vdash \sigma(e) : \sigma(t_e); \varepsilon \uplus \varepsilon'$$

Now, to get the right affinity qualifier on the result, we must show $p((A_x, A), \varepsilon) = p(A, \varepsilon \uplus \varepsilon')$. But, Notice that $(Ax, A) \neq \cdot \wedge fx \uplus fx' \neq \cdot$. Thus, we have $p((Ax, A), \varepsilon) = p(A, \varepsilon \uplus \varepsilon') = ¡$

Sub-sub-case $x \notin A_x$: In this case, we apply Proposition 7 to get either (i) $\Gamma \vdash t_x{::}\mathtt{A}$ or (ii) $\Gamma; \cdot \vdash_{\varphi} v : t_x, \cdot$.

In case (i), since $k = \mathtt{A}$ and $x \notin Ax, A$, we must have $x \notin FV(e)$. Thus, $\sigma(e) = e$ and we can use the induction hypothesis to construct

$$(P3') \quad \Gamma, \sigma(\Gamma', y{:}t_y); A, a(y,k) \vdash e : \sigma(t_e); \varepsilon$$

Now, the affinity qualifier on the result is $p(A, \varepsilon) = p((A_x, A), \varepsilon)$, as required.

In case (ii), we can use $\Gamma; \cdot \vdash_{\varphi} v : t_x, \cdot$ to reduce to the first subcase with $\varepsilon' = \cdot$.

Case (T-APP): We have two possible ways of inverting (T-APP), depending on whether A_x is used in the first or the second premise. Let $A_x, A = A', A''$.

$$\frac{\begin{array}{c}\Gamma, x{:}t_x, \Gamma'; A' \vdash e : !((y{:}t') \to t); \varepsilon_1 \\ \Gamma, x{:}t_x, \Gamma'; A'' \vdash e' : t'; \varepsilon_2\end{array}}{\Gamma, x{:}t_x, \Gamma'; A_x, A \vdash e\ e' : [y \mapsto e']t; \varepsilon_1 \uplus \varepsilon_2} \text{ (T-APP)}$$

We can apply the induction hypothesis to each of the two premises and obtain

$$\Gamma, \sigma(\Gamma'); A' \setminus A_x \vdash \sigma(e) : !(y{:}\sigma(t')) \to \sigma(t); \varepsilon_1 \uplus \varepsilon_1'$$

and

$$\Gamma, \sigma(\Gamma'); A'' \setminus A_x \vdash \sigma(e') : \sigma(t'); \varepsilon_2 \uplus \varepsilon_2'$$

Sub-case $\varepsilon' = \cdot$: In this case we have $\varepsilon_1' = \varepsilon_2'$ and the conclusion is straightforward.

Sub-case $\varepsilon' \neq \cdot$: Here, we proceed as in (T-ABS) and apply Proposition 7 to split into two subcases.

Sub-sub-case $\Gamma \vdash t_x :: \mathtt{A}$: In this case, since x is in only either A' or A'' (not both), we can conclude that x is free in either e or in e' only (not both). Thus, either $\varepsilon'_1 = \cdot \vee \varepsilon'_2 = \cdot$ and $\varepsilon_1 \uplus \varepsilon_2 \uplus \varepsilon'_1 \uplus \varepsilon'_2$.

Sub-sub-case $\Gamma; \cdot \vdash v : t_x, \cdot$: In this case, we can reduce to the first subcase and contruct the premises appropriately to construct $\Gamma, \sigma(\Gamma') \vdash_\varphi \sigma(e\ e') : \sigma(t), \varepsilon_1 \uplus \varepsilon_2$. Finally, we can conclude with (T-AFN) to introduce ε' and establish the name constraint $\varepsilon_1 \uplus \varepsilon_2 \uplus \varepsilon'$.

Case (T-TAB):

$$\frac{\begin{array}{c}\Gamma, x{:}t_x, \Gamma', \alpha{::}k; A_x, A \vdash e : t_e; \varepsilon_0 \uplus \varepsilon_1 \\ \alpha \notin \varepsilon_0 \qquad \varepsilon_1 \in \{\cdot, \alpha\} \qquad q = p((A_x, A), \varepsilon_0)\end{array}}{\Gamma, x{:}t_x, \Gamma'; A_x, A \vdash \Lambda\alpha{::}k.e : q(\forall \alpha{::}k \xrightarrow{\varepsilon_1} t_e); \varepsilon_0}$$

Applying the induction hypothesis to the first premise, we get

$$\Gamma, \sigma(\Gamma', \alpha{::}k); A \vdash \sigma(e) : \sigma(t_e); \varepsilon_0 \uplus \varepsilon_1 \uplus \varepsilon'$$

, where we get ε' disjoint from ε_1 by α-renaming.

Since $\alpha \notin \mathrm{dom}(\sigma)$ we can rewrite this as

$$\Gamma, \sigma(\Gamma'), \alpha{::}k; A \vdash \sigma(e) : \sigma(t_e); \varepsilon_0 \uplus \varepsilon' \uplus \varepsilon_1$$

This allows us to conclude with

$$\Gamma, \sigma(\Gamma'); A \vdash \sigma(e) : q'(\sigma(\forall \alpha{::}k \xrightarrow{\varepsilon_1} t_e)); \varepsilon_0 \uplus \varepsilon'$$

where $q' = p(A, \varepsilon_0 \uplus \varepsilon')$. In order to show that $q = q' = p((A_x, A), \varepsilon_0)$ we follow the same argument as in (T-ABS).

Case (T-TAP):

$$\frac{\Gamma, x{:}t_x, \Gamma'; A_x, A \vdash_\varphi e : q(\forall \alpha{::}k \xrightarrow{\varepsilon_1} t'; \varepsilon_0 \qquad \Gamma, x{:}t_x, \Gamma' \vdash t :: k}{\Gamma, x{:}t_x, \Gamma'; A_x, A \vdash_\varphi e\ [t] : [\alpha \mapsto t]t'; \varepsilon_0 \uplus ([\alpha \mapsto t]\varepsilon_1)}\ \text{(T-TAP)}$$

From the induction hypothesis on the first premise we get

$$\Gamma, \sigma(\Gamma'); A \vdash \sigma(e) : \forall \alpha{::}k \xrightarrow{\varepsilon_1} \sigma(t'); \varepsilon_0 \uplus \varepsilon'$$

For the second premise, from the mutual induction hypothesis of Lemma 10 we get

$$\Gamma, \sigma(\Gamma') \vdash \sigma(t) :: k$$

This is sufficient to establish the conclusion

$$\Gamma, \sigma(\Gamma'); A \vdash \sigma(e)\ [\sigma(t)] : [\alpha \mapsto \sigma(t)]\sigma(t'); \varepsilon_0 \uplus ([\alpha \mapsto t]\varepsilon_1) \uplus \varepsilon'$$

Case (T-CAP): Similar to (T-APP), we have two cases depending on whether A_x is used in the first or second premise. As previously, we establish that if $A_x = x$ then $\varepsilon' = \varepsilon'_1 \uplus \varepsilon'_2$ and in the conclusion we get $\varepsilon_1 \uplus \varepsilon_2 \uplus \varepsilon'_1 \uplus \varepsilon'_2$.

Case (T-CASE): Induction hypothesis on the first, third, fourth and fifth premise. Lemma 10 on the second premise. As with (T-CAP) and (T-APP), we use affinity to ensure that if $t_x :: A$, then x is free only in e or in e' and e''. If it is the latter, then since the effects ε' and ε'' are allowed to overlap, we can establish the conclusion.

Case (T-AFN), (T-DROP), (T-B1), (T-B2), (T-B3): All follow from the induction hypothesis.

Case (T-CONV): Here, we must show that type equivalence is preserved under substitution. The only interesting case here is (TE-RED). But, here we restrict reduction to closed type-level expressions e and e'. Thus $\sigma(e) = e$ and $\sigma(e') = e'$ and so $M \vdash e \longrightarrow e' \iff M \vdash \sigma(e) \longrightarrow \sigma(e')$. By axiomatizing that the type of e is preserved under reduction to e', we can use the induction hypothesis on the first and third premises to establish the conclusion. □

Lemma 10 (Substitution for kinding judgment)**.** *Given well-formed* $\Gamma; \cdot$ *well-formed, and* Γ' *such that:*

(A1) Γ, Γ' *is well-formed.*

(A2) $\Gamma, x : t_x, \Gamma'$ *also well-formed.*

(A3) $\Gamma, x : t_x, \Gamma' \vdash t :: k$

(A4) $\Gamma; \cdot \vdash v : t_x$

Then, for $\sigma = x \mapsto v$*,*

$$\Gamma, \sigma(\Gamma') \vdash \sigma(t) :: k$$

Proof. By mutual induction on the structure of (A3), together with Lemma 9.

Case (K-A): $\alpha \in \Gamma, \Gamma'$ and $FV(k) = \emptyset$.

Case (K-N): Induction hypothesis on the first premise gives us $\Gamma, \sigma(\Gamma') \vdash \sigma(t) :: A$. The second premise is trivial since $\alpha \in \Gamma, \Gamma'$ and $\alpha \notin \mathrm{dom}(\sigma)$.

Case (K-TC): S is unchanged, and $FV(K) = \emptyset$.

Case (K-AFN): Induction hypothesis.

Case (K-ALL): Induction hypothesis gives us $\Gamma, \sigma\Gamma', \alpha{::}k \vdash \sigma(t) :: k'$.

Case (K-FUN): Induction hypothesis on the first premise gives us $\Gamma, \sigma(\Gamma') \vdash \sigma(t) :: k$, and on the second premise $\Gamma, \sigma(\Gamma', x{:}t) \vdash \sigma(t') :: k'$. Finally, for the third premise, $\varepsilon \cap \mathrm{dom}(\sigma) = \emptyset$.

Case (K-TAP): Induction hypothesis on each premise.

Case (K-DEP): This is the interesting case, where we must rely on mutual recursion with Lemma 9 for the second premise to establish $\Gamma, \sigma(\Gamma'); \cdot \vdash_{\mathsf{type}} \sigma(e) : \sigma(t'); \varepsilon'$. We can conclude with (T-NC-type) and use the phase distinction to establish $\Gamma, \sigma(\Gamma'); \cdot \vdash_{\mathsf{type}} \sigma(e) : \sigma(t'); \cdot$ as required. For the first premise, we use the induction hypothesis to get $\Gamma, \sigma(\Gamma') \vdash t{::}t' \to K$, and from well-formedness of kinds we have that $FV(t') = \emptyset$. Thus, $\sigma(t') = t'$ and we have the conclusion $\Gamma, \sigma(\Gamma') \vdash t\ e :: K$.

Case (K-CON): Induction hypothesis on each premise.

□

Lemma 11 (Type substitution)**.** *Given well-formed* Γ *well-formed, and* Γ' *such that:*

(A1) $\Gamma, \alpha :: k, \Gamma'$ *well-formed.*

(A2) $\Gamma, \alpha :: k, \Gamma'; \cdot \vdash e : t; \varepsilon$ *or* $\Gamma, \alpha :: k, \Gamma' \vdash t :: k$

(A3) $\Gamma \vdash t' :: k$

(A5) $\sigma = \alpha \mapsto t'$

Then,

$$\Gamma, \sigma(\Gamma') \vdash_{\varphi} \sigma e : \sigma t; \sigma\varepsilon$$

and

$$\Gamma, \sigma(\Gamma') \vdash \sigma t :: K$$

Proof. By induction on the structure of (A2). □

Proposition 12 (Unification respects types)**.** *If all of the following are true*

(A1) $\cdot; \cdot \vdash v : t$

(A2) $\overrightarrow{x : t}; A \vdash e_p : t$*, where* $\cdot; \cdot \vdash t_i :: A \ \Rightarrow\ x \in A$

(A3) $v \succ e_p : \sigma$

Then, $dom(\sigma) = \overrightarrow{x}$, and $\cdot;\cdot \vdash \sigma(x_i) : t_i$.

Proposition 13 (Inversion of type abstractions)**.** *Given $\Gamma;\cdot$ well formed, and $\Gamma;\cdot \vdash_\varphi \Lambda\alpha{::}k.e : \forall\alpha{::}k \xrightarrow{\varepsilon_1} t_e;\varepsilon_0$. Then, $\Gamma,\alpha{::}k;\cdot \vdash_\varphi e : t_e : \varepsilon_0 \uplus \varepsilon_1$.*

Proposition 14 (Inversion of abstractions)**.** *Given $\Gamma;\cdot$ well formed, and $\Gamma;\cdot \vdash_\varphi \lambda x{:}t.e : (x{:}t) \to t_e;\varepsilon$. Then, for some k, $\Gamma,x{:}t;a(x,k) \vdash_\varphi e : t_e;\varepsilon$.*

Theorem 15 (Preservation)**.** *Given $\Gamma = \alpha_1{::}N,\ldots,\alpha_n{::}N$ and $(A1)\Gamma;\cdot \vdash_{\mathsf{term}} e : t;\varepsilon$ and the interpretation M and S are type-consistent, then if $(A2)M \vdash e \xrightarrow{l} e'$, $\Gamma;\cdot \vdash e' : t;\varepsilon$.*

Proof. By induction on the structure of the derivation (A1).

Case (T-X), (T-XA), (T-X-type), (T-NC-type): Inapplicable.

Case (T-B): Inversion on (A2) gives (E-DEL). Establish the conclusion by using (T-B3) with (A1) in the premise.

Case (T-BOT): Inversion on (A2) gives (E-INF). Conclusion is trivial from (A1).

Case (T-NEW): Inversion on (A2) gives (E-CTX) or (E-BOT). In the first case, the induction hypothesis applied to the first premise suffices. In the second case, (T-BOT) suffices.

Case (T-DROP): Apply (T-DROP) with the induction hypothesis in the premise.

Case (T-WKN): Apply (T-WKN) with the induction hypothesis in the premise.

Case (T-TAB), (T-ABS): In both cases, e is value.

Case (T-TAP): If e is not a value, inversion of (A2) gives (E-CTX) of (E-BOT) and apply the induction hypothesis or (T-BOT) to conclude.

If e is a value, the inversion of (A2) gives (E-TAP), (E-B3) or (E-B4).

Sub-case (E-TAP): From the inversion lemma, Proposition 13 applied to the first premise, we get

$$(A1.1) \qquad \Gamma,\alpha{::}k;\cdot \vdash e : t_e;\varepsilon_0 \uplus \varepsilon_1$$

Applying the type substitution lemma, Lemma 11, to (A1.1) using the second premise of (A1) we get the desired result:

$$\Gamma;\cdot \vdash_{\alpha\mapsto t} e : [\alpha \mapsto t]t';\varepsilon_0 \uplus [\alpha \mapsto t]\varepsilon_1$$

Sub-case (E-B3): The result follows from type consistency of M and S.

Sub-case (E-B4): Apply (T-B4) using (A1) in the premise.

Case (T-APP): If either e or e' are not values, inversion of (A2) gives (E-CTX) or (E-BOT) and we apply the induction hypothesis of (T-BOT) to conclude.

If both are values, then inversion of (A2) gives (E-APP), (E-B1) or (E-B2).

Sub-case (E-APP): From the inversion lemma, Proposition 14 applied to the first premise, we get

$$(A1.1) \qquad \Gamma,x{::}t';a(x,k) \vdash e : t;\varepsilon$$

Applying the substitution lemma, Lemma 9, to (A1.1) using the second premise of (A1) we get the desired result:

$$\Gamma;\cdot \vdash_{x\mapsto v} e : [\alpha \mapsto t]t';\varepsilon_0 \uplus [\alpha \mapsto t]\varepsilon_1$$

Sub-case (E-B1): The result follows from type consistency of M and S.

Sub-case (E-B2): Apply (T-B1) using (A1) in the premise.

Case (T-CAP): Inversion of (A1) gives (E-CTX) or (E-BOT). Result follows from induction hypothesis or (T-BOT).

Case (T-CASE): If the discriminant e is not a value, inversion of (A2) gives (E-CTX) or (E-BOT) which we handle as in the other cases.

If e is a value, then inversino of (A2) gives (E-CASE). There are two sub-cases, depending on which branch is taken.

Sub-case (else-branch): The final premise of (T-CASE) gives:

$$\Gamma;\cdot \vdash e'' : t;\varepsilon''$$

where $\varepsilon'' \subseteq \varepsilon$. However, we can always use (T-WKN) to introduce additional effects in the conclusion.

Sub-case (case-branch): From the premises of (A1) we have

$$\Gamma;\cdot \vdash_{\varphi} v : t_v;\varepsilon_v$$

and

$$\Gamma,\overrightarrow{x:t};A'' \vdash_{\varphi} e' : t$$

From Lemma 12 we can show that in $v \succ e_p : \sigma$, $\forall x_i \in \mathrm{dom}(\sigma).\Gamma;\cdot \vdash \sigma(x_i) : t_i$, then by repeated application of Lemma 9, we can conclude

$$\Gamma;\cdot \vdash_{\varphi} \sigma(e') : t;\varepsilon'$$

Finally, as in the else case, we can always expand the effects ε' using (T-AFN), if necessary.

Case (T-B1), (T-B2), (T-B3): All of these are values.

Case (T-CONV): Induction hypothesis on the premise. □

A concise syntax for AIR

$$\begin{array}{lrcl} \text{policy} & \pi & ::= & (id, P, \overrightarrow{\sigma}, \overrightarrow{\mathsf{R}_r}, \overrightarrow{\mathsf{R}_t}) \\ \text{rule} & \mathsf{R} & ::= & (id, x, d, \overrightarrow{\exists x{:}t.C}, e, A) \\ \text{judgment index} & \rho & ::= & r \mid t \end{array}$$

$\pi \models S$ Translation from a policy π to a signature S

$$\frac{S = id{:}Id, B{:}Class, S_0 \quad S \models \overrightarrow{\sigma} : S_\sigma \quad S, S_\sigma \models_r \overrightarrow{\mathsf{R}_r} : S_r \quad S, S_\sigma, S_r \models_t \overrightarrow{\mathsf{R}_t} : S_t}{(id, P, \overrightarrow{\sigma}, \overrightarrow{\mathsf{R}_r}, \overrightarrow{\mathsf{R}_t}) \models S, S_\sigma, S_r, S_t} \qquad \frac{S \models_\rho \mathsf{R}_0 : S_0 \quad S, S_0 \models_\rho \overrightarrow{\mathsf{R}} : S'}{S \models_\rho \mathsf{R}_0, \overrightarrow{\mathsf{R}} : S_0, S'}$$

$S \models \overrightarrow{\sigma} : S'$ Translation from states $\overrightarrow{\sigma}$ to signature S'

$$\frac{S \models \sigma_0 : S_0 \quad S, S_0 \models \overrightarrow{\sigma} : S'}{S \models \sigma_0, \overrightarrow{\sigma} : S_0, S'} \qquad \frac{C \notin \mathrm{dom}(S) \quad t = Class \Rightarrow Instance}{S \models C : (C{:}t)} \qquad \frac{i \in 1 \ldots n \quad \cdot \vdash_{S_0} t_i :: U \quad C \notin \mathrm{dom}(S)}{S \models C \text{ of } \overrightarrow{t} : (C{:}Class \Rightarrow t_1 \Rightarrow \ldots \Rightarrow t_n \Rightarrow Instance)}$$

$S \models_\rho \overrightarrow{\mathsf{R}} : S'$ Translation of a rule.

$$\frac{\begin{array}{l} \Gamma = src{::}N, dst{::}N, \alpha{::}U, s{:}!Instance^{src}, x{:}Protected\ \alpha\ src, d{:}¡Instance^{dst}, d'{:}Class \quad src; dst; S; \Gamma \models_\rho \overrightarrow{\exists x{:}t.C}; e : t' \\ id \notin \mathrm{dom}(S) \qquad t_r = \forall src{::}N, dst{::}N, \alpha{::}U.(s{:}!Instance^{src}) \to (ClassOf\ src\ B) \to (x{:}Protected\ \alpha\ src) \to \\ B{:}Class \in S \qquad\qquad (d{:}!Instance^{dst}) \to (d'{:}Class) \to (ClassOf\ dst\ d') \to t' \end{array}}{S \models_\rho (id, x, d, \overrightarrow{\exists x{:}t.C}, e, A) : (id{:}t_r)}\ \text{(S-RULE)}$$

$src; dst; S; \Gamma \models_\rho \overrightarrow{\exists x{:}t.C}; e : t'$ Translation of a rule body.

$$\frac{\Gamma \vdash t :: k \quad src; dst; S; \Gamma, x{:}t \models C : t' \quad src; dst; S; \Gamma, x{:}t \models_\rho \overrightarrow{\exists x{:}t.C}; e : t''}{src; dst; S; \Gamma \models_\rho \exists x{:}t.C, \overrightarrow{\exists x{:}t.C}; e : (x{:}t) \to t' \to t''}\ \text{(TR-COND)}$$

$$\frac{\Gamma, s'{::}N; \cdot \vdash e : (!Instance^{s'} \times !Instance^{d}); \varepsilon}{s; d; S; \Gamma \models_t \cdot; e : (!Instance^{s} \times !Instance^{d})}\ \text{(T-BODY)} \qquad \frac{\Gamma, s'{::}N; \cdot \vdash e : (¡Instance^{s'} \times ¡Instance^{d} \times Protected\ \alpha\ s); \varepsilon}{s; d; S; \Gamma \models_r \cdot; e : (¡Instance^{s} \times ¡Instance^{d} \times Protected\ \alpha\ d)}\ \text{(R-BODY)}$$

$src; dst; S; \Gamma \models C : t$ Translation of a condition expression to a witness type.

$$\frac{s; d; S; \Gamma \models A_1 : (e_1, Class) \quad s; d; S; \Gamma \models A_2 : (e_2, Class)}{s; d; S; \Gamma \models A_1\ IsClass\ A_2 : IsClass\ e_1\ e_2}\ \text{(C-CLS)} \qquad \frac{s; d; S; \Gamma \models A_1 : (e_1, Prin) \quad s; d; S; \Gamma \models A_2 : (e_2, Prin)}{s; d; S; \Gamma \models A_1\ ActsFor\ A_2 : ActsFor\ e_1\ e_2}\ \text{(C-ACTS)}$$

$$\frac{s; d; S; \Gamma \models A_1 : (e_1, !Instance) \quad s; d; S; \Gamma \models A_2 : (e_2, Instance)}{s; d; S; \Gamma \models A_1\ InState\ A_2 : InState\ e_1\ e_2}\ \text{(C-STATE)} \qquad \frac{s; d; S; \Gamma \models A_1 : (e_1, Int) \quad s; d; S; \Gamma \models A_2 : (e_2, Int)}{s; d; S; \Gamma \models A_1\ LessThan\ A_2 : LEQ\ e_1\ e_2}\ \text{(C-LEQ)}$$

$src; dst; S; \Gamma \models A : (e, t)$ Translation of an atom A to an expression e of type t

$$\frac{\Gamma; \cdot \vdash x : t; \cdot}{s; d; S; \Gamma \models x : (x, t)}\ \text{(A-X)} \qquad \frac{\Gamma; \cdot \vdash B : t; \cdot}{s; d; S; \Gamma \models B : (B, t)}\ \text{(A-B)} \qquad \frac{(B{:}Class) \in S \quad s; d; S; \Gamma \models A_i : (e_i, t_i)}{s; d; S; \Gamma \models C(\overrightarrow{A}) : (C\ \{B\}\ \{e_1\}\ \{\ldots\}\ \{e_n\}, Instance^{\alpha})}\ \text{(A-ST)}$$

$$\frac{\Gamma; \cdot \vdash x : ¡Instance^{s}; \cdot}{s; d; S; \Gamma \models Self : (x, ¡Instance^{s})}\ \text{(A-SELF)} \qquad \frac{s; d; S; \Gamma \models A : (e, Class)}{s; d; S; \Gamma \models Principal(A) : (principal\ e, Prin)}\ \text{(A-PRIN)}$$

$$\frac{s; d; S; \Gamma \models A : (e, ¡Instance^{src}) \quad (B : Class) \in S}{s; d; S; \Gamma \models Class(A) : (B, Class)}\ \text{(A-SCLS)} \qquad \frac{s; d; S; \Gamma \models A : (e, ¡Instance^{dst}) \quad (d : Class) \in \Gamma}{s; d; S; \Gamma \models Class(A) : (d, Class)}\ \text{(A-DCLS)}$$

Figure 9. Translating an AIR policy to a λ_{AIR} signature.

$A;\pi \models T;A'$ An automaton instance in state A, accepts the trace T and transitions to A' according to the policy π

$$\frac{A;\pi \models T;A'}{A;\pi \models \cdot,T;A'}\ \text{(L-DOT)} \qquad \frac{A;\pi \models l;A' \quad A';\pi \models T;A''}{A;\pi \models l,T;A''}\ \text{(L-TR)} \qquad \frac{(B,\ldots) \notin \pi.\overrightarrow{\mathsf{R}}}{A;\pi \models (B:\mathscr{D});A}\ \text{(L-NOTX)} \qquad \frac{A;\pi \models (B:\mathscr{D});A'}{A;\pi \models (B:t,\mathscr{D});A'}\ \text{(L-SKIP-t)}$$

$$\frac{src::N;\cdot \vdash v_{inst} : \text{¡}Instance^{src} \quad src::N;\cdot \vdash v_{ev} : ClassOf\ src\ (\mathsf{Class}\ \{\pi.id\}\ \{v_{prin}\}) \quad v_{inst};A;\pi \overset{\mathscr{D}}{\Rightarrow} A';B}{A;\pi \models (B:v_{inst},v_{ev},\mathscr{D});A'}\ \text{(L-ENTER)}$$

$v_{src};A;\pi \overset{\mathscr{D}}{\Rightarrow} A';B$ Given trace event $\mathscr{D}$ automaton in state A transitions to A' using rule B of policy π

$$\frac{\begin{array}{c}(B,x,d,\overrightarrow{\mathsf{RC}},e,A') \in \pi.\overrightarrow{\mathsf{R}} \qquad v_{src} = \mathsf{new}\ v'_{src} \qquad \pi \models S \qquad s;d;S;\cdot \models A : (v'_{src},t) \\ \sigma_c = ((\mathsf{Class}\ d) \mapsto v_{cd}) \qquad \sigma = (\mathsf{Self} \mapsto v_{src}, x \mapsto v_x, d \mapsto v_{dst}) \qquad \sigma(\sigma_c(\overrightarrow{\mathsf{RC}})) \models \mathscr{D}\end{array}}{v_{src};A;\pi \overset{v_x,v_{dst},v_{cd},v_{ev},\mathscr{D}}{\Rightarrow} A';B}\ \text{(TX)}$$

$\overrightarrow{\mathsf{RC}} \models \mathscr{D}$ Check that all obligations in the release conditions are justified by evidence in $\mathscr{D}$

$$\frac{\cdot;\cdot \vdash v_{ev} : \mathscr{T}(\mathscr{R})\ e_1\ e_2 \qquad \sigma_0 = (x \mapsto v_x) \qquad \sigma_0 A_1 \succ e_1 : \sigma_1 \qquad \sigma_1\sigma_0 A_2 \succ e_2 : \sigma_2 \qquad \sigma = \sigma_0,\sigma_1,\sigma_2 \qquad \sigma(\overrightarrow{\mathsf{RC}}) \models \mathscr{D}}{\exists x{:}t.A_1\ \mathscr{R}\ A_2, \overrightarrow{\mathsf{RC}} \models v_x, v_{ev}, \mathscr{D}}\ \text{(CERT)}$$

where $\mathscr{T}(\text{ActsFor}) = ActsFor \qquad \mathscr{T}(\text{InState}) = InState \quad \mathscr{T}(\text{IsClass}) = IsClass \quad \mathscr{T}(\leq) = LEQ$

$A \succ e : \sigma$ Unification of atom A with expression e

$$\frac{}{x \succ e : x \mapsto e} \qquad \frac{}{e \succ e :} \qquad \frac{A \succ e : \sigma}{\text{Principal}\ A \succ principal\ e : \sigma} \qquad \frac{A \succ e : \sigma}{\text{Class}\ A \succ class\_of\_inst\ e : \sigma}$$

Figure 10. Trace acceptance condition defined by an AIR class.

B Proof of correct API usage

In this section, we first define a translation that produces a λ_{AIR} signature from an AIR policy. We then define a semantics for an AIR policy as a regular language—i.e., an automaton that accepts strings. Finally, we show that the execution trace generated by a λ_{AIR} program is a member of the language accepted by the AIR automaton.

Definition 16 (Consistency of a model)**.** *A model M and a policy π are consistent if, and only if,*

$$\pi \models S \quad \Rightarrow S \text{ and } M \text{ are type consistent}$$

And, the state transitions specified by M are in accordance with the policy π. I.e., if all of the following are true

1. $(B,x,d,\overrightarrow{\exists x{:}t.C},e,A) \in \pi$
2. $$\Gamma = src::N, dst::N, \alpha::U, s:!Instance^{src}, \\ x{:}Protected\ \alpha\ src, d{:}\text{¡}Instance^{dst}, d'{:}Class, \overrightarrow{x{:}t}$$
3. $src;dst;S;\Gamma \models A : (e_A,t)$
4. $B : \mathscr{D} \rightsquigarrow v \in M$
5. $\sigma = (src \mapsto \mathscr{D}_1, dst \mapsto \mathscr{D}_2, \ldots, x_n \mapsto \mathscr{D}_m)$

Then, $v = \mathsf{Pair}\ \{v_{src}\}\ \{v'\}$, and $v = \sigma(e_A)$.

Theorem 17 (Security)**.** *Given all of the following:*

(A1) An AIR *declaration π of a class with identifier* C *owned by principal P.*

(A2) A signature S such that $\pi \models S$.

(A3) M *and* π *are consistent.*

(A4) $\Gamma = src::N, dst::N, s : ¡Instance^{src}$.

(A5) $\Gamma; s \vdash e : t; dst$, *and* e *is* π*-free.*

(A6) $v = \mathsf{new}\ \mathsf{Init}\ \{\mathsf{Class}\ \{\mathsf{C}\}\ \{P\}\}$.

(A7) $M \vdash ((s \mapsto v)e) \xrightarrow{l_1} e_1 \ldots \xrightarrow{l_n} e_n$.

Then

$$Init; \pi \models (l_1, \ldots, l_n); A'$$

Proof. By induction on the length of the string $l_1, \ldots, l_n$. We strengthen the induction hypothesis to actually prove $A; \pi \models (l_1, \ldots, l_n); A'$, where

(A6') $s; d; S; \cdot \models A : (v', (Instance^{src}))$ and $v = \mathsf{new}\ v'$, such that $src::N; \cdot \vdash v : ¡Instance^{src}$.

Case $(l_1 = \cdot)$: Apply (L-DOT) and use the induction hypothesis in the premise, since the length of the trace is reduced by one.

Case $(l_1 = B : \mathscr{D})$ AND $B \notin \pi.\overrightarrow{\mathsf{R}}$: Apply (L-TR) with (L-NOT-X) in the first premise, and the induction hypothesis in the second premise.

Case $(l_1 = B : \mathscr{D})$ AND $(B, \ldots) \in \pi.\overrightarrow{\mathsf{R}}$: This is the interesting case. We will use (L-TR) with the induction hypothesis in the second premise. Our first goal is to show that the first premise can be satisfied.

By premise (A2), and since we have $(B, ldots) \in \pi.\overrightarrow{\mathsf{R}}$, we have $(B : t_B) \in S$, where $S' \models_\rho (B, \ldots) : (B : t_B)$ from (S-RULE). Additionally,

$$\begin{aligned} t_B = \forall src, dst::N, \alpha::U.(s{:}!Instance^{src}) \rightarrow (ClassOf\ src\ C_B) \rightarrow \\ (x{:}Protected\ \alpha\ src) \rightarrow (d{:}!Instance^{dst}) \rightarrow \\ (d'{:}Class) \rightarrow (ClassOf\ dst\ d') \rightarrow t'_B \end{aligned}$$

where, $C_B = \mathsf{Class}\pi.idP$, is the representation of the class of π. Next, from type consistency of M and S we have

$$\mathscr{D} = t_{src}, t_{dst}, t_\alpha, v_{src}, v_{ev}, v_x, v_{dst}, v_{cd}, v_{ev}\mathscr{D}'$$

Thus, for the first premise, we apply (L-SKIP-t) three times to skip past the first three types in $\mathscr{D}$. We must now show

$$A; \pi \models B : v_{src}, v_{ev}, v_x, v_{dst}, v_{cd}, v_{ev}, \mathscr{D}' : A'$$

Again, from type consistency, we can satisfy the first two premises of (L-ENTER). We now must show

$$v_{src}; A; \pi \overset{v_x, v_{dst}, \mathscr{D}'}{\Rightarrow} A'; B$$

The first premise of this judgment (TX) is satisfied by the assumptions of this case. For the second premise, we use inversion on the first premise of (L-ENTER) to establish that $v_{src} = \mathsf{new}\ v_{src'}$. The third premise is given by (A2). The fourth premise is interesting. We have to show that the current state of the automaton instance v_{src} in the program is the same as the current state A in the trace-acceptance relation. However, from our strengthened induction hypothesis, we have that $s; d; S; \cdot \models A : (v', Instance^{src})$. Additionally, from the second premise of (L-ENTER) we have that v_{src} is an instance of the π class and from (A5) we have that e is π-free. Thus, we can conclude that $v_{src} = v = \mathsf{new}\ v'$. The fifth and sixth premises are constructed to enable showing that the final premise of (TX) can be satisfied. In particular, a straightforward induction on the length of $\mathscr{D}'$, and by relying on type consistency, we can easily show that (TX) is satisfiable.

The more interesting goal is to show to re-establish the premises of this theorem so as to be able to apply the induction hypothesis for the second premise of (L-TR). This will follow from the type-soundness result for λAIR.
From inversion of the reduction relation, we have :

$$\frac{B : \mathscr{D}, v \rightsquigarrow v' \in M \qquad l_1 = B : \mathscr{D}, v}{M \vdash \mathscr{E} \cdot [\![B]\!]^{\mathscr{D}}\ v \xrightarrow{l_1} \mathscr{E} \cdot v'} \text{(E-B1)}$$

From type consistency of M and S, we have that the type of e is determined by either (T-BODY) or (R-BODY). I.e., e has type

$$(¡Instance^{src} \times ¡Instance^{dst})$$

or

$$(¡Instance^{src} \times ¡Instance^{dst} \times Protected\ t_x\ dst).$$

I.e., $v' = Pair\ \{v'_{src}\}\ \{v''\}$. From (A5) we have

$$src{::}N, dst{::}N, s : ¡Instance^{src} \vdash e : t; dst$$

We can use the substitution lemma, relying on the fact that s is an affine assumption, to establish

$$src{::}N, dst{::}N \vdash (s \mapsto v)e : t; dst \uplus src$$

.
Now, from subject reduction, we have

$$src{::}N, dst{::}N \vdash \mathscr{E} \cdot v' : t; dst \uplus src$$

where

$$src{::}N \vdash v'_{src} : ¡Instance^{src}; src$$

Now, applying the converse of the substitution lemma, we can also establish

$$\Gamma; s \vdash \mathscr{E} \cdot Pair\ \{s\}\ \{v''\} : t; dst$$

which is the form of (A5) necessary to apply the induction hypothesis. To conclude, we must also establish that assumption (A6')—i.e., that $v'_{src} = \mathsf{new}\ v''_{src}$ where $s; d; S; \cdot \models A' : (v''_{src}, t'_A)$. However, this follows from the semantic consistency of M with respect to π, as established by (A3). □