

# Distributed Programming with Distributed Authorization \*

Kumar Avijit

School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA 15213  
kavijit@cs.cmu.edu

Anupam Datta

CyLab  
Carnegie Mellon University  
Pittsburgh, PA 15213  
danupam@cmu.edu

Robert Harper

School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA 15213  
rwh@cs.cmu.edu

## Abstract

We propose a programming language, called PCML<sub>5</sub>, for building distributed applications with distributed access control. Target applications include web-based systems in which programs must compute with stipulated resources at different sites. In such a setting, access control policies are *decentralized* (each site may impose restrictions on access to its resources without the knowledge of or cooperation with other sites) and *spatially distributed* (each site may store its policies locally). To enforce such policies PCML<sub>5</sub> employs a distributed proof-carrying authorization framework in which sensitive resources are governed by reference monitors that authenticate principals and demand logical proofs of compliance with site-specific access control policies. The language provides primitive operations for *authentication*, and *acquisition* of proofs from local policies. The type system of PCML<sub>5</sub> enforces locality restrictions on resources, ensuring that they can only be accessed from the site at which they reside, and enforces the authentication and authorization obligations required to comply with local access control policies. This ensures that a well-typed PCML<sub>5</sub> program cannot incur a runtime access control violation at a reference monitor for a controlled resource.

**Categories and Subject Descriptors** D.3.1 [Programming Languages]: Formal Definitions and Theory; F.3.3 [Logics and Meanings of Programs]: Studies of Program Constructs—Type structure; D.4.6 [Operating Systems]: Security and Protection—Access Controls, Authentication

**General Terms** Languages, Design, Theory, Security

**Keywords** Distributed programming, proof-carrying authorization, authorization logic, logical frameworks, phase distinction

## 1. Introduction

We present a programming language for writing distributed programs that compute with resources spread across multiple sites on

\*The material is based on work supported in part by NSF grants 0716469, and CCR-0424422 (TRUST). The second author was also supported in part by the Army Research Office through grant DAAD19-02-1-0389 at Carnegie Mellon CyLab.

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

TLDI'10, January 23, 2010, Madrid, Spain.

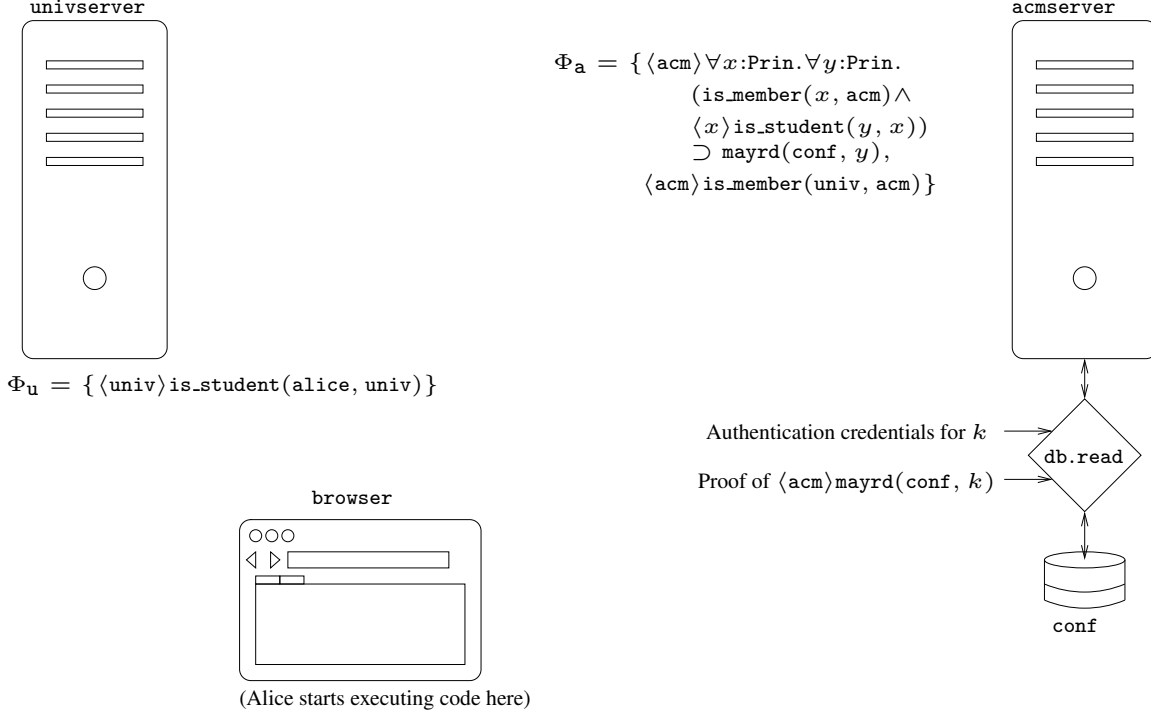
Copyright © 2010 ACM 978-1-60558-891-9/10/01...\$5.00

a network. Each site on the network is assumed to be administered by a principal who wishes to control access to all the resources at that site.

We consider an approach based on distributed proof-carrying authorization (DPCA), which was first introduced by Appel and Felten (Appel and Felten 1999), and developed further in subsequent work (Bauer et al. 2001, 2005; Garg 2009). There are two key ingredients to the DPCA approach. First, the policies specifying the conditions under which a principal may access a resource are formulated as theories in a formal logic. The exact form of the logic varies from one setting to another, but must be expressive enough to admit reasoning from the perspective of a principal using judgments of the form “*k* affirms *p*”. This expressivity is required when not all resources are controlled by the same principal. The second component of DPCA is the use of proof-carrying reference monitors to control access to resources. With each access request to a resource, its reference monitor demands a formal proof stating that the controlling principal at the site affirms that the accessing principal is permitted to access the resource. The burden of producing such a proof from the local policies is thus on the accessing principal, and not on the reference monitor. The reference monitor simply verifies the proof before granting access.

The security of the DPCA approach relies on the integrity of reference monitors and the soundness of the authorization logic. It is assumed that an attacker cannot access the resource in any manner other than having to go through the reference monitor. It is further assumed that the proof-checker of the reference monitor is sound in the sense that it accepts only those proofs that are valid under the host’s policies. The effectiveness of the reference monitor depends, in turn, on the expressiveness of the authorization logic. In particular, the affirms judgment is essential to stating distributed authorization policies. For example, if *k* controls access to a resource *r*, then an access control entry for *r* has to be of the form “*k* affirms mayrd(*r*, *k*’)”, stating that the principal *k*’ is permitted by the principal *k* to access *r*. The logic should not allow for deductions of “*k* affirms *p*” for propositions *p* that were not intended by *k*. Such properties have been investigated by Garg and Pfenning in context of a particular authorization logic (Garg and Pfenning 2006).

To facilitate writing non-malicious code in such a setting, we propose an extension to ML<sub>5</sub> (Murphy 2008; Murphy et al. 2007), which is a distributed programming language, with support for distributed proof-carrying authorization. The ML<sub>5</sub> type system supports distributed programming by using a modal interpretation to distribution of resources. In particular, the ML<sub>5</sub> type system ensures that resources on a network can only be accessed at the site where they reside, even though references to these resources may be distributed freely. This is achieved through the typing judgment *m* : *A*@*w* stating that the computation *m* may be executed at the site *w*. ML<sub>5</sub> itself only supports distribution, but not authorization.



**Figure 1.** An example scenario for distributed PCA

The contribution of this paper is an extension of the  $\text{ML}_5$  type system, to statically enforce the proof requirements imposed by proof-carrying reference monitors that control the resources on the network. There are three main ingredients in this extension:

1. **Authentication:**  $\text{PCML}_5$  provides a primitive operation that allows the executing program to authenticate as a principal, obtaining credentials that are presented to the reference monitor in order to identify the accessing principal.
2. **Distributed policy acquisition:** We assume that access control policies are distributed on the network, and must be acquired dynamically to construct the proofs demanded by the reference monitors at a site. As a simple example, a program must consult the access control list for a resource at a site to determine whether a specified principal is permitted to access a specified resource in a particular manner. We enrich  $\text{PCML}_5$  with a primitive for doing deductions using local authorization policies at the sites.
3. **Specification of proof obligations imposed by reference monitors:** Each resource is governed by an API whose type specifies all the proof demands imposed by the reference monitor. In particular, the type for a resource  $r$  specifies that the reference monitor must be presented with the following:
  - A credential that witnesses authentication of the accessing principal  $k$ .
  - A proof in the authorization logic that  $k$  is permitted by the governing authority to access  $r$ .
  - Any other arguments appropriate to the resource and the particular mode of access.

The main result of the paper is the authorization theorem which states that well-typed programs cannot incur authorization failure at the reference monitor governing the resource. This theorem holds because the type system ensures that the accessing principal

must provide the proof required by the reference monitor, which therefore cannot fail at execution time. This imposes the obligation on the programmer to collect sufficient credentials and derivations prior to arriving at the API call-site. This process of collecting credentials might itself fail. For instance, the program may query the host policy to determine if the principal  $k$  is allowed access to a resource  $r$ , i.e., attempt to prove that  $k'$  affirms  $\text{mayrd}(r, k)$ , where  $k'$  is the controlling authority of  $r$ . This query can fail, if the indicated access is not permitted. However, if the query succeeds, the program acquires evidence that can be used in further deductions that are presented to the reference monitor. An analogy with array bounds check may be helpful here. One may envision a language in which the ability to perform a subscript operation on an array demands a proof that the given index is within bounds. It is incumbent on the program to acquire, through a combination of run-time checks and logical deductions, a proof that the proposed access is legitimate. The run-time checks used to assemble the ingredients of these proofs may fail, but the language can ensure that the proof presented at the access site will always be valid, and hence proof-checking at reference monitor cannot fail.

Figure 1 illustrates the use of  $\text{PCML}_5$  for authorized, distributed computation. There are three sites on the network, viz., browser, acmserver and univserver, and acmserver hosts a database conf. Access to conf at acmserver is governed by the principal acm. The site univserver is administered by univ. The local policy of each site  $w$  is specified as a logical theory  $\Phi_w$ , which is a set of propositions of the form  $\langle k \rangle P$ , where  $k$  is the governing authority at  $w$ . We use the modal proposition  $\langle k \rangle P$  as the internalization of the judgment  $k$  affirms  $P$ . The database conf is protected by a reference monitor that demands an authentication credential, and a proof of the proposition  $\langle \text{acm} \rangle \text{mayrd}(\text{conf}, k)$  with  $k$  being the principal authenticated.

We wish to write a distributed program that starts execution at browser, travels to the site acmserver, and accesses the database conf with appropriate authentication and authorization proofs. The

first step is to authenticate the program as running on behalf of a principal using PCML<sub>5</sub>'s authentication primitive. This generates a credential that will be used to prove the identity of accessing principal at the reference monitor. The rest of the program can be written as if executing on behalf of this principal, say  $k$ . The next step is assembling a proof for  $\langle \text{acm} \rangle \text{mayrd}(\text{conf}, k)$ . The program does this by first proving that  $k$  is a student of  $\text{univ}$ , i.e., by obtaining a proof for  $\langle \text{univ} \rangle \text{is\_student}(k, \text{univ})$ , and then using this proof to deduce that  $\text{acm}$  permits  $k$  to read  $\text{conf}$ . First of all, the program moves to  $\text{univ}$  to do a local proof search for  $\langle \text{univ} \rangle \text{is\_student}(k, \text{univ})$ . For the particular authorization theories in the example, this proof search fails unless  $k$  is  $\text{alice}$ . In case of failure, the program simply aborts. However if the query is successful, the program moves to  $\text{acmserver}$  to deduce

$$\langle \text{univ} \rangle \text{is\_student}(k, \text{univ}) \supset \langle \text{acm} \rangle \text{mayrd}(\text{conf}, k)$$

from the local policy. This proof check succeeds under the theory  $\Phi_a$ . The two locally deduced proofs are moved to  $\text{acmserver}$  where they are combined using an  $\supset$ -elimination to deduce

$$\langle \text{acm} \rangle \text{mayrd}(\text{conf}, k)$$

This proof, along with the authentication credential for  $k$  is passed to the API for reading  $\text{conf}$ . Provided that the previous local deductions at  $\text{univserver}$  and  $\text{acmserver}$  succeed, it is guaranteed that the proof is accepted by the reference monitor.

## 2. Authorization logic

The presentation of PCML<sub>5</sub> in this paper uses the GP authorization logic (Garg and Pfenning 2006), which we summarize in this section. However, we emphasize that we pick a specific logic for illustration only. Since the logic is integrated into the type system using higher-order abstract syntax (as described in Section 5), the PCML<sub>5</sub> language can be easily coupled with other authorization logics that can be encoded in the framework.

Sorts	$s$	::=	$\text{prin} \mid \text{db}$
First-order terms	$a$	::=	$\alpha \mid \mathbf{a}$
Predicate symbols	$p$		
Propositions	$P$	::=	$P_1 \supset P_2 \mid P_1 \wedge P_2$ $\mid \forall \alpha:s.P \mid \langle a \rangle P \mid p(a_1, \dots, a_n)$
Basic judgments	$J$	::=	$a:s \mid P \text{ true} \mid a \text{ affirms } P$
Proposition context	$\Phi$	::=	$\cdot \mid \Phi, P \text{ true}$
Term context	$\Delta$	::=	$\cdot \mid \Delta, \alpha:s$
Signature	$\Sigma$	::=	$\cdot \mid \Sigma, \mathbf{a}:s \mid \Sigma, p:(s_1, \dots, s_n)$

The logic is parameterized by a signature  $\Sigma$  that fixes the arity of the predicates;  $p:(s_1, \dots, s_n)$  denotes that  $p$  is a predicate of arity  $n$  with  $s_i$  being the sort of its  $i$ th argument. There are at least two sorts, viz.,  $\text{prin}$  and  $\text{db}$ , representing principals and databases respectively. The signature also fixes principal and database constants  $\mathbf{a}$ .

Apart from the standard propositional connectives, the logic features a family of modalities indexed by principals, to express propositions affirmed by principals. The modal proposition  $\langle a \rangle P$  (read as “ $a$  says  $P$ ”) expresses that proposition  $P$  is endorsed by principal  $a$ . There are three forms of basic judgments: the sorting judgment  $a:s$ , which means that the first-order term  $a$  is of the sort  $s$ ; the truth judgment  $P \text{ true}$ ; and the affirmation judgment  $a \text{ affirms } P$ , which means that  $a$  (a term of the sort  $\text{prin}$ ) endorses  $P$ .

We use  $\Delta; \Phi \vdash_{\Sigma}^{\mathcal{L}} J$  to mean the hypothetical judgment  $J$  from assumptions  $\Delta$  and  $\Phi$ , under the signature  $\Sigma$ . The superscript  $\mathcal{L}$  differentiates entailment in the logic from that in hypothetical judgments of PCML<sub>5</sub> which we shall introduce later. We collect hypotheses about sorting judgments  $\alpha:s$  under the context  $\Delta$ ; hy-

$$\frac{\Delta; \Phi \vdash_{\Sigma}^{\mathcal{L}} P \text{ true} \quad \Delta; \Phi \vdash_{\Sigma}^{\mathcal{L}} a : \text{prin}}{\Delta; \Phi \vdash_{\Sigma}^{\mathcal{L}} a \text{ affirms } P} (\text{truaff})$$

$$\frac{\Delta; \Phi \vdash_{\Sigma}^{\mathcal{L}} a \text{ affirms } P}{\Delta; \Phi \vdash_{\Sigma}^{\mathcal{L}} \langle a \rangle P \text{ true}} (\langle \rangle\text{-I})$$

$$\frac{\Delta; \Phi \vdash_{\Sigma}^{\mathcal{L}} \langle a \rangle P_1 \text{ true} \quad \Delta; \Phi, P_1 \text{ true} \vdash_{\Sigma}^{\mathcal{L}} a \text{ affirms } P_2}{\Delta; \Phi \vdash_{\Sigma}^{\mathcal{L}} a \text{ affirms } P_2} (\langle \rangle\text{-E})$$

**Figure 2.** Selected rules from natural deduction for the authorization logic. (Reproduced from (Garg and Pfenning 2006)).

potheses of truth judgments  $P \text{ true}$  are written as  $\Phi$ . We do not need to form hypotheses about affirmation judgments in this logic.

In order to illustrate the *says* modality, we present selected rules from a natural deduction for the logic in Figure 2. Rule ( $\text{truaff}$ ) says that true propositions are affirmed by all principals. Rule ( $\langle \rangle$ -I) internalizes the judgment  $a \text{ affirms } P$  as the proposition  $\langle a \rangle P$ . Rule ( $\langle \rangle$ -E) eliminates the modality in  $\langle a \rangle P_1$  by using a proof of  $\langle a \rangle P_1$  to cut a hypothesis  $P_1 \text{ true}$  in a derivation of affirmation  $a \text{ affirms } P_2$  by the same principal  $a$ .

## 3. Distributed policies and proof-checking

In our setup, each site is governed by an administrative principal who decides the authorization policy at that site. We shall use the notation  $\bar{w}$  to denote the governing principal at site  $w$ . We model the local policy as a context of propositions declared to be true by the governing principal. From a global perspective, a proposition  $P$  declared by the principal  $\bar{w}$  has the force of the proposition  $\langle \bar{w} \rangle P$ . Thus the local policy at  $w$ , called  $\Phi_w$ , is a context of the form:

$$\Phi_w ::= \langle \bar{w} \rangle P_1 \text{ true}, \dots, \langle \bar{w} \rangle P_n \text{ true}$$

A proof presented to a reference monitor may contain facts from local policies at various sites. Logically, such a proof is valid under the union of local policies at all sites. We refer to this union as the *amalgamated* authorization policy and denote it by

$$\Phi = \bigcup_{w \in \text{Wld}} \Phi_w$$

where  $\text{Wld}$  is the set of all sites. We use the same notation for amalgamated policy as a proposition context from Section 2 in order to highlight that a policy is just a context in the logic. In implementations, evidence for the local policy at a site is provided by digitally signed certificates, which can be verified by reference monitors at any site. A certificate for proposition  $P$  signed by a principal  $a$  is unforgeable and establishes the judgment  $\langle a \rangle P \text{ true}$  irrefutably. Verification of a proof constructed from such evidences requires checking the validity of these certificates. We view these evidences as implementing the amalgamated authorization policy. Therefore, our formalism directly uses the full policy  $\Phi$  for the purpose of proof-checking at reference monitors.

In practical distributed authorization systems for trust management (Blaze et al. 1996) and other applications (Bauer et al. 2005) policy statements made by an administrator at its site  $w$  are often cached at another site  $w'$  in order to aid distributed construction. For example, consider the authorization policy at  $\text{univserver}$  from Figure 1. The evidence for  $\langle \text{univ} \rangle \text{is\_student}(\text{alice}, \text{univ})$  may be cached at Alice's machine browser. In this paper, we do not consider caching of proofs

at sites. The policy available at a site is restricted to only consist of declarations made by the principal governing that site.

## 4. Overview of PCML<sub>5</sub>

We shall now informally discuss the key ideas behind PCML<sub>5</sub>. There are three main aspects of PCML<sub>5</sub>:

1. PCML<sub>5</sub> is a language for writing distributed programs that interact with resources which are distributed across different administrative domains. The type system tracks the site where a computation is executed. This allows the language to enforce the requirement that a resource can be accessed only from the site where it is situated.
2. PCML<sub>5</sub> uses an authorization logic as a part of its type system. The logic is used to express within the language authorization proof obligations that are imposed by reference monitors.
3. PCML<sub>5</sub> allows construction of proofs by combining the results of a distributed proof search with proof constructors from authorization logic. The type system statically ensures that proofs passed to a reference monitor satisfy its proof specifications. This ensures that a call made to a reference monitor from a well-typed PCML<sub>5</sub> program never incurs a runtime error due to failure of verification of proof.

We now elaborate on each of these aspects.

### 4.1 Distribution of computation

PCML<sub>5</sub> is based on ML<sub>5</sub> (Murphy 2008; Murphy et al. 2004), which is a distributed programming language based on propositions-as-types interpretation of the modal logic Intuitionistic S5 where modal worlds are interpreted as sites in a distributed system. The central idea that PCML<sub>5</sub> inherits from ML<sub>5</sub> is that terms are classified using types *relative* to sites. The typing judgment  $\Gamma \vdash m : A@w$  means that under the hypotheses about variables in  $\Gamma$ , the term  $m$  is typed as  $A$  for the site  $w$  where it is to be evaluated.

The idea of typing relative to worlds is central to writing distributed programs since it allows the type system to track where different pieces of code are meant to execute. A term  $m$  for the site  $w'$  can be remotely executed from a world  $w$  by doing a `get[w']m` at  $w$ . This causes computation to move to  $w'$  to execute  $m$ , and the result is brought back to  $w$ , provided type of the result is mobile. We discuss mobility later in Section 5.3 after having introduced the formalism.

### 4.2 Authorization logic

PCML<sub>5</sub> incorporates an authorization logic as static constructors using higher-order abstract syntax. Constructors are classified by kinds. Propositions belong to the kind `Prop` and proofs of a proposition  $P$  are kinded as `Prf(A)` where  $A$  is the representation of the proposition  $P$ . Proof of an affirmation judgment  $a$  affirms  $P$  is represented as a constructor of kind `Affirms(A1, A2)`, where  $A_1$  represents the principal  $a$ , and  $A_2$  is the representation of  $P$ . As we shall see later, this translation provides a compositional bijection between propositions, proofs, and affirmations, on one hand, and constructors of kind families `Prop`, `Prf`, and `Affirms`, on the other.

The main requirement from this translation is to be able to reflect the consequence relation of logic in an adequate manner using PCML<sub>5</sub> as is done in LF (Harper et al. 1993). Thus if there is a proof of  $P_2$  hypothetical in a proof of  $P_1$  in the logic, one can write a constructor of kind `Prf(A2)` with a free variable of kind `Prf(A1)` in the language, and vice versa, where  $A_i$ 's are the representations of propositions  $P_i$ 's. We shall make this relation precise once we introduce the formalism.

In order to ensure that adequacy is without doubt, we follow a phase distinction (Harper et al. 1990) between static and dynamic phases of PCML<sub>5</sub>. The constructor level is constrained to be pure in the sense that constructors do not depend on runtime terms. This way terms can be arbitrarily effectful, but they do not affect the types in any way.

### 4.3 Representation of principals and resources

Principals and resources appear in propositions and proofs, and also occur as runtime values. In order to maintain a phase distinction, we differentiate between the runtime and the compile-time representations of principals and resources, yielding a dual representation for them:

1. They appear as static constructors of kinds `Prin` and `Db` resp. These constructors appear in types and propositions, such as `mayrd(d, p)`.
2. They are also represented as runtime values that mediate access to them. The type system tracks the identities of these runtime values by linking them to their static counterparts using their types. The type `db(A)` is a singleton type that represents runtime database values indexed by  $A$ , which has the kind `Db`. We refer to the static representations, like  $A$ , as indices.

The two representations are created together. For databases, this is done using a primitive operation `db.open`:  $\exists\alpha::\text{Db.db}(\alpha)$ .

In the case of principals, we need runtime representations of only those principals that are authenticated. The type `Iam(A)` represents the type of authentication tokens for  $A::\text{Prin}$ .

### 4.4 Runtime acquisition of proofs

PCML<sub>5</sub> provides a primitive operation `acquire[A]{α.m1 | m2}` to enable querying for proofs using local security policies. When executed at a site  $w$ , the primitive does a local proof search for a proof of  $A$  using only  $\Phi_w$ , the policy at  $w$ . The language itself does not specify or depend on any particular proof search strategy. It is however reasonable to assume that any strategy used in an implementation should at least succeed in finding a proof when one is present as an atomic fact in the policy.

The key idea behind `acquire` is that it enables acquiring facts from security policies and incorporating them in the program at runtime. The results of the runtime query, if successful, discharge the hypothesis  $\alpha$  bound in the branch  $m_1$ . In case the query succeeds, the result is substituted for  $\alpha$  in  $m_1$  and execution proceeds with  $m_1$ . In case the query fails,  $m_2$  is executed.

## 5. Syntax and static semantics

Figure 3 presents the syntax of PCML<sub>5</sub>. We divide the syntax into three levels: runtime terms, constructors, and kinds. Runtime terms are classified by types, which are a subset of constructors, and constructors are classified by kinds. The language is parameterized by two forms of signatures:

1.  $\Sigma_c$ : is the signature that introduces constructor constants,  $\mathbf{a} \Rightarrow K$ . Besides constant principals and databases, this signature introduces the encoding of authorization logic.
2.  $\Sigma_t$ : is the term signature and it gives types to externally implemented API constants  $\mathbf{c} : \forall(\alpha_1 \Rightarrow K_1, \dots, \alpha_n \Rightarrow K_n). A @ w$ .

In order to avoid normalization at the constructor level, we present only the normal (canonical) forms in the style of Canonical LF (Harper and Licata 2007). The idea is to exclude  $\beta$  and  $\eta$  redexes altogether by rearranging the constructors as *neutral* ( $N$ ) and *normal* ( $A$ ) forms so that all occurrences of eliminations appear before any introduction. Since we admit only canonical constructors,

Kinds	$K$	$::=$	Type   Wld   Prop   Prin   Prf( $A$ )   Affirms( $A_1, A_2$ )   Db   $\Pi\alpha::K_1.K_2$
Constructors (Neutral forms)	$N$	$::=$	$\alpha$   <b>a</b>   $A_1 \rightarrow A_2$   $A_1 \times A_2$   $A_1 + A_2$   <b>unit</b>   $\exists\alpha::K.A$   <b>Iam</b> ( $A$ )   ( $N A$ )
(Normal forms)	$A, w, k$	$::=$	$N$   $\lambda\alpha::K.A$
Polytypes	$\tau$	$::=$	$\forall(\alpha_1:K_1, \dots, \alpha_n:K_n)A$
Terms	$m$	$::=$	$x$   $\lambda x:A.m$   ( $m_1 m_2$ )   $\langle m_1, m_2 \rangle$   $\pi_1 m$   $\pi_2 m$   <b>inl</b> $m$   <b>inr</b> $m$   <b>case</b> $m$ of $x.(m_1   m_2)$   $\langle \rangle$   <b>let</b> $x = m_1$ in $m_2$   <b>get</b> [ $w$ ] $m$   <b>acquire</b> [ $A$ ]   <b>authenticate</b>   <b>c</b> [ $A_1, \dots, A_n$ ]( $m$ )   $\{\alpha = A; m : A'\}$   <b>open</b> $\{\alpha, x\} = m_1$ in $m_2$   <b>iam</b> [ <b>a</b> ]
Constructor signature	$\Sigma_c$	$::=$	$\cdot$   $\Sigma_c, \mathbf{a}:K$
Term signature	$\Sigma_t$	$::=$	$\cdot$   $\Sigma_t, \mathbf{c}:\tau@w$
Constructor context	$\Delta$	$::=$	$\cdot$   $\Delta, \alpha \Rightarrow K$
Term context	$\Gamma$	$::=$	$\cdot$   $\Gamma, x:A@w$

**Figure 3.** PCML<sub>5</sub> syntax: Shaded parts are PCML<sub>5</sub>'s additions to ML<sub>5</sub>

we use hereditary substitution (Watkins et al. 2002) to keep constructors canonical upon substitution, details of which are omitted here.

### 5.1 Judgment forms

The basic judgment forms used in static semantics are:

$K$ kind	$K$ is a well-formed kind
$A \Leftarrow K$	$A$ is checked for the kind $K$
$N \Rightarrow K$	$K$ is synthesized as the kind of $N$
$m : A@w$	$m$ is of type $A$ at the world $w$

We maintain a phase-distinction between static and dynamic phases in PCML<sub>5</sub>. This means that runtime terms do not appear in constructors. Accordingly, we divide the context into a static part,  $\Delta$ , which gives kinds to constructor variables  $\alpha$ , and a dynamic part  $\Gamma$  which types term variables  $x$ . Because of the phase-distinction, the dynamic context is needed only for term typing judgments. Selected rules defining well-formed kinds are shown here:

$$\frac{\Delta \vdash_{\Sigma_c} A \Leftarrow \text{Prop}}{\Delta \vdash_{\Sigma_c} \text{Prf}(A) \text{ kind}} (\text{Proof})$$

$$\frac{\Delta \vdash_{\Sigma_c} A_1 \Leftarrow \text{Prin} \quad \Delta \vdash_{\Sigma_c} A_2 \Leftarrow \text{Prop}}{\Delta \vdash_{\Sigma_c} \text{Affirms}(A_1, A_2) \text{ kind}} (\text{Affirms})$$

The kind  $\text{Prf}(A)$  classifies proofs of proposition  $A$ , and the kind  $\text{Affirms}(A_1, A_2)$  classifies evidences for the judgment  $A_1$  affirms  $A_2$  from Section 2.

All types except one in PCML<sub>5</sub> are inherited from ML<sub>5</sub>. PCML<sub>5</sub> adds an abstract type  $\text{Iam}(A)$  of authentication credentials of the principal  $A$ . We shall return to this type in Section 5.4 when we discuss authentication.

$$\frac{\Delta \vdash_{\Sigma_c} A \Leftarrow \text{Prin}}{\Delta \vdash_{\Sigma_c} \text{Iam}(A) \Rightarrow \text{Type}} (\text{auth})$$

### 5.2 Representing authorization logic

The kind level of PCML<sub>5</sub> is sufficiently rich to encode the authorization logic from Section 2 using higher-order abstract syntax. For instance, the proposition  $P_1 \supset P_2$  can be encoded as the constructor  $\text{imp } A_1 A_2$ , where  $\text{imp} \Rightarrow \text{Prop} \rightarrow \text{Prop} \rightarrow \text{Prop}$  is a constructor constant. Propositions in the logic become constructors of kind  $\text{Prop}$  in PCML<sub>5</sub>. Proofs, i.e., derivations of the judgment  $\Delta; \Phi \vdash_{\Sigma}^c P \text{ true}$  get translated into open constructors of kind  $\text{Prf}(A)$ , where  $A$  is translation of  $P$ , under the context  $\Delta$  that represents  $\Delta, \Phi$ . Affirmations, i.e., derivations of judgment  $\Delta; \Phi \vdash_{\Sigma}^c a \text{ affirms } P$  are translated into open constructors of kind  $\text{Affirms}(A_1, A_2)$  where  $A_1$  and  $A_2$  are representations of  $a$  and  $P$  respectively.

The central principle behind the translation of logic is that hypothetical reasoning in PCML<sub>5</sub> represents the consequence relation of the logic. This enables constructing proofs in a program that are hypothetical in proofs of certain propositions (such as those where a runtime policy query is involved), and later discharging those hypotheses using facts from authorization policies. In order to state an adequacy theorem, we first need some terminology: We use the notation  $\Sigma_c^+$  to denote the signature consisting of constants in  $\Sigma_c$  and the constants defining the embedding of authorization logic in PCML<sub>5</sub>. The following judgments show representation of an entity from the logic into a construct in the language (rules are omitted here but can be found in our full report (Avijit et al. 2009)):

$$\begin{aligned} \Delta \vdash_{\Sigma_c^+}^c P \text{ prop} &\gg \Delta \vdash_{\Sigma_c^+} A \Leftarrow \text{Prop} && \text{Propositions} \\ \Delta; \Phi \vdash_{\Sigma_c^+}^c P \text{ true} &\gg \Delta \vdash_{\Sigma_c^+} A_1 \Leftarrow \text{Prf}(A_2) && \text{Proofs} \\ &\Delta; \Phi \vdash_{\Sigma_c^+}^c a \text{ affirms } P \gg && \\ \Delta \vdash_{\Sigma_c^+} A &\Leftarrow \text{Affirms}(A_1, A_2) && \text{Affirmations} \end{aligned}$$

We lift the translation to contexts and signatures in the straightforward manner, representing a hypothesis  $P \text{ true}$  as  $\alpha \Rightarrow \text{Prf}(A)$  where  $A$  is the representation of  $P$ , and  $\alpha$  is a fresh variable. In the

following theorem, we shall assume that  $\Sigma_c$  is a representation of  $\Sigma$  and  $\Delta_1, \Delta_2$  represent the contexts  $\Delta$  and  $\Phi$  resp..

**Theorem 5.1** (Adequacy). *Let  $\Delta \vdash_{\Sigma}^c P_1 \text{ prop} \gg \Delta_1 \vdash_{\Sigma_c^+} A_1 \Leftarrow \text{Prop}$ , and  $\Delta \vdash_{\Sigma}^c P_2 \text{ prop} \gg \Delta_1 \vdash_{\Sigma_c^+} A_2 \Leftarrow \text{Prop}$ . Then  $\Delta; \Phi, P_1 \text{ true} \vdash_{\Sigma}^c P_2 \text{ true}$  iff there exists  $A_3$  such that  $\Delta_1, \Delta_2, \alpha \Rightarrow \text{Prf}(A_1) \vdash_{\Sigma_c^+} A_3 \Leftarrow \text{Prf}(A_2)$ .*

In the rest of the paper, we freely use logical notation instead of its translation for readability. For instance, we use  $A_1 \supset A_2$  to mean  $\text{imp } A_1 \ A_2$ . The notation  $\hat{\Phi}$  denotes the representation of the authorization policy  $\Phi$  in the language.

### 5.3 Situated typing and distributed computation

PCML<sub>5</sub> inherits the notion of typing relative to worlds from ML<sub>5</sub>. Here we briefly recap the central ideas from ML<sub>5</sub>: distribution of computation and mobility of data.

The main motivation behind situated typing in ML<sub>5</sub> is to express locality of resources, and use it to restrain where a piece of code can run, with the idea being that a resource can only be accessed by code running at the same location. Distribution of execution is achieved by moving result of a computation from one location to another using a `get`. This is expressed in the following typing rule:

$$\frac{\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t} m : A@w' \quad \Delta \vdash_{\Sigma_c} A \text{ mobile}}{\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t} \text{get}[w']m : A@w} (\text{get})$$

If  $m$  is well-typed for  $w'$ , then it may be executed using a remote call from another world  $w$ . ML<sub>5</sub> however restricts the remote invocation to terms of *mobile* types only. The first requirement for a type to be mobile is that all values of that type should be typable at all worlds. This ensures that when the result of  $m$  is brought back from  $w'$  to  $w$ , the value makes sense at  $w$ . Base types such as `string` are mobile as their values can be typed at any site. Types representing local resources such as reference cells are not mobile. Function types are also not mobile in ML<sub>5</sub> because, by design, computations are never shipped across sites. Instead code meant to be run at a site is compiled and stored at that site. Only the locus of execution shifts from one site to another, as in a remote procedure call.

PCML<sub>5</sub> adds a new type  $\text{Iam}(A)$  for typing authentication credentials to the types inherited from ML<sub>5</sub>. The type  $\text{Iam}(A)$  is mobile because authentication credentials are typable at all worlds. An existential type  $\exists \alpha :: K.A$  is mobile if the type  $A$  of the packed value is mobile, regardless of  $\alpha$ . Thus, for instance, the type  $\exists \alpha :: \text{Prin.Iam}(\alpha)$  is mobile. All constructors are considered mobile in the sense that their kinding is independent of their location. In particular, principal and resource indices, and proofs are typable at all sites.

### 5.4 Authentication

For each constant principal  $\mathbf{a}$  declared in the constructor signature  $\Sigma_c$ , we use a constant `iam[a]` as its authentication token which is typed as  $\text{Iam}(\mathbf{a})$ .

$$\frac{\mathbf{a} \Rightarrow \text{Prin} \in \Sigma_c}{\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t} \text{iam}[\mathbf{a}] : \text{Iam}(\mathbf{a})@w} (\text{Iam-I})$$

PCML<sub>5</sub> provides an abstract operation `authenticate` to authenticate the program on behalf of a principal. In an implementation, this could be achieved through well-established protocols for authentication, e.g. Kerberos (Neuman and Ts'o 1994). The operation returns a principal together with an authentication token that serves as a proof of the principal having been authenticated.

The operation `authenticate` is typed as:

$$\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t} \text{authenticate} : (\exists \alpha :: \text{Prin.Iam}(\alpha)) \text{ option}@w$$

The primitive is typed using an option type because it may fail to return any meaningful value. Let us consider the case when the primitive is successful. The result is typed as  $\exists \alpha :: \text{Prin.Iam}(\alpha)$  in this case. Notice that the authentication credential is typed as  $\text{Iam}(\alpha)$  linking it to the abstract component  $\alpha$  of the abstraction. The runtime term of type  $\text{Iam}(\alpha)$  and the static proxy  $\alpha$  can be viewed as dual representations of an authenticated principal. Because of the abstraction, the static component uniquely represents a particular instance of opening up the package, since it is assumed to be different from everything else. Thus a credential of type  $\text{Iam}(\alpha)$  represents a particular instance of doing authentication in the program.

### 5.5 Authorization

An authorization proof is constructed by doing proof search using local security policies at various sites. This is done using a primitive operation `acquire[A]`. The type system tracks the kinds of proofs starting with `acquire[A]` to API calls where the proofs are used.

#### 5.5.1 Distributed proof acquisition

`acquire[A]` optionally returns a proof of  $A$ . We use an existential type to model this.

$$\frac{\Delta \vdash_{\Sigma_c} A \Leftarrow \text{Prop}}{\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t} \text{acquire}[A] : (\exists \alpha :: \text{Prf}(A).\text{unit}) \text{ option}@w}$$

The success case is typed as  $\exists \alpha :: \text{Prf}(A).\text{unit}$ . When this existential type is eliminated as `open { $\alpha, x$ } =  $m_1$  in  $m_2$` , an assumption  $\alpha \Rightarrow \text{Prf}(A)$  is introduced in the scope of  $m_2$  during type-checking. We have mentioned before that all kinds are mobile meaning that constructors are typable at all sites. Thus the variable  $\alpha$  stands for a globally-valid proof of proposition  $A$  in  $m_2$  even though it is discharged using a proof obtained locally at a site.

#### 5.5.2 Proof checking at API call sites

API constants,  $\mathbf{c}$ , are introduced using the term signature  $\Sigma_t$ . In order to call an API, its polymorphic arguments have to be fully instantiated. The typing rule statically reflects the proof-verification that happens at the reference monitor during such a call. All the constructor arguments are typed statically according to the type of the API constant.

$$\frac{\Sigma_t(\mathbf{c}) = \forall \langle \alpha_1 : K_1, \dots, \alpha_n : K_n \rangle (A \rightarrow A') \quad \Delta \vdash_{\Sigma_c} A_k \Leftarrow [A_1/\alpha_1] \dots [A_{k-1}/\alpha_{k-1}] K_k \ (k = 1..n) \quad \Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t} m : [A_1/\alpha_1] \dots [A_n/\alpha_n] A@w}{\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t} \mathbf{c}[A_1, \dots, A_n](m) : [A_1/\alpha_1] \dots [A_n/\alpha_n] A'@w}$$

Notice that for an API call to be well-typed, the constructor arguments should be of the proper kind under the context  $\Delta$  as per the type specification of the API. At runtime, the hypotheses are discharged using closed constructors. The hypothetical judgment ensures that typing of arguments is preserved upon discharging these hypotheses. This in turn ensures success of runtime verification of these proofs at reference monitors. We elaborate on proof verification at reference monitors when we give the dynamics of API calls in Section 6.6.

An important technical detail here is the dependency among the constructor arguments. Consider the type of an API constant:

$$\forall \langle \alpha_1 : K_1, \dots, \alpha_n : K_n \rangle (A \rightarrow A')$$

The kind of an argument potentially depends on all previous arguments. This is manifest in the second premiss of the typing rule:  $\Delta \vdash_{\Sigma_c} A_k \Leftarrow [A_1/\alpha_1] \dots [A_{k-1}/\alpha_{k-1}] K_k$  that substitutes arguments  $A_1$  through  $A_{k-1}$  in  $K_k$  while checking the kind of  $A_j$ .

## 6. Runtime semantics

Term evaluation in PCML<sub>5</sub> has the following aspects, in addition to distribution:

1. **Principals, databases and APIs:** The transition system is parameterized by a constructor signature  $\Sigma_c$ , and an API signature  $\Sigma_t$ . The signatures remain fixed throughout the evaluation of the program. We assume a fixed set of principal and database indices, introduced through the constructor signature  $\Sigma_c$ . In addition to principals and databases,  $\Sigma_c$  also introduces representation of an authorization logic.

2. **Execution under a security policy:** Each location has a fixed authorization policy associated with it. The policy at site  $w$  consists of assertions made by the principal  $\bar{w}$  who governs that site. We formulate the local policy at site  $w$  as:

$$\hat{\Phi}_w ::= \alpha_1 \Rightarrow \text{Prf}(\langle \bar{w} \rangle A_1), \dots, \alpha_n \Rightarrow \text{Prf}(\langle \bar{w} \rangle A_n)$$

where  $\alpha_1, \dots, \alpha_n$  are chosen fresh.

We assume that all local policies are well-formed with respect to the signature  $\Sigma_c$  that defines the authorization logic. As mentioned before, the amalgamated authorization policy  $\hat{\Phi}$  is the union of all local policies.

3. **Authorization checks at reference monitors:** An API call  $\mathbf{c}[A_1, \dots, A_n](m)$  results in the reference monitor verifying  $A_1$  through  $A_n$  (as shown by Rule (api-reduce) later). Some of these  $A_i$ 's may be proofs constructed by a distributed acquisition of local policy assertions from various sites. As discussed in Section 3, local assertions make sense globally and are implemented in an unforgeable and irrefutable manner by signing the asserted proposition with the principal's signing key. In the formalism, we directly use the global security policy  $\hat{\Phi}$  to serve as the context for type-checking the proofs at API call-sites at runtime.

4. **Authentication and active principals:** As evaluation progresses, principals may be authenticated using the primitive `authenticate`. Principals who have been authenticated during a program run are called *active* principals. We keep a record of all active principals as a set of principals  $\mathcal{A} \subseteq \{\mathbf{a} \mid \mathbf{a} \Rightarrow \text{Prin} \in \Sigma_c\}$ . The program starts execution with  $\mathcal{A} = \{\}$ , i.e. no principal is assumed to be authenticated at the beginning of the program.

### 6.1 Judgment forms

We describe the dynamic semantics using a transition relation between terms. We use the following judgment forms:

1.  $m; \mathcal{A} \mapsto_{w; \Sigma_c; \Sigma_t; \hat{\Phi}} m'; \mathcal{A}'$ : Under the signatures  $\Sigma_c; \Sigma_t$  and the amalgamated authorization policy  $\hat{\Phi}$ , the term  $m$ , executed at world  $w$ , steps to the term  $m'$  in a single transition; the set of active principals changes from  $\mathcal{A}$  to  $\mathcal{A}'$ . Since  $\Sigma_c; \Sigma_t$  and  $\hat{\Phi}$  remain fixed during evaluation, we often omit them while presenting the transition system.
2.  $m \text{ val}_{\mathcal{A}}$ : means that the term  $m$  is a value under the runtime record  $\mathcal{A}$ , and is not evaluated further.

### 6.2 The value judgment

The value judgment  $m \text{ val}_{\mathcal{A}}$  defines values with respect to the runtime state  $\mathcal{A}$ . The rules for determining values of the types such as function, product, sum and existential types are straightforward; we adopt eager evaluation for products, sums and existential packages.

The crucial part of this definition is the case for values of type  $\text{Iam}(A)$ . An authentication token  $\text{iam}[\mathbf{a}]$  is a value only if  $\mathbf{a}$  is an authenticated principal, i.e. if  $\mathbf{a}$  is part of the active set  $\mathcal{A}$ . The need

for the requirement that  $\mathbf{a}$  be in  $\mathcal{A}$  is explained in Section 6.6, where the value  $\text{iam}[\mathbf{a}]$  is used at an API call-site.

$$\frac{\mathbf{a} \in \mathcal{A}}{\text{iam}[\mathbf{a}] \text{ val}_{\mathcal{A}}} (\text{Iam-V})$$

### 6.3 Distribution

A remote call  $\text{get}[w']m$  is evaluated at a world  $w$  as follows:

$$\frac{m; \mathcal{A} \mapsto_{w'} m'; \mathcal{A}'}{\text{get}[w']m; \mathcal{A} \mapsto_w \text{get}[w']m'; \mathcal{A}'} (\text{get-eval})$$

$$\frac{m \text{ val}_{\mathcal{A}}}{\text{get}[w']m; \mathcal{A} \mapsto_w m; \mathcal{A}} (\text{get-reduce})$$

Rule (get-eval) expresses the remote execution of  $m$  at the world  $w'$ . In case  $m$  is a value, it is brought to  $w$  from  $w'$  (Rule (get-reduce)). This transfer is safe, i.e.,  $m$  is well-typed at  $w$  because the type system guarantees that  $m$  has a mobile type.

### 6.4 Local policy acquisition

The `acquire[A]` primitive does a proof search for the proposition  $A$  using the local policy  $\hat{\Phi}_w$  at the site where it is executed. This proof search may fail. We do not stipulate the procedure used for local theorem-proving. For this reason, we model this primitive using a non-deterministic step: the Rule (acq-succ) non-deterministically chooses a proof  $A'$  such that  $A'$  has the kind  $\text{Prf}(A)$  under the assumptions  $\hat{\Phi}_w$ .

$$\frac{\hat{\Phi}_w \vdash_{\Sigma_c} A' \Leftarrow \text{Prf}(A)}{\text{acquire}[A]; \mathcal{A} \mapsto_w \text{SOME } \{\alpha = A'; \langle \rangle : \text{unit}\}; \mathcal{A}} (\text{acq-succ})$$

The failure case is modeled by a transition to the value NONE:

$$\frac{}{\text{acquire}[A]; \mathcal{A} \mapsto_w \text{NONE}; \mathcal{A}} (\text{acq-fail})$$

### 6.5 Authentication

Authentication, if successful, results in the production of a token for the authenticated principal. We use a non-deterministic rule which may result in a token for any principal that has been declared in the signature  $\Sigma_c$ .

$$\frac{\mathbf{a}::\text{Prin} \in \Sigma_c}{\text{authenticate}; \mathcal{A} \mapsto_w \text{SOME } \{\alpha = \mathbf{a}; \text{iam}[\mathbf{a}] : \text{Iam}(\alpha)\}; \mathcal{A} \cup \{\mathbf{a}\}} (\text{auth-succ})$$

The record  $\mathcal{A}$  is augmented with  $\mathbf{a}$  to note this authentication.

In case of failure, `authenticate` returns NONE, and  $\mathcal{A}$  is left unchanged:

$$\frac{}{\text{authenticate}; \mathcal{A} \mapsto_w \text{NONE}; \mathcal{A}} (\text{auth-fail})$$

An important technical detail to note here is that we require  $\Sigma_c$  to contain all possible principals. This is manifest in the Rule (auth-succ) where `authenticate` does not generate a new principal identity but simply picks one up from  $\Sigma_c$ .

### 6.6 API calls to reference monitors

API constants represent externally implemented functions. We model their behavior using a non-deterministic transition to an arbitrary term of the appropriate type.

$$\frac{m; \mathcal{A} \mapsto_w m'; \mathcal{A}'}{\mathbf{c}[A_1, \dots, A_n](m); \mathcal{A} \mapsto_w \mathbf{c}[A_1, \dots, A_n](m'); \mathcal{A}'} (\text{api-eval})$$

$$\frac{\Sigma_t(\mathbf{c}) = \forall (\alpha_1::K_1, \dots, \alpha_n::K_n) A \rightarrow A' @w}{\forall i \in [0..n-1] \hat{\Phi} \vdash_{\Sigma_c} A_{i+1} \Leftarrow [A_1/\alpha_1] \dots [A_i/\alpha_i] K_{i+1}} (\text{api-reduce})$$

$$\frac{m_1 \text{ val}_{\mathcal{A}}}{\mathbf{c}[A_1, \dots, A_n](m_1); \mathcal{A} \mapsto_w \Sigma_c; \Sigma_t; \hat{\Phi} m_2; \mathcal{A}} (\text{api-reduce})$$

The Rule (`api-reduce`) illustrates the central actions that take place during a call to a reference monitor:

1. **Authorization checks** The central point of PCA is that reference monitors for APIs check proofs. Proofs are passed as constructor arguments to the API. In order to verify them, the reference monitor needs the amalgamated policy  $\hat{\Phi}$ .

Consider the Rule (`api-reduce`). Before evaluating the API call, the reference monitor type-checks all the constructor arguments  $A_1 \dots A_n$  under  $\hat{\Phi}$ . For well-typed programs, the type system (cf. Section 5.5.2) guarantees that all type checks succeed.

2. **Authentication checks** In addition to verifying proofs, the reference monitor checks authentication credentials. APIs for such monitors are polymorphic in the authenticated principal  $\alpha$ , and require an extra parameter of type  $\text{Iam}(\alpha)$ . In order to verify an authentication credential  $\text{iam}[a]$ , the reference monitor simply checks whether  $a \in \mathcal{A}$ , where  $\mathcal{A}$  is the active set at the time of API call. In the formalism, authentication checks are reflected implicitly in the value judgment. An authentication token  $\text{iam}[a]$  is considered a value under an active set  $\mathcal{A}$  only if  $a \in \mathcal{A}$  (Rule `Iam-V`). This way, we reduce authentication checking to checking that the argument to the API call ( $m_1$  in Rule (`api-reduce`)) is a value.

## 6.7 PCA runtime errors

A term may incur a runtime fault at a reference monitor in the following two ways:

1. Either the proofs passed to the monitor API do not type-check,
2. An authentication credential is not valid, in the sense that the purported principal does not appear in the authentication record  $\mathcal{A}$ .

We formalize these errors explicitly using the judgment  $m \uparrow_{\mathcal{A}}$ . This judgment formalizes the intuition that  $m$  is not a value, and  $m; \mathcal{A} \not\vdash_w$ .

$$\frac{\mathbf{a} \notin \mathcal{A}}{\text{iam}[\mathbf{a}] \uparrow_{\mathcal{A}}} (\text{iam}^\dagger)$$

$$\frac{\hat{\Phi} \not\vdash_{\Sigma_c} A_i \Leftarrow [A_1/\alpha_1] \dots [A_{i-1}/\alpha_{i-1}] K_i}{\mathbf{c}[A_1, \dots, A_n](m) \uparrow_{\mathcal{A}}} (\mathbf{c}^\dagger)$$

In addition, we have rules to propagate errors through evaluation of other terms, of which we present only a sample here:

$$\frac{m_1 \uparrow_{\mathcal{A}}}{\langle m_1, m_2 \rangle \uparrow_{\mathcal{A}}} (\text{pair}_1^\dagger) \quad \frac{m_1 \text{val}_{\mathcal{A}} \quad m_2 \uparrow_{\mathcal{A}}}{\langle m_1, m_2 \rangle \uparrow_{\mathcal{A}}} (\text{pair}_2^\dagger)$$

The rules for the error judgment follow the evaluation order. A pair  $\langle m_1, m_2 \rangle$  can incur an error at two instances: either when  $m_1$  which is evaluated first incurs an error (as in Rule (`pair1†`)), or when  $m_1$  has been evaluated to a value and  $m_2$  incurs an error (as shown in Rule (`pair2†`)).

## 7. Metatheory

We now prove the central result of this paper: a class of well-typed PCML<sub>5</sub> programs do not incur runtime faults at reference monitors. We have already summarized a notion of *error* states (Section 6.7) that formalizes exactly the runtime failures we are trying to avoid.

We start by proving progress and preservation for PCML<sub>5</sub>. This means that well-typed terms do not get stuck. However, static typing alone does not rule out authentication errors, i.e. failures that happen when the reference monitor gets a token  $\text{iam}[a]$  and  $a$  is not

an active principal. This happens because the type system does not track runtime authentication events, i.e. calls to `authenticate`.

Throughout this section we shall assume that  $\Sigma_c; \Sigma_t$  are well-formed signatures and  $\hat{\Phi}$  is an amalgamated security policy which is well-formed under  $\Sigma_c$ .

### 7.1 Type safety

**Theorem 7.1** (Progress). *Let  $\mathcal{A}$  be a set of active principals from  $\Sigma_c$ . If  $\Delta; \cdot \vdash_{\Sigma_c; \Sigma_t} m : A@w$ , then*

1. either  $m \text{val}_{\mathcal{A}}$ ,
2. or  $\exists m', \mathcal{A}'. m; \mathcal{A} \mapsto_w^{\Sigma_c; \Sigma_t; \hat{\Phi}} m'; \mathcal{A}'$ ,
3. or  $m \uparrow_{\mathcal{A}}$  without using the Rule (`c†`).

**Theorem 7.2** (Preservation of typing). *Assume that all API constants declared in  $\Sigma_t$  preserve the typing. If  $m; \mathcal{A} \mapsto_w^{\Sigma_c; \Sigma_t; \hat{\Phi}} m'; \mathcal{A}'$  and  $\Delta; \cdot \vdash_{\Sigma_c; \Sigma_t} m : A@w$ , then  $\Delta, \hat{\Phi}; \cdot \vdash_{\Sigma_c; \Sigma_t} m' : A@w$ .*

In order to guarantee that terms do not incur authentication errors in addition to not getting stuck, we define a notion of *authentication safety* for terms with respect to a set of active principals. A well-typed term is regarded as authentication safe if it never evaluates to an unauthenticated token. We prove that an authentication safe term does not evaluate to an error state. In addition, we prove that authentication safety is preserved by evaluation; these two theorems together ensure that authentication safe terms do not incur runtime access violations at reference monitors. Finally we show how authentication safety for the case where the set of active principals is empty can be enforced using the type system.

### 7.2 Authentication history

The runtime semantic ensures that `authenticate` is the only way in which a new authentication token can be generated. We further wish to enforce that a token  $\text{iam}[a]$  appears in a term only when the associated history  $\mathcal{A}$  mentions  $a$  as one of the authenticated principals. We use a judgment of the form  $\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t}^{\mathcal{A}} m : A@w$  to denote that, in addition to the term being well-typed, all authentication tokens in  $m$  have the corresponding principal recorded in  $\mathcal{A}$ . This judgment resembles the static typing judgment  $\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t} m : A@w$  except for the case of  $\text{iam}[a]$ , where we demand that the principal  $a$  be present in  $\mathcal{A}$ . We present a sample of rules here:

$$\frac{\mathbf{a}::\text{Prin} \in \Sigma_c \quad \mathbf{a} \in \mathcal{A}}{\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t}^{\mathcal{A}} \text{iam}[\mathbf{a}] : \text{Iam}(\mathbf{a})@w}$$

$$\frac{\Delta; \Gamma, x:A@w \vdash_{\Sigma_c; \Sigma_t}^{\mathcal{A}} m : A'@w}{\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t}^{\mathcal{A}} \lambda x:A. m : A \rightarrow A'@w}$$

$$\frac{\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t}^{\mathcal{A}} m_1 : A_1 \rightarrow A_2@w \quad \Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t}^{\mathcal{A}} m_2 : A_1@w}{\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t}^{\mathcal{A}} (m_1 m_2) : A_2@w}$$

**Theorem 7.3.** *If  $\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t}^{\mathcal{A}} m : A@w$  then  $\Delta; \Gamma \vdash_{\Sigma_c; \Sigma_t} m : A@w$ .*

We assume that API calls do not directly introduce spurious authentication tokens in their results. The following assumption about API calls summarizes this:

**Definition 7.4** (Authentication safety for APIs). *Let  $\Sigma_t(\mathbf{c}) = \forall(\alpha_1::K_1, \dots, \alpha_n::K_n) A_1 \rightarrow A_2@w$ . Let  $\mathcal{P}$  be the set of constants  $\mathbf{a}$  s.t.  $\mathbf{a}::\text{Prin} \in \Sigma_c$ . The API  $\mathbf{c}$  is authentication safe if the following holds:*

*For all constructors  $B_1, \dots, B_n$  such that  $\cdot \vdash_{\Sigma_c} B_{i+1} \Leftarrow [B_1/\alpha_1] \dots [B_i/\alpha_i] K_{i+1}$ . Let  $v$  be well-typed,*



and be safe for  $\mathcal{A}$  as  $\Delta; \cdot \vdash_{\Sigma_c; \Sigma_t}^A v : [B_1/\alpha_1] \dots [B_n/\alpha_n] A_1 @w$ .  
 If  $\mathbf{c}[B_1, \dots, B_n](v); \mathcal{A} \mapsto_w m; \mathcal{A}$ , then  $m$  is safe for  $\mathcal{A}$ , i.e.  
 $\Delta; \cdot \vdash_{\Sigma_c; \Sigma_t}^A m : [B_1/\alpha_1] \dots [B_n/\alpha_n] A_2 @w$ .

This restriction forces the resulting term to be safe w.r.t. all the possible  $\mathcal{A}$ 's for which the argument  $v$  is safe, thus ensuring that all authentication tokens in the result of the call are inherited from the argument  $v$ , i.e. no new  $\mathbf{iam}[a]$  are introduced in the result.

If all APIs are authentication safe, then authentication safety is preserved by evaluation:

**Theorem 7.5** (Preservation of authentication safety). *Let all API constants declared in  $\Sigma_t$  be safe in the sense of Def. 7.4.*

*If  $m; \mathcal{A} \mapsto_w^{\Sigma_c; \Sigma_t; \hat{\Phi}} m'; \mathcal{A}'$  and  $\Delta; \cdot \vdash_{\Sigma_c; \Sigma_t}^A m : A @w$ , then  $\Delta, \hat{\Phi}; \cdot \vdash_{\Sigma_c; \Sigma_t}^{A'} m' : A @w$ .*

Furthermore, authentication safety ensures that reference monitors cannot reject calls from well-typed and authentication-safe programs.

**Theorem 7.6** (Progress under authentication safety). *Let  $\mathcal{A}$  be a set of active principals from  $\Sigma_c$ . If  $\Delta; \cdot \vdash_{\Sigma_c; \Sigma_t}^A m : A @w$ , then either  $m \mathbf{val}_{\mathcal{A}}$ , or  $\exists m', \mathcal{A}'. m; \mathcal{A} \mapsto_w^{\Sigma_c; \Sigma_t; \hat{\Phi}} m'; \mathcal{A}'$ .*

Notice that in contrast to Theorem 7.1 which has a failure clause ( $m \uparrow_{\mathcal{A}}$ ), Theorem 7.6 excludes the possibility of an authentication-safe term getting stuck and ensures that it is either a value, or that it takes a step to another term. The progress theorem (Theorem 7.6), together with preservation (Theorem 7.5) ensures that an authentication safe, well-typed term never incurs a runtime fault at reference monitors.

### 7.3 Initializing evaluation

A closed well-typed program  $m$  begins evaluation under an empty set of authenticated principals. By virtue of the progress and preservation theorems for authentication safety, in order to ensure that evaluation never incurs a runtime error, it is sufficient to ensure that  $\cdot; \cdot \vdash_{\Sigma_c; \Sigma_t}^{\{\}} m : A @w$ . This can be enforced easily using the typing rules by simply disallowing all constants  $\mathbf{iam}[c]$  as well-typed terms. This is possible because while checking for authentication safety of a term under  $\mathcal{A}$ , all the sub-terms are checked under the same  $\mathcal{A}$ .

## 8. Example

We revisit the example scenario from Figure 1. The program shown in Figure 4 illustrates how a distributed program running at `browser` can access the database `conf` at `acmserver` using a proof constructed using distributed proof-search at `univserver` and `acmserver`.

### 8.1 External API and other declarations

The program in Figure 4 first declares principal, database and world constants as `extern` declarations. The `extern` principal declaration runs an initialization code to acquire the runtime representation of the principal (using the name supplied) and binds it to the variable in the declaration. The classifiers `bytecode` and `javascript` in world declarations determine the language of the generated code for the respective worlds. Lines 10-12 declare the addresses for the three sites.

Lines 14-16 declare predicates using the `extern prop` declaration. A declaration `extern prop (s1, ..., sn) p` declares  $p$  to be a predicate with  $s_1$  through  $s_n$  being the sorts of its arguments.

Next we declare types for the database API at `acmserver`. The constructor `dbhandle` has the kind `Db`  $\rightarrow$  `Type`. The API `db.open` takes a database name of type `string` and returns the runtime

identity of the database packaged with its static proxy, as discussed in Section 4.3. We use the syntax  $\{a:K, A\}$  for the existential type  $\exists a::K.A$ . For simplicity, databases in this paper are simply (*key, value*) pairs, where both the *key* and the *value* are strings. Line 21 introduces the proof-carrying API for reading databases: `db.read` is typed polymorphically in the accessing principal  $p$ , the database  $d$ , and the proof  $f$  of `acm` affirming that  $p$  is allowed to read  $d$ . The argument to `db.read` is a triple consisting of (1) a value of type `Iam(p)` which forms an evidence of  $p$ 's authentication, (2) the runtime structures associated with the database  $d$ , and (3) the key to be read. Both the database operations `db.open`, and `db.read` are local to `acmserver`.

### 8.2 Distributed proof construction

We now discuss the steps involved in reading `conf` at `acmserver`. This is done in Lines 28-57 by the function `readpaper` which is typed at `acmserver`.

The first operation (Lines 29-31) is to authenticate the principal. In case authentication fails, the program raises the `Abort` exception and halts. In case of a successful authentication, `authenticate` returns a value of type  $\exists \alpha::\text{Prin}.\text{Iam}(\alpha)$ , which is bound to `authpkg`.

Let us consider the case when authentication succeeds. In this case `authpkg` gets bound to a package  $\{\alpha = k; \mathbf{iam}[k] : \text{Iam}(\alpha)\}$ , where  $k$  is the authenticated principal. The side-effect of this successful authentication is that  $k$  gets added to the set of active principals.

Assume that the amalgamated policy is given as the following context  $\hat{\Phi}$ , with the variables  $\alpha_1, \alpha_2, \alpha_3$  chosen fresh. The whole program is run under this context. In an implementation, these variables would be bound to digital certificates.

$$\hat{\Phi} = \left\{ \begin{array}{l} \alpha_1 \Rightarrow \langle \text{univ} \rangle \text{is\_student}(\text{alice}, \text{univ}), \\ \alpha_2 \Rightarrow \langle \text{acm} \rangle \text{is\_member}(\text{univ}, \text{acm}), \\ \alpha_3 \Rightarrow \langle \text{acm} \rangle \forall x:\text{prin}. \forall y:\text{prin}. \\ \quad \text{is\_member}(x, \text{acm}) \wedge \langle x \rangle \text{is\_student}(y, x) \\ \quad \supset \text{mayrd}(\text{conf}, y) \\ \vdots \end{array} \right\}$$

The task now is to construct a proof of

$$\langle \text{acm} \rangle \text{mayrd}(\text{conf}, k).$$

This task depends on the identity of the authenticated principal. The static index of the authenticated principal is obtained by opening the package `authpkg` (Line 33), binding `me` to  $k$ , and `cookie` to  $\mathbf{iam}[k]$ . The rest of the code in the function executes in the scope of this `open`.

We begin by first acquiring the certificate for studentship for  $k$ . This is done by moving evaluation to the site `univserver` and searching for a proof of  $\langle \text{univ} \rangle \text{is\_student}(k, \text{univ})$  there. The query `acquire` on Line 36 succeeds if  $k$  is `alice` because in this case there is a direct proof,  $\alpha_1$ , in the policy. In the absence of any other proofs in `univserver`'s local policy, this query fails for other principals  $k$ . A successful `acquire` on Line 36 binds `studentcert` to the optional package `SOME`  $\{\alpha = \alpha_1; \langle \rangle : \text{unit}\}$ , and the result is brought back to `acmserver` by the `get`. Notice that the result of an `acquire` is of a mobile type and therefore can be brought over from `univserver` to `acmserver`. The program raises the exception `abort` and terminates in case this `acquire` fails.

Next we need to acquire a proof that  $k$  is allowed access to the database `conf`. This proof search depends on the static identity of the database `conf`. This identity is obtained using `db.open` at Line 40, which, if successful binds the proxy to  $d$ , and the runtime handle to  $h$  upon opening the package returned by `db.open`. At this point, we have already acquired a proof of studentship for `me`. We now use this fact to guide our proof search. We proceed (Line 44) with a query for the following proposition at `acmserver`:

```

1  unit
2  import "std.mlh"
3
4  extern principal acm = "acm"
5  extern principal univ = "univ"
6  extern bytecode world acmserver
7  extern bytecode world univserver
8  extern javascript world browser
9
10 extern val acmaddr ~ acmserver addr
11 extern val univaddr ~ univserver addr
12 extern val browseraddr ~ browser addr
13
14 extern prop (prin, prin) is_member
15 extern prop (db, prin) mayrd
16 extern prop (prin, prin) is_student
17
18 extern type (d:db) dbhandle
19 extern val db.open :
20   string -> {d:db, dbhandle(d)} option @ acmserver
21 extern val (p: prin, d: db, f: acm says mayrd(d, p))
22   db.read :
23     lam(p) * dbhandle(d) * string -> string
24     @ acmserver
25
26 exception Abort of string
27
28 fun readpaper () =
29   let val authpkg = case authenticate of
30     NONE => raise Abort "authentication failed"
31     | SOME p => p
32   in
33     open authpkg as {me: prin, cookie} in
34       let val studcert =
35         case from univaddr
36           get acquire [univ says is_student(me, univ)] of
37         NONE => raise Abort "studentship search failed"
38         | SOME c => c
39       in
40         case db.open "conf" of
41         NONE => raise Abort "db.open failed"
42         | SOME dbase =>
43           open dbase as {d:db, h} in
44             case acquire [univ says is_student(me, univ)
45               implies
46                 acm says mayrd(d, me)] of
47             NONE => raise Abort "mayrd search failed"
48             | SOME prf => open studcert as {studpf, _} in
49               open prf as {pf, _} in
50                 db.read[me, d, impE pf studpf]
51                   (cookie, h, "paper.pdf")
52               end
53             end
54           end
55         end
56       end
57     end
58   do from acmaddr get readpaper ()
59 end
60
61 end

```

**Figure 4.** PCML<sub>5</sub> code for accessing a PCA-enabled resource under a distributed authorization policy. The code corresponds to the setup introduced in Figure 1.

```

univ says is_student(me, univ)
implies
acm says mayrd(d, me)

```

This particular proof search illustrated the idea of combining proofs from various sites. The proof search is done at `acmserver` even though the antecedents of the implication represent non-local facts. The hypothesis is discharged (Line 50) using proofs  $\alpha_1$  acquired separately at `univserver`. This discharge of assumptions using non-local facts to obtain a globally valid proof is possible because constructors are typeable at all locations.

Notice that in order to construct the accessibility proof, it is critical to assemble `univ`'s policy from `univserver`. The policy of `acm` alone does not suffice for the required proof. That is, one cannot hope for the direct query

```

acquire[acm says mayrd (d, me)]

```

to be successful at `acmserver`. The proof query in Lines 44-46 can be viewed as temporarily extending the policy at `acmserver` to include the non-local fact `(univ)is_student(k, univ)`, whereupon the accessibility proof becomes derivable at `acmserver`.

### 8.3 Database access using required proofs

Lines 48-53 show the code that makes an API call to read `conf` using the proof constructed above. First the proof packages returned by `acquire`'s are opened binding the proof constructors to `studpf` and `pf`. The accessibility proof is constructed as `impE pf studpf`. `impE` is a constructor from the embedding of the authorization logic representing  $\supset$ -elimination; it has the kind

$$\Pi\alpha::\text{Prop}.\Pi\beta::\text{Prop}.\text{Prf}(\alpha \supset \beta) \rightarrow \text{Prf}(\alpha) \rightarrow \text{Prf}(\beta)$$

The constructor `pf` is a proof of

$$\langle \text{univ} \rangle \text{is\_student}(k, \text{univ}) \supset \langle \text{acm} \rangle \text{mayrd}(\text{conf}, k)$$

under  $\hat{\Phi}_{\text{acmserver}}$ , and the constructor `studpf` is a proof of

$$\langle \text{univ} \rangle \text{is\_student}(k, \text{univ})$$

under  $\hat{\Phi}_{\text{univserver}}$ . Both proofs are therefore well-typed under  $\hat{\Phi}$ . Thus the proof `impE pf studpf` proves `(acm)mayrd(conf, k)` under the amalgamated policy  $\hat{\Phi}$ .

Apart from the authorization check, the monitor checks if the credential `cookie` represents an authenticated principal by consulting the set of active principals. This check also succeeds because `authenticate` done at Line 29 augments the active set with `k`. Since both checks succeed, the API call is successful. The result, which is of the mobile type `string`, is returned back to `browser` on Line 59.

## 9. Related work

The problem of specifying and enforcing policies in a scenario with decentralized and spatially distributed policies was first studied in the context of authentication in the Taos operating system (Wobber et al. 1994), and later in the context of trust management systems (Blaze et al. 1996; Clarke et al. 2001; Li et al. 2002). Such policies call for logics that are expressive enough to distinguish assertions made by different principals. A number of logics, beginning with the seminal work by Lampson et al. (Lampson et al. 1992), provide support for such assertions.

Recently, a number of languages for writing programs compliant with authorization policies have been developed. We review those that we consider to be the closest to PCML<sub>5</sub>, in terms of objectives and/or techniques.

Aura (Jia et al. 2008; Vaughan et al. 2008) is a language for enforcing authorization policies. It is based on DCC (Abadi et al. 1999). Our language differs from Aura in terms of the domain of

use because ours is a distributed programming language for distributed authorization policies. Aura is neither a distributed programming language, nor does it handle distributed policies.

As a matter of technique, PCML<sub>5</sub> differs from Aura in the way they incorporate the authorization logic. While PCML<sub>5</sub> uses a higher-order encoding of an authorization logic as static constructors, with a phase distinction between static and dynamic phases, the proof level in Aura is a Curry-Howard interpretation of an authorization logic, and is based on DCC (Abadi 2006).

Also, Aura does not have a notion of authentication of the principal executing the program. A special principal identifier called `self` is used to refer to the executing principal. In contrast, PCML<sub>5</sub> uses authentication tokens as indicators of the fact that a certain principal had been authenticated. This also allows for a program to acquire multiple authenticities during its evaluation.

RCF (Bengtson et al. 2008; Fournet et al. 2007) uses refinement types together with dependent types to express pre- and post-conditions. The proof obligations are represented as preconditions in the API. Thus a function for reading databases may be typed as

$$\text{read} : \text{file:string}\{\text{mayrd}(\text{file})\} \rightarrow \text{string}$$

where `file:string{mayrd(file)}` is the refinement type of strings `f` for which `mayrd(f)` holds. Refinement types are introduced using a term of the form `assume C`, which is typed as `..unit{C}`. The typing context can be thought of as defining a theory which is the set of all the formulae appearing in it. The most important difference between PCML<sub>5</sub> and RCF is in the respective problems being addressed by the two languages. While PCML<sub>5</sub> is a language for programming with PCA-enabled resources that demand proofs of accessibility along with each access request, the motivation of RCF is how to verify that a program conforms to a security policy. As a matter of technique, PCML<sub>5</sub> uses explicit proof terms unlike RCF. In absence of proof terms, the type-checking algorithm of RCF uses an SMT solver to verify if the typing context proves a particular logic formula. Another difference is that RCF does not have a phase distinction since runtime values can appear inside types, because formulas need to mention runtime entities.

PCAL (Chaudhuri and Garg 2009) is another language that relies on external SMT solvers during compile-time to construct PCA proofs. Users annotate program points with propositions that they expect to hold there. The compiler first checks that the annotation at a point is sufficient to guarantee access to the command executed at that point. Then it attempts to construct a proof statically as per the annotation. In case it cannot construct a proof statically, the compiler produces code to dynamically construct the required proof. Using a combination of both methods, the compiler ensures compliance with the PCA interface.

Fable(Swamy et al. 2008) is another language that provides statically enforced compliance with security policies. The idea is to have an abstract type of tagged program values that can only be manipulated using trusted policy functions. A program is divided into two fragments: the policy fragment that provides the abstraction, and the application fragment that functions as a client for the abstraction, treating tagged values abstractly. Different kinds of security properties can be expressed by having different interpretations for the tags. Thus instead of designing a language around a particular form of policies, such as is PCML<sub>5</sub>, Fable attempts to provide a general framework in which different kinds of policies can be expressed. This however comes at a cost: the language itself does not guarantee any security property itself (other than type safety); the relevant properties need to be proved separately for every policy fragment.

The idea of statically checking the permission for accessibility has been used in a completely different setting as compared to ours by Krishnaswami *et al.* (Krishnaswami and Aldrich 2005). They

use the notion of *domains* with inter-domain accessibility permissions to statically enforce that code from a domain may access an element of another domain only if there is a chain of access permissions from the former to the latter domain. Analogous to the proof-carrying APIs in PCML<sub>5</sub>, their proposal allows specification of access permissions associated with a domain. They use domains to encapsulate stateful parts of modules for which it is desirable to restrict access from outside domains. The type system enforces compliance of the module and its interface to a high-level policy of accessibility, i.e., protected parts of the module are not leaked out by the interface.

## 10. Conclusion and future work

We have presented a language-based approach for enforcing distributed authorization policies. We are currently working on a prototype implementation of PCML<sub>5</sub>, building upon the implementation of ML<sub>5</sub> (Murphy 2008). We plan to implement distributed applications using PCML<sub>5</sub>. Currently the programmer constructs all the proofs by himself. In future, we would like to provide support for automated proof construction. We also plan to mechanize the metatheory of PCML<sub>5</sub>. In this paper, we have assumed a very simple model of spatial distribution of policies: the local policy at a site *w* contains only assertions made by the principal who governs *w*. In practice, however, assertions made by one principal may be cached at other sites, as is done in trust management systems (Blaze et al. 1996). In future work, we plan to support such richer forms of distribution of authorization policies and distributed proof search algorithms.

## Acknowledgments

We would like to thank Deepak Garg, Limin Jia and Dan Licata for many helpful comments and suggestions.

## References

- Martín Abadi. Access control in a core calculus of dependency. In *ICFP*, pages 263–273, 2006.
- Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. A core calculus of dependency. In *Proceedings 26th ACM Symp. on Principles of Programming Languages (POPL)*, pages 147–160, 1999.
- Andrew W. Appel and Edward W. Felten. Proof-carrying authentication. In *CCS '99: Proceedings of the 6th ACM conference on Computer and communications security*, pages 52–62, 1999.
- Kumar Avijit, Anupam Datta, and Robert Harper. Distributed programming with distributed authorization, 2009. Available at <http://www.cs.cmu.edu/~kavijitu/papers/pcml5-full.pdf>.
- Lujo Bauer, Michael A. Schneider, and Edward W. Felten. A proof-carrying authorization system. Technical Report TR-638-01, Princeton University, April 2001. URL <http://www.ece.cmu.edu/~lbauer/papers/pcaprototr.pdf>.
- Lujo Bauer, Scott Garriss, Jonathan M. McCune, Michael K. Reiter, Jason Rouse, and Peter Rutenbar. Device-enabled authorization in the grey system. In *Proceedings of the 8th Information Security Conference (ISC05)*, pages 431–445, 2005.
- Jesper Bengtson, Karthikeyan Bhargavan, Cédric Fournet, Andrew D. Gordon, and Sergio Maffei. Refinement types for secure implementations. In *Proceedings of the 2008 21st IEEE Computer Security Foundations Symposium*, pages 17–32, 2008.
- M. Blaze, J. Feigenbaum, and J. Lacy. Decentralized trust management. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 164–173, May 1996.
- Avik Chaudhuri and Deepak Garg. PCAL: Language support for proof-carrying authorization systems. In *Proceedings of the European Symposium on Research in Computer Security*, pages 184–199, 2009.
- Dwayne E. Clarke, Jean-Emile Elien, Carl M. Ellison, Matt Fredette, Alexander Morcos, and Ronald L. Rivest. Certificate chain discovery in SPKI/SDSI. *Journal of Computer Security*, 9(4):285–322, 2001.
- Cedric Fournet, Andrew Gordon, and Sergio Maffei. A type discipline for authorization in distributed systems. In *CSF '07: Proceedings of the 20th IEEE Computer Security Foundations Symposium*, pages 31–48, 2007.
- Deepak Garg. *Proof Theory for Authorization Logic and its Application to a Practical File System*. PhD thesis, School of Computer Science, Carnegie Mellon University, 2009. Available as Technical Report CMU-CS-09-123.
- Deepak Garg and Frank Pfenning. Non-interference in constructive authorization logic. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop (CSFW 19)*, pages 283–296, 2006.
- Robert Harper and Daniel R. Licata. Mechanizing metatheory in a logical framework. *Journal of Functional Programming*, 17(4–5):613–673, July 2007.
- Robert Harper, John C. Mitchell, and Eugenio Moggi. Higher-order modules and the phase distinction. In *POPL '90: Proceedings of the 17th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 341–354, 1990.
- Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
- Limin Jia, Jeffrey A. Vaughan, Karl Mazurak, Jianzhou Zhao, Luke Zarko, Joseph Schorr, and Steve Zdancewic. Aura: a programming language for authorization and audit. In *ICFP '08: Proceeding of the 13th ACM SIGPLAN International Conference on Functional programming*, pages 27–38, 2008.
- Neel Krishnaswami and Jonathan Aldrich. Permission-based ownership: encapsulating state in higher-order typed languages. *SIGPLAN Notices*, 40(6):96–106, 2005.
- Butler W. Lampson, Martín Abadi, Michael Burrows, and Edward Wobber. Authentication in distributed systems: Theory and practice. *ACM Transactions on Computer Systems*, 10(4):265–310, 1992.
- Ninghui Li, John C. Mitchell, and William H. Winsborough. Design of a role-based trust-management framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 114–130, 2002.
- Tom Murphy, VII. *Modal Types for Mobile Code*. PhD thesis, Carnegie Mellon, January 2008. Available as technical report CMU-CS-08-126.
- Tom Murphy, VII, Karl Cray, Robert Harper, and Frank Pfenning. A symmetric modal lambda calculus for distributed computing. In *Proceedings of the 19th IEEE Symposium on Logic in Computer Science (LICS 2004)*, pages 286–295, 2004.
- Tom Murphy, VII, Karl Cray, and Robert Harper. Type-safe distributed programming with ML5. In *Trustworthy Global Computing 2007*, November 2007.
- B. C. Neuman and T. Ts'o. Kerberos: An authentication service for computer networks. *Communications Magazine, IEEE*, 32(9):33–38, 1994.
- Nikhil Swamy, Brian J. Corcoran, and Michael Hicks. Fable: A language for enforcing user-defined security policies. In *Proceedings of the IEEE Symposium on Security and Privacy (Oakland)*, May 2008.
- Jeffrey A. Vaughan, Limin Jia, Karl Mazurak, and Steve Zdancewic. Evidence-based audit. In *CSF '08: Proceedings of the 2008 21st IEEE Computer Security Foundations Symposium*, pages 177–191, 2008.
- Kevin Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. A concurrent logical framework I: Judgments and properties. Technical Report CMU-CS-02-101, Carnegie Mellon University, School of Computer Science, 2002.
- Edward Wobber, Martín Abadi, Michael Burrows, and Butler Lampson. Authentication in the taos operating system. *ACM Transactions on Computer Systems*, 12:256–269, 1994.