

AURA: A programming language for authorization and audit

Limin Jia Jeffrey A. Vaughan Karl Mazurak Jianzhou Zhao Luke Zarko Joseph Schorr
Steve Zdancewic

University of Pennsylvania

{liminjia, vaughan2, mazurak, jianzhou, zarko, jschorr, stevez}@seas.upenn.edu

Abstract

This paper presents AURA, a programming language for access control that treats ordinary programming constructs (e.g., integers and recursive functions) and authorization logic constructs (e.g., principals and access control policies) in a uniform way. AURA is based on polymorphic DCC and uses dependent types to permit assertions that refer directly to AURA values while keeping computation out of the assertion level to ensure tractability. The main technical results of this paper include a proof of decidability for AURA's type system, a fully mechanically verified proof of soundness, and a prototype typechecker and interpreter.

1. Introduction

There can be no universal definition of security. Every piece of confidential data and every sensitive resource may have specialized access control requirements. At the same time, almost every modern computer system stores some private information or provides a service intended only for certain clients. To ensure that only allowed principals—human users or other computer systems—can reach the protected resources, these access-control requirements must be carefully defined and enforced. An *authorization policy* specifies whether a request by a principal to access a resource should be granted, and a *reference monitor* mediates all access to the resource, ensuring that handling of requests complies with the authorization policy.

One significant challenge in building secure systems that enforce access control is that, as the number of resources and principals grows, specifying the authorization policy becomes more difficult. The situation is further complicated in decentralized or distributed settings, where resources may have different owners and the principals may have non-trivial trust relationships. Once the policies become sufficiently complex, understanding which principals may access which resources is itself a daunting problem. Consequently, reference monitors that enforce such policies also become complex, which is not a desired situation when (as in a conventional access control scheme) the reference monitor is part of the trusted computing base.

To help mitigate this complexity, researchers have proposed *authorization logics* that facilitate reasoning about principals, requests, and policy assertions [4, 13, 19, 1, 2]. Several of these log-

ics have been concerned with specifying access-control policies in distributed settings [43, 5, 10, 20, 19]. Part of the appeal of authorization logics is that proofs of propositions in the logic can act as *capabilities* that provide the reference monitor with evidence that a given request should be granted. As proposed by Appel and Felten [5], such a *proof-carrying authorization* approach places the burden of validating the authorization decision on the principal requesting access. Moreover, the explicit proofs can be logged for future auditing, which can help track down bugs in the authorization policy [38].

Authorization logics provide rich and concise languages for specifying access-control policies, abstracting from low-level details like authentication and cryptography. Unfortunately, these logics are rather removed from the languages used to write software that must respect the access-control policies; typecheckers and other tools that help the programmer write correct programs will not necessarily help the programmer correctly make use of an authorization logic. This is especially problematic in the case of the reference monitor, which has the task of enforcing policies written in the authorization logic and must still be considered part of the trusted computing base.

This paper presents the design of AURA, a domain-specific programming language that incorporates a constructive authorization logic based on DCC [3, 2] as part of its type system. Rather than mediate between programs and policy statements written in two distinct languages, AURA uses *dependent types* to permit policy statements that refer directly to AURA values (like integers or datatype constructors). For example, a function *playFor* that acts as a reference monitor for playing MP3 files might have the following type, which requires a proof that principal p is permitted to access the song:

$$(s : \text{Song}) \rightarrow (p : \text{prin}) \rightarrow \text{pf}(\text{self says } \text{MayPlay } p \ s) \rightarrow \text{Unit}.$$

As indicated by this type, AURA programs may construct and manipulate authorization proofs just as they might other program values, and the AURA programming model provides notions of principals (p), authority (self), and policy assertions (*MayPlay*) in addition to standard functional language features like higher-order functions, polymorphism, and recursive algebraic datatypes. In addition, security-relevant implementation details—like the creation of audit trails or the cryptographic interpretation of certain logical statements—can be handled automatically with little to no programmer intervention.

Because policy assertions are part of AURA's type system, deciding whether to grant access amounts to typechecking a proof object. This can be encapsulated in AURA's runtime, removing individual reference monitors from the trusted computing base. Moreover, any program written in AURA benefits from the immediate availability of the authorization logic; many misbehaving programs can now be ruled out at compile time. Finally, DCC, on which AURA is based, has been shown to be useful in representing other

forms of language-based security, such as the type-based enforcement of information-flow properties as found in Jif [29] or Flow-Caml [33]; AURA thus represents a promising avenue for further work in connecting these concepts.

The main contributions of this paper can be summarized as follows:

- We present the design of core AURA, a language with support for first-class, dependent authorization policies.
- We give a fully machine-checked proof of type soundness for the core language.
- We prove decidability of AURA’s type system.
- We describe a prototype implementation of a typechecker, interpreter, and sample programs.

AURA represents a relatively unexplored facet of language design. Typical dependently typed languages (see Section 6) use types to encode precise program specifications. Our goal is different; AURA uses dependent types to naturally connect data with proofs for run-time policy enforcement. Compared with a conventional dependently type language AURA adds some features—assertion types, digitally signed objects as proofs, the *says* monad and *pf* modality—and restricts or removes others—only values may appear in dependent types. The result is a system tuned for dynamic authorization but unsuitable for, e.g., static program verification.

Our proof of soundness is implemented in Coq and encompasses all of AURA’s features, including: higher order and polymorphic types, mutually recursive data types and propositions, a restricted form of dependent types, and authorization proofs. We believe that the mechanized proof is of independent value, because parts of the proof may be reused in other settings.

The rest of this paper focuses on the novel core features of AURA. The next section overviews AURA’s programming model and illustrates the novel features by example. Section 3 gives a formal account of AURA’s core language, its type system, operational semantics, and the main technical results (soundness and decidability of type checking). Section 4 describes our prototype implementation. Section 5 gives a larger scale example demonstrating how AURA’s features work in concert. Section 6 situates AURA with respect to related work, especially prior work on authorization logics and languages with dependent types. Finally, Section 7 concludes with a discussion of future avenues for extending AURA.

AURA as we present it is intended to be suitable as a compilation target for a more convenient surface syntax. As such, we defer the important (and practical) issues of type inference, pattern-match compilation, and the like to future work. Additional topics for future study include authentication, credential revocation, the interpretation of AURA values in cryptography, and integration with mixed language (e.g. C or .Net) systems.

2. Programming in AURA

AURA is intended to be used to implement reference monitors [11] for access control in security sensitive settings. A reference monitor must first mediate access by allowing and denying requests to a resource (based, in this case, on policy specified in an authorization logic) and second log accesses to enable *ex post facto* audit. This latter point we have covered in detail elsewhere [38] (although we discuss logging briefly in Section 2.3); in this paper we concentrate on the details integrating programming with an authorization logic.

The potential design space of dependently-typed languages is quite large, and there are many challenges in striking a good balance between expressiveness and tractability of type checking. AURA’s design strives for simplicity, even at the cost of expres-

siveness. This section describes AURA’s design, concentrating on the features specific to authorization policies.

As alluded to by the function *playFor* in the introduction, we use an AURA implementation of a musical jukebox server as a running example throughout this paper. The full example is given in Section 5; the rest of this section will illustrate *playFor* in more detail.

2.1 AURA as an authorization logic

We first turn our attention to AURA’s assertions, which are based on the polymorphic core calculus of dependency (DCC) by [3] and in particular on DCC’s interpretation as an authorization logic [2]. In both DCC and AURA, an indexed monad *says* associates propositions with principals. The statement *a says P* holds when the principal *a* has actively affirmed the proposition *P*, when a direct proof for *P* is known, or when *a says P* logically follows from monad operations that we will describe shortly—it is critical to note, however, that *a says P* does not imply *P*. We augment DCC with dependent types, allowing principals to assert propositions about data, and with the constructs *say* and *sign*, which we will describe shortly.

Principals in AURA, written *a*, *b*, etc. and having type *prin*, represent distinct components of a software system. They may correspond to human users, system components such as an operating system kernel, a particular server, etc. Formally, principals are treated as special values in AURA; they are characterized by their ability to index the family of *says* monads.

As ‘*a says*’ is a *monad* [40], we can construct a term of type *a says P* from a proof *p* of *P* using the operation *return a p*. A proof encapsulated in a *says* monad cannot be used directly; rather, the monad’s bind operation, written (*bind p (λx:P. q)*) allows *x* to stand in for the proof inside *p* and appear in the expression *q*.

For example, consider the principals *a* and *b*, the song *freebird*, and the assertion *MayPlay* introduced earlier. The statements

$$\begin{aligned} ok & : a \text{ says } (MayPlay \ a \ freebird) \\ delegate & : b \text{ says } ((p:prin) \rightarrow (s:Song) \rightarrow \\ & \quad (a \text{ says } (MayPlay \ p \ s)) \rightarrow \\ & \quad (MayPlay \ p \ s)) \end{aligned}$$

assert that *a* gives herself permission to play *freebird* and *b* delegates to *a* the authority to make any variety of *MayPlay* statement on his behalf. These two terms may be used to create a proof of *b says (MayPlay a freebird)* as follows:

$$\begin{aligned} bind \ delegate \ (\lambda d: ((p: prin) \rightarrow (s: Song) \rightarrow \\ & \quad (a \text{ says } (MayPlay \ p \ s)) \rightarrow \\ & \quad (MayPlay \ p \ s)). \\ & \quad return \ b \ (d \ a \ freebird \ ok)) \end{aligned}$$

Such a proof might have direct utility—it could be passed to the *playFor* function if self is *b*—or it might become part of a larger chain of reasoning.

In addition to *return*, AURA allows for the introduction of proofs of *a says P* without corresponding proofs of *P*. We provide a pair of constructs, *say* and *sign*, that represent a principal’s active affirmation of a proposition. The value *sign(a, P)* has type *a says P*; intuitively we may think of it as a digital signature by *a*’s private key on proposition *P*. Such a value is intended to have a stable meaning as it is passed throughout a distributed systems.

A principal should only be able to create a term of the form *sign(a, P)* if it is—or, at least, has access to the private key of—*a*. We thus prohibit such terms from appearing in source programs and introduce the related term (*say P*), which represents an effectful computation that uses the runtime’s current authority—that is, its private key—to sign proposition *P*. When executed, *say P* generates

a fresh value $\text{sign}(\text{self}, P)$, where self is a distinguished principal representing the current runtime authority.

It is worth noting that a principal can assert any proposition, even *False*. Because assertions are confined to the monad—thanks to the non-interference property of DCC—a false assertion can do little harm other than making the principal’s own assertions inconsistent. In practice, it is useful to restrict the kinds of assertions that various principals can make, but *a priori*, AURA requires no such constraints.

The concept of a program’s runtime authority already has a natural analog in the operating system world—a UNIX process, for example, has an associated user ID that often, but not always, corresponds to the user who started the process. In a more distributed setting, running under the authority of a can indeed be represented by possession of a ’s private key. In such a setting objects of the form $\text{sign}(a, P)$ can be represented by actual digital signatures, and principal identifiers—which, in AURA, are first class values of type prin —can be thought of as public keys.

The restriction of authority to a single principal is only for simplicity’s sake—although syntax would need to be changed, nothing in our development would conflict with a more complex notion of authority. AURA currently provides no means of *transferring* authority, in effect disallowing programs from directly manipulating private keys; this prevents AURA programs from creating new principals (i.e., key pairs) at runtime but also trivially disallows the accidental disclosure of private keys. Were AURA to be extended with support for dynamically generated principals, the addition of information flow tracking could assist in ensuring that private keys stay sufficiently private.

2.2 Authorization proofs and dependent types

By treating assertions as types and proofs as expressions we are taking advantage of the well-known Curry-Howard Isomorphism [18, 23] between logic and programming languages. Yet while the ability to manipulate proofs just as one can manipulate other datatypes is quite useful, we cannot hide from the fact that useful assertions must somehow refer to objects within our language. In the above example, for instance, *freebird* is data that appears at the assertion (i.e., type) level; the function *playFor* in the introduction also a clearly dependent type.

AURA incorporates dependent types directly—in contrast to, for example, using GADTs [32] or static equality proofs [35] to simulate the required dependencies. Such an approach allows straightforward use of data at the type level and avoids replicating the same constructs in both static and dynamic form, but unconstrained use of dependent types can quickly lead to an undecidable typing judgment. Moreover, care must be taken to separate effectful computations from pure proof objects.

Much like CIC [17], AURA has separate universes Type and Prop , with Type and Prop themselves being classified by Kind . The previously mentioned assertion *MayPlay*, for instance, would be given the assertion type $\text{Prin} \rightarrow \text{Song} \rightarrow \text{Prop}$. Unlike CIC, both types of kind Type and propositions of kind Prop describe data that may be available at runtime. Propositions, however, are required to be completely computation-free: propositions never reduce and AURA does not employ type-level reduction during typechecking, meaning that only dependencies on values (i.e., well-formed normal forms) for which equality comparison is available can be used non-trivially. This turns out to be enough to ensure decidability of AURA’s type system.

AURA offers a type-refining equality test on *atomic* values—for instance, principals and booleans—as well as a dynamic cast between objects of equivalent types, which prove necessary for certain equalities that arise only at runtime. For example, when typechecking if $\text{self} = a$ then e_1 else e_2 , the fact that $\text{self} = a$ is

automatically made available while typechecking e_1 (due to the fact that prin is an atomic type), and hence proofs of type $\text{self says } P$ can be cast to type $a \text{ says } P$ and vice-versa.

The distinction between Type and Prop is also illustrated by the previously introduced *say* and *sign*. On the one hand, *say* P certainly belongs in Type ’s universe—not only do we intend it to be reduced by our operational semantics, this reduction is an effectful (if trivial) computation dependent on a program’s runtime authority. On the other hand, $\text{sign}(a, P)$ should be of type $a \text{ says } P$, which, like P , is of kind Prop . To solve this dilemma we introduce the modality $\text{pf} : \text{Prop} \rightarrow \text{Type}$, allowing us to give *say* P the type $\text{pf}(\text{self says } P)$ of kind Type . The pf modality also comes equipped with its own *bind* and *return* operations, much like *says*, thus allows proofs to be manipulated by computations while keeping the worlds of computations and assertions separate.

AURA’s dependent types also address something that might have seemed odd about our cryptographic interpretation of the *says* monad, namely that one most often thinks of digitally signing *data*, whereas $\text{sign}(a, P)$ signs only an assertion. With dependent types, however, this issue evaporates, as an assertion can refer to whatever data might be endorsed. We find this design compelling, because a digital signature on raw data does not necessarily have a sensible meaning; signing only propositions ensures that the signed data is attributed with some semantics, just as, for example, a physical signature on a contract will indicate whether the signer is party to the contract or merely a witness.

2.3 Auditing in AURA

Passing proofs at runtime is also useful for after the fact auditing of AURA programs. The full details are given elsewhere [38] but we note that, when full proofs are logged for every resource access, it becomes possible to determine *how* access was granted at a very fine granularity. This is of great importance when the intent of some institutional policy is not properly reflected in the actual rules enforced by a software system—for example, an auditor can examine the proof that allowed an unwanted access to take place and determine whether and where authority was improperly delegated.

These guarantees can be made as long as the interface to the resources of interest is sufficiently rich: we can simply decree that every interface function—that is, a function that wraps a lower level operating system call—writes its arguments to the log. There are no constraints on what the rest of the reference monitor may do other than that it must respect this interface—it is not possible to inadvertently add a path through the program that causes insufficient information to be logged. This is in keeping with AURA’s general philosophy of resilience toward mistakes on the part of the programmer.

Returning to *playFor*, let us assume that there exists a native function $\text{rawPlayFor} : \text{Song} \rightarrow \text{Unit}$ that is not security-aware and hence is not available to the programmer. We define the interface function *playFor* as simply

$$\lambda s : \text{Song}. \lambda p : \text{prin}. \lambda \text{proof} : \text{pf}(\text{self says } \text{MayPlay } p \ s). \\ \text{rawPlayFor } s.$$

Because *playFor* is an interface function—i.e., because it has access to *rawPlayFor*—its arguments will automatically be logged, and because the access control policy is entirely encoded in *playFor*’s signature, the log will automatically contain everything an auditor needs to determine precisely how any song was authorized to be played.

3. The AURA Core Language

This section presents the main technical contributions of this paper, namely a formal description of the AURA core language, its type

system, operational semantics, and the corresponding proofs of type soundness and decidability of type checking.

We adopt the design philosophy of elaboration-style semantics (as used, for example, by Lee *et. al* [27]): the AURA intermediate language is intended to make type checking as explicit as possible. Following this principle, our design eschews complex pattern matches, equality tests over complex values, and implicit casts. Our goal was to cleanly divide the compiler into two parts: an elaboration phase that uses inference, possibly with heuristics and programmer-supplied hints, to construct an internal representation that makes all type information explicit; and a compilation phase that processes the fully elaborated intermediate representation into executable code.

3.1 AURA core syntax

As described above, AURA is a call-by-value polymorphic lambda calculus. It consists of a “term-level” programming language (whose expressions are classified by types of kind Type) for writing algorithms and manipulating data and a “proof-level” assertion language (whose expressions are classified by propositions of kind Prop) for writing proofs of access-control statements. These two languages share many features (λ -abstraction, application, constructors, etc.) and, due to the dependent types, propositions and types may mention terms. To simplify the presentation of AURA, it makes sense to unify as many of these constructs as possible. We thus adopt a lambda-cube style presentation [9] that uses the same syntactic constructs for terms, proofs, types, and propositions. Different categories are distinguished by the type system as necessary. This approach also has the appeal of greatly reducing the number of objects in the language, which simplifies both the metatheory and implementation. Our design was significantly influenced by the Henk intermediate language [24], which also adopts this compact representation.

The lambda-cube terms of the AURA core syntax are given by:

$$\begin{array}{l} \text{Terms } t ::= x \mid ctr \mid \dots \\ \quad \quad \quad \mid \lambda x:t_1. t_2 \mid t_1 t_2 \mid (x:t_1) \rightarrow t_2 \\ \quad \quad \quad \mid \text{match } t_1 t_2 \text{ with } \{b\} \mid \langle t_1 : t_2 \rangle \\ \text{Branches } b ::= \cdot \mid ctr \Rightarrow t \mid b \end{array}$$

Here, x ranges over variables, and ctr ranges over programmer-defined constructors created using data type declarations as described below. In addition to the standard lambda abstraction, application, and dependent arrows, AURA also has a pattern matching construct and an explicit type cast. In the expression $\text{match } t t_1 \text{ with } \{b\}$, t is the term that is being analyzed, t_1 is the return type, and b is a list of branches t_1 is matched against; the type annotation t_1 in the syntax of the pattern-matching expression ensures that type checking is straightforward even when the set of branches is empty. The explicit cast $\langle t_1 : t_2 \rangle$ ensures (safely) that t_1 be considered at type t_2 .

To express and reason about access control, AURA extends the core syntax above with additional terms. Here, and throughout the rest of the paper, we use metavariable conventions that make it easier to recall constraints placed on a term by the type system: a ranges over principals, P ranges over propositions, p ranges over proofs, e ranges over program expressions, and v stands for values. All of these metavariables are synonymous with t , which we use to indicate syntactic objects of any flavor. The AURA-specific syntax

is given by¹:

$$\begin{array}{l} t ::= \dots \mid \text{Type} \mid \text{Prop} \mid \text{Kind} \\ \quad \quad \quad \mid \text{prin } a \text{ says } P \mid \text{pf } P \\ \quad \quad \quad \mid \text{self } \mid \text{sign}(a, P) \mid \text{say } P \\ \quad \quad \quad \mid \text{return}_s a p \mid \text{bind}_s e_1 e_2 \\ \quad \quad \quad \mid \text{return}_p p \mid \text{bind}_p e_1 e_2 \\ \quad \quad \quad \mid \text{if } v_1 = v_2 \text{ then } e_1 \text{ else } e_2 \end{array}$$

3.2 Type checking AURA

AURA’s type system contains the following judgments:

$$\begin{array}{ll} \text{Well-formed signatures} & S \vdash \diamond \\ \text{Well-formed typing environments} & S \vdash E \\ \text{Well-formed terms} & S; E \vdash t : s \\ \text{Well-formed match branches} & S; E; s; \text{args} \vdash \text{branches} : t \end{array}$$

Figure 1 shows the term type checking rules. We omit the rules for typechecking signatures and branches, though we describe their salient features below. The full type system can be found in the Coq implementation.

In these judgments, S is a signature that declares types, propositions, and assertions (described in more detail below). Typing environments E map variables to their types as usual, but they also record the hypothetical equalities among atomic run-time values. In the definition of environments, E , below, a binding $x \sim (v_1 = v_2) : t$ indicates that v_1 and v_2 have type t , and that the run-time values of v_1 and v_2 are equal.

$$\text{Environments } E ::= \cdot \mid E, x:t \mid E, x \sim (v_1 = v_2) : t$$

3.3 Signatures: data declarations and assertions

Programmers can define bundles of mutually recursive data types and propositions in AURA, just as in other programming languages. A signature S collects together these data definitions and, as a consequence, a well-formed signature can be thought of as map from constructor identifiers to their types. We omit the formal grammar and typing rules for signatures, as they are largely straightforward. Instead we explain signatures via examples.

Data definitions may be parameterized. For example, the familiar polymorphic list declaration is written:

$$\begin{array}{l} \text{data List : Type} \rightarrow \text{Type} \{ \\ \quad \mid \text{nil} : (t:\text{Type}) \rightarrow \text{List } t \\ \quad \mid \text{cons} : (t:\text{Type}) \rightarrow t \rightarrow \text{List } t \rightarrow \text{List } t \\ \quad \} \end{array}$$

AURA’s type system rules out data declarations that require nontrivial equality constraints at the type level. For example, the following GADT-like declaration is ruled out, since $\text{Bad } t u$ would imply $t = u$:

$$\begin{array}{l} \text{data Bad : Type} \rightarrow \text{Type} \rightarrow \text{Type} \{ \\ \quad \mid \text{bad} : (t:\text{Type}) \rightarrow \text{Bad } t t \\ \quad \} \end{array}$$

Logical connectives like conjunction and disjunction can be encoded using dependent propositions, as in Coq and other type-based provers. For example:

$$\begin{array}{l} \text{data And : Prop} \rightarrow \text{Prop} \rightarrow \text{Prop} \{ \\ \quad \mid \text{both} : (p1:\text{Prop}) \rightarrow (p2:\text{Prop}) \rightarrow p1 \rightarrow p2 \rightarrow \text{And } p1 p2 \\ \quad \} \end{array}$$

AURA’s type system conservatively constrains Prop definitions to be inductive by disallowing negative occurrences of Prop constructors. Such a restriction is essential for consistency of the logic,

¹ In the Coq development, these constructs are represented using constants and term application.

$$\begin{array}{c}
\frac{S \vdash E}{S; E \vdash \text{Type} : \text{Kind}} \text{WF-TM-TYPE} \quad \frac{S \vdash E}{S; E \vdash \text{Prop} : \text{Kind}} \text{WF-TM-PROP} \\
\frac{S \vdash E \quad S(\text{ctr}) = t}{S; E \vdash \text{ctr} : t} \text{WF-TM-CTR} \quad \frac{S \vdash E \quad E(x) = t}{S; E \vdash x : t} \text{WF-TM-FV} \quad \frac{S; E, x : t_1 \vdash t_2 : k_2 \quad k_2 \in \{\text{Type}, \text{Prop}, \text{Kind}\}}{S; E \vdash (x : t_1) \rightarrow t_2 : k_2} \text{WF-TM-ARR} \\
\frac{S; E \vdash t : k \quad S; E, x : t \vdash u : k_1 \quad S; E \vdash (x : u) \rightarrow k_1 : k_2 \quad k \in \{\text{Type}, \text{Prop}, \text{Kind}\} \quad k_2 \in \{\text{Type}, \text{Prop}\}}{S; E \vdash \lambda x : t. u : (x : t) \rightarrow k_1} \text{WF-TM-ABS} \\
\frac{S; E \vdash t_1 : (x : u_2) \rightarrow u \quad S; E \vdash t_2 : u_2 \quad \text{val}(t_2) \text{ or } x \notin \text{fv}(u)}{S; E \vdash t_1 t_2 : \{x/t_2\}u} \text{WF-TM-APP} \\
\frac{S; E \vdash e : s \quad \text{fully_applied } s \text{ ctr args } k \quad S(\text{ctr}) = k \quad \text{branches_cover } S \text{ branches } \text{ctr} \quad S; E; s; \text{args} \vdash \text{branches} : t \quad S; E \vdash s : u \quad S; E \vdash t : u \quad u \in \{\text{Type}, \text{Prop}\}}{S; E \vdash \text{match } e \text{ t with } \{\text{branches}\} : t} \text{WF-TM-MATCHES} \\
\frac{S \vdash E}{S; E \vdash \text{prin} : \text{Type}} \text{WF-TM-PRIN} \quad \frac{S \vdash E}{S; E \vdash \text{self} : \text{prin}} \text{WF-TM-SELF} \\
\frac{S; E \vdash a : \text{prin} \quad S; E \vdash P : \text{Prop}}{S; E \vdash a \text{ says } P : \text{Prop}} \text{WF-TM-SAYS} \quad \frac{S; E \vdash a : \text{prin} \quad \text{val}(a) \quad S; E \vdash p : P \quad S; E \vdash P : \text{Prop}}{S; E \vdash \text{return}_s a p : a \text{ says } P} \text{WF-TM-SAYS-RET} \\
\frac{S; E \vdash e_1 : a \text{ says } P \quad S; E \vdash e_2 : (x : P) \rightarrow a \text{ says } Q \quad x \notin \text{fv}(Q)}{S; E \vdash \text{bind}_s e_1 e_2 : a \text{ says } Q} \text{WF-TM-SAYS-BIND} \\
\frac{S; \cdot \vdash a : \text{prin} \quad S; \cdot \vdash P : \text{Prop}}{S; E \vdash \text{sign}(a, P) : a \text{ says } P} \text{WF-TM-SIGN} \quad \frac{S; E \vdash P : \text{Prop}}{S; E \vdash \text{say } P : \text{pf self says } P} \text{WF-TM-SAY} \\
\frac{S; E \vdash P : \text{Prop}}{S; E \vdash \text{pf } P : \text{Type}} \text{WF-TM-PF} \quad \frac{S; E \vdash p : P \quad S; E \vdash P : \text{Prop}}{S; E \vdash \text{return}_p p : \text{pf } P} \text{WF-TM-PF-RET} \\
\frac{S; E \vdash e_1 : \text{pf } P \quad S; E \vdash e_2 : (x : P) \rightarrow \text{pf } Q \quad x \notin \text{fv}(Q)}{S; E \vdash \text{bind}_p e_1 e_2 : \text{pf } Q} \text{WF-TM-PF-BIND} \\
\frac{S; E \vdash v_1 : k \quad S; E \vdash v_2 : k \quad \text{atomic } S k \quad \text{val}(v_1) \quad \text{val}(v_2) \quad S; E, x \sim (v_1 = v_2) : k \vdash e_1 : t \quad S; E \vdash e_2 : t}{S; E \vdash \text{if } v_1 = v_2 \text{ then } e_1 \text{ else } e_2 : t} \text{WF-TM-IF} \\
\frac{S; E \vdash e : s \quad \text{converts } E s t}{S; E \vdash \langle e : t \rangle : t} \text{WF-TM-CAST}
\end{array}$$

Figure 1. AURA typing rules

since otherwise it would be possible to write loops that inhabit any proposition, including *False*. *False* itself is definable: it is a proposition with no constructors:

```
data False : Prop { }
```

Assertions, like the *MayPlay* proposition from above define uninhabited constants that construct Props:

```
assert MayPlay : Prin → Song → Prop
```

While assertions are similar in flavor to datatypes with no constructors, there is a key difference. When an empty datatype is scrutinized by a match expression, the match may be assigned any type. Hence if we were to define *MayPlay* as an empty inductive type, *A says False* would follow from *A says MayPlay A freebird*. In contrast, there is no elimination form for assertions. This means that principals may sign assertion without compromising their says monad’s consistency.

3.4 Core term typing

Type is the type for computation expressions, and Prop is the type for propositions. Constant Kind classifies both Type and Prop as shown in rules WF-TM-TYPE and WF-TM-PROP. (Here and elsewhere, we use the lowercase word “type” to mean a classifier in the type system—Prop and Type are both “types” in this sense.)

The typechecking rules for constructors declared in the signature and free variables are completely standard (see WF-TM-CTR and WF-TM-FV). More interesting is WF-TM-ARR, which says that the type of an arrow is the type of arrow’s output type. The latter is required to be one of Type, Prop, or Kind, which rules out non-sensical arrow forms. For example, $(x : \text{Type}) \rightarrow \text{Type}$ is legal whereas $(x : \text{Type}) \rightarrow \text{self}$ is not—the former could be the type of the polymorphic list constructor while the latter doesn’t make sense since *self* is a computation-level value.

The WF-TM-ABS rule for introducing functions is standard except that, as in other lambda-cube like languages, AURA restricts what sorts of values may be abstracted in others. The argument to a function can be a term value, a proof, a type or a proposition. The

resulting lambda must be typable with an arrow that itself has type `Type` or `Prop`. These restrictions imply that all the lambda abstractions are either a computation or a proof term for a proposition. AURA does not support Type-level lambdas as in F_ω because doing so would require support for β -reduction at the type level. Such reductions, while useful for verification, appear superfluous here.

The interesting part of the `WF-TM-APP` rule is the side condition that either t_2 is a value ($val(t_2)$), or u does not depend on x ($x \notin fv(u)$). This restriction reveals that even though AURA seems to be quite liberal with respect to the dependencies allowed by well-formed $(x : s) \rightarrow t$ terms, the actual dependencies admitted by the type system are quite simple. For instance, although the type system supports singleton types like `S(0)`, it cannot check `S(1+2)` because the latter type depends on a non-value.

The upshot of these restrictions is that AURA requires that types may only depend on values (i.e. terms that cannot reduce). This decision limits the applicability of dependent types for program verification tasks, but greatly simplifies the metatheory, since there is no possibility of effectful computations appearing in a type.

Typechecking pattern match expressions is fairly standard (`WF-TM-MATCHES`), though it is a bit intricate because AURA supports a rich class of parameterized recursive datatypes. Only expressions that have saturated (fully applied) types can be matched against. The types of the branches must exhaustively match the constructors declared in the signature, and any parameters to the datatype being analyzed are also made available inside the branches. Each branch must return an expression of the same type, which is the result type of the entire match expression. Since data types and propositions in AURA may be nullary (have zero constructors), typechecking without inference requires the match expression to carry an annotation. For lack of space, we omit the auxiliary definitions and the judgment used for typechecking the branches themselves.

3.5 Principals and proofs

Principals are an integral part of access control logics and AURA treats principals as first-class objects with type `prin`. The only built-in principal is `self`, which represents the identity of the currently running process (see `WF-TM-PRIN` and `WF-TM-SELF`); additional principal identifier constants could be accommodated simply by adding them with type `prin`, but we omit such a rule for simplicity's sake.

As described above, AURA uses the principal-indexed `says` monad to express access control policies. Proposition $a \text{ says } P$ means that principal a has asserted proposition P (either directly or indirectly). Expression $\text{return}_s a p$ is the return operation for the $a \text{ says}$ monad, and $\text{bind}_s e_1 e_2$ is the corresponding bind operation. These constraints are shown in rules `WF-TM-SAYS`, `WF-TM-SAYS-RET` and `WF-TM-SAYS-BIND`. The rules are adapted from DCC [2], with the exception that AURA eschews DCC's label lattice in favor of explicit delegation among principals. (Abadi has called a similar DCC fragment CDD, standing for cut-down DCC).

The expression $\text{sign}(a, P)$ witnesses the assertion of proposition P made by principal a (`WF-TM-SIGN`). Since $\text{sign}(a, P)$ is intended to model evidence manufactured by a without justification, it should never appear in a source program. Moreover, since signed propositions are intended to be distributed and thus may escape the scope of the running AURA program, they are required to be closed. Note, however, that the declaration signature S must be available in whatever context the signature is to be ascribed meaning. In practice, this means that two distributed AURA programs that wish to exchange proofs need to agree on the signatures used to construct those proofs.

Creating $\text{sign}(a, P)$ requires a 's authority. AURA models the authority vested in a running program using the principal constant `self`. The `say P` operation creates an object of type `pf self says P`.

Intuitively, this operation creates the signed assertion $\text{sign}(\text{self}, P)$ and injects it as a proof term for further manipulation (see `WF-TM-SAY`).

AURA uses the constant `pf : Prop \rightarrow Type` to wrap the access-control proofs that witness propositions as program values, as shown by the rule `WF-TM-PF`. The `pf` type operates monadically: the return operation $\text{return}_p p$ injects a proof p into the term level and the corresponding bind operation bind_p allows a computation to compose proofs (rules `WF-TM-PF-RET` and `WF-TM-PF-BIND`). This separation between proofs and computations is necessary to prevent effectful program expressions from appearing in a proof term. For example, if `say P` was given type `self says P` rather than `pf self says P`, it would be possible to create a bogus ‘‘proof’’ $\lambda x : \text{Prop}. \text{say } x$; the meaning of this ‘‘proof’’ would depend on the authority (`self`) of the program that applied the proof object.

3.6 Equality and conversion

Some typing rules (e.g. `WF-TM-APP`) require checking that two terms can be given the same type. Satisfying such constraints in a dependently type language requires deciding when two terms are equal—a difficult static analysis problem in the best case.

In AURA we address this with a conditional construct. Dynamically, if $v_1 = v_2$ then e_1 else e_2 steps to e_1 when v_1 and v_2 are equal, otherwise the expression steps to e_2 . Statically (rule `WF-TM-IF`), the `then` branch is typed in an environment containing the static constraint ($v_1 = v_2$). As we will see shortly, the constraint may be used to perform safe typecasts. This is an instance of the type refinement problem, well known from pattern matching in languages such as Coq [16], Agda [31], and Epigram [28].

AURA limits its built-in equality tests to inhabitants of *atomic* types. First, the built in `prin` type is atomic. Second, a type is also atomic when it is defined by a non-parameterized `Type` declaration each of whose constructors take no arguments. The `list` type above is not atomic, nor is `list nat` (since `cons` takes an argument). However, the following `Song` type is atomic:

```
data Song: Type { | freebird: Song | ironman: Song }
```

Our definition of atomic type is limiting, but we believe it can be naturally extended to first-order datatypes.

With equalities over atomic types in the context, we can now consider the issue of general type equality. As in standard presentations of the Calculus of Constructions [9] we address type equality in two acts.

Two types in AURA are considered equivalent when they are related by *converts*. The conversion relation, defined in Figure 2, is reflexive, symmetric, and transitive. The key rule is `CONV-AXIOM`; it uses equality assumptions in the environment. For instance, under assumption $x = \text{self}$, term $x \text{ says } P$ converts with `self says P`. As equalities only mention atomic values, conversion will only alter the ‘‘value’’ parts of a type—convertible types always have the same shape up to embedded data values.

AURA contains explicit, safe typecasts. As specified in rule `WF-TM-CAST`, term $\langle e : T \rangle$ is assigned type T whenever e 's type is convertible with T . Standard presentations of dependently type languages instead use implicit conversions which may occur anywhere in a type derivation. Using the explicit cast is appealing because it gives an algorithmic type system. Casts have no runtime effect and are simply discarded by our operational semantics.

3.7 Evaluation rules

Figure 3 defines AURA's operational semantics using a call-by-value small-step evaluation relation.

Most of the evaluation rules are straightforward. The rule `PF-BIND` is a standard beta reduction for monads. Term `say P` creates a proof that principal `self` has asserted that proposition P is

$$\begin{array}{c}
\frac{}{\text{converts } E t t} \text{ CONV-REFL} \quad \frac{\text{converts } E t s}{\text{converts } E s t} \text{ CONV-SYMM} \\
\frac{\text{converts } E s u \quad \text{converts } E u t}{\text{converts } E s t} \text{ CONV-TRANS} \\
\frac{x \sim (s = t) : k \in E}{\text{converts } E s t} \text{ CONV-AXIOM} \\
\frac{\text{converts } E s_1 t_1 \quad \text{converts } E s_2 t_2}{\text{converts } E (s_1 s_2) (t_1 t_2)} \text{ CONV-APP} \\
\frac{\text{converts } E s_1 t_1 \quad \text{converts } E s_2 t_2}{\text{converts } E (\lambda x : s_1. s_2) (\lambda x : t_1. t_2)} \text{ CONV-ABS} \\
\frac{\text{converts } E s_1 t_1 \quad \text{converts } E s_2 t_2}{\text{converts } E ((x : s_1) \rightarrow s_2) ((x : t_1) \rightarrow t_2)} \text{ CONV-ARR}
\end{array}$$

Figure 2. Conversion

$$t \mapsto t'$$

$$\begin{array}{c}
\frac{\text{val}(v)}{(\lambda x : t. e) v \mapsto \{v/x\}e} \text{ APP} \\
\frac{}{\text{bind}_p (\text{return}_p e_1) e_2 \mapsto e_2 e_1} \text{ PF-BIND} \\
\frac{}{\text{say } P \mapsto \text{return}_p (\text{sign}(\text{self}, P))} \text{ SAY} \\
\frac{v_1 = v_2}{\text{if } v_1 = v_2 \text{ then } e_1 \text{ else } e_2 \mapsto e_1} \text{ IF-EQ} \\
\frac{v_1 \neq v_2}{\text{if } v_1 = v_2 \text{ then } e_1 \text{ else } e_2 \mapsto e_2} \text{ IF-NEQ} \\
\frac{\text{val}(v)}{\langle v : t \rangle \mapsto v} \text{ CAST} \\
\frac{(v, \text{branches}) \mapsto_b e}{\text{match } v t \text{ with } \{\text{branches}\} \mapsto e} \text{ MATCH}
\end{array}$$

$$(v, b) \mapsto_b e$$

$$\begin{array}{c}
\frac{(v, c, \text{body}) \mapsto_c (e, 0)}{(v, \text{brn } c \text{ body } \{\text{rest}\}) \mapsto_b e} \text{ B-HERE} \\
\frac{(v, \text{rest}) \mapsto_b e}{(v, \text{brn } c \text{ body } \{\text{rest}\}) \mapsto_b e} \text{ B-EARLIER}
\end{array}$$

$$(v, c, \text{body}) \mapsto_c (e, n)$$

$$\begin{array}{c}
\frac{}{(c, n), (c, n), \text{body}) \mapsto_c (\text{body}, n)} \text{ CTR-BASE} \\
\frac{\text{val}(v_2) \quad m > 0}{(v_1, (c, n), \text{body}) \mapsto_c (\text{body}, m)} \\
\frac{}{(v_1 v_2, (c, n), \text{body}) \mapsto_c (\text{body}, m - 1)} \text{ CTR-PARAM} \\
\frac{(v_1, (c, n), \text{body}) \mapsto_c (e, 0)}{(v_1 v_2, (c, n), \text{body}) \mapsto_c (e v_2, 0)} \text{ CTR-ARG}
\end{array}$$

Figure 3. Reduction Rules

true; therefore, it evaluates to an assertion “signed” by principal self. There are two possibilities in the evaluation of the if $v_1 = v_2$ then e_1 else e_2 statement: when v_1 equals to v_2 , it evaluates to e_1 ; otherwise it evaluate to e_2 . The reduction rule for pattern matching is most complicated, and we need to define two auxiliary reduction relations to implement it. We write $(v, b) \mapsto_b e$ to denote the evaluation of value v against a set of branches. These evaluation rules search through the list of branches until v matches with the constructor of one of the branches, at which point the rules focus on the branch and supply the body of the branch with the arguments in v . The tricky part lies in correctly identifying the arguments in v and discarding the type parameters. We write $(v, c, \text{body}) \mapsto_c (e, n)$ to denote the evaluation of the body of the branch where v matches with the constructor c in the branch. Here, n is the number of parameters that should be discarded before the first argument of v is found. Note that the semantics represents constructors as a pair of the constructor name c and its number of type parameters. For instance, in the definition of polymorphic lists shown previously, the representation of *cons* is $(\text{cons}, 1)$.

3.8 Metatheory

We have proved the following progress and preservation theorems for AURA. The soundness proofs are fully mechanized in the Coq proof assistant.

Theorem 1 (Preservation). *If $S; \cdot \vdash e : t$ and $e \mapsto e'$, then $S; \cdot \vdash e' : t$.*

Theorem 2 (Progress). *If $S; \cdot \vdash e : t$ then either $\text{val}(e)$ or exists e' such that $e \mapsto e'$.*

We have also proved that type checking in AURA is decidable by giving a constructive proof of the following theorem:

Theorem 3 (Type Checking is Decidable).

- If $S \vdash \diamond$ and $S \vdash E$, then $\forall e, \forall t$, either there exists a derivation such that $S; E \vdash e : t$ or there doesn't exist a derivation such that $S; E \vdash e : t$.
- If $S \vdash \diamond$ then $\forall E$ either there exists a derivation such that $S \vdash E$ or doesn't exist a derivation such that $S \vdash E$.
- Either there exists a derivation such that $S \vdash \diamond$ or doesn't exist a derivation such that $\text{sig} \vdash \diamond$.

We have mechanized all of the decidability proofs except for the decidability of the *converts* relation, which is proved on paper. A sketch of the latter is given below.

Lemma 4 (Converts is Decidable). $\forall E, \forall t, \forall s$, it is decidable whether there exists a derivation such that *converts* $E t s$.

Proof (sketch):

1. Define an algorithmic version of the *converts* relation as follows. Given $E t$ and s , apply CONV-APP, CONV-ABS, and CONV-ARR rules until the t 's and s 's in all subgoals are atomic (variables, constructors, or constants). If there exists a subgoal *converts* $E t' s'$ such that one of t' and s' is atomic, and the other is not, then there does not exist a derivation such that *converts* $E t s$. Then we do a graph search using E as the graph definition to see if t' can reach s' . If the graph search succeeds on all subgoals, then *converts* $E t s$; otherwise there does not exist a derivation such that *converts* $E t s$.
2. Prove that the algorithmic version of the *converts* is sound and complete with respect to the original definition.
3. The algorithmic version of *converts* is obviously decidable since the graph algorithm is decidable (E is finite).
4. According to the sound and completeness argument, the *converts* relation is decidable.

□

While defining the graph search algorithm is easy on paper, defining such a function in Coq is non-trivial; we must explicitly declare termination for Coq functions. Furthermore, given the graph search function in Coq, proving the soundness and completeness of the algorithm with regard to the inductively defined converts relation also requires a significant amount of engineering; we leave it for future work.

4. Validation and prototype implementation

Mechanized Proofs The judgments and rules presented in this paper are a close approximation of the formal Coq definitions of AURA. For instance, in order to prove the preservation of pattern-matching, we have to take the parameters and arguments supplied to the constructor in the pattern-matching evaluation rules. In order to prove the decidability of type checking, we strengthened the typing judgments to take two signature arguments: one contains the type declarations of the top-level constructors that can appear in mutually recursively defined data types, and the other is used for looking up the constructors of the data types.

AURA has 20 reduction rules, 40 typing judgments including the well-formedness of terms, environments and signatures, and numerous other relations such as atomic equality types to constraint the type system. For a system of this size, implementing a fully mechanized version of the soundness proofs of the entire language is challenging.

We formalized the soundness proofs of AURA in the Coq proof assistant². We use a variant of the locally nameless representation [8] to formalize the metatheory of the language. Well documented definitions of AURA including typing-rules, reduction rules, and other related relations are about 1400 lines of Coq code. The progress and preservation proofs take about 6000 lines of Coq code. Though the automation used in these Coq proofs is relatively rudimentary, we did not devote much time to writing automation tactics by ourselves.

The most intricate parts of the language design are the invariants of the inductive data types, the dependent types, atomic equality types, and the conversion relations. This complexity is reflected in the Coq proof development in two ways: one is in the number of lemmas stating the invariants of the signatures of the data types; the other is in the number of revisions to the Coq proofs due to design changes caused by failure to prove the soundness. We found that for such a complicated system, mechanized proofs are definitely better suited for dealing with iterative revisions of the language design, since Coq could easily identify which proofs require modification when the language design changes.

Because AURA is a superset of system F plus inductively defined data types we conjecture that without much difficulty, we could extract mechanized soundness proofs of other related type systems from the Coq proofs of AURA.

Typechecker and Interpreter The prototype AURA type checker and interpreter together implement the language as it is formalized in Coq with only minor differences. The typechecker recognizes a small number of additional types and constants that are not present in the formal definition. These types include literal 32-bit integers, literal strings and tuples. Although it is derivable in AURA, we include a *fix* constant for defining recursive functions. By using this constant together with tuples, mutually recursive functions can be defined more succinctly than is possible in the formal definition. To allow for code reuse, we have added an *include* statement that performs textual substitution from external files. The software sorts included files in dependency order and copies each only once. Finally,

while the formal definition allows for implicit coercion, the prototype typechecker requires that all coercion be made explicit. The interpreter directly implements the formally defined single-step operational semantics.

AURA is not meant for general-purpose application development; instead, we intend for it to be used synergistically with existing production programming languages. To simplify the technical demands for reaching this goal, we intend to eventually target the .NET runtime, as the CLR encourages language intermingling (see Section 7). We plan to expose authorization policies written in AURA to the .NET common type system by providing libraries for interacting at runtime with propositions. We will also explore the possibilities of rewriting annotated methods in compiled .NET code to make implicit calls to these libraries. This approach should allow any language that uses the common type system to interoperate with AURA.

5. An Extended Example

In this section, we illustrate the key features of AURA's type system by explaining a program implementing a simple streaming music server.

The extended code sample is listed in Figures 4 and 5. The example program typechecks in the prototype AURA interpreter and uses some of the language extensions discussed in Section 4. At the very beginning (Line 1) the program imports library code that defines unit (with other tuple types), list, and maybe types.

We imagine that the server implements the following policy. Every song may have one or more owners, corresponding to principals who have purchased rights to play the song. Additionally, song owners may delegate their listening rights to other principals.

The rights management policy is defined over predicates *Owns* and *MayPlay*, which are declared as assertions in Lines 5 and 6. Recall that assertions are appropriate because we cannot expect to find closed proofs of ownership and delegation in pure type theory.

The main policy rule, *shareRule* (see Line 12) is defined using a say expression. The type of *shareRule* is an implication wrapped in two monads. The outer *pf* monad is required because *say* accesses a private key and must be treated effectfully. The inner *self* monad is required to track the provenance of the policy. The implication encodes the delegation policy above. The *shareRule* provides a way to build up a proof of *pf self says (MayPlay A s)*, which is required before *A* can play song *s*.

The exact form of *shareRule* is somewhat inconvenient. We derive two more convenient rules, *shareRule'* and *shareRule''* (see lines 53 and 76). These use monadic bind and return operations to change the placement of *pf* and *says* type constructors relative to *shareRule*'s type. The resulting type of *shareRule''* shows that one can obtain a proof term of *pf self says (MayPlay A s)* by a simple application of *shareRule'* to various arguments, as shown in Line 101.

The key functionality of the music server is provided by a function stub, *playFor*, which is intended to model an effectful function that streams a provided song to a specified principal. Its type is given by the annotation on line 20. The *playFor* function takes the song to be played and the principal it should play to as the first two arguments. The third argument is a proof of the proposition *self says (MayPlay A s)* demonstrating the requesting principal's capability to play the song, which is required by the server's policy. As modeling an audio API would clutter the example, *playFor* simply returns a unit value. In a real implementation, *playFor* would call into the trusted computing base, which would also log appropriate proofs for future auditing.

The remaining code implements the application's main computation. The *handleRequest* function takes a delegation request and, using a provided database of owner information, attempts to

²Code available at: <http://www.cis.upenn.edu/~stevez/so1/>


```

include "tuple.core" include "list.core" include "maybe.core"
2
data Song : Type { | freebird : Song | ironman : Song }
4
assert Owns : prin → Song → Prop;
6 assert MayPlay : prin → Song → Prop;

8 data OwnerRecord : Type {
| ownerRecord : (p : prin) → (s : Song) →
10   (pf (self says (Owns p s))) → OwnerRecord }

12 let shareRule :
  pf (self says ((o : prin) → (r : prin) → (s : Song) →
14   (Owns o s) → (o says (MayPlay r s)) → (MayPlay r s))) →
  say ((o : prin) → (r : prin) → (s : Song) →
16   (Owns o s) → (o says (MayPlay r s)) → (MayPlay r s))
  in
18   (* A real implementation would do something here *)
20 let playFor : (s : Song) → (p : prin) →
  (pf (self says (MayPlay p s))) → Unit =
22   λs : Song . λp : prin . λproof : (pf (self says (MayPlay p s))) . unit
  in
24 let notFound : (p : prin) → (s : Song) →
26   (Maybe (pf (self says (Owns p s)))) =
  λp : prin . λs : Song . Nothing (pf (self says Owns p s))
28 in

30 let getOwnerProof : (s : Song) → (p : prin) →
  (List OwnerRecord) → (Maybe (pf (self says (Owns p s)))) =
32   λs : Song . λp : prin . λownerRecords : List OwnerRecord .
  fix (λrec : (List OwnerRecord) →
34   (Maybe (pf (self says (Owns p s)))) .
  λl : (List OwnerRecord) .
36   match l with (Maybe (pf (self says Owns p s))) {
| nil → notFound p s
38   | cons → λx : OwnerRecord . λxs : List OwnerRecord .
  match x with (Maybe (pf (self says Owns p s))) {
40   | ownerRecord → λp' : prin . λs' : Song .
  λproof : pf (self says (Owns p' s')).
42     if p = p'
  then if s = s'
44       then
  Just (pf (self says (Owns p s)))
46     (proof : (pf (self says (Owns p s))))
  else rec xs
48     else rec xs
  } } )
50   ownerRecords
  in
52 let shareRule' :
54   (pf ((o : prin) → (r : prin) → (s : Song) →
  (self says (Owns o s)) → (o says (MayPlay r s)) →
56   (self says (MayPlay r s))) =
  bind shareRule (λsr : (self says
58     ((o : prin) → (r : prin) →
  (s : Song) → (Owns o s) →
60   (o says (MayPlay r s)) →
  (MayPlay r s))) .
62   return (λo : prin . λr : prin . λs : Song .
  λowns : (self says (Owns o s)) .
64   λmay : (o says (MayPlay r s)) .
  bind sr (λsr' : ((o' : prin) → (r' : prin) → (s' : Song) →
66   (Owns o' s') → (o' says (MayPlay r' s')) →
  (MayPlay r' s')) .
68   bind owns (λowns' : (Owns o s) .
  return self (sr' o r s owns' may))))))
70 in

```

Figure 4. AURA code for a music store (cont. in Figure 5).

```

76 let shareRule'' : (o : prin) → (p : prin) → (s : Song) →
  (pf self says (Owns o s)) →
78   (pf (o says (MayPlay p s))) →
  (pf self says (MayPlay p s)) =
80   λo : prin . λp : prin . λs : Song .
  λownsPf : pf (self says (Owns o s)) .
82   λplayPf : pf (o says (MayPlay p s)) .
  bind ownsPf (λopf : (self says (Owns o s)) .
84   bind playPf (λppf : (o says (MayPlay p s)) .
  bind shareRule' (λsr' :
86   ((o' : prin) → (r' : prin) → (s' : Song) →
  (self says (Owns o' s')) →
88   (o' says (MayPlay r' s')) →
  (self says (MayPlay r' s')) .
90   (return (sr' o p s opf ppp))))))
  in
92 let handleRequest : (s : Song) → (p : prin) → (o : prin) →
94   (List OwnerRecord) →
  (delPf : pf (o says (MayPlay p s))) → Unit =
96   λs : Song . λp : prin . λo : prin . λl : List OwnerRecord .
  λdelPf : pf (o says (MayPlay p s)) .
98   match (getOwnerProof s o l) with Unit {
| Nothing → unit
100   | Just → λx : (pf (self says (Owns o s))) .
  playFor s p (shareRule'' o p s x delPf)
102   }
  in unit

```

Figure 5. AURA code for a music store (cont. from Figure 4).

construct an appropriate self says *MayPlay* proof. If it succeeds *playSong* is invoked.

The implementation of *handleRequest* (line 93) is straight forward. There are two interesting things to note. First, *handleRequest* takes a database of owner information expressed as a list of *OwnerRecords*. *OwnerRecord* (line 8) is an inductive type whose single constructor has a dependent type. Because *ownerRecord*'s third argument depends on its first two, *OwnerRecord* encodes an existential type. Second, the match expression on line 98 relies on the fact that (*getOwnerProof s o l*) returns an object of type *Maybe (pf (self says (Owns p s)))*. Getting such a type is possible because when *getOwnerProof* pulls a proof from the list, its type is refined so that the existentially bound principal and song are identified with *p* and *s*.

GetOwnerProof (line 30) performs this type refinement in several steps. It uses the fixpoint combinator (line 33) to perform a list search. After each *OwnerRecord* is decomposed, we must check its constituent parts to determine if it is the correct record and, if so, refine the types appropriately. The action occurs between lines 42 and 48. At runtime the first if expression tests for dynamic equality between the principal we're searching for, *p*, and the principal store in the current record, *p'*. A similar check is performed for between *Songs s* and *s'*. If both checks succeed then we cast *proof*: pf (self says *Owns p' s'*) to type pf (self says *Owns p s*) and return it packaged as a *Maybe*. If either dynamic check fails we repeat again and, if no match is found, eventually return *Nothing*.

6. Related Work

We have published related results on AURA₀, a language closely related to the Prop fragment of AURA [38]. This includes soundness and sound normalization proofs for AURA₀, as well as discussion and examples of audit in the presence of authorization proofs.

One intended semantics for AURA implements objects of form *sign(A, P)* as digital signatures. All cryptography occurs at a lower level of abstraction than the language definition. This approach

has previously be used to implement declarative information flow policies [39]. An alternative approach is to treat keys as types or first class objects and to provide encryption or signing primitives in the language [6, 14, 34, 26, 25]. Such approaches typically provide the programmer with additional flexibility, but complicate the programming model.

Authorization logics Many logics and languages [4, 5, 13, 10, 20, 2, 19] have tracked authorization using *says*. We follow the approach of DCC [2], the logic in which *says* was first defined as an indexed monad. This is compelling for several reasons. First, DCC proofs are lambda-terms, a fact we exploit to closely couple the Prop and Type universes. Second, DCC is a strong logic and important authorization concepts, such as the *acts-for* relation and the hand-off rule ($A \text{ says } B \text{ acts-for } A \rightarrow (B \text{ acts-for } A)$), can be defined or derived. Third, DCC is known to enjoy a non-interference property: in the absence of delegation, statements in the *A says* monad will not effect the *B says* monad. In our setting this means that a given program cannot be tricked by what an untrusted program says. AURA modifies DCC in several ways. In addition to adding dependent types, AURA omits DCC’s *protects* relation. The *protects* relation strengthens monadic bind, making propositions $A \text{ says } (B \text{ says } P)$ and $B \text{ says } (A \text{ says } P)$ interderivable. While useful in other settings, such equivalences appear incorrect for access control. Additionally AURA’s use of signatures changes some meta-theoretic properties of the DCC leading to, for example, a more subtle proof of normalization [38].

Fournet, Gordon and Maffeis [20, 21] discuss authorization logic in the context of distributed systems. They use a limited form dependent pairs to associate propositions with data. Unlike in AURA proofs are erased at runtime. Consequently, their type discipline is best suited for closed systems that do not require high-assurance logging.

The Grey project [10] uses proof carrying authorization in manner similar in propose to AURA. In Grey, mobile phone handsets build authorization proofs that unlock doors. While AURA is a unified authorization logic and computation language, Grey’s logic is not integrated with a computation language.

DeYoung, Garg, and Pfenning [19] describe a constructive authorization logic that is parameterized by a notation of time. Propositions and proofs are annotated with time intervals during which they may be judged valid. This allows revocation to be modeled as credential expiration.

The trust management system PolicyMaker [12] treats access control decisions as a distributed programming problem. A PolicyMaker *assertion* is a pair containing a function and (roughly speaking) a principal. In general, assertion functions may communicate with each other, and each function’s output is tagged by the associated principal. PolicyMaker checks if a request complies with policy by running all assertions functions and seeing if they produce an output in which distinguished principal POLICY says approve. Principal tags appear similar is purpose, but not realization, to *says* modalities in AURA. Note also that expressing security properties via term-level computation is fundamentally different than expressing them as types, the approach followed in most other work discussed here. The ideas in PolicyMaker have been refined in KeyNote [12] and REFEREE [15].

The Fable language [36] associates security labels with data values. Labels may be used to encode information flow, access control, and other policies. Technically, labels are terms which may be referred to at the type level; *colored* judgments are used to separate the data and label worlds. The key security property is that standard computations (i.e. application computations described with color *app*) are parametric in their labeled inputs. Unlike AURA proofs, the label sub-language (i.e. policy computations described with color *pol*) admits arbitrary recursion. Hence the color separa-

tion may restrict security sensitive operations to a small trusted computing base, but does not give rise to a logical soundness property.

Dependent type theory The AURA language design was influenced by dependent type systems like the Calculus of Constructions (CoC) [9, 17], and proof carrying authorization logics, especially Dependency Core Calculus (DCC) [2]. Both CoC and AURA contain dependent types and a unified syntax encompassing both types and terms. However there are several important differences between CoC and AURA. Most critically, CoC quotients type equality by beta-equivalence but AURA does not. Type-level beta reduction, while convenient for verification, is unnecessary for expressing authorization predicates, and greatly complicates language design and use.

As realized in the Coq Proof Assistant [16], CoC can contain inductive types and different universes for computation and logic types—AURA universes Prop and Type correspond to Prop and Set in Coq. However, because Set is limited to pure computations, Coq does not wrap Props in a pf monad. In Coq all inductive declarations are subject to a complex *positivity* constraint which ensures inductive types have a well-defined logical interpretation. In contrast AURA uses a simpler positivity constraint in Prop and no constraint in Type. Additionally, AURA performs less type refinement than Coq for GADTs/type indices. When compared with Coq, AURA is strictly weaker for defining logical predicates, but can define certain stronger algebra datatypes for use in computation.

Several other projects have combined dependent types and pragmatic language design. Ynot (an embedding of Hoare Type Theory [30] in Coq), Agda [31], and Epigram [28] are intended to support general purpose program verification and usually require that the programmer construct proofs interactively. In contrast Dependent ML [45], ATS [45, 44], and RSP1 [42] provide distinguished dependency domains and can only express constraints on objects from these domains. These dependency domains are intended to be amenable to automated analysis. Cayenne [7] extends Haskell with general purpose dependent types. In Cayenne type equality is checked by normalizing potentially divergent Haskell terms—a strategy which may cause type checking itself to diverge. Hancock and Setzer [22] present a core calculus for interactive programming in dependent type theory. Their language uses an IO monad to encapsulate stateful computations. Inhabitants of the monad are modeled as imperative programs and type equality is judged up to a bisimulation on (imperative) program text.

Peyton Jones and Meijer describe the Henk typed intermediate language [24]. Henk is an extension of the lambda cube family of programming languages that includes CoC. Like AURA, Henk is intended to be a richly-typed compiler intermediate language. Unlike AURA, Henk has not been proved sound. Additionally, its lack of a pf monad (or equivalent technique for isolating computations from proofs) makes it unsuitable for programming in the presence of both dependent types and effects.

7. Future Work

Future work: Theory As discussed in Section 3, we have proved that the degree to which AURA restricts dependency allows for tractable typechecking. This does not, however, rule out the possibility that this tractability could be preserved if less restrictive sorts of dependency were added to AURA. In particular, it may well be useful to look at the simulation of dependency with GADTs [32] in search of examples that, while relatively simple, cannot be handled directly in AURA, and to look for ways to extend AURA with support for such features without losing decidability.

Section 2 describes the correspondence between $a \text{ says } P$ and objects digitally signed by a ’s private key. It is natural,

then, to wonder about the possibility of an analog for public key encryption—perhaps terms of type T for a could be constructed from objects of type T encrypted with a 's public key. It is unclear how precisely to integrate such an additional monad with AURA, however, not least because, while the says monad makes complete sense operating only at the Prop level, we almost certainly want to encrypt data of kind Type. Additionally, our use of dependent types means that the type of an AURA term will often reference part of the term itself, which may well be unacceptable for data that is meant to be encrypted.

The tracking of information flow was one of the first uses proposed for DCC [3], and even without encryption AURA's assertions are sometimes reminiscent of confidentiality tracking; were encryption to be added, the similarities would be even more pronounced. It may be possible to take advantage of this by equipping AURA with a more general notion of information flow—which does not necessarily have as straightforward a cryptographic interpretation—for use internal to a single well-typed application while reverting to the coarser-grained says (and possibly for) when communication with the outside world is desired. The challenge, of course, is to make this change of granularities as fluid and unencumbering as possible.

Even without information flow it may still be useful to have a better idea of which proofs may come from the outside world. After all, operations on digital signatures are not trivial, but since proofs are defined to be computation-free, a purely local proof could be given a more efficient but less portable representation, and certain proofs might be completely elided at runtime. For the first case, we would first need to extend our formalism with some notion of network communication; inference could be performed backwards from communication points to ascertain which proofs need not be represented in portable form. As an initial step towards recognizing the second case, we might consider an additional form of abstraction, with an argument that cannot be used in certain ways but is guaranteed to be necessary at compile time only; ideally, however, we would want to infer these abstractions as part of compilation.

It is clear from our examples that AURA is fairly verbose. As it is meant to be an intermediate language, this is not a pressing usability issue. We hope that a higher-level language that generates AURA will be able to cut down on this verbosity using inference techniques. Our proof-passing style also suggests the use of some variety of proof inference. Of course, this very quickly becomes undecidable, but that does not rule out practical partial solutions.

Finally, although AURA emphasizes the security aspects of programming with an embedded authorization logic, there might be other applications of this idea. In particular, one of the challenges of making program verification via dependent types practical is the need to construct and otherwise manipulate proof objects. Of course, one can always add axioms to the logic, but doing so can easily compromise its consistency. Failures due to a poor choice of axioms might be hard to isolate when debugging. The says monads of DCC provide a possible intermediate ground: One could imagine associating a principal with each module of the program and then allowing modules to make assertions. Explicit trust delegations would then be required when importing axioms from one module to another; such delegations would document the module dependencies and help the type checker isolate uses of faulty axioms. We speculate that it is even possible that blame (in the style similar to that proposed by Wadler and Findler [41]) can be appropriately assigned to offending modules whenever a run-time error caused by incorrect assertions is encountered.

Future work: Practice The single-step interpreter is useful as a tool for checking the correctness of small examples; however, it is infeasible to use it to run code in a production environment. As

such, we are extending the implementation to generate CIL compatible with both Microsoft's .NET CLR and the open-source Mono runtime. Don Syme's work on ILX aids us greatly in this effort. ILX, described in [37], is a group of extensions to the CIL that facilitates the use of higher-order functions, discriminated unions and parametric polymorphism. By compiling for this existing standard execution environment, we will gain access to the ecosystem of .NET software and libraries. Most notably, we should be able to make use of existing code for cryptography and cross-platform networking. We will also be free from having to worry about lower-level issues like efficient machine code generation and garbage collection, both of which are well outside of the AURA project's scope.

Additionally, there remain practical issues that AURA must address in order to fully express policies likely to be found in its intended problem domain. Chief among these is the demand for the signatures that *expire*, either due to explicit revocation or simply the passage of time. This stands in contrast to our current formalism—and, indeed, most formalisms of programming languages, as a term that successfully typechecks is generally seen as valid for regardless of the time or the state of the world. It would, of course, be possible to define the operational semantics of AURA such that every operation has a chance to fail at runtime due to digital signature expiration, but this would quickly make programming quite cumbersome. Instead, we hope to find a solution that allows time and revocation to be referenced by AURA in an intuitive way; one possibility is, explored by Garg and Pfenning [19], is the use of *linear* logic, which is naturally suited to describing resources that can, in some sense, be used up.

References

- [1] Martín Abadi. Logic in access control. In *Proceedings of the 18th Annual Symposium on Logic in Computer Science (LICS'03)*, pages 228–233, June 2003.
- [2] Martín Abadi. Access control in a core calculus of dependency. *Electronic Notes in Theoretical Computer Science*, 172:5–31, April 2007.
- [3] Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon Riecke. A core calculus of dependency. In *Proc. 26th ACM Symp. on Principles of Programming Languages (POPL)*, pages 147–160, San Antonio, TX, January 1999.
- [4] Martín Abadi, Michael Burrows, Butler W. Lampson, and Gordon D. Plotkin. A calculus for access control in distributed systems. *Transactions on Programming Languages and Systems*, 15(4):706–734, September 1993.
- [5] 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, New York, NY, USA, 1999. ACM.
- [6] Aslan Askarov, Daniel Hedin, and Andrei Sabelfeld. Cryptographically masked information flows. In *Proceedings of the International Static Analysis Symposium, LNCS*, Seoul, Korea, August 2006.
- [7] Lennart Augustsson. Cayenne—a language with dependent types. In *Proc. 3rd ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 239–250, September 1998.
- [8] Brian E. Aydemir, Arthur Charguraud, Benjamin C. Pierce, Randy Pollack, and Stephanie Weirich. Engineering formal metatheory. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, pages 3–15. ACM, 2008.
- [9] Henk P. Barendregt. Lambda calculi with types. In Samson Abramsky, Dov M. Gabbay, and Thomas S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 2, pages 117–309. Clarendon Press, Oxford, 1992.
- [10] Lujao Bauer, Scott Garriss, Jonathan M. McCune, Michael K. Reiter, Jason Rouse, and Peter Rutenbar. Device-enabled authorization in the

- Grey system. In *Information Security: 8th International Conference, ISC 2005*, volume 3650 of *Lecture Notes in Computer Science*, pages 431–445, September 2005.
- [11] Matt Bishop. *Computer Security: Art and Science*. Addison-Wesley Professional, 2002.
- [12] Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis. KeyNote: Trust management for public-key infrastructures (position paper). *Lecture Notes in Computer Science*, 1550:59–63, 1999.
- [13] J.G. Cederquist, R. Corin., M.A.C. Dekker, S. Etalle, and J.J. den Hartog. An audit logic for accountability. In *The Proceedings of the 6th IEEE International Workshop on Policies for Distributed Systems and Networks*, 2005.
- [14] Tom Chothia, Dominic Duggan, and Jan Vitek. Type based distributed access control. In *Proceedings of the 16th IEEE Computer Security Foundations Workshop (CSFW'03)*, Asilomar, Ca., USA, July 2003.
- [15] Yang-Hua Chu, Joan Feigenbaum, Brian LaMacchia, Paul Resnick, and Martin Strauss. REFEREE: Trust management for web applications. *Computer Networks and ISDN Systems*, 29:953–964, 1997.
- [16] The Coq Development Team, LogiCal Project. *The Coq Proof Assistant Reference Manual*, 2006.
- [17] T. Coquand and G. Huet. The calculus of constructions. *Information and Computation*, 76, 1988.
- [18] Haskell B. Curry, Robert Feys, and William Craig. *Combinatory Logic*, volume 1. North-Holland, Amsterdam, 1958.
- [19] Henry DeYoung, Deepak Garg, and Frank Pfenning. An authorization logic with explicit time. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF-21)*, Pittsburgh, June 2008. To appear.
- [20] Cédric Fournet, Andrew D. Gordon, and Sergio Maffeis. A type discipline for authorization policies. In *Proc. of the 14th European Symposium on Programming*, volume 3444 of *LNCS*, pages 141–156, Edinburgh, Scotland, April 2005. Springer-Verlag.
- [21] Cédric Fournet, Andrew D. Gordon, and Sergio Maffeis. A type discipline for distributed systems. In *Proc. of the 20th IEEE Computer Security Foundations Symposium*, pages 31–45, Venice, Italy, July 2007.
- [22] Peter Hancock and Anton Setzer. Interactive programs in dependent type theory. In *Proceedings of the 14th Annual Conference of the EACSL on Computer Science Logic*, pages 317–331, London, UK, 2000. Springer-Verlag.
- [23] W. A. Howard. The formulae-as-types notion of construction. In J. P. Seldin and J. R. Hindly, editors, *To H. B. Curry: Essays on Combinatory Logic, Lambda-Calculus, and Formalism*, pages 479–490. Academic Press, New York, 1980.
- [24] Simon Peyton Jones and Erik Meijer. Henk: A typed intermediate language. In *Proceedings of the Types in Compilation Workshop*, Amsterdam, June 1997.
- [25] Peeter Laud. On the computational soundness of cryptographically masked flows. *SIGPLAN Not.*, 43(1):337–348, 2008.
- [26] Peeter Laud and Varmo Vene. A type system for computationally secure information flow. In *Proceedings of the 15th International Symposium on Fundamentals of Computational Theory*, volume 3623, pages 365–377, Lbeck, Germany, 2005.
- [27] Daniel K. Lee, Karl Crary, and Robert Harper. Towards a mechanized metatheory of Standard ML. In *POPL '07: Proceedings of the 34th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 173–184, New York, NY, USA, 2007. ACM.
- [28] Conor McBride. *The Epigram Prototype: a nod and two winks*, April 2005. Available from <http://www.e-pig.org/downloads/epigram-system.pdf>.
- [29] Andrew C. Myers, Stephen Chong, Nathaniel Nystrom, Lantian Zheng, and Steve Zdancewic. Jif: Java information flow. 1999.
- [30] A. Nanevski, G. Morrisett, and L. Birkedal. Polymorphism and separation in Hoare Type Theory. In *Proc. 11th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, 2006.
- [31] Ulf Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, Department of Computer Science and Engineering, Chalmers University of Technology, SE-412 96 Göteborg, Sweden, September 2007.
- [32] Simon Peyton Jones, Dimitrios Vytiniotis, Stephanie Weirich, and Geoffrey Washburn. Simple unification-based type inference for GADTs. In *ICFP '06: Proceedings of the Eleventh ACM SIGPLAN International Conference on Functional Programming*, pages 50–61, New York, NY, USA, 2006. ACM Press.
- [33] François Pottier and Vincent Simonet. Information flow inference for ML. *ACM Trans. Program. Lang. Syst.*, 25(1):117–158, 2003.
- [34] Geoffrey Smith and Rafael Alparz. Secure information flow with random assignment and encryption. In *Proceedings of The 4th ACM Workshop on Formal Methods in Security Engineering: From Specifications to Code (FSME'06)*, pages 33–43, Alexandria, Virginia, USA, November 2006.
- [35] Martin Sulzmann, Manuel M. T. Chakravarty, Simon Peyton Jones, and Kevin Donnelly. System F with type equality coercions. In *TLDI '07: Proceedings of the 2007 ACM SIGPLAN international workshop on Types in languages design and implementation*, pages 53–66, New York, NY, USA, 2007. ACM.
- [36] 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. To appear.
- [37] Don Syme. ILX: Extending the .NET Common IL for functional language interoperability. *Electronic Notes in Theoretical Computer Science*, 59(1), 2001.
- [38] Jeffrey A. Vaughan, Limin Jia, Karl Mazurak, and Steve Zdancewic. Evidence-based audit. In *Proc. of the IEEE Computer Security Foundations Symposium*, 2008. To appear. Extended version available as U. Pennsylvania Technical Report MS-CIS-08-09.
- [39] Jeffrey A. Vaughan and Steve Zdancewic. A cryptographic decentralized label model. In *IEEE Symposium on Security and Privacy*, pages 192–206, Berkeley, California, 2007.
- [40] Philip Wadler. Monads for functional programming. In J. Jeuring and E. Meijer, editors, *Advanced Functional Programming*, volume 925 of *LNCS*. Springer Verlag, 1995. Some errata fixed August 2001.
- [41] Philip Wadler and Robert Bruce Findler. Well-typed programs can't be blamed. In *Workshop on Scheme and Functional Programming*, pages 15–26, 2007.
- [42] E. Westbrook, A. Stump, and I. Wehrman. A Language-based Approach to Functionally Correct Imperative Programming. In B. Pierce, editor, *10th ACM SIGPLAN International Conference on Functional Programming*, Tallinn, Estonia, 2005.
- [43] Edward Wobber, Martín Abadi, Michael Burrows, and Butler Lampson. Authentication in the Taos operating system. *ACM Trans. Comput. Syst.*, 12(1):3–32, 1994.
- [44] Hongwei Xi. Applied Type System (extended abstract). In *post-workshop Proceedings of TYPES 2003*, pages 394–408. Springer-Verlag LNCS 3085, 2004.
- [45] Hongwei Xi and Frank Pfenning. Dependent types in practical programming. In *Proc. 26th ACM Symp. on Principles of Programming Languages (POPL)*, San Antonio, Texas, September 1998.