# Evidence-based Audit, Technical Appendix

Jeffrey A. Vaughan     Limin Jia     Karl Mazurak     Steve Zdancewic

University of Pennsylvania
Department of Compture and Information Science
Technical Report Number MS-CIS-08-09

## Abstract

*Authorization logics provide a principled and flexible approach to specifying access control policies. One of their compelling benefits is that a proof in the logic is evidence that an access-control decision has been made in accordance with policy. Using such proofs for auditing reduces the trusted computing base and enables the ability to detect flaws in complex authorization policies. Moreover, the proof structure is itself useful, because proof normalization can yield information about the relevance of policy statements. Untrusted, but well-typed, applications that access resources through an appropriate interface must obey the access control policy and create proofs useful for audit.*

*This paper presents* AURA$_0$, *an authorization logic based on a dependently-typed variant of DCC and proves the metatheoretic properties of subject-reduction and normalization. It shows the utility of proof-based auditing in a number of examples and discusses several pragmatic issues that must be addressed in this context.*

## 1 Introduction

*Logging*, *i.e.* recording for subsequent audit significant events that occur during a system's execution, has long been recognized as a crucial part of building secure systems. A typical use of logging is found in a firewall, which might record the access control decisions that it makes when deciding whether to permit connection requests. In this case, the log might consist of a sequence of time stamped strings written to a file where each entry indicates some information about the nature of the request (IP addresses, port numbers, *etc.*) and whether the request was permitted. Other scenarios place more stringent requirements on the log. For example, a bank server's transactions log should be tamper resistant, and log entries should be authenticated and not easily forgeable. Logs are useful because they can help administrators audit the system both to identify sources of

unusual or malicious behavior and to find flaws in the authorization policies enforced by the system.

Despite the practical importance of auditing, there has been surprisingly little research into what constitutes good auditing procedures.[1] There has been work on cryptographically protecting logs to prevent or detect log tampering [29, 11], efficiently searching confidential logs [32], and experimental research on effective, practical logging [6, 26]. But there is relatively little work on *what* the contents of an audit log should be or how to ensure that a system implementation performs appropriate logging (see Wee's paper on a logging and auditing file system [33] for one approach to these issues, however).

In this paper, we argue that audit log entries should constitute *evidence* that justifies the authorization decisions made during the system's execution. Following an abundance of prior work on authorization logic [4, 24, 17, 1, 27, 2, 21], we adopt the stance that log entries should contain *proofs* that access should be granted. Indeed, the idea of logging such proofs is implicit in the proof-carrying authorization literature [5, 7, 10], but, to our knowledge, the use of proofs for auditing purposes has not been studied outright.

There are several compelling reasons why it is advantageous to include proofs of authorization decisions in the log. First, by connecting the contents of log entries directly to the authorization policy (as expressed by a collection of rules stated in terms of the authorization logic), we obtain a principled way of determining what information to log. Second, proofs contain structure that can potentially help administrators find flaws or misconfigurations in the authorization policy. Third, storing verifiable evidence helps reduce the size of the trusted computing base; if every access-restricting function automatically logs its arguments and re-

---

[1]Note that the term auditing can also refer to the practice of *statically* validating a property of the system. Code review, for example, seeks to find flaws in software before it is deployed. Such auditing is, of course, very important, but this paper focuses on *dynamic* auditing mechanisms such as logging.

sult, the reasoning behind any particular grant of access cannot be obscured by a careless or malicious programmer.

The impetus for this paper stems from our experience with the (ongoing) design and implementation of a new security-oriented programming language called AURA [25]. The primary goal of this work is to find mechanisms that can be used to simplify the task of manipulating authorization proofs and to ensure that appropriate logging is always performed regardless of how a reference monitor is implemented. Among other features intended to make building secure software easier, AURA provides a built-in notion of principals, and its type system treats authorization proofs as first-class objects; the authorization policies may themselves depend on program values.

This paper focuses on the use of proofs for logging purposes and the way in which we envision structuring AURA software to take advantage of the authorization policies to minimize the size of the trusted computing base. The main contributions of this paper can be summarized as follows.

Section 2 proposes a system architecture in which logging operations are performed by a trusted kernel, which can be thought of as part of the AURA runtime system. Such a kernel accepts proof objects constructed by programs written in AURA and logs them while performing security-relevant operations.

To illustrate AURA more concretely, Section 3 develops a dependently typed authorization logic based on DCC [2] and similar to that found in the work by Gordon, Fournet, and Maffeis [19, 20]. This language, $\text{AURA}_0$, is intended to model the fragment of AURA relevant to auditing. We show how proof-theoretic properties such as subject reduction and normalization can play a useful role in this context. Of particular note is the normalization result for $\text{AURA}_0$ authorization proofs.

Section 4 presents an extended example of a file system interface; as long as a client cannot circumvent this interface, any reference monitor code is guaranteed to provide appropriate logging information. This example also demonstrates how additional domain-specific rules can be built on top of the general kernel interface, and how the logging of proofs can be useful when it isn't obvious which of these rules are appropriate.

Of course, there are many additional engineering problems that must be overcome before proof-enriched auditing becomes practical. Although it is not our intent to address all of those issues here, Section 5 highlights some of the salient challenges and sketches future research directions. Section 6 discusses related work. The appendix contains definitions and proofs elided from the accompanying conference paper [31].

## 2 Kernel Mediated Access Control

A common system design idiom protects a resource with a *reference monitor*, which takes requests from (generally) untrusted clients and decides whether to allow or deny access to the resource [12]. Ideally a reference monitor should be configured using a well-specified set of *rules* that define the current access-control policy and mirror the intent of some institutional policy.

Unfortunately, access-control decisions are not always made in accordance with institutional intent. This can occur for a variety of reasons including the following:

1. The reference monitor implementation or rule language may be insufficient to express institutional intent. It this case, the rules must necessarily be too restrictive or too permissive.

2. The reference monitor may be configured with an incorrect set of rules.

3. The reference monitor itself may be buggy. That is, it may reach an incorrect decision even when starting from correct rules.

The first and second points illustrate an interesting tension: rule language expressiveness is both necessary and problematic. While overly simple languages prevent effective realization of policy intent, expressive languages make it more likely that a particular rule set has unintended consequences. The latter issue is particularly acute in light of Harrison and colleagues' observation that determining the ultimate effect of policy changes—even in simple systems—is generally undecidable [23]. The third point recognizes that reference monitors may be complex and consequently vulnerable to implementation flaws.

The AURA programming model suggests a different approach to protecting resources, illustrated in Figure 1. There are three main components in the system: a trusted kernel, an untrusted application, and a set of rules that constitute the formal policy. The kernel itself contains a log and a resource to be protected. The application may only request resource access through kernel interface $I_K$. This interface (made up of the $\text{op}_i$s in the figure) wraps each of the resource's native operations (the raw-$\text{op}_i$s) with new operations taking an additional argument—a proof that access is permitted. $\Sigma_K$ and $\Sigma_{ext}$ contain constant predicate symbols that may be occur in these proofs.

Unlike in the standard reference monitor model, an AURA kernel forwards every well-typed request to its underlying resource. Each $\text{op}_i$ function takes as an additional argument a proof that the operation is permitted and returns a corresponding proof that the operation was performed, so the well-typedness of a call ensures that the requested access is permitted. Proofs can be typechecked dynamically
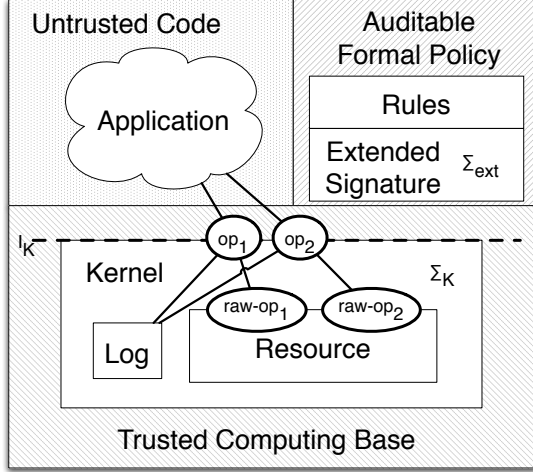
**Figure 1. A monolithic application decomposed into several components operating with various degrees of trust.**

in time linearly proportional to the size of the proof should the request not come from a well-typed application. Moreover, logging these proofs is enough to allow an auditor to ascertain precisely why any particular access was allowed.

We define a language $\text{AURA}_0$ to provide an expressive policy logic for writing rules and kernel interface types. It is a cut-down version of full AURA [25], which itself is a polymorphic and dependent variant of Abadi's Dependency Core Calculus [3, 2]. In $\text{AURA}_0$, software components may be explicitly associated with one or more principals. Typically, a trusted kernel is identified with principal K, and an untrusted application may work on behalf of several principals: A, B, etc. Principals can make assertions; for instance, the (inadvisable) rule "the kernel asserts that all principals may open any file," is written as proposition K **says** $((A{:}\textbf{prin}) \rightarrow (f{:}\textbf{string}) \rightarrow \text{OkToOpen } A \ f)$. Evidence for this rule contains one or more signature objects—possibly implemented as cryptographic signatures—that irrefutably tie principals to their utterances.

The above design carries several benefits. Kernels log the reasoning used to reach access control decisions; if a particular access control decision violates policy intent but is allowed by the rules, audit can reveal which rules contributed to this failure. Additionally, because all resource access is accompanied by a proof, the trusted computing base is limited to the proof checker and kernel. As small, standard programs these components are less likely to suffer from software vulnerabilities than than traditional, full-scale reference monitors.

A key design principle is that kernels should be small and general; this is realized by removing complex, specialized reasoning about policy (e.g. proof search) from the trusted

computing base. In this sense, AURA systems are to traditional reference monitors as operating system microkernels are to monolithic kernels.

## 2.1 The formal system description

We model a system consisting of a trusted kernel $K$ wrapping a security-oblivious resource $R$ and communicating with an untrusted application. The kernel is the trusted system component that mediates between the application and resource by checking and logging access control proofs; we assume that applications are prevented from accessing resources directly by using standard resource isolation techniques deployed in operating systems or type systems.

A resource $R$ is a stateful object with a set of operators that may query and update its state. Formally, $R = (\sigma, \textit{States}, I_R, \delta)$ where $\sigma \in \textit{States}$ and

$$I_R = \text{raw-op}_1 : T_1 \Rightarrow S_1, \ldots, \text{raw-op}_n : T_n \Rightarrow S_n$$

The current state $\sigma$ is an arbitrary structure that representing $R$'s current state, and $\textit{States}$ is the set of all possible resource states. $I_R$ is the resource's interface; each $\text{raw-op}_i : T_i \Rightarrow S_i$ is an operator with its corresponding type signature. The transition function $\delta$ describes how the raw operations update state, as well as their input-output behavior. For instance, $(u, \sigma') = \delta(\text{raw-op}_i, v, \sigma)$ when raw operation $i$—given input $v$ and initial resource state $\sigma$—produces output $u$ and updates the resource state to $\sigma'$.

We formalize a trusted kernel $K$ as a tuple $(L, R, \Sigma_K, I_K)$; the authority of the kernel is denoted by the constant principal K. The first component, $L$, is a list of proofs representing the log. The second component is the resource encapsulated by the kernel. Signature $\Sigma_K$ contains pairs of predicates, $\text{OkToOp}_i : T_i \rightarrow \textbf{Prop}$ and $\text{DidOp}_i : T_i \rightarrow S_i \rightarrow \textbf{Prop}$ for each $\text{raw-op}_i$ of type $T_i \Rightarrow S_i$ in $I_R$. These predicates serve as the core lexicon for composing access control rules: a proof of K **says** $\text{OkToOp } t$ signifies that an operation raw-op is permitted with input $t$, and a proof of K **says** $\text{DidOp } t \ s$ means that raw-op was run with input $t$ and returned $s$. Lastly, the kernel exposes an application programming interface $I_K$, which contains a security-aware wrapper operation

$$\text{op}_i : (x : T_i) \Rightarrow K \ \textbf{says} \ (\text{OkToOp}_i \ x) \Rightarrow$$
$$\{y{:}S_i; K \ \textbf{says} \ \text{DidOp}_i \ x \ y\}$$

for each $\text{raw-op}_i$ in $I_R$. Applications must access $R$ through $I_K$ rather than $I_R$.

The type of $\text{op}_i$ shows that the kernel requires two arguments before it will provide access to $\text{raw-op}_i$. The first argument is simply $\text{raw-op}_i$'s input; the second is a proof that the kernel approves the operation, typically a composition of policy rules (globally known statements signed by

K) and statements made by other relevant principals. The return value of $op_i$ is a pair of $\mathsf{raw\text{-}op}_i$'s output with a proof that acts as a receipt, affirming that the kernel called $\mathsf{raw\text{-}op}_i$ and linking the call's input and output. Note that $\mathsf{OkToOp}_i$ and $\mathsf{DidOp}_i$ depend on the arguments $x$ and $y$.

The final components in the model are the application, the rule set, and the extended signature. We assume either that the application is well-typed—and thus that it respects $I_K$—or, equivalently, that the kernel performs dynamic typechecking on incoming untrusted arguments. The rule set is simply a well-known set of proofs intended to represent some access control policy; the extended signature ($\Sigma_{ext}$ in Figure 1) defines predicate symbols that these rules may use in addition to those defined in $\Sigma_K$.

**Remote procedure call example**  Consider a simple remote procedure call resource with only the single raw operation, $\mathsf{raw\text{-}rpc} : \textbf{string} \Rightarrow \textbf{string}$. The kernel associated with this resource exposes the following predicates:

$$\Sigma_K = \; \mathsf{OkToRPC} : \textbf{string} \to \textbf{Prop},$$
$$\mathsf{DidRPC} : \textbf{string} \to \textbf{string} \to \textbf{Prop}$$

and the kernel interface

$$I_K = \mathsf{rpc} : (x : \textbf{string}) \Rightarrow \mathsf{K} \; \textbf{says} \; \mathsf{OkToRPC} \; x \Rightarrow$$
$$\{y{:}\textbf{string}; \mathsf{K} \; \textbf{says} \; \mathsf{DidRPC} \; x \; y\}.$$

A trivial policy could allow remote procedure call. This policy is most simply realized by the singleton rule set $Rules = \{r_0 : \mathsf{K} \; \textbf{says} \; ((x{:}\textbf{string}) \to \mathsf{OkToRPC} \; x)\}$.

## 2.2  State transition semantics

While the formalism presented thus far is sufficient to describe what $\mathrm{AURA}_0$ systems look like at one instant in time, it is much more interesting to consider an evolving system. Here we describe variant operational semantics of the $\mathrm{AURA}_0$ system at a semi-formal level, with emphasis on logging. The full $\mathrm{AURA}$ language includes a computation fragment capable of expressing the ideas in this section by way of a standard monadic state encoding, although its analysis by Jia and colleagues [25] does not address logging directly.

To demonstrate the key components of authorization and auditing in $\mathrm{AURA}_0$, we consider evaluations from three perspectives listed as follows. In each we will consider updating states according to the transition relations defined in Figure 2.

1. Resource evaluation, written with $-\{\}\mapsto^r$, models the state transition for raw resources. This relation does no logging and does not consider access control.

2. Logged evaluation, written with $-\{\}\mapsto^l$, models state transitions of an $\mathrm{AURA}_0$ system implementing logging as described in this paper. All proofs produced or consumed by the kernel are recorded in the log.

3. Semi-logged evaluation, written with $-\{\}\mapsto^s$, models the full system update with weaker logging. While proofs are still required for access control, the log contains only operation names, not the associated proofs.

Resource evaluation is the simplest evaluation system. A transition $I_R; \delta \vdash \sigma -\{\mathsf{raw\text{-}op}_i v\}\mapsto^r \sigma'$ may occur when $v$ is a well typed input for $\mathsf{raw\text{-}op}_i$ according to resource interface $I_R$ and $\delta$ specifies that $\mathsf{raw\text{-}op}_i$, given $v$ and starting with a resource in state $\sigma$, returns $u$ and updates the resource state to $\sigma'$. (In the following we will generally omit the $\vdash$ and objects to its left, as they are constant and can be inferred from context.)

The logged evaluation relation is more interesting: instead of simply updating resource states, it updates configurations. A configuration $C$, is a triple $(L, \sigma, \mathcal{S})$, where $L$ is a list of proofs representing a log, $\sigma$ is an underlying resource state, and $\mathcal{S}$ is a set of proofs of the form $\textbf{sign}(A, P)$ intended to track all assertions made by principals. There are two logged evaluation rules, L-SAY and L-ACT.

Intuitively, L-SAY allows principals other than the kernel K to add objects of the form $\textbf{sign}(A, P)$ to $\mathcal{S}$, corresponding to the ability of clients to sign arbitrary propositions, as long as all of signatures found within $P$ already appear in $\mathcal{S}$. This last condition is written $\mathcal{S} \vDash P$ and prevents principals from forging evidence—in particular, from forging evidence signed by K. $\mathcal{S} \vDash P$ holds when all signatures embedded in $P$ appear in $\mathcal{S}$.

Rule L-ACT models the use of a resource through it's public interface. The rules ensure that both of the operation's arguments—the data component $v$ and the proof $p$—are well typed, and all accepted access control proofs are appended to the log. After the resource is called through its raw interface, the kernel signs a new proof term, $q$, as a receipt; it is both logged and added to $\mathcal{S}$. Again, the premise $\mathcal{S} \vDash p$ guarantees the unforgeability of $\textbf{sign}$ objects.

The semi-logged relation functions similarly (see rules S-SAY and S-ACT), although it logs only the list of operations performed rather than any proofs.

By examining the rules in Figure 2, we can see that the kernel may only sign $\mathsf{DidOp}$ receipts during evaluation. Since statements signed by any other principal may be added to $\mathcal{S}$ at any time, we may identify the initial set of sign objects in $\mathcal{S}$ with the system's policy rules.

**Audit and access control**  The three transition relations permit different operations and record different information about allowed actions. Resource evaluation allows all well-typed calls to the raw interface, and provides no information

4

Resource evaluation relation $\quad \cdot;\cdot \vdash \cdot \dashv\{\cdot\}\mapsto^r \cdot$

$$\frac{\cdot;\cdot \vdash v:T \qquad \text{raw-op}_i : T \Rightarrow S \in I_R \qquad (\_, \sigma') = \delta(\text{raw-op}_i, v, \sigma)}{I_R; \delta \vdash \sigma \dashv\{\text{raw-op}_i\ v\}\mapsto^r \sigma'}\ \text{R-ACT}$$

Semi-logged evaluation relation $\quad \cdot;\cdot;\cdot \vdash \cdot \dashv\{\cdot\}\mapsto^s \cdot$

$$\frac{\begin{array}{c} \text{op}_i : (x:T) \Rightarrow \mathsf{K}\ \textbf{says}\ \mathsf{OkToOp}_i\ x \Rightarrow \{y{:}S; \mathsf{K}\ \textbf{says}\ \mathsf{DidOp}_i\ x\ y\} \in I_{\mathsf{K}} \qquad \mathcal{S} \vDash p \\ \cdot;\cdot \vdash v:T \qquad \Sigma_{ext};\cdot \vdash p : \mathsf{K}\ \textbf{says}\ \mathsf{OkToOp}_i\ v \qquad (u,\sigma') = \delta(\text{raw-op}_i, v, \sigma) \qquad q = \textbf{sign}(\mathsf{K}, \mathsf{DidOp}_i\ v\ u) \end{array}}{\Sigma_{ext}; I_{\mathsf{K}}; \delta \vdash (L, \sigma, \mathcal{S}) \dashv\{\text{op}_i, v, p\}\mapsto^s (\text{op}_i :: L, \sigma', \mathcal{S} \cup \{q\})}\ \text{S-ACT}$$

$$\frac{\Sigma_{ext};\cdot \vdash P : \textbf{Prop} \qquad A \neq \mathsf{K} \qquad \mathcal{S} \vDash P}{\Sigma_{ext}; I_{\mathsf{K}}; \delta \vdash (L, \sigma, \mathcal{S}) \dashv\{\text{assert:}A\ \textbf{says}\ P\}\mapsto^s (L, \sigma, \mathcal{S} \cup \{\textbf{sign}(A, P)\})}\ \text{S-SAY}$$

Proof-logged evaluation relation $\quad \cdot;\cdot;\cdot \vdash \cdot \dashv\{\cdot\}\mapsto^l \cdot$

$$\frac{\begin{array}{c} \text{op}_i : (x:T) \Rightarrow \mathsf{K}\ \textbf{says}\ \mathsf{OkToOp}_i\ x \Rightarrow \{y{:}S; \mathsf{K}\ \textbf{says}\ \mathsf{DidOp}_i\ x\ y\} \in I_{\mathsf{K}} \qquad \mathcal{S} \vDash p \\ \cdot;\cdot \vdash v:T \qquad \Sigma_{ext};\cdot \vdash p : \mathsf{K}\ \textbf{says}\ \mathsf{OkToOp}_i\ v \qquad (u,\sigma') = \text{raw-op}_i(v, \sigma) \qquad q = \textbf{sign}(\mathsf{K}, \mathsf{DidOp}_i\ v\ u) \end{array}}{\Sigma_{ext}; I_{\mathsf{K}}; \delta \vdash (L, \sigma, \mathcal{S}) \dashv\{\text{op}_i, v, p\}\mapsto^l (q :: p :: L, \sigma', \mathcal{S} \cup \{q\})}\ \text{L-ACT}$$

$$\frac{\Sigma_{ext};\cdot \vdash P : \textbf{Prop} \qquad A \neq \mathsf{K} \qquad \mathcal{S} \vDash P}{\Sigma_{ext}; I_{\mathsf{K}}; \delta \vdash (L, \sigma, \mathcal{S}) \dashv\{\text{assert:}A\ \textbf{says}\ P\}\mapsto^l (L, \sigma, \mathcal{S} \cup \{\textbf{sign}(A, P)\})}\ \text{L-SAY}$$

**Figure 2. Operational semantics**

to auditors. Semi-logged evaluation allows only authorized access to the raw interface via access control, and provides audit information of the list of allowed operations. Logged evaluation, like semi-logged evaluation, allows only authorized access to the raw interface; it also produces a more informative log of the proofs of the authorization decisions. Intuitively, semi-logged and logged evaluation, which deploy access control, allow strictly fewer operations than resource evaluation. Logged evaluation provides more information than the semi-logged evaluation for auditing, and semi-logged evaluation provides more information than resource evaluation.

The rest of this section sketches a technical framework in which the above claims are formalized and verified. The main result, Lemma 2.1, states that logged evaluation provides more information during audit than resource evaluation; similar results hold when comparing the logged and semi-logged relations or the semi-logged and resource relations. Before we present the formal statement of this lemma, we define a few auxiliary concepts.

Each of the three relations can be lifted to define *traces*. For instance, a resource trace is a sequence of the form

$$\tau = \sigma_0 \dashv\{\text{raw-op}_1\ v_1\}\mapsto^r \sigma_1 \cdots \dashv\{\text{raw-op}_n\ v_n\}\mapsto^r \sigma_n$$

Logged and semi-logged traces are defined similarly.

The following meta-function, pronounced "erase", shows how a logged trace is implemented in terms of its encapsulated resource:

$$\lfloor (L, \sigma, \mathcal{S}) \rfloor_{l/r} = \sigma$$
$$\lfloor C \dashv\{\text{assert: }\_\}\mapsto^l \tau \rfloor_{l/r} = \lfloor \tau \rfloor_{l/r}$$
$$\lfloor C \dashv\{\text{op}, v, \_\}\mapsto^l \tau \rfloor_{l/r} = \lfloor C \rfloor_{l/r} \dashv\{(\text{raw-op}, v)\}\mapsto^r \lfloor \tau \rfloor_{l/r}$$

For a set of traces, $\lfloor (H) \rfloor_{l/r}$ is defined as $\{\lfloor \tau \rfloor_{l/r} \mid \tau \in H\}$. Analogous functions can be defined to relate other pairs of evaluation schemes.

The $\sigma_0, \mathcal{S}_0$-*histories* of a configuration $C$, written $H^l(\sigma_0, \mathcal{S}_0, C)$, is defined as the set of all traces that terminate at configuration $C$ and begin with an initial state of the form $(nil, \sigma_0, \mathcal{S}_0)$. The $\sigma_0$-histories of a resource state $\sigma$, written $H^r(\sigma_0, \sigma)$, is defined as the set of all resource traces that terminate at $\sigma$.

The following lemma makes precise the claim that logged evaluation is strictly more informative, for audit, than resource evaluation. It describes a thought experiment where an auditor looks at either a logged evaluation configuration or its erasure as a resource state. In either case the auditor can consider the histories leading up to his ob-

servation. The lemma shows that there are histories consistent with resource evaluation that are not consistent with logged evaluation. Intuitively, this means logged evaluation makes more distinctions than—and is more informative than—resource evaluation.

**Lemma 2.1.** *There exists a kernel $K$, extended signature $\Sigma_{ext}$, configuration $C = (L, \sigma, \mathcal{S})$, rule set $\mathcal{S}_0$, initial trace $\sigma_0$ and resource trace $\tau$ such that $\tau \in H^r(\sigma_0, \sigma)$, but $\tau \notin \lfloor (H^l(\sigma_0, \mathcal{S}_0, C)) \rfloor_{l/r}$.*

*Proof Sketch.* By construction. Let *States* $= \{up, down\}$, with initial state $up$. Pick a configuration $C$ whose log contains six proofs and reflects a trace of the form $(\_, up, \_)\dashv\{\}\mapsto^l(\_, down, \_)\dashv\{\}\mapsto^l(\_, up, \_)$. Now consider trivial resource trace $\tau = up$. Observe that $\tau \in H^r(up, \lfloor C \rfloor_{l/r})$, but $\tau \notin H^l(C)$. □

Not surprisingly, it is possible to make similar distinctions between logged and semi-logged histories, as logged histories can ensure that a particular L-ACT step occurred, but this is not possible in the semi-logged case. As we will see in Section 3.3, this corresponds to the inability of the semi-logged system to distinguish between different proofs of the same proposition and thus to correctly assign blame.

## 3 The Logic

This section defines AURA$_0$, a language for expressing access control. AURA$_0$ is a higher-order, dependently typed, cut-down version of Abadi's Dependency Core Calculus [3, 2], Following the Curry-Howard isomorphism [16], AURA$_0$ types correspond to propositions relating to access control, and expressions correspond to proofs of these propositions. Dependent types allow propositions to be parameterized by objects of interest, such as principals or file handles. The interface between application and kernel code is defined using this language.

After defining the syntax and typing rules of AURA$_0$ and illustrating its use with a few simple access-control examples, this section gives the reduction rules for AURA$_0$ and discusses the importance of normalization with respect to auditing. It concludes with proofs of subject reduction, strong normalization and confluence for AURA$_0$; details may be found in the appendix.

### 3.1 Syntax

Figure 3 defines the syntax of AURA$_0$, which features two varieties of terms: access control proofs $p$, which are classified by corresponding propositions $P$ of kind **Prop**, and conventional expressions $e$, which are classified by types $T$ of the kind **Type**.[2] For ease of the subsequent pre-

[2]Our use of several syntactic categories in Figure 3 is purely for illustrative purposes.

| $t, s$ | ::= | $k \mid T \mid e$ | Terms |
| | | | |
| $k$ | ::= | $\mathbf{Kind^P} \mid \mathbf{Kind^T}$ | Sorts |
| | $\mid$ | $\mathbf{Prop} \mid \mathbf{Type}$ | Base kinds |
| | | | |
| $T, P$ | ::= | $\mathbf{string} \mid \mathbf{prin}$ | Base types |
| | $\mid$ | $x \mid a$ | Variables and constants |
| | $\mid$ | $t \; \mathbf{says} \; t$ | Says modality |
| | $\mid$ | $(x{:}t) \to t$ | Logical implication |
| | $\mid$ | $(x{:}t) \Rightarrow t$ | Computational arrows |
| | $\mid$ | $\{x{:}t; t\}$ | Dependent pair type |
| | | | |
| $e, p$ | ::= | $\texttt{"a"} \mid \texttt{"b"} \mid \ldots$ | String literals |
| | $\mid$ | $\mathsf{A} \mid \mathsf{B} \mid \mathsf{C} \ldots$ | Principal literals |
| | $\mid$ | $\mathbf{sign}(A, t)$ | Signature |
| | $\mid$ | $\mathbf{return}@[t] \; t$ | Injection into **says** |
| | $\mid$ | $\mathbf{bind} \; x = t \; \mathbf{in} \; t$ | Reasoning under **says** |
| | $\mid$ | $\lambda x{:}t. \, t \mid t \, t$ | Abstraction, application |
| | $\mid$ | $\langle t, t \rangle$ | Pair |

**Figure 3. Syntax of** AURA$_0$

sentation of the typing rules, we introduce two sorts, $\mathbf{Kind^P}$ and $\mathbf{Kind^T}$, which classify **Prop** and **Type** respectively. The base types are **prin**, the type of principals, and **string**; we use $x$ to range over variables, and $a$ to range over constants. String literals are `" "`–enclosed ASCII symbols; A, B, C etc. denote literal principals, while principal variables are written $A, B, C$.

In addition to the standard constructs for the functional dependent type $(x{:}t_1) \to t_2$, dependent pair type $\{x{:}t_1; t_2\}$, lambda abstraction $\lambda x{:}t_1. \, t_2$, function application $t_1 \; t_2$, and pair $\langle t_1, t_2 \rangle$, AURA$_0$ includes a special computational function type $(x{:}t_1) \Rightarrow t_2$. Intuitively, $(x{:}t_1) \to t_2$ is used for logical implication and $(x{:}t_1) \Rightarrow t_2$ describes kernel interfaces; Section 3.2 discusses this further. We will sometimes write $t_1 \to t_2$, $t_1 \Rightarrow t_2$, and $\{t_1; t_2\}$ as a shorthand for $(x{:}t_1) \to t_2$, $(x{:}t_1) \Rightarrow t_2$, and $\{x{:}t_1; t_2\}$, respectively, when $x$ does not appear free in $t_2$.

As in DCC, the modality **says** associates claims relating to access control with principals. The term $\mathbf{return}@[A] \; p$ creates a proof of $A$ **says** $P$ from a proof of $P$, while $\mathbf{bind} \; x = p_1 \; \mathbf{in} \; p_2$ allows a proof of $A$ **says** $P_1$ to be used as a proof of $P_1$, but only within the scope of a proof of $A$ **says** $P_2$. Finally, expressions of the form $\mathbf{sign}(\mathsf{A}, P)$ represent assertions claimed without proof. Such an expression is indisputable evidence that $P$ was asserted by $A$—rather than, for example, someone to whom $A$ has delegated authority. Such signed assertions must be verifiable, binding (i.e. non-repudiable), and unforgeable; signature implementation strategies are discussed in Section 5.

$$\boxed{\Sigma; \Gamma \vdash t : t}$$

$$\frac{\Sigma \vdash \Gamma}{\Sigma; \Gamma \vdash \mathbf{Prop} : \mathbf{Kind^P}} \text{ T-PROP} \qquad \frac{\Sigma \vdash \Gamma}{\Sigma; \Gamma \vdash \mathbf{Type} : \mathbf{Kind^T}} \text{ T-TYPE} \qquad \frac{\Sigma \vdash \Gamma \qquad T \in \{\mathbf{string}, \mathbf{prin}\}}{\Sigma; \Gamma \vdash T : \mathbf{Type}} \text{ T-BASE}$$

$$\frac{\Sigma \vdash \Gamma \qquad x : t \in \Gamma}{\Sigma; \Gamma \vdash x : t} \text{ T-VAR} \qquad \frac{\Sigma \vdash \Gamma \qquad a : t \in \Sigma}{\Sigma; \Gamma \vdash a : t} \text{ T-CONST} \qquad \frac{\Sigma; \Gamma \vdash t_1 : \mathbf{prin} \qquad \Sigma; \Gamma \vdash t_2 : \mathbf{Prop}}{\Sigma; \Gamma \vdash t_1 \ \mathbf{says} \ t_2 : \mathbf{Prop}} \text{ T-SAYS}$$

$$\frac{\Sigma; \Gamma \vdash t_1 : k_1 \qquad \Sigma; \Gamma, x : t_1 \vdash t_2 : k_2 \qquad k_1 \in \{\mathbf{Kind^P}, \mathbf{Type}, \mathbf{Prop}\} \qquad k_2 \in \{\mathbf{Type}, \mathbf{Prop}\}}{\Sigma; \Gamma \vdash (x{:}t_1) \rightarrow t_2 : k_2} \text{ T-ARR}$$

$$\frac{\Sigma; \Gamma \vdash t_1 : k_1 \qquad \Sigma; \Gamma, x : t_1 \vdash t_2 : k_2 \qquad k_1, k_2 \in \{\mathbf{Type}, \mathbf{Prop}\}}{\Sigma; \Gamma \vdash \{x{:}t_1; t_2\} : k_2} \text{ T-PAIRTYPE}$$

$$\frac{\Sigma \vdash \Gamma \qquad A \in \{\mathsf{A}, \mathsf{B}, \ldots\} \qquad \Sigma; \cdot \vdash t : \mathbf{Prop}}{\Sigma; \Gamma \vdash \mathbf{sign}(A, t) : A \ \mathbf{says} \ t} \text{ T-SIGN}$$

$$\frac{\Sigma \vdash \Gamma \qquad s \in \{\texttt{"a"}, \texttt{"b"}, \ldots\}}{\Sigma; \Gamma \vdash s : \mathbf{string}} \text{ T-LITSTR} \qquad \frac{\Sigma; \Gamma \vdash t_1 : \mathbf{prin} \quad \Sigma; \Gamma \vdash t_2 : s_2 \quad \Sigma; \Gamma \vdash s_2 : \mathbf{Prop}}{\Sigma; \Gamma \vdash \mathbf{return}@[t_1] \ t_2 : t_1 \ \mathbf{says} \ s_2} \text{ T-RETURN}$$

$$\frac{\Sigma \vdash \Gamma \qquad A \in \{\mathsf{A}, \mathsf{B} \ldots\}}{\Sigma; \Gamma \vdash A : \mathbf{prin}} \text{ T-LITPRIN} \qquad \frac{\Sigma; \Gamma \vdash e_1 : t \ \mathbf{says} \ P_1 \quad \Sigma; \Gamma, x : P_1 \vdash e_2 : t \ \mathbf{says} \ P_2 \quad x \notin fv(P_2)}{\Sigma; \Gamma \vdash \mathbf{bind} \ x \ = \ e_1 \ \mathbf{in} \ e_2 : t \ \mathbf{says} \ P_2} \text{ T-BIND}$$

$$\frac{\Sigma; \Gamma, x : t \vdash p : P \qquad \Sigma; \Gamma \vdash (x{:}t) \rightarrow P : \mathbf{Prop}}{\Sigma; \Gamma \vdash \lambda x{:}t. \ p : (x{:}t) \rightarrow P} \text{ T-LAM} \qquad \frac{\Sigma; \Gamma \vdash t_1 : (x{:}P_2) \rightarrow P \qquad \Sigma; \Gamma \vdash t_2 : P_2}{\Sigma; \Gamma \vdash t_1 \ t_2 : \{t_2/x\}P} \text{ T-APP}$$

$$\frac{\Sigma; \Gamma \vdash t_1 : s_1 \qquad \Sigma; \Gamma \vdash t_2 : \{t_1/x\}s_2 \qquad \Sigma; \Gamma, x : s_1 \vdash s_2 : k}{\Sigma; \Gamma \vdash \langle t_1, t_2 \rangle : \{x{:}s_1; s_2\}} \text{ T-PAIR}$$

$$\boxed{\Sigma; \Gamma \vdash t}$$

$$\frac{\Sigma; \Gamma \vdash t_1 : t_2 \quad t_2 \in \{\mathbf{Kind^P}, \mathbf{Kind^T}, \mathbf{Prop}, \mathbf{Type}\}}{\Sigma; \Gamma \vdash t_1} \text{ T-C} \qquad \frac{\Sigma; \Gamma \vdash t_1 : k \quad k \in \{\mathbf{Type}, \mathbf{Prop}\} \quad \Sigma; \Gamma, x : t_1 \vdash t_2}{\Sigma; \Gamma \vdash (x{:}t_1) \Rightarrow t_2} \text{ T-ARR-C}$$

**Figure 4. The typing relation**

### 3.2 Type system

AURA$_0$'s type system is defined in terms of constant signatures $\Sigma$, and variable typing contexts $\Gamma$, which associate types to global constants and local variables, respectively, and are written:

$$\Gamma ::= \cdot \mid \Gamma, x : t \qquad \Sigma ::= \cdot \mid \Sigma, a : t.$$

Typechecking consists of four judgments:

| | |
|---|---|
| $\Sigma \vdash \diamond$ | Signature $\Sigma$ is well-formed |
| $\Sigma \vdash \Gamma$ | Context $\Gamma$ is well formed |
| $\Sigma; \Gamma \vdash t_1 : t_2$ | Term $t_1$ has type $t_2$ |
| $\Sigma; \Gamma \vdash t$ | Computation type $t$ is well-formed |

The signature $\Sigma$ is well-formed if $\Sigma$ maps constants to types of sort $\mathbf{Kind^P}$—in other words, all AURA$_0$ constants construct propositions. The context $\Gamma$ is well-formed with respect to signature $\Sigma$ if $\Gamma$ maps variables to well-formed types. A summary of the typing rules for terms can be found in Figure 4. Most of the rules are straightforward, and we explain only a few key rules.

Rule T-SIGN states that a signed assertion created by the principal $A$ signing a proposition $P$ has type $A$ **says** $P$; here, $P$ can be any proposition, even false. More interesting, however, is when $P$ contains a constant symbol defined in the signature $\Sigma$; as there is no introduction form for constants, there can be no proof of $P$ within the logic, but the

existence of signatures allows for terms of type $A$ **says** $P$. These signed assertions are an essential part of encoding access control. The premises of T-SIGN typechecks $A$ and $P$ in the empty variable context, as signatures are intended to have unambiguous meaning in any scope—a signature with free variables is inherently meaningless.

The rule T-RETURN states that if we can construct a proof term $p$ for proposition $P$, then the term **return**@$[A]$ $p$ is a proof term for proposition $A$ **says** $P$—in other words, all principals believe what can be independently verified. The T-BIND rule is a standard bind rule for monads and ensures that what principal $A$ believes can only be used when reasoning from $A$'s perspective.

The rule for the functional dependent type T-ARR restricts the kinds of dependencies allowed by the type system, ruling out functional dependencies on objects of kind **Type**. Note that, in the T-LAM rule, the type of the lambda abstraction must be of kind **Prop**. With such restrictions in place, it is rather straightforward to observe that these two rules allow us to express flexible access control rules while at the same time ruling out type level computations and preserving decidability of type checking.[3]

The interfaces between the application code and the kernel also requires a type description. For this reason, AURA$_0$ introduces a special computational arrow type, $(x{:}t_1) \Rightarrow t_2$. Computations cannot appear in proofs or propositions. This decouples AURA$_0$ proof reduction from effectful computation, and simplifies the interpretation of propositions. While AURA [25] demonstrates how to achieve similar results using a single arrow type and restrictions on applications, computation types simplify the exposition of AURA$_0$.

The typing rule T-PAIRTYPE for dependent pairs is standard and permits objects of kinds **Type** and **Prop** to be freely mixed; for simplicity we prohibit types and propositions themselves from appearing in a pair. Notice that AURA$_0$ features an introduction proof for pairs but no corresponding elimination form. While full AURA does, of course, feature such terms, AURA$_0$ uses dependent pairs only when associating proofs with the data on which they depend, and hence the elimination forms for pairs are unnecessary and have been elided for brevity.

## 3.3 Examples

The combination of dependent types and the **says** modality in AURA$_0$ can express many interesting policies. For instance, Abadi's encoding of speaks-for [2] is easily adopted:

$A$ speaksfor $B \triangleq B$ **says** $((P{:}\textbf{Prop}) \to A$ **says** $P \to P)$

---

[3]Using two sorts, **Kind**$^{\textbf{T}}$ and **Kind**$^{\textbf{P}}$, makes it easy to state these restrictions on function types. Full AURA [25] implements a similar restriction using only a single sort; this makes some of its typing rules slightly heavier, but the two approaches appear largely equivalent.

Adding dependency allows for more fine grained delegation. For example, we can encode partial delegation:

$B$ **says** $((x{:}\textbf{string}) \to A$ **says** $\mathsf{Good}\ x \to \mathsf{Good}\ x)$

Here $A$ speaks for $B$ only when certifying that a string is "good." Such fine-grained delegation is important for real applications where the full speaks-for relation may be too permissive.

Recall also the Remote Procedure Call example from Section 2.1. While an application might use $r_0$ (of type $\mathsf{K}$ **says** $((x{:}\textbf{string}) \to \mathsf{OkToRPC}\ x))$ directly when building proofs, it could also construct a more convenient derived rule by using AURA$_0$'s **bind** to reason from $\mathsf{K}$'s perspective. For instance:

$$r_0' \quad : \quad (x{:}\textbf{string}) \to \mathsf{K}\ \textbf{says}\ \mathsf{OkToRPC}\ x$$
$$r_0' \quad = \quad \lambda x{:}\textbf{string}.\ \textbf{bind}\ y\ =\ r_0\ \textbf{in}\ \textbf{return}@[\mathsf{K}]y\ x$$

Rules like $r_0$ and its derivatives, however, are likely too trivial to admit interesting opportunities for audit; a more interesting policy states that any principal may perform a remote procedure call so long as that principal signs the input string. One encoding of this policy uses the extended context

$$\Sigma_{ext} = \mathsf{ReqRPC} : \textbf{string} \to \textbf{Prop}, \Sigma_K$$

and singleton rule set

$Rules = \{r_1 = \ \textbf{sign}(\mathsf{K}, (x{:}\textbf{string}) \to (A{:}\textbf{prin}) \to$
$\qquad\qquad\qquad (A\ \textbf{says}\ \mathsf{ReqRPC}\ x) \to \mathsf{OkToRPC}\ x)\}.$

Given this rule, an auditor might find the following proofs in the log:

$p_1 = \ \textbf{bind}\ x\ =\ r_1\ \textbf{in}$
$\qquad\quad \textbf{return}@[\mathsf{K}](x\ \texttt{"hi"}\ \mathsf{A}\ \textbf{sign}(\mathsf{A}, \mathsf{ReqRPC}\ \texttt{"hi"}))$
$p_2 = \ (\lambda x{:}\mathsf{K}\ \textbf{says}\ \mathsf{OkToRPC}\ \texttt{"ab"}.$
$\qquad\quad \lambda y{:}\mathsf{C}\ \textbf{says}\ \mathsf{ReqRPC}\ \texttt{"cd"}.\,x)$
$\qquad\quad (\textbf{bind}\ z\ =\ r_1\ \textbf{in}$
$\qquad\qquad \textbf{return}@[\mathsf{K}](z\ \texttt{"ab"}\ \mathsf{B}\ \textbf{sign}(\mathsf{B}, \mathsf{ReqRPC}\ \texttt{"ab"}))$
$\qquad\quad (\textbf{sign}(\mathsf{C}, \mathsf{ReqRPC}\ \texttt{"cd"})).$

As $p_1$ contains only $\mathsf{A}$'s signature, and as signatures are unforgeable, the auditor can conclude that $\mathsf{A}$ is responsible for the request—the ramifications of this depend on the real-world context of in question. Proof $p_2$ is more complicated; it contains signatures from both $\mathsf{B}$ and $\mathsf{C}$. An administrator can learn several things from this proof.

We can simplify the analysis of $p_2$ by reducing it as discussed in the following section. Taking the normal form of $p_2$ (i.e., simplifying it as much as possible) yields

$p_2' = \ \textbf{bind}\ z\ =\ r_1$
$\qquad\quad \textbf{in return}@[\mathsf{K}](z\ \texttt{"ab"}\ \mathsf{B}\ \textbf{sign}(\mathsf{B}, \mathsf{ReqRPC}\ \texttt{"ab"}).$

$$\boxed{\vdash t \to t'}$$

$$\frac{x \notin \mathit{fv}(t_2)}{\vdash \mathbf{bind}\ x\ =\ t_1\ \mathbf{in}\ t_2 \to t_2}\ \text{R-BINDS}$$

$$\frac{}{\vdash \mathbf{bind}\ x\ =\ \mathbf{return}@[t_0]\ t_1\ \mathbf{in}\ t_2 \to \{t_1/x\}t_2}\ \text{R-BINDT}$$

$$\frac{\vdash t_2 \to t_2'}{\vdash \mathbf{return}@[t_1]\ t_2 \to \mathbf{return}@[t_1]\ t_2'}\ \text{R-SAYS}$$

$$\frac{y \notin \mathit{fv}(t_3)}{\vdash \mathbf{bind}\ x\ =\ (\mathbf{bind}\ y\ =\ t_1\ \mathbf{in}\ t_2)\ \mathbf{in}\ t_3 \to}$$
$$\mathbf{bind}\ y\ =\ t_1\ \mathbf{in}\ \mathbf{bind}\ x\ =\ t_2\ \mathbf{in}\ t_3$$
$$\text{R-BINDC}$$

**Figure 5. Selected reduction rules**

This term contains only B's signature, and hence B may be considered accountable for the action. This is exactly the ruling out of histories discussed in Section 2.2.

Proofs $p_2$ and $p_2'$ illustrate a tension inherent to this computation model. A configuration whose log contains $p_2$ will be associated with fewer histories (i.e. those in which C make no assertions) than an otherwise similar configuration containing $p_2'$. While normalizing proofs inform policy analysis, it can also discard interesting information. To see this, consider how C's signature may be significant on an informal level. If the application is intended to pass normalized proofs to the kernel, then this is a sign that the application is malfunctioning. If principals are only supposed to sign certain statements, C's apparently spurious signature may indicate an violation of that policy, even if the signature was irrelevant to actual access control decisions.

## 3.4 Formal language properties

**Subject reduction** As the preceding example illustrates, proof simplification is a useful tool for audit. Following the Curry-Howard isomorphism, proof simplification corresponds to $\lambda$-calculus reductions on proof terms.

Most of the reduction rules for $\text{AURA}_0$ are standard; selected rules can be seen in Figure 5, and the entire reduction relation can be found in the appendix. For **bind**, in addition to the standard congruence, beta reduction, and commute rules as found in monadic languages, we also include a special beta reduction rule R-BINDS. The R-BINDS rule eliminates bound proofs that are never mentioned in the **bind**'s body. Rule R-BINDS permits simplification of terms like **bind** $x\ =\ \mathbf{sign}(\mathsf{A}, P)$ **in** $t$, which are not subject to R-BINDT reductions. $\text{AURA}_0$ disallows reduction under **sign**, as signatures are intended to represent fixed objects realized, for example, via cryptographic means.

The following lemma states that the typing of an expression is preserved under reduction rules:

**Lemma 3.1** (Subject Reduction). *If $\vdash t \to t'$ and $\Sigma; \Gamma \vdash t : s$ then $\Sigma; \Gamma \vdash t' : s$.*

*Proof Sketch.* The proof proceeds by structural induction on the reduction relation and depends on several standard facts. Additionally, the R-BINDS cases requires a nonstandard lemma observing that we may remove a variable $x$ from the typing context when $x$ is not used elsewhere in the typing judgment. □

**Proof normalization** An expression is in *normal form* when it has no applicable reduction rules; as observed in Section 3.3, reducing a proof to its normal form can be quite useful for auditing. Proof normalization is most useful when the normalization process always terminates and every term has a unique normal form.

An expression $t$ is *strongly normalizing* if application of any sequence of reduction rules to $t$ always terminates. A language is strongly normalizing if all the terms in the language are strongly normalizing. We have proved that $\text{AURA}_0$ is strongly normalizing, which implies that any algorithm for proof normalization will terminate. The details of the proofs are presented in the appendix

**Lemma 3.2** (Strong Normalization). *$\text{AURA}_0$ is strongly normalizing.*

*Proof Sketch.* We prove that $\text{AURA}_0$ is strongly normalizing by translating $\text{AURA}_0$ to the Calculus of Constructions extended with dependent pairs, which is known to be strongly normalizing [22], in a way that preserves both types and reduction steps. The interesting cases are the translations of terms relating to the **says** monad: **return** expressions are dropped, **bind** expressions are translated to to lambda application, and a term $\mathbf{sign}(t_1, t_2)$ is translated to a variable whose type is the translation of $t_2$. One subtle point is the tracking of dependency in the types of these newly introduced variables, which must be handled delicately. □

We have also proved that $\text{AURA}_0$ is confluent—i.e., that two series of reductions starting from the same term can always meet at some point. Let $t \to^* t'$ whenever $t = t'$ or $t$ reduces to $t'$ in one or more steps.

**Lemma 3.3** (Confluence). *If $t \to^* t_1$, and $t \to^* t_2$, then there exists $t_3$ such that $t_1 \to^* t_3$ and $t_2 \to^* t_3$.*

*Proof Sketch.* We first prove that $\text{AURA}_0$ is weakly confluent, which follows immediately from inspection of the reduction rules. We then apply the well-known fact that strong normalization and weak confluence imply confluence. □

A direct consequence of these properties is that every $\text{AURA}_0$ term has a unique normal form; any algorithm for proof normalization will yield the same normal form for a given term. This implies that the set of relevant evidence—i.e., signatures—in a given proof term is also unique, an important property to have when assigning blame.

## 4  File System Example

As a more substantial example, we consider a file system in which file access is authorized using $\text{AURA}_0$ and log entries consist of authorization proofs. In a traditional file system, authorization decisions regarding file access are made when a file is opened, and thus we begin by considering only the open operation and only briefly consider additional operations. Our open is intended to provide flexible access control on top of a system featuring a corresponding raw-open and associated constants:

$$\text{Mode} : \textbf{Type} \qquad \text{FileDes} : \textbf{Type}$$

$$\text{RDONLY} : \text{Mode} \qquad \text{WRONLY} : \text{Mode}$$
$$\text{APPEND} : \text{Mode} \qquad \text{RDWR} : \text{Mode}$$

$$\text{raw-open} : \{\text{Mode}; \textbf{string}\} \Rightarrow \text{FileDes}$$

We can imagine that raw-open is part of the interface to an underlying file system with no notion of per-user access control or $\text{AURA}_0$ principals; it, of course, should not be exposed outside of the kernel. Taking inspiration from Unix, we define RDONLY, WRONLY, APPEND, and RDWR (the inhabitants of Mode), which specify whether to open a file for reading only, overwrite only, append only, or unrestricted reading and writing, respectively. Type FileDes is left abstract; it classifies file descriptors—unforgeable capabilities used to access the contents of opened files.

Figure 6 shows the interface to open, the extended signature of available predicates, and the rules used to construct the proofs of type $\mathsf{K}$ **says** $\text{OkToOpen} \langle m, f \rangle$ (for some file $f$ and mode $m$) that open requires. OkToOpen and DidOpen are as specified in Section 2, and the other predicates have the obvious readings: $\text{Owns}\ A\ f$ states that the principal $A$ owns the file $f$, $\text{ReqOpen}\ m\ f$ is a request to open file $f$ with mode $m$, and $\text{Allow}\ A\ m\ s$ states that $A$ should be allowed to open $f$ with mode $m$. (As we are not modeling authentication we will take it as given that all proofs of type $A$ **says** $\text{ReqOpen}\ m\ f$ come from $A$; we discuss ways of enforcing this in Section 5.)

We assume, for each file $f$, the existence of a rule $\text{owner}f$ of type $\mathsf{K}$ **says** $\text{Owns}\ A\ f$ for some constant principal A—as only one such rule exists for any $f$ and no other means are provided to generate proofs of this type, we can be sure that each file will always have a unique owner. Aside from such statements of ownership, the only rule a

---

Kernel Signature $\Sigma_K$

$$\text{OkToOpen} : \{\text{Mode}; \textbf{string}\} \rightarrow \textbf{Prop}$$
$$\text{DidOpen} : (x : \{\text{Mode}; \textbf{string}\}) \rightarrow$$
$$\text{FileDes} \rightarrow \textbf{Prop}$$

Kernel Interface $I_K$

$$\text{open} : (x : \{\text{Mode}; \textbf{string}\}) \Rightarrow$$
$$\mathsf{K}\ \textbf{says}\ \text{OkToOpen}\ x \Rightarrow$$
$$\{h{:}\text{FileDes}; \mathsf{K}\ \textbf{says}\ \text{DidOpen}\ x\ h\}$$

Additional Types in Extended Signature $\Sigma_{ext}$

$$\text{Owns} : \textbf{prin} \rightarrow \textbf{string} \rightarrow \textbf{Prop}$$
$$\text{ReqOpen} : \text{Mode} \rightarrow \textbf{string} \rightarrow \textbf{Prop}$$
$$\text{Allow} : \textbf{prin} \rightarrow \text{Mode} \rightarrow \textbf{string} \rightarrow \textbf{Prop}$$

Rule Set $R$:

$$\text{owner}f : \mathsf{K}\ \textbf{says}\ \text{Owns}\ \mathsf{A}\ f$$

$$\begin{aligned}
\text{delegate} : \mathsf{K}\ \textbf{says}\ ((A : \textbf{prin}) \rightarrow (B : \textbf{prin}) \rightarrow \\
(m : \text{Mode}) \rightarrow (f : \textbf{string}) \rightarrow \\
A\ \textbf{says}\ \text{ReqOpen}\ m\ f \rightarrow \\
\mathsf{K}\ \textbf{says}\ \text{Owns}\ B\ f \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ m\ f \rightarrow \\
\text{OkToOpen}\ \langle m, f \rangle)
\end{aligned}$$

$$\begin{aligned}
\text{owned} : \mathsf{K}\ \textbf{says}\ ((A : \textbf{prin}) \rightarrow (m : \text{Mode}) \rightarrow \\
(f : \textbf{string}) \rightarrow \\
A\ \textbf{says}\ \text{ReqOpen}\ m\ f \rightarrow \\
\mathsf{K}\ \textbf{says}\ \text{Owns}\ A\ f \rightarrow \\
\text{OkToOpen}\ \langle m, f \rangle)
\end{aligned}$$

$$\begin{aligned}
\text{readwrite} : \mathsf{K}\ \textbf{says}\ ((A : \textbf{prin}) \rightarrow (B : \textbf{prin}) \rightarrow \\
(f : \textbf{string}) \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ \text{RDONLY}\ f \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ \text{WRONLY}\ f \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ \text{RDWR}\ f)
\end{aligned}$$

$$\begin{aligned}
\text{read} : \mathsf{K}\ \textbf{says}\ ((A : \textbf{prin}) \rightarrow (B : \textbf{prin}) \rightarrow \\
(f : \textbf{string}) \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ \text{RDWR}\ f \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ \text{RDONLY}\ f)
\end{aligned}$$

$$\begin{aligned}
\text{write} : \mathsf{K}\ \textbf{says}\ ((A : \textbf{prin}) \rightarrow (B : \textbf{prin}) \rightarrow \\
(f : \textbf{string}) \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ \text{RDWR}\ f \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ \text{WRONLY}\ f)
\end{aligned}$$

$$\begin{aligned}
\text{append} : \mathsf{K}\ \textbf{says}\ ((A : \textbf{prin}) \rightarrow (B : \textbf{prin}) \rightarrow \\
(f : \textbf{string}) \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ \text{RDWR}\ f \rightarrow \\
B\ \textbf{says}\ \text{Allow}\ A\ \text{APPEND}\ f)
\end{aligned}$$

**Figure 6. Types for the file system example**

file system absolutely needs is delegate, which states that the kernel allows anyone to access a file with a particular mode if the owner of the file allows it.

The other rules, however, are of great convenience. The rule owned relieves the file owner $A$ from the need to create signatures of type $A$ **says** Allow $A\ m\ f$ for files $A$ owns, while readwrite allows a user who has acquired read and write permission for a file from different sources to open the file for reading and writing simultaneously. The rules read, write, and append do the reverse, allowing a user to drop from RDWR mode to RDONLY, WRONLY, or AP-PEND. These last four rules simply reflect semantic facts about constants of type Mode.

With the rules given in Figure 6 and the other constructs of our logic it is also easy to create complex chains of delegation for file access. For example, Alice (A) may delegate full authority over any files she can access to Bob (B) with a signature of type

$$A\ \textbf{says}\ (C : \textbf{prin} \to m : \textsf{Mode} \to f : \textbf{string} \to$$
$$B\ \textbf{says}\ \textsf{Allow}\ C\ m\ f \to A\ \textbf{says}\ \textsf{Allow}\ C\ m\ f),$$

or she may restrict what Bob may do by adding further requirements on $C$, $m$, or $f$. She might restrict the delegation to files that she owns, or replace $C$ with B to prevent Bob from granting access to anyone but himself. She could do both with a signature of type

$$A\ \textbf{says}\ (m : \textsf{Mode} \to f : \textbf{string} \to K\ \textbf{says}\ \textsf{Owns}\ A\ f$$
$$B\ \textbf{says}\ \textsf{Allow}\ B\ m\ f \to A\ \textbf{says}\ \textsf{Allow}\ B\ m\ f).$$

As described in Section 2, the kernel logs the arguments to our interface functions whenever they are called. So far we have only one such function, open, and logging its arguments means keeping a record every time the system permits a file to be opened. Given the sort of delegation chains that the rules allow, it should be clear that the reason why an open operation is permitted can be rather complex, which provides a strong motivation for the logging of proofs.

One can easily imagine using logged proof terms—and in particular proof terms in normal form, as described in Section 3.3—to assist in assigning the blame for an unusual file access to the correct principals. For example, a single principal who carelessly delegates RDWR authority might be blamed more severely than two unrelated principals who unwittingly delegate RDONLY and WRONLY authority to someone who later makes use of readwrite. Examining the structure of proofs can once again allow an auditor to, in the terminology of Section 2.2, rule out certain histories.

We can also see how logging proofs might allow a system administrator to debug the rule set. The rules in Fig-

ure 6 might well be supplemented with, for example

$$\textsf{surely} : K\ \textbf{says}\ ((A : \textbf{prin}) \to (B : \textbf{prin}) \to$$
$$(f : \textbf{string}) \to$$
$$B\ \textbf{says}\ \textsf{Allow}\ A\ \textsf{RDONLY}\ f \to$$
$$B\ \textbf{says}\ \textsf{Allow}\ A\ \textsf{APPEND}\ f \to$$
$$B\ \textbf{says}\ \textsf{Allow}\ A\ \textsf{RDWR}\ f)$$

$$\textsf{maybe} : K\ \textbf{says}\ ((A : \textbf{prin}) \to (B : \textbf{prin}) \to$$
$$(f : \textbf{string}) \to$$
$$B\ \textbf{says}\ \textsf{Allow}\ A\ \textsf{WRONLY}\ f \to$$
$$B\ \textbf{says}\ \textsf{Allow}\ A\ \textsf{APPEND}\ f)$$

Rule surely is clearly erroneous, as it allows a user with only permission to read from and append to a file to alter its existing content, but such a rule could easily be introduced by human error when the rule set is created. Since any uses of this rule would be logged, it would not be possible to exploit such a problematic rule without leaving a clear record of how it was done, allowing a more prudent administrator to correct the rule set.

Rule maybe, on the other hand, is a bit more subtle— it states that the ability to overwrite a file is strictly more powerful than the ability to append to that file, even in the absence of any ability to read. Whether such a rule is valid depends on the expectations of the system's users: maybe is clearly unacceptable if users desire to allow others to overwrite but not to append to files; otherwise, maybe may be seen as quite convenient, allowing, for examples, easy continuation of long write operations that were prematurely aborted. Examining the proofs in the log can help the administrator determine whether the inclusion of maybe best suits the needs of most users.

We have so far discussed only open, but there is still much $\textsc{Aura}_0$ has to offer a file system, even in the context of operations that do not involve authorization.

**Reading and writing**   While access permission is granted when a file is opened, it is worth noting that, as presented, the type FileDes conveys no information about what sort of access has been granted; consequently, attempting, for example, to write to a read-only file descriptor will result in a run-time error. Since we already have a system with dependent types, this need not be the case; while it is somewhat orthogonal to our concerns of authorization, FileDes could easily be made to depend on the mode with which a file has been opened, and operations could expect file descriptors equipped with the correct Mode arguments. This would, however, require some analog to the subsumption rules read, write, and append—and perhaps also readwrite—along with, for pragmatic reasons, a means of preventing the kernel from logging file contents being read or written, as discussed in Section 5.

**Close** At first glance it seems that closing a file, like reading or writing, is an operation that requires only a valid file descriptor to ensure success, yet there is something more the type system can provide. For example, if we require a corresponding DidOpen when constructing proofs of type OkToClose, we can allow a user to share an open file descriptor with other processes secure in the knowledge that those processes will be unable to prematurely close the file. In addition, logging of file close operations can help pinpoint erroneous double closes, which, while technically harmless, may be signs of deeper logic errors in the program that triggered them.

**Ownership** File creation and deletion are certainly operations that should require authorization, and they are especially interesting due to their interaction with the Owns predicate. The creation of file $f$ by principal A should introduce a rule owner $f$ : Owns A $f$ into the rule set, while the deletion of a file should remove said rule; a means of transferring file ownership would also be desirable. This can amount to treating a subset of our rules as a protected resource in its own right, with a protected interface to these rules and further rules concerning the circumstances under which access to this new resource should be granted. An alternate approach is to dispense with ownership rules completely and instead use signed objects and signature revocation, discussed further in Section 5, to represent ownership.

## 5 Discussion

**Signature implementation** Thus far we have treated signatures as abstract objects that may only be created by principals or programs with sufficient authority. This suggests two different implementation strategies.

The first approach is cryptographic: a **sign** object can be represented by a digital signature in public key cryptography. Each principal must be associated with a well known public key and in possession of its private counterpart; implementing rule T-SIGN reduces to calling a digital signature verification function. The cryptographic scheme is well suited for distributed systems with mutual distrust.

A decision remains to be made, however: we can interpret **sign**(A, $P$) either as a tuple containing the cryptographic signature along with A and $P$ in plaintext, or as the cryptographic signature alone. In the latter case signatures are small (potentially constructed from a hash of the contents), but recovering the text of a proposition from its proof (i.e., doing type inference) may not be possible. In the former case, inference is trivial, but proofs are generally large. Note that proof checking of **sign**s in either case involves validating digital signatures, a polynomial time operation.

An alternative implementation of signatures assumes that all principals trust some *moderator*, who maintains a table of signatures as abstract data values; each **sign** may then be represented as an index into the moderator's table. Such indices can be small while still allowing for easy type inference, but such a scheme requires a closed system with a mutually trusted component. In a small system, the moderator can be the kernel itself, but a larger system might contain several kernels protecting different resources and administered by disparate organizations, in which case finding a suitable moderator may be quite difficult.

**Temporary signatures** Real-world digital signature implementations generally include with each signature an interval of time outside of which the signature should not be considered valid. In addition, there is often some notion of a revocation list to which signatures can be added to ensure their immediate invalidation. Both of these concepts could be useful in our setting, as principals might want to delegate authority temporarily and might or might not know in advance how long this delegation should last. Potentially mutable rules—which could be very important in a truly distributed setting—can even be represented by digital signatures in the presence of a revocation list.

The question remains, however, how best to integrate these concepts with AURA$_0$. One possible answer is to change nothing in the logic and simply allow for the possibility that any proof might be declared invalid at runtime due to an expired signature. Following this strategy requires operations to dynamically validate the timestamps in the signatures before logging, thereby making all kernel operations partial (i.e., able to fail due to expired proofs). In such a setting, it seems appealing to incorporate some kind of transaction mechanism so that clients can be guaranteed that their proofs are current before attempting to pass them to the kernel. While easy to implement, this approach may be unsatisfying in that programmers are left unable to reason about or account for such invalid proofs.

Signatures might also be limited in the number of times they may be used, and this seems like a natural application for *linear* or *affine* types (see Bauer *et al.* for an authorization logic with linearity constraints [8]). Objects of a linear or affine type must be used exactly or at most once, respectively, making such types appropriate for granting access to a resource only a set number of times. They can also be used to represent protocols at the type level, ensuring, for example, that a file descriptor is not used after it is closed.

Garg, deYoung, and Pfenning [18] are studying a constructive and linear access control logic with an explicit time intervals. Their syntax includes propositions of the form $P@[T_1, T_2]$, meaning "$P$ is valid between times $T_1$ and $T_2$." To handle time, the judgment system is parameterized by an interval; the interpretation of sequent $\Psi; \Gamma; \Delta \implies P[I]$ is, "given assumptions $\Psi, \Gamma$, and $\Delta$, $P$ is valid during interval $I$." Adopting this technique could allow AURA$_0$ to address the problems of temporal policies, though it is currently un-

clear what representations of time and revocation might best balance concerns of simplicity and expressive power.

**Proof normalization**   Proofs in normal form are useful for audit because they provide a particularly clear view of authorization decisions. Normalization, however, is an expensive operation—even for simply typed lambda calculus, the worst-case lower-bound on the complexity of the normalization is on the order of $exp(2, n)$, where $exp(2, 0) = 1$, $exp(2, n) = 2^{exp(2, (n-1))}$, etc., and $n$ is the size of the term [30]. Furthermore, the size of a normalized proof can grow to $exp(2, n)$ as well. On the other hand, checking whether a proof is in normal form is linear to the size of the proof, and, in practice, non-malicious proof producers will likely create proofs that are simple to normalize. Consequently, where the normalization process should be carried out depends on the system in question.

A kernel operating in a highly untrusted environment might require all submitted proofs to be in normal form, shifting the computational burden to potentially malicious clients (as is commonly done to defend against denial of service attacks). By contrast, a kernel providing services to a "smart dust" network might normalize proofs itself, shifting work away from computationally impoverished nodes and onto a more robust system, again a standard design. Server side normalization might be done online as proofs come in (to amortized computation cost) or offline during audit (to avoid latency). Ultimately, the AURA programming model naturally accommodates these approaches and others; an implementation should allow programmers to select whatever normalization model is appropriate.

**Authentication**   In Section 4 we assumed that signatures of type $A$ **says** ReqOpen $m$ $f$ are always sent from $A$. Such an assumption is necessary because we are not currently modeling any form of authentication—or even the association of a principal with a running program—but a more realistic solution is needed when moving beyond the scope of this paper. For example, communication between programs running on behalf of different principals could take place over channel endpoints with types that depend on the principal on the other end of the channel.

Of course, when this communication is between different machines on an inherently insecure network, problems of secure authentication become non-trivial, as we must implement a secure channel on top of an insecure one. In practice this is done with cryptography, and one of the long-term goals of the AURA project is to elegantly integrate cryptographic methods with the type system, following the work of, for example, Fournet, *et al.* [20].

**Pragmatics**   We are in the process of implementing AURA, in part to gain practical experience with the methodology proposed in this paper. Besides the issues with temporal policies and authentication described above, we anticipate several other concerns that need to be addressed.

In particular, we will require efficient log operations and compact proof representations. Prior work on proof compression for proof-carrying code [28] should apply in this setting, but until we have experience with concrete examples, it is not clear how large the authorization proofs may become in practice. A related issue is tool support for browsing querying the audit logs: tools should allow system administrators to issue queries against the log and analyze the evidence that is present and rules that have been used.

For client developers, we expect that it will often prove useful to log information beyond what is logged by the kernel. A simple means of doing this is to treat the log itself as a resource protected by the kernel. The kernel interface could expose a generic "log" operation

$$\text{log} : (x : \textbf{string}) \rightarrow \text{K } \textbf{says } \text{OkToLog } x \rightarrow$$
$$\text{K } \textbf{says } \text{DidLog } x$$

with (hopefully permissive) rules for constructing OkToLog proofs. It might be especially useful to log failed attempts at proof construction. For example, users of the file system presented in Section 4 might repeatedly attempt to construct proofs for APPEND access given only the privileges necessary for WRONLY access, indicating that the rule maybe might be appropriate for their needs.

Conversely, some operations take arguments that should not be logged, perhaps due to security or space constraints. Section 4 mentions the possibility of logging file read and write operations, which touches on both these issues—even if it were practical to log all data read from and written to each file, many users would likely prefer that their file contents not be included in the system logs. Terms that must be excluded from the log, however, limit not just the scope of auditing but also the dependencies that may occur within propositions, as it would hardly suffice for data excised from the log to appear inside a type annotation.

## 6   Related Work

Earlier work on proof-carrying access control [4, 5, 14, 9, 10, 19] recognized the importance of **says** and provided a variety of interpretations for it. Garg and Pfenning [21] and, later, Abadi [2] introduced the treatment of **says** as an indexed monad. Both systems [21, 3] also enjoy the crucial noninterference property: in the absence of delegation, nothing B says can cause A to say false. AURA$_0$ builds on this prior work, especially Abadi's DCC, in several ways. The addition of dependent types enhances the expressiveness of DCC, and the addition of **sign** allows for a robust distributed interpretation of **says**. AURA$_0$'s treatment of principals as terms, as opposed to members of a special index set, enables quantification over principals. Lastly,

AURA$_0$ eliminates DCC's built-in acts-for lattice (which can be encoded as described in Section 3.3) along with the protects relation (which allows additional commutation and simplification of **says** with with regards to that lattice).

Our work is closely related to Fournet, Gordon and Maffeis's research on authorization in distributed systems. [19, 20] Fournet *et al.* work with an explicit $\pi$-calculus based model of computation. Like us, they use dependent types to express access control properties. Fournet and colleagues focus on the security properties that are maintained during execution, which are reflected into the type system using static theorem proving and a type constructor Ok. The inhabitants of Ok, however, do not contain dynamic information and cannot be logged for later audit. Additionally, while AURA$_0$ treats signing abstractly, Fournet and colleagues' type system (and computation model) can explicitly discuss cryptographic operations.

Trust management systems like PolicyMaker and Keynote [13] are also related to our work. Trust management systems are meant to determine whether a set of credentials proves that the request complies with a security policy, and they use general purpose compliance checkers to verify these credentials. In PolicyMaker, proofs are programs—written in a safe language—that operate on strings; a request $r$ is allowed when the application can combine proofs such that the result returns true on input $r$. While validity of AURA$_0$ propositions is tested by type checking, validity in PolicyMaker is tested by *evaluation*; this represents a fundamentally different approaches to logic. Similar to this paper, trust management systems intend for proof checking to occur in a small and application-independent trusted computing base; proof search may be delegated to untrusted components.

Proof carrying access control has been field tested by Bauer and colleagues in the Grey project [9, 10]. In their project, smart phones build proofs which can be used to open office doors or log into computer systems. The Grey architecture shares structural similarities with the model discussed in this paper: in Grey, proof generating devices, like our applications, need not reside in the trusted computing base, and both systems use expressive foundational logics to define policies (Grey uses higher-order logic [15]). In order to make proof search effective, Bauer suggests using cut-down fragments of higher order logic for expressing particular rule sets and using a distributed, tactic-based proof search algorithm.

Wee implemented the Logging and Auditing File System (LAFS) [33], a practical system which shares several architectural elements with AURA$_0$. LAFS uses a lightweight daemon, analogous to our kernel, to wrap NFS file systems; like our kernel, the LAFS daemon forwards all requests to the underlying resources. Both systems also configure policy using sets of rules defined outside the trusted computing base. The systems differ in three key respects. First, the LAFS policy language is too weak to express many AURA$_0$ policies. Second, AURA$_0$ requires some privileged K **says** rules to bootstrap a policy, while LAFS can be completely configured with non-privileged policies. Third, the LAFS interface is designed to be transparent to application code and does not provide any access control properties; instead LAFS logs—but does not prevent—rule violations.

Cederquist and colleagues describe a distributed system architecture with discretionary logging and no reference monitor [14]. In this system agents—i.e. principals—may choose to enter proofs (written in a first-order natural deduction style logic) into a a trusted log when performing actions. Cederquist *et al.* formalize accountability such that agents are guilty until proved innocent—that is, agents use log entries to reduce the quantity of actions for which they can be held accountable. This relies on the ability of some authority to independently observe certain actions; such observations are necessary to begin the audit process.

## 7 Conclusion

This paper has argued for evidence-based auditing, in which audit log entries contain proofs about authorization; such proofs are useful for minimizing the trusted computing base and provide information that can help debug policies. This paper has presented an architecture for structuring systems in terms of trusted kernels whose interfaces require proofs. As a concrete instance of this approach, this paper has developed AURA$_0$, a dependently-typed authorization logic that enjoys subject reduction and strong normalization properties. Several examples using AURA$_0$ have demonstrated how we envision applying these ideas in practice.

## 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. In *ICFP '06: Proc. of the 11th International Conference on Functional Programming*, pages 263–273, 2006.

[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 dis-

tributed 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] S. Axelsson, U. Lindqvist, U. Gustafson, and E. Jonsson. An approach to UNIX security logging. In *Proc. 21st NIST-NCSC National Information Systems Security Conference*, pages 62–75, 1998.

[7] Lujo Bauer. *Access Control for the Web via Proof-Carrying Authorization*. PhD thesis, Princeton U., November 2003.

[8] Lujo Bauer, Kevin D. Bowers, Frank Pfenning, and Michael K. Reiter. Consumable credentials in logic-based access control. Technical Report CMU-CYLAB-06-002, Carnegie Mellon University, February 2006.

[9] Lujo 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*, pages 431–445, 2005.

[10] Lujo Bauer, Scott Garriss, and Michael K. Reiter. Distributed proving in access-control systems. In *Proc. of the 2005 IEEE Symposium on Security & Privacy*, pages 81–95, May 2005.

[11] Mihir Bellare and Bennet Yee. Forward integrity for secure audit logs. Technical report, Computer Science and Engineering Department, U. California at San Diego, 1997.

[12] Matt Bishop. *Computer Security: Art and Science*. Addison-Wesley Professional, 2002.

[13] Matt Blaze, Joan Feigenbaum, John Ioannidis, and Angelos D. Keromytis. The role of trust management in distributed systems security. In *Secure Internet programming: security issues for mobile and distributed objects*, pages 185–210. Springer-Verlag, London, UK, 1999.

[14] 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.

[15] Alonzo Church. A formulation of the simple theory of types. *The Journal of Symbolic Logic*, 5(2):56–68, June 1940.

[16] Haskell B. Curry, Robert Feys, and William Craig. *Combinatory Logic*, volume 1. North-Holland, Amsterdam, 1958.

[17] John DeTreville. Binder, a logic-based security language. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 105–113, May 2002.

[18] Henry deYoung, Deepak Garg, and Frank Pfenning. An authorization logic with explicit time. In *Proc. of the 21st IEEE Computer Security Foundations Symposium*, 2008.

[19] Cédric Fournet, Andrew D. Gordon, and Sergio Maffeis. A type discipline for authorization policies. In *Proc. of the 14th European Symposium on Programming*, 2005.

[20] Cédric Fournet, Andrew D. Gordon, and Sergio Maffeis. A type discipline for authorization in distributed systems. In *Proc. of the 20th IEEE Computer Security Foundations Symposium*, 2007.

[21] Deepak Garg and Frank Pfenning. Non-interference in constructive authorization logic. In *Proc. of the 19th IEEE Computer Security Foundations Workshop*, pages 283–296, 2006.

[22] Herman Geuvers. A short and flexible proof of strong normalization for the calculus of constructions. In *TYPES '94: Selected papers from the International Workshop on Types for Proofs and Programs*, pages 14–38, London, 1995.

[23] M. A. Harrison, W. L Ruzzo, and J. D. Ullman. Protection in operating systems. *Comm. ACM*, 19(8):461–471, 1976.

[24] Sushil Jajodia, Pierangela Samarati, and V. S. Subrahmanian. A logical language for expressing authorizations. In *SP '97: Proceedings of the 1997 IEEE Symposium on Security and Privacy*, page 31. IEEE Computer Society, 1997.

[25] Limin Jia, Jeffrey A. Vaughan, Karl Mazurak, Jianzhou Zhao, Luke Zarko, Joseph Schorr, and Steve Zdancewic. Aura: A programming language for authorization and audit, preliminary technical results. Technical Report MS-CIS-08-10, U. Pennsylvania, 2008.

[26] C. Ko, M. Ruschitzka, and K. Levitt. Execution monitoring of security-critical programs in distributed systems: a specification-based approach. In *SP '97: Proceedings of the 1997 IEEE Symposium on Security and Privacy*, page 175, Washington, DC, USA, 1997. IEEE Computer Society.

[27] Ninghui Li, Benjamin N. Grosof, and Joan Feigenbaum. Delegation logic: A logic-based approach to distributed authorization. *ACM Trans. Inf. Syst. Secur.*, 6(1), 2003.

[28] G. C. Necula and P. Lee. Efficient representation and validation of proofs. In *LICS '98: Proceedings of the 13th Annual IEEE Symposium on Logic in Computer Science*, page 93, Washington, DC, USA, 1998. IEEE Computer Society.

[29] Bruce Schneier and John Kelsey. Cryptographic support for secure logs on untrusted machines. In *Proc. of the 7th on USENIX Security Symposium.*, pages 53–62, January 1998.

[30] Helmut Schwichtenberg. Normalization. Lecture Notes for Marktoberdorf Summer School, 1989.

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

[32] B. Waters, D. Balfanz, G. E. Durfee, and D. K. Smetters. Building an encrypted and searchable audit log. In *11th Annual Network and Distributed Security Symposium*, 2004.

[33] Christopher Wee. LAFS: A logging and auditing file system. In *Annual Computer Security Applications Conference*, pages 231–240, New Orleans, LA, USA, December 1995.

# A   Appendix

This appendix, not included in the conference version of the paper [31], provides definitions and proofs for $\text{AURA}_0$.

## A.1   Definitions

The substitution and free variable functions are defined as usual.

**Definition A.1** (Metafunctions over syntax)**.**

*Substitution of term s for variable z in term t, $\{s/z\}t$:*

$$
\begin{aligned}
\{s/z\}z &= s \\
\{s/z\}x &= x \quad\ \text{if } z \neq x \\
\{s/z\}(x{:}t_1) \to t_2 &= (x{:}\{s/z\}t_1) \to \{s/z\}t_2 \quad \text{if } z \neq x \\
\{s/z\}\{x{:}t_1; t_2\} &= \{x{:}\{s/z\}t_1; \{s/z\}t_2\} \quad\ \text{if } z \neq x \\
\{s/z\}\textbf{sign}(t_1, t_2) &= \textbf{sign}(\{s/z\}t_1, \{s/z\}t_2) \\
\{s/z\}\textbf{return}@[t_1]\ t_2 &= \textbf{return}@[\{s/z\}t_1]\ \{s/z\}t_2 \\
\{s/z\}\textbf{bind } x = t_1 \textbf{ in } t_2 &= \textbf{bind } x = \{s/z\}t_1 \textbf{ in } \{s/z\}t_2 \quad \text{if } z \neq x \\
\{s/z\}\lambda x{:}t_1.\, t_2 &= \lambda x{:}\{s/z\}t_1.\, \{s/z\}t_2 \quad \text{if } z \neq x \\
\{s/z\}t_1\ t_2 &= (\{s/z\}t_1)\ (\{s/z\}t_2) \\
\{s/z\}\langle t_1, t_2\rangle &= \langle\{s/z\}t_1, \{s/z\}t_2\rangle \\
\{s/z\}t &= t \quad \text{otherwise}
\end{aligned}
$$

*Substitution of term s for variable z in context $\Gamma$, $\{s/z\}\Gamma$:*

$$
\begin{aligned}
\{s/z\}\cdot &= \cdot \\
\{s/z\}\Gamma, x : t &= (\{s/z\}\Gamma), x : \{s/z\}t
\end{aligned}
$$

*Free variables of term t, $fv(t)$:*

$$
\begin{aligned}
fv(x) &= \{x\} \\
fv(t_1 \textbf{ says } t_2) &= fv(t_1) \cup fv(t_2) \\
fv((x{:}t_1) \to t_2) &= fv(t_1) \cup fv(t_2) \setminus \{x\} \\
fv(\{x{:}t_1; t_2\}) &= fv(t_1) \cup fv(t_2) \setminus \{x\} \\
fv(\textbf{sign}(t_1, t_2)) &= fv(t_1) \cup fv(t_2) \\
fv(\textbf{return}@[t_1]t_2) &= fv(t_1) \cup fv(t_2) \\
fv(\textbf{bind } x = t_1 \textbf{ in } t_2) &= fv(t_1) \cup fv(t_2) \setminus \{x\} \\
fv(\lambda x{:}t_1.\, t_2) &= fv(t_1) \cup fv(t_2) \setminus \{x\} \\
fv(t_1\ t_2) &= fv(t_1) \cup fv(t_2) \\
fv(\langle t_1, t_2\rangle) &= fv(t_1) \cup fv(t_2)
\end{aligned}
$$

*Free variables of context $\Gamma$, $fv(\Gamma)$:*

$$
\begin{aligned}
fv(\cdot) &= \cdot \\
fv(\Gamma, x : t) &= fv(\Gamma) \cup fv(t)
\end{aligned}
$$

$$\boxed{\Sigma \vdash \diamond}$$

$$\frac{}{\cdot \vdash \diamond} \text{ S-EMPTY} \qquad\qquad \frac{\Sigma \vdash \diamond \qquad \Sigma; \cdot \vdash t : \mathbf{Kind}^{\mathbf{P}}}{\Sigma, a : t \vdash \diamond} \text{ S-CONS}$$

$$\boxed{\Sigma \vdash \Gamma}$$

$$\frac{\Sigma \vdash \diamond}{\Sigma \vdash \cdot} \text{ E-EMPTY} \qquad \frac{\Sigma \vdash \Gamma \qquad \Sigma; \Gamma \vdash t : k \qquad \text{if } x \notin \mathit{fv}(\Gamma) \qquad k \in \{\mathbf{Kind}^{\mathbf{P}}, \mathbf{Kind}^{\mathbf{T}}, \mathbf{Prop}, \mathbf{Type}\}}{\Sigma \vdash \Gamma, x : t} \text{ E-CONS}$$

**Figure 7. Well formed signature and environment judgments (defined mutually with typing relation)**

$$\boxed{\vdash t \to t'}$$

$$\frac{}{\vdash (\lambda x{:}t_1.\, t_2)\; t_3 \to \{t_3/x\}t_2} \text{ R-BETA} \qquad\qquad \frac{x \notin \mathit{fv}(t_2)}{\vdash \mathbf{bind}\; x\; =\; t_1\; \mathbf{in}\; t_2 \to t_2} \text{ R-BINDS}$$

$$\frac{}{\vdash \mathbf{bind}\; x\; =\; \mathbf{return}@[t_0]\; t_1\; \mathbf{in}\; t_2 \to \{t_1/x\}t_2} \text{ R-BINDT} \qquad \frac{\vdash t_2 \to t_2'}{\vdash \mathbf{return}@[t_1]\; t_2 \to \mathbf{return}@[t_1]\; t_2'} \text{ R-SAYS}$$

$$\frac{y \notin \mathit{fv}(t_3)}{\vdash \mathbf{bind}\; x\; =\; (\mathbf{bind}\; y\; =\; t_1\; \mathbf{in}\; t_2)\; \mathbf{in}\; t_3 \to \mathbf{bind}\; y\; =\; t_1\; \mathbf{in}\; \mathbf{bind}\; x\; =\; t_2\; \mathbf{in}\; t_3} \text{ R-BINDC}$$

$$\frac{\vdash t_2 \to t_2'}{\vdash \lambda x{:}t_1.\, t_2 \to \lambda x{:}t_1.\, t_2'} \text{ R-LAM} \qquad\qquad \frac{\vdash t_1 \to t_1'}{\vdash \mathbf{bind}\; x\; =\; t_1\; \mathbf{in}\; t_2 \to \mathbf{bind}\; x\; =\; t_1'\; \mathbf{in}\; t_2} \text{ R-BIND1}$$

$$\frac{\vdash t_2 \to t_2'}{\vdash \mathbf{bind}\; x\; =\; t_1\; \mathbf{in}\; t_2 \to \mathbf{bind}\; x\; =\; t_1\; \mathbf{in}\; t_2'} \text{ R-BIND2} \qquad \frac{\vdash t_1 \to t_1'}{\vdash t_1\; t_2 \to t_1'\; t_2} \text{ R-APP1} \qquad \frac{\vdash t_2 \to t_2'}{\vdash t_1\; t_2 \to t_1\; t_2'} \text{ R-APP2}$$

$$\frac{\vdash t_1 \to t_1'}{\vdash \langle t_1, t_2 \rangle \to \langle t_1', t_2 \rangle} \text{ R-PAIR1} \qquad\qquad \frac{\vdash t_2 \to t_2'}{\vdash \langle t_1, t_2 \rangle \to \langle t_1, t_2' \rangle} \text{ R-PAIR2}$$

**Figure 8. Reduction relation**

## A.2  Subject Reduction

**Lemma A.1** (Weakening). *If $\Sigma; \Gamma, \Gamma' \vdash t_1 : t_2$ and $\Sigma \vdash \Gamma, x : t_3, \Gamma'$ then $\Sigma; \Gamma, x : t_3, \Gamma' \vdash t_1 : t_2$.*

*Proof.* By structural induction on the typing derivation. □

**Lemma A.2** (Inversion Var—Same). *If $\Sigma; \Gamma, z : u, \Gamma' \vdash z : s$ then $u = s$.*

*Proof.* Inverting the typing derivation (which ends in T-VAR) yields $z : y \in \Gamma, z : u, \Gamma'$ and $\Sigma \vdash \Gamma, z : u, \Gamma'$. Proof precedes by a trivial induction on the environment's well-formedness. □

**Lemma A.3** (Inversion Var—Different). *If $\Sigma; \Gamma, z : u, \Gamma' \vdash x : s$ and $x \neq z$ then $x : s \in \Gamma$ or $x : s \in \Gamma'$.*

*Proof.* By inverting the typing derivation, than structural induction on the environment's well formedness. □

**Lemma A.4** (Variables closed in context). *If $\Sigma; \Gamma \vdash s : t$ and $x \in \mathit{fv}(t) \cup \mathit{fv}(s)$ then $x \in \mathit{dom}(\Gamma)$.*

*Proof.* Proof by induction on the typing derivation. □

**Lemma A.5** (Context Ordering). *If* $\Sigma \vdash \Gamma, z : u$ *and* $x : s \in \Gamma$ *then* $z \notin fv(s)$.

*Proof.* By the definition of $\in$, we know $\Gamma = \Gamma_1, x : s, \Gamma_2$. The well-formedness of $\Gamma_1, x : s, \Gamma_2, z : u$ shows that (1) $dom(z) \notin \Gamma_1$ and (2) $\Sigma; \Gamma_1 \vdash s : k$ for some $k$. From these and Lemma A.4 we can conclude $z \notin fv(s)$. □

**Lemma A.6** (Well-formedness). *(1) If* $\Sigma; \Gamma \vdash t : s$ *then* $s \in \{\mathbf{Kind^P}, \mathbf{Kind^T}\}$ *or there exists* $k$ *such that* $\Sigma; \Gamma \vdash s : k$.
  *(2) If* $\Sigma \vdash \Gamma$ *and* $x : s \in \Gamma$ *then there exists* $k$ *such that* $\Sigma; \Gamma \vdash s : k$.
  *(3) If* $\Sigma \vdash \diamond$ *and* $a : s \in \Gamma$ *then there exists* $k$ *such that* $\Sigma; \cdot \vdash s : k$.

*Proof.* By mutual induction on the typing, well-formed signature, and well-formed environment judgments. □

**Lemma A.7** (Substitution, strong form for induction). *Assume* $\Sigma; \Gamma \vdash t_u : u$. *Then*

*(1)* $\Sigma; \Gamma, z : u, \Gamma' \vdash t : s$ *implies* $\Sigma; \Gamma, \{t_u/z\}\Gamma' \vdash \{t_u/z\}t : \{t_u/z\}s$. *And*

*(2)* $\Sigma \vdash \Gamma, z : u, \Gamma'$ *implies* $\Sigma \vdash \Gamma, \{t_u/z\}\Gamma'$.

*Proof.* By mutual structural induction over the typing and well-formed environment derivations. Proceed with inversion on the form of the last typing or well-formedness rule.

Case T-PROP. We have $t = \mathbf{Prop}$ and $s = \mathbf{Kind^P}$. So it suffices to show $\Sigma \vdash \Gamma, \{t_u/z\}\Gamma'$. This follows immediately from the induction hypothesis.

Case T-VAR. Suppose $t = z$. Then, by Lemma A.2 $s = u$. Therefore it suffices to show $\Sigma; \Gamma, \{t_u/z\}\Gamma' \vdash t_u : u$, which we get by applying weakening (finitely many times) to the assumption $\Sigma; \Gamma \vdash t_u : u$. Instead suppose $t = x \neq z$. Then we must show $\Sigma; \Gamma, \{t_u/z\}\Gamma' \vdash x : \{t_u/z\}s$. By Lemma A.3, either $x : s \in \Gamma$ or $x : s \in \Gamma'$. Suppose $x : s \in \Gamma$. Then by Lemma A.5 we find $z \notin fv(s)$ so $s = \{t_u/z\}s$. Thus it suffices to show $\Sigma; \Gamma, \{t_u/z\}\Gamma' \vdash x : s$, which follows from T-VAR, the induction hypothesis and the form of $\Gamma$. Lastly, consider the case that $x : s \in \Gamma'$. Then $x : \{t_u/z\}s \in \{t_u/z\}\Gamma'$, and we conclude using this, T-VAR, and the induction hypothesis.

Case T-LAM. We have $t = \lambda x{:}t_1. t_2$ and $s = (x{:}t_1) \rightarrow P$. Without loss of generality assume $x \neq z$. The induction hypothesis yields $\Sigma; \Gamma, \{t_u/z\}(\Gamma', x : t_1) \vdash \{t_u/z\}t_2 : \{t_u/z\}P$ and $\Sigma; \Gamma, \{t_u/z\}\Gamma' \vdash \{t_u/z\}(x{:}t_1) \rightarrow P : \mathbf{Prop}$. We conclude by applying T-LAM and the following facts about substitution: $\{t_u/z\}(\Gamma', x : t_1) = (\{t_u/z\}\Gamma'), x : (\{t_u/z\}t_1)$ and $\{t_u/z\}((x{:}t_1) \rightarrow t_2) = (x{:}\{t_u/z\}t_1) \rightarrow \{t_u/z\}t_2$. (The latter holds because $x \neq z$.)

Case T-BIND. This case is similar to T-LAM, but uses the additional fact that, for all $t_1$ and $t_2$, $\{t_u/z\}(t_1 \ \mathbf{says} \ t_2) = (\{t_u/z\}t_1) \ \mathbf{says} \ (\{t_u/z\}t_2)$.

Case T-SIGN. We have $t = \mathbf{sign}(t_1, t_2)$ and $s = t_1 \ \mathbf{says} \ t_2$. From Lemma A.4, we find $fv(t_1) = fv(t_2) = \emptyset$. Hence $\{t_u/z\}t = t$ and $\{t_u/z\}s = s$, so to use T-SIGN, we need only show $\Sigma \vdash \Gamma, \{t_u/z\}\Gamma'$. This follows immediately from the induction hypothesis.

Case T-PAIR. We have $s = \{s_1{:}x; s_2\}$. Assume without loss of generality $x \neq z$. This case follows from the induction hypothesis and the fact $\{t_u/z\}(\{t_1/x\}s_2) = \{(\{t_u/z\}t_1)/x\}(\{t_u/z\}t_2)$.

The remaining cases are similar to T-PROP (T-TYPE, T-STRING, T-CONST, T-PRIN, T-LITSTR, T-LITPRIN), or T-LAM (T-ARR, T-PAIRTYPE), or are trivial (T-SAYS, T-RETURN, T-APP). □

Subject reduction will need both substitution (above) and the following strengthening lemma (below). Note that strengthening is *not* a special case of of substitution, as strengthening works even when $u$ is uninhabited.

**Lemma A.8** (Strengthening). *If* $\Sigma; \Gamma, z : u, \Gamma' \vdash t : s$ *and* $\Sigma \vdash \Gamma, \Gamma'$ *and* $z \notin fv(t) \cup fv(s)$ *then* $\Sigma; \Gamma, \Gamma' \vdash t : s$.

*Proof.* Proof by structural induction on the typing relation.

Case T-VAR. Then $t$ is a variable, $x$. By the definition of $fv(\cdot)$, $x \neq z$. Inverting the typing relation yields $x : s \in \Gamma, z : u, \Gamma'$. Thus $z \in \Gamma, \Gamma'$, and we conclude with T-VAR.

Case T-PAIR. We have $s = \{x{:}s_1; s_2\}$ for some $x$, $s_1$, and $s_2$. By A.6 and a simple case analysis, $\Sigma; \Gamma \vdash s : k$ where $k \in \{\mathbf{Prop}, \mathbf{Type}, \mathbf{Kind^P}, \mathbf{Kind^T}\}$. Thus $z \notin fv(k)$. Thus the case follows from the induction hypothesis and T-PAIR.

All other cases follow directly from the induction hypothesis. □

**Lemma A.9** (Subject Reduction). *If* $\vdash t \rightarrow t'$ *and* $\Sigma; \Gamma \vdash t : s$ *then* $\Sigma; \Gamma \vdash t' : s$.

*Proof.* Proof is by structural induction on the reduction relation. Proceed by case analysis on the last rule used.

Case R-BETA. We have $t = (\lambda x{:}t_1.\, t_2)\, t_3$ and $t' = \{t_3/x\}t_2$. Term $t$ could only have been typed by a derivation ending in

$$
\text{T-APP} \cfrac{\text{T-LAM} \cfrac{\vdots \quad \Sigma;\Gamma, x:t_1 \vdash t_2 : s_2 \quad \dots}{\Sigma;\Gamma \vdash \lambda x{:}t_1.\, t_2 : (x{:}t_1) \to s_2} \qquad \cfrac{\vdots}{\Sigma;\Gamma \vdash t_3 : s_3}}{\Sigma;\Gamma \vdash (\lambda x{:}t_1.\, t_2)\, t_3 : \{s_3/x\}s_2}
$$

for some $s_2$ and $s_3$. So $s = \{t_3/x\}t_2$. That $\Sigma;\Gamma \vdash t' : s$ holds follows directly from Lemma A.7 and the judgments written in the above derivation.

Case R-BINDS. We have $t = \mathbf{bind}\ x\ =\ t_1\ \mathbf{in}\ t_2$ and $t' = t_2$. Term $t$ could only be typed by T-BIND, and inverting this rule gives $s = a\ \mathbf{says}\ s_2$ and $\Sigma;\Gamma, x:s_1 \vdash t_2 : a\ \mathbf{says}\ s_2$. Before concluding with Lemma A.8, we must show $x \notin a\ \mathbf{says}\ s_2$. This is a consequence of Lemma A.4, and the hypothesis that $a\ \mathbf{says}\ s_2$ is a type assignment in $\Gamma$.

Case R-BINDT. We have $t = \mathbf{bind}\ x\ =\ \mathbf{return}@[t_0]\ t_1\ \mathbf{in}\ t_2$ and $t' = \{t_1/x\}t_2$ and $s = t_0\ \mathbf{says}\ s_2$. Term $t$ can only be typed by a derivation ending with, for some $s_1$,

$$
\text{T-BIND} \cfrac{\text{T-APP} \cfrac{\vdots \quad \cfrac{\vdots}{\Sigma;\Gamma \vdash t_1 : s_1}}{\Sigma;\Gamma \vdash \mathbf{return}@[t_0]\ t_1 : t_0\ \mathbf{says}\ s_1} \qquad \cfrac{\vdots}{\Sigma;\Gamma, x:s_1 \vdash t_2 : t_0\ \mathbf{says}\ s_2}}{\Sigma;\Gamma \vdash \mathbf{bind}\ x\ =\ \mathbf{return}@[t_0]\ t_1\ \mathbf{in}\ t_2 : t_0\ \mathbf{says}\ s_2}
$$

By Lemma A.7, we find $\Sigma;\Gamma \vdash \{t_1/x\}t_2 : \{t_1/x\}(t_0\ \mathbf{says}\ s_2)$. The contrapositive of Lemma A.4 shows $x \notin t_0\ \mathbf{says}\ s_2$, so we can rewrite the above to $\Sigma;\Gamma \vdash t' : s$.

Case R-BINDC. We have $t = \mathbf{bind}\ x\ =\ (\mathbf{bind}\ y\ =\ t_1\ \mathbf{in}\ t_2)\ \mathbf{in}\ t_3$ and $s = u\ \mathbf{says}\ s_3$. Following the Barendregt variable convention, assume $y \notin fv(\Gamma) \cup \{x\}$. Inverting the typing derivation twice shows, for some $s_1$ and $s_2$, that $\Sigma;\Gamma \vdash t_1 : u\ \mathbf{says}\ s_1$, $\Sigma;\Gamma, y:s_1 \vdash t_2 : u\ \mathbf{says}\ s_2$, $\Sigma;\Gamma, x:s_2 \vdash t_3 : u\ \mathbf{says}\ s_3$, and $x \notin fv(s_3)$. With Lemma A.1 we find $\Sigma;\Gamma, y:s_1, x:s_2 \vdash t_3 : u\ \mathbf{says}\ s_3$. With Lemma A.4, $y \notin fv(s_3)$. We conclude using T-BINDC twice.

The remaining cases follow directly from the induction hypothesis. $\qquad\square$

## A.3 Proof of Strong Normalization

We prove $\text{AURA}_0$ is strongly normalizing by translating $\text{AURA}_0$ to the Calculus of Construction extended with product dependent types (CC).

The main property of the translation, which we will prove later in this section, is that the translation has to preserves both the typing relation and the reduction relation. The translation of terms has the form: $[\![t]\!]_\Delta = (s, \Delta')$, where context $\Delta$ is a typing context for variables. To translate a $\text{AURA}_0$ term, we take in a context $\Delta$, and produce a new context $\Delta'$ together with a term in CC.

Before we present the formal definitions of the translation, we define the following auxiliary definitions.

**Definitions**

- *unique*$(\Delta)$ if for all $fvar_1, fvar_2 \in dom(\Delta)$, $\Delta(fvar_1) \neq \Delta(fvar_2)$.

- *wf*$(\Gamma)$:

$$
\cfrac{}{wf(\cdot)} \qquad\qquad \cfrac{\Gamma \vdash^{CC} t : s \qquad s \in \{*, \square\} \qquad v \notin dom(\Gamma)}{wf(\Gamma, v : t)}
$$

The translation of $\text{AURA}_0$ terms to CC terms is defined in Figure A.3. The translation collapses $\mathbf{Kind^P}$ and $\mathbf{Kind^T}$ to the kind $\square$ in CC, and $\mathbf{Prop}$, $\mathbf{Type}$ to $*$. We translate all base types to $\mathbf{unit}$, and constants to $(\ )$. The interesting cases are the translation of DCC terms. The translation drops the monads, and translates the $\mathbf{bind}$ expression to lambda application. The

term $\textbf{sign}(t_1, t_2)$ has type $t_1$ **says** $t_2$; therefore, it has to be translated to a term whose type is the translation of $t_2$. One way to find such a term is to generate a fresh variable and assign its type to be the translation of $t_2$. The context $\Delta$ is used to keep track of the type mapping of those fresh variables generated. There are two cases in translation $\textbf{sign}(t_1, t_2)$. In the first case, the variable we need has already been generated. In the second case, we need to generate a fresh variable and append its type binding to $\Delta_1$ as the output context. To make proofs easier, we assume that the fresh variables are denoted by a *fvar*, not to be confused with the variable $x$.

Both of $\text{AURA}_0$'s signature $\Sigma$ and context $\Gamma$ are translated into CC's typing context. The translation of $\Sigma$ has the form $[\![\Sigma]\!] = \Sigma'$. The translation of $\Gamma$ context has the form $[\![\Gamma]\!]_\Sigma = (\Gamma', \Delta')$. The context $\Delta'$ contains all the fresh variables generated while translating the types in $\Gamma$. One subtlety of the context translation is that it has "weakening" built-in. Notice that in the translation of $\Gamma, v : t$, the translation of $\Gamma$ yields $\Sigma_1, \Delta_1$, but $t$ is translated in the larger context $\Sigma_1, \Delta_1, \Delta_2$. This also means that the translation of context $\Gamma$ is not unique. The judgment $[\![\Gamma]\!]_\Sigma = (\Gamma', \Delta')$ is more precisely read as $(\Gamma', \Delta')$ is a legitimate translation of $\Gamma$ given $\Sigma$. This is good enough for our proof because we only need to show that for any well-typed $\text{AURA}_0$ term $t$, there is a typing derivation for the translation of $t$ in CC.

**Lemma A.10** (Translation Weakening). *If $[\![t]\!]_{\Delta_1} = (s, \Delta_2)$, unique($\Delta$), and $(\Delta_1, \Delta_2) \subseteq \Delta$, then $[\![t]\!]_\Delta = (s, \cdot)$.*

*Proof.* By induction on the structure of $t$. The key is when t is $\textbf{sign}(t_1, t_2)$.

case: $t = \textbf{sign}(t_1, t_2)$.
    By assumptions,

$$unique(\Delta) \tag{1}$$
$$[\![\textbf{sign}(t_1, t_2)]\!]_{\Delta_1} = (x, \Delta_2)$$
$$\text{and } [\![t_2]\!]_{\Delta_1} = (s, \Delta_2), \quad (\Delta_1, \Delta_2)(fvar) = s \tag{2}$$
$$(\Delta_1, \Delta_2) \subseteq \Delta \tag{3}$$

    By I.H. on $t_2$,

$$[\![t_2]\!]_\Delta = (s, \cdot) \tag{4}$$

    By (2), (1), (3),

$$\Delta(fvar) = s \tag{5}$$

    By the rules for translation,

$$[\![\textbf{sign}(t_1, t_2)]\!]_\Delta = (fvar, \cdot) \tag{6}$$

case: $t = \textbf{sign}(t_1, t_2)$.
    By assumptions,

$$unique(\Delta) \tag{1}$$
$$[\![\textbf{sign}(t_1, t_2)]\!]_{\Delta_1} = (fvar_1, (\Delta_2, fvar_1 : s))$$
$$\text{and } [\![t_2]\!]_{\Delta_1} = (s, \Delta_2),$$
$$\nexists fvar \in dom(\Delta_2) \text{ s.t. } (\Delta, \Delta_2)(fvar) = s \tag{2}$$
$$(\Delta_1, \Delta_2, fvar : s) \subseteq \Delta \tag{3}$$

    By I.H. on $t_2$,

$$[\![t_2]\!]_\Delta = (s, \cdot) \tag{4}$$

    By (1), (3),

$$\Delta(fvar_1) = s \tag{5}$$

    By the rules for translation,

$$[\![\textbf{sign}(t_1, t_2)]\!]_\Delta = (fvar_1, \cdot) \tag{6}$$

$\square$

**Lemma A.11** (CC Typing Weakening). *If $\Gamma_1, \Gamma_2 \vdash^{CC} t : s$, and wf($\Gamma_1, \Gamma', \Gamma_2$), then $\Gamma_1, \Gamma', \Gamma_2 \vdash^{CC} t : s$.*

*Proof.* By induction on structure of the derivation $\mathcal{E} :: \Gamma_1, \Gamma_2 \vdash^{CC} t : s$. $\square$

**Lemma A.12** (CC well-formed term gives well-formed environment). *If $\Gamma \vdash^{CC} t : s$ then wf($\Gamma$).*

*Proof.* By induction on the typing derivation. $\square$

**Lemma A.13** (CC well-formed term gives well-formed type). *If $\Gamma \vdash^{CC} t : s$ then either $s = \square$ or exists $k$ such that $\Gamma \vdash^{CC} s : k$.*

$\boxed{[\![t]\!]_\Delta = (s, \Delta')}$

$$\frac{\text{if } t \in \{\mathbf{Kind^P}, \mathbf{Kind^T}\}}{[\![t]\!]_\Delta = (\square, \cdot)} \qquad \frac{\text{if } t \in \{\mathbf{Prop}, \mathbf{Type}\}}{[\![t]\!]_\Delta = (*, \cdot)} \qquad \frac{\text{if } t \in \{\texttt{"a"}, \ldots, \mathsf{A} \ldots\}}{[\![t]\!]_\Delta = ((), \cdot)} \qquad \frac{\text{if } t \in \{\mathbf{string}, \mathbf{prin}\}}{[\![t]\!]_\Delta = (\mathbf{unit}, \cdot)}$$

$$[\![a]\!]_\Delta = (a, \cdot) \qquad\qquad [\![x]\!]_\Delta = (x, \cdot) \qquad\qquad \overline{[\![t_1 \text{ says } t_2]\!]_\Delta = [\![t_2]\!]_\Delta}$$

$$\frac{[\![t_1]\!]_\Delta = (s_1, \Delta_1) \qquad [\![t_2]\!]_{\Delta, \Delta_1} = (s_2, \Delta_2)}{[\![(x{:}t_1) \to t_2]\!]_\Delta = ((x{:}s_1) \to s_2, (\Delta_1, \Delta_2))} \qquad \frac{[\![t_1]\!]_\Delta = (s_1, \Delta_1) \qquad [\![t_2]\!]_{\Delta, \Delta_1} = (s_2, \Delta_2)}{[\![(x : t_1) \Rightarrow t_2]\!]_\Delta = ((x{:}s_1) \to s_2, (\Delta_1, \Delta_2))}$$

$$\frac{[\![t_1]\!]_\Delta = (s_1, \Delta_1) \qquad [\![t_2]\!]_{\Delta, \Delta_1} = (s_2, \Delta_2)}{[\![\{x{:}t_1; t_2\}]\!]_\Delta = (\{x{:}s_1; s_2\}, (\Delta_1, \Delta_2))} \qquad \frac{[\![t_2]\!]_\Delta = (s, \Delta_1) \qquad (\Delta, \Delta_1)(\textit{fvar}) = s}{[\![\mathbf{sign}(t_1, t_2)]\!]_\Delta = \textit{fvar}, \Delta_1}$$

$$\frac{[\![t_2]\!]_\Delta = (s, \Delta_1) \qquad \text{not exists } \textit{fvar} \in \textit{dom}(\Delta, \Delta_1) s.t. (\Delta, \Delta_1)(\textit{fvar}) = s \qquad y \text{ is fresh}}{[\![\mathbf{sign}(t_1, t_2)]\!]_\Delta = (\textit{freshvar}, (\Delta_1, y : s))}$$

$$\frac{}{[\![\mathbf{return}@[t_1] \; t_2]\!]_\Delta = [\![t_2]\!]_\Delta} \qquad \frac{[\![t_0]\!]_\Delta = (s, \Delta_1) \qquad [\![t_1]\!]_{\Delta, \Delta_1} = (s_1, \Delta_2) \qquad [\![t_2]\!]_{\Delta, \Delta_1, \Delta_2} = (s_2, \Delta_3)}{[\![\mathbf{bind} \; x{:}t_0 \; = \; t_1 \; \mathbf{in} \; t_2]\!]_\Delta = ((\lambda x{:}s. \; s_2) \; s_1, (\Delta_1, \Delta_2, \Delta_3))}$$

$$\frac{[\![t_1]\!]_\Delta = (s_1, \Delta_1) \qquad [\![t_2]\!]_{\Delta, \Delta_1} = (s_2, \Delta_2)}{[\![\lambda x{:}t_1. \; t_2]\!]_\Delta = (\lambda x{:}s_1. \; s_2, (\Delta_1, \Delta_2))} \qquad \frac{[\![t_1]\!]_\Delta = (s_1, \Delta_1) \qquad [\![t_2]\!]_{\Delta, \Delta_1} = (s_2, \Delta_2)}{[\![t_1 \; t_2]\!]_\Delta = (s_1 \; s_2, (\Delta_1, \Delta_2))}$$

$$\frac{[\![t_1]\!]_\Delta = (s_1, \Delta_1) \qquad [\![t_2]\!]_{\Delta, \Delta_1} = (s_2, \Delta_2)}{[\![\langle t_1, t_2 \rangle]\!]_\Delta = (\langle s_1, s_2 \rangle, (\Delta_1, \Delta_2))}$$

$\boxed{[\![\Sigma]\!] = \Sigma'}$

$$\frac{}{[\![\cdot]\!] = \cdot} \qquad\qquad \frac{[\![\Sigma]\!] = \Sigma' \qquad [\![t]\!]_\Sigma = (s, \Delta_2)}{[\![\Sigma, v : t]\!] = (\Sigma', \Delta_2, v : s)}$$

$\boxed{[\![\Gamma]\!]_\Sigma = (\Gamma', \Delta)}$

$$\frac{}{[\![\cdot]\!]_\Sigma = (\cdot, \cdot)} \qquad\qquad \frac{[\![\Gamma]\!]_\Sigma = (\Gamma', \Delta_1) \qquad \textit{wf}(\Sigma, \Delta_1, \Delta_2) \qquad \textit{unique}(\Sigma, \Delta_1, \Delta_2) \qquad [\![t]\!]_{\Sigma, \Delta_1, \Delta_2} = (s, \Delta_3)}{[\![\Gamma, v : t]\!]_\Sigma = ((\Gamma', v : s), (\Delta_1, \Delta_2, \Delta_3))}$$

**Figure 9. Translation of** $\textsc{Aura}_0$**'s terms and contexts to CC**

*Proof.* By induction on the typing derivation. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

**Lemma A.14** (Substitution). *If* $\Sigma; \Gamma \vdash t_1 : k$, *unique*$(\Delta)$, $[\![t_1]\!]_\Delta = (s_1, \cdot)$ *and* $[\![t_2]\!]_\Delta = (s_2, \cdot)$, *then* $[\![\{t_2/x\}t_1]\!]_\Delta = (\{s_2/x\}s_1, \cdot)$.

*Proof.* By induction on the structure of $t_1$.

case: $t_1 = \mathbf{sign}(t, p)$.
    By assumption,

$$[\![t_2]\!]_\Delta = (s_2, \cdot) \tag{1}$$
$$\Sigma; \Gamma \vdash \mathbf{sign}(t, p) : k \tag{2}$$
$$[\![\mathbf{sign}(t, p)]\!]_\Delta = (\textit{fvar}, \cdot) \tag{3}$$

By the definition of translation, (3),

$$\llbracket p \rrbracket_\Delta = (s, \cdot) \tag{4}$$

$$\text{and } \Delta(\textit{fvar}) = s \tag{5}$$

By inversion on (2),

$$\Sigma; \cdot \vdash p : \textbf{Prop} \tag{6}$$

$$x \notin \textit{fv}(p) \tag{7}$$

$$\{t_2/x\}p = p \tag{8}$$

By (8), (4), (5),

$$\llbracket \{t_2/x\}(\textbf{sign}(t, p)) \rrbracket_\Delta = (\textit{fvar}, \cdot) = (\{s_2/x\}\textit{fvar}, \cdot) \tag{9}$$

$\square$

**Lemma A.15** (Correctness of Translation)**.**

1. *If* $\Sigma \vdash \diamond$, $\llbracket \Sigma \rrbracket = \Sigma_1$, *then* $\textit{wf}(\Sigma_1)$ *and* $\textit{unique}(\Sigma_1)$.

2. *If* $\Sigma \vdash \Gamma$, $\llbracket \Sigma \rrbracket = \Sigma_1$, $\llbracket \Gamma \rrbracket_{\Sigma_1} = (\Gamma_1, \Delta)$, *then* $\textit{wf}(\Sigma_1, \Delta, \Gamma_1)$ *and* $\textit{unique}(\Sigma_1, \Delta)$.

3. *If* $\mathcal{E} :: \Sigma; \Gamma \vdash t : s$, $\llbracket \Sigma \rrbracket = \Sigma_1$, $\llbracket \Gamma \rrbracket_{\Sigma_1} = (\Gamma_1, \Delta_1)$, $\textit{wf}(\Sigma_1, \Delta_1, \Delta_2, \Gamma_1)$, $\textit{unique}(\Sigma_1, \Delta_1, \Delta_2)$, $\llbracket t \rrbracket_{\Sigma_1, \Delta_1, \Delta_2} = (t_1, \Delta_3)$, *then* $\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1 \vdash^{CC} t_1 : s_1$, *and* $\llbracket s \rrbracket_{(\Sigma_1, \Delta_2, \Delta_3)} = (s_1, \cdot)$ *and* $\textit{unique}(\Sigma_1, \Delta_1, \Delta_2, \Delta_3)$.

4. *If* $\mathcal{E} :: \Sigma; \Gamma \vdash t$, $\llbracket \Sigma \rrbracket = \Sigma_1$, $\llbracket \Gamma \rrbracket_{\Sigma_1} = (\Gamma_1, \Delta_1)$, $\textit{wf}(\Sigma_1, \Delta_1, \Delta_2, \Gamma_1)$, $\textit{unique}(\Sigma_1, \Delta_1, \Delta_2)$, $\llbracket t \rrbracket_{\Sigma_1, \Delta_1} = (t_1, \Delta_3)$, *then* $\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1 \vdash^{CC} t_1 : */\square$ *(read as $t_3$ is classified by $*$ or $\square$), and* $\textit{unique}(\Sigma_1, \Delta_1, \Delta_2, \Delta_3)$.

*Proof.* In the proof of 1 and 2, we use 3 only when $\Sigma$ or $\Gamma$ is smaller.
1. By induction on the structure of $(\Sigma)$.

case: $\Sigma = \Sigma', a : t$
    By assumption,

$$\Sigma', a : t \vdash \diamond \tag{1}$$

$$\llbracket \Sigma', a : t \rrbracket = (\Sigma_1, \Delta, a : s) \tag{2}$$

$$\text{where } \llbracket \Sigma' \rrbracket = \Sigma_1 \tag{3}$$

$$\text{and} \llbracket t \rrbracket_{\Sigma_1} = (s, \Delta) \tag{4}$$

By inversion of (1),

$$\Sigma' \vdash \diamond \tag{5}$$

$$\Sigma'; \cdot \vdash t : \textbf{Kind}^{\textbf{P}} \tag{6}$$

By I.H. on $\Sigma'$,

$$\textit{wf}(\Sigma_1) \text{ and } \textit{unique}(\Sigma_1) \tag{7}$$

By 3, (6), (7), (4),

$$\Sigma_1, \Delta \vdash^{CC} s : \square \text{ and } \textit{unique}(\Sigma_1, \Delta) \tag{8}$$

By definition of *wf* and *unique*, and (8),

$$\textit{wf}(\Sigma_1, \Delta, a : s), \text{ and } \textit{unique}(\Sigma_1, \Delta, a : s) \tag{9}$$

2. By induction on the structure of $\Gamma$.

case: $\Gamma = \Gamma', x : t$
    By assumption,

$$\Sigma \vdash \Gamma', x : t \tag{1}$$

$$\llbracket \Sigma \rrbracket = \Sigma_1 \tag{2}$$

$$\llbracket \Gamma', x : t \rrbracket_{\Sigma_1} = ((\Gamma_1, x : s), (\Delta_1, \Delta_2, \Delta_3)) \tag{3}$$

$$\text{where } \llbracket \Gamma' \rrbracket_{\Sigma_1} = (\Gamma_1, \Delta_1) \tag{4}$$

$$\text{and } \textit{wf}(\Sigma_1, \Delta_1, \Delta_2), \textit{unique}(\Sigma_1, \Delta_1, \Delta_2) \tag{5}$$

$$\text{and } \llbracket t \rrbracket_{\Sigma_1, \Delta_1, \Delta_2} = (s, \Delta_3) \tag{6}$$

By inversion of (1),

$$\Sigma \vdash \Gamma' \tag{7}$$

$$\Sigma; \Gamma' \vdash t : \textbf{Kind}^{\textbf{P}}/\textbf{Kind}^{\textbf{T}}/\textbf{Prop}/\textbf{Type} \tag{8}$$

22

By 3, (2), (3), (4), (5), (6), (8),
$$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1 \vdash^{CC} s : \Box/* \tag{9}$$
$$\text{and } unique(\Sigma_1, \Delta_1, \Delta_2, \Delta_3) \tag{10}$$
By definition of *wf*, and (9),
$$wf(\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1, x : s) \tag{11}$$

3. By induction on the structure of the derivation $\mathcal{E}$.

case: $\mathcal{E}$ ends in T-PROP.

By assumption,
$$\mathcal{E} = \frac{\mathcal{E}' :: \Sigma \vdash \Gamma}{\Sigma; \Gamma \vdash \textbf{Prop} : \textbf{Kind}^\textbf{P}} \tag{1}$$
$$\llbracket \Sigma \rrbracket = \Sigma_1 \tag{2}$$
$$\llbracket \Gamma \rrbracket_{\Sigma_1} = (\Gamma_1, \Delta_1) \tag{3}$$
$$\llbracket \textbf{Prop} \rrbracket_{\Sigma_1, \Delta_1} = (*, \cdot) \tag{4}$$
By *ax* rule,
$$\cdot \vdash^{CC} * : \Box \tag{5}$$
By 2, $\mathcal{E}'$, (2), (3),
$$wf(\Sigma_1, \Delta_1, \Gamma_1) \text{ and } unique(\Sigma_1, \Delta_1) \tag{6}$$
By Lemma weakening,
$$\Sigma_1, \Delta_1, \Gamma_1 \vdash^{CC} * : \Box \tag{7}$$

case: $\mathcal{E}$ ends in T-ARR.

By assumption,
$$\mathcal{E} = \frac{\mathcal{E}_1 :: \Sigma; \Gamma \vdash t_1 : (\textbf{Kind}^\textbf{P}, \textbf{Type}, \textbf{Prop}) \qquad \mathcal{E}_2 :: \Sigma; \Gamma, x : t_1 \vdash t_2 : k_2 \qquad k_2 \in \{\textbf{Kind}^\textbf{P}, \textbf{Prop}\}}{\Sigma; \Gamma \vdash (x{:}t_1) \to t_2 : k_2} \tag{1}$$
$$\llbracket \Sigma \rrbracket = \Sigma_1 \tag{2}$$
$$\llbracket \Gamma \rrbracket_{\Sigma_1} = (\Gamma_1, \Delta_1) \tag{3}$$
$$wf(\Sigma_1, \Delta_1, \Delta_2) \text{ and } unique(\Sigma_1, \Delta_1, \Delta_2) \tag{4}$$
$$\llbracket (x{:}t_1) \to t_2 \rrbracket_{\Sigma_1, \Delta_1, \Delta_2} = ((x{:}s_1) \to s_2, (\Delta_3, \Delta_4)) \tag{5}$$
$$\text{where } \llbracket t_1 \rrbracket_{\Sigma_1, \Delta_1, \Delta_2} = (s_1, \Delta_3) \tag{6}$$
$$\text{and } \llbracket t_2 \rrbracket_{\Sigma_1, \Delta_1, \Delta_2, \Delta_3} = (s_2, \Delta_4) \tag{7}$$
By I.H. on $\mathcal{E}_1$,
$$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1 \vdash^{CC} s_1 : (*, \Box) \tag{8}$$
$$\text{and } unique(\Sigma_1, \Delta_1, \Delta_2, \Delta_3) \tag{9}$$
By Definition of the translation of $\Gamma$, (3), (4),
$$\llbracket \Gamma, x : t_1 \rrbracket_{\Sigma_1} = ((\Gamma_1, x : s_1), (\Delta_1, \Delta_2, \Delta_3)) \tag{10}$$
By Lemma A.12, (8),
$$wf(\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1) \tag{11}$$
By I.H. on $\mathcal{E}_2$, (7), (10), (11), (9),
$$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4, \Gamma_1, x : s_1 \vdash^{CC} s_2 : (*/\Box) \tag{12}$$
$$\text{and } unique(\Sigma_1, \Delta_1, \Delta_3, \Delta_4) \tag{13}$$
By $\Pi$, (8), (12),
$$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4, \Gamma_1 \vdash^{CC} (x{:}s_1) \to s_2 : (*/\Box) \tag{14}$$

case: $\mathcal{E}$ ends in T-SIGN.

By assumption,
$$\mathcal{E} = \frac{\mathcal{E}_1 :: \Sigma \vdash \Gamma \qquad \Sigma; \cdot \vdash t_1 : \textbf{prin} \qquad \mathcal{E}_2 :: \Sigma; \cdot \vdash t_2 : \textbf{Prop}}{\Sigma; \Gamma \vdash \textbf{sign}(t_1, t_2) : t_1 \textbf{ says } t_2} \tag{1}$$
$$\llbracket \Sigma \rrbracket = \Sigma_1 \tag{2}$$
$$\llbracket \Gamma \rrbracket_{\Sigma_1} = (\Gamma_1, \Delta_1) \tag{3}$$
$$wf(\Sigma_1, \Delta_1, \Delta_2), unique(\Sigma_1, \Delta_1, \Delta_2) \tag{4}$$
$$\llbracket \textbf{sign}(t_1, t_2) \rrbracket_{\Sigma_1, \Delta_1, \Delta_2} = (fvar, \Delta_3) \tag{5}$$
$$\text{where } \llbracket t_2 \rrbracket_{\Sigma_1, \Delta_1, \Delta_2} = (s_2, \Delta_3) \tag{6}$$

and $(\Sigma_1, \Delta_1, \Delta_2, \Delta_3)(\textit{fvar}) = s_2$ (7)

By I.H. on $\mathcal{E}_2$, (2), (4), (6),

$\Sigma_1, \Delta_1, \Delta_2, \Delta_3 \vdash^{CC} s_2 : *$ (8)

and $\textit{unique}(\Sigma_1, \Delta_1, \Delta_2, \Delta_3)$ (9)

By (7), (8), Lemma A.11 weakening,

$\Sigma_1, \Delta_1, \Delta_2, \Delta_3 \vdash^{CC} \textit{fvar} : s_2$ (10)

By 2 on $\mathcal{E}_1$, (2), (3),

$\textit{wf}(\Sigma_1, \Delta_1, \Gamma_1)$ (11)

By Weakening and the domain of $\Gamma_1$ and $\Delta_2 \; \Delta_3$ are disjoint,

$\textit{wf}(\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1)$ (12)

By Lemma A.11 weakening, (12), (10),

$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1 \vdash^{CC} \textit{fvar} : s_2$ (13)

case: $\mathcal{E}$ ends in T-SIGN.

By assumption,

$$\mathcal{E} = \dfrac{\mathcal{E}_1 :: \Sigma \vdash \Gamma \qquad \Sigma; \cdot \vdash t_1 : \mathbf{prin} \qquad \mathcal{E}_2 :: \Sigma; \cdot \vdash t_2 : \mathbf{Prop}}{\Sigma; \Gamma \vdash \mathbf{sign}(t_1, t_2) : t_1 \; \mathbf{says} \; t_2}$$ (1)

$[\![\Sigma]\!] = \Sigma_1$ (2)

$[\![\Gamma]\!]_{\Sigma_1} = (\Gamma_1, \Delta_1)$ (3)

$\textit{wf}(\Sigma_1, \Delta_1, \Delta_2), \textit{unique}(\Sigma_1, \Delta_1, \Delta_2)$ (4)

$[\![\mathbf{sign}(t_1, t_2)]\!]_{\Sigma_1, \Delta_1, \Delta_2} = (\textit{fvar}_1, (\Delta_3, \textit{fvar}_1 : s_2))$ (5)

where $[\![t_2]\!]_{\Sigma_1, \Delta_1, \Delta_2} = (s_2, \Delta_3)$ (6)

and $\not\exists \textit{fvar}$ such that $(\Sigma_1, \Delta_1, \Delta_2, \Delta_3)(\textit{fvar}) = s_2$, and $\textit{fvar}$ is fresh (7)

By I.H. on $\mathcal{E}_2$, (2), (4), (6),

$\Sigma_1, \Delta_1, \Delta_2, \Delta_3 \vdash^{CC} s_2 : *$ (8)

and $\textit{unique}(\Sigma_1, \Delta_1, \Delta_2, \Delta_3)$ (9)

By $\textit{var}$ rule, (8),

$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \textit{fvar} : s_2 \vdash^{CC} \textit{fvar} : s_2$ (10)

By (9), (7),

$\textit{unique}(\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \textit{fvar} : s_2)$ (11)

By 2 on $\mathcal{E}_1$, (2), (3),

$\textit{wf}(\Sigma_1, \Delta_1, \Gamma_1)$ (12)

By Weakening and the domain of $\Gamma_1$ and $\Delta_2 \; (\Delta_3, \textit{fvar}_1 : s_2)$ are disjoint,

$\textit{wf}(\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \textit{fvar}_1 : s_2, \Gamma_1)$ (13)

By Lemma A.11 weakening, (13), (10),

$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \textit{fvar}_1 : s_2, \Gamma_1 \vdash^{CC} \textit{fvar} : s_2$ (14)

case: $\mathcal{E}$ ends in T-APP RULE

By assumption,

$$\mathcal{E} = \dfrac{\mathcal{E}_1 :: \Sigma; \Gamma \vdash t_1 : (x{:}u_2) \rightarrow u \qquad \mathcal{E}_2 :: \Sigma; \Gamma \vdash t_2 : u_2}{\Sigma; \Gamma \vdash t_1 \; t_2 : \{t_2/x\}u}$$ (1)

$[\![\Sigma]\!] = \Sigma_1$ (2)

$[\![\Gamma]\!]_{\Sigma_1} = (\Gamma_1, \Delta_1)$ (3)

$\textit{wf}(\Sigma_1, \Delta_1, \Delta_2), \textit{unique}(\Sigma_1, \Delta_1, \Delta_2)$ (4)

$[\![t_1 \; t_2]\!]_{\Sigma_1, \Delta_1} = (s_1 \; s_2, (\Delta_2, \Delta_3, \Delta_4))$ (5)

and $[\![t_1]\!]_{\Sigma_1, \Delta_1, \Delta_2} = (s_1, \Delta_3)$ (6)

and $[\![t_2]\!]_{\Sigma_1, \Delta_1, \Delta_2, \Delta_3} = (s_2, \Delta_4)$ (7)

By I.H. on $\mathcal{E}_1$,

$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1 \vdash^{CC} s_1 : k$ (8)

and $(k, \cdot) = [\![(x{:}u_2) \rightarrow u]\!]_{\Sigma_1, \Delta_1, \Delta_2, \Delta_3}$ (9)

and $\textit{unique}(\Sigma_1, \Delta_1, \Delta_2, \Delta_3)$ (10)

By (8), and Lemma A.12,

$\textit{wf}(\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Gamma_1)$ (11)

24

By I.H. on $\mathcal{E}_2$, (10), (11), (7),
$$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4, \Gamma_1 \vdash^{CC} s_2 : k_2 \qquad (12)$$
$$\text{where } (k_2, \cdot) = [\![u_2]\!]_{\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4} \qquad (13)$$
$$\text{and } unique\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4 \qquad (14)$$
By translation weakening Lemma A.10 and (9),
$$\text{and } (k, \cdot) = [\![(x{:}u_2) \to u]\!]_{\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4} \qquad (15)$$
By definition of translation and (15), (13),
$$k = (x{:}k_2) \to ku \text{ and } (ku, \cdot) = [\![u]\!]_{\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4} \qquad (16)$$
By *app* rule, (8), (12), (16),
$$\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4, \Gamma_1 \vdash^{CC} s_1 \ s_2 : \{s_2/x\}ku \qquad (17)$$
By translation weakening Lemma A.10 and (7),
$$[\![t_2]\!]_{\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4} = (s_2, \cdot) \qquad (18)$$
By Lemma A.14, (16), (18),
$$[\![\{t_2/x\}u]\!]_{\Sigma_1, \Delta_1, \Delta_2, \Delta_3, \Delta_4} = (\{s_2/x\}ku, \cdot) \qquad (19)$$

4. By induction on the structure of the derivation $\mathcal{E}$. □

The following $\beta'$ reduction rule mirrors the commute reduction rule in AURA$_0$.

**Special Reduction Rule:**
$(\lambda x{:}t. \, t_1)((\lambda y{:}s. \, t_2)u) \to_{\beta'} (\lambda y{:}s. \, ((\lambda x{:}t. \, t_1)t_2))u$

Calculus of Construction extended with product dependent types is know to be strongly normalizing [22]. We use $\mathbf{SN}(\beta)$ to denote the set of terms that are strongly normalizing under $\beta$ reductions in CC; similarly, $\mathbf{SN}(\beta\beta')$ is the set of terms that are strongly normalizing under the $\beta$ and $\beta'$ reduction rules. We demonstrate that CC augmented with $\beta'$ is also strongly normalizing.

**Lemma A.16** (Strong normalization of $\beta\beta'$-reduction in CC). *For all term* $t \in \mathbf{SN}(\beta)$, $t \in \mathbf{SN}(\beta\beta')$.

*Proof.* We assign an ordering between terms as the dictionary order of a pair $(\beta(t), \delta(t))$, where $\beta(t)$ is the maximum *beta*-reduction steps of $t$, and $\delta(t)$ is defined as follows. $\delta(x) = 1, \delta(\lambda x{:}t. \, s) = \delta(s), \delta(t_1 \, t_2) = \delta(t_1) + 2\delta(t_2)$. We then prove that if $t \to_{\beta'} t'$ then $\beta(t') \leq \beta(t)$, by examining all possible $\beta$-reductions of $t'$, and showing that $t$ has an corresponding reduction that takes at least the same number of $\beta$-reduction steps as $t'$. Now $\delta((\lambda y{:}s. \, ((\lambda x{:}t. \, t_1)t_2))u) = \delta(t_1) + 2\delta(t_2) + 2\delta(u)$, and $\delta(\lambda x{:}t. \, t_1)((\lambda y{:}s. \, t_2)u) = \delta(t_1) + 2\delta(t_2) + 4\delta(u)$. Therefore, when $t \to_{\beta'} t'$, $(\beta(t'), \delta(t')) < (\beta(t), \delta(t))$. Consequently, for all term $t \in \mathbf{SN}(\beta)$, $t \in \mathbf{SN}(\beta\beta')$. □

Now we prove that the reductions in CC augmented with the $\beta'$ reduction rule simulates the reduction in AURA$_0$.

**Lemma A.17** (Simulation). *If* $t \to t'$, *and and* $[\![t]\!]_\Delta = (s, \Delta)$, $[\![t']\!]_\Delta = (s', \Delta)$, *then* $s \to^+_{\beta, \beta'} s'$.

*Proof.* By examining all the reduction rules. □

**Lemma A.18** (Strong normalization). AURA$_0$ *is strongly normalizing.*

*Proof.* By Lemma A.17, and Lemma A.16. A diverging path in AURA$_0$ implies a diverging path in CC. Since CC is strongly normalizing, AURA$_0$ is also strongly normalizing. □

**Lemma A.19.** *If* $s \to s'$, *then* $\{t/x\}s \to^* \{t/x\}s'$.

*Proof.* By induction on the structure of $s$. □

**Lemma A.20.** *If* $t \to t'$, *then* $\{t/x\}s \to^* \{t'/x\}s$.

*Proof.* By induction on the structure of $s$. □

**Lemma A.21** (Weak Confluence). *If* $t \to t_1$, $t \to t_2$, *then exists* $t_3$ *such that* $t_1 \to^* t_3$, *and* $t_2 \to^* t_3$.

*Proof.* By induction on the structure of $t$. We invoke induction hypothesis directly in most of the cases. We show a few key cases below.

Case : $t = \lambda x{:}u.\, s$

By assumption,

$$t_1 = \lambda x{:}u.\, s_1 \text{ where } s \to s_1 \tag{1}$$

$$t_2 = \lambda x{:}u.\, s_2 \text{ where } s \to s_2 \tag{2}$$

By I.H. on $s$,

$$\exists s_3 \text{ such that } s_1 \to^* s_3, \text{ and } s_2 \to^* s_3 \tag{3}$$

By reduction rules,

$$t_3 = \lambda x{:}u.\, s_3 \text{ such that } t_1 \to^* t_3 \text{ and } t_2 \to^* t_3 \tag{4}$$

Case: $t = (\lambda x{:}u.\, s_1)s_2$, $t_1 = \{s_2/x\}s_1$, and
$\quad t_2 = (\lambda x{:}u.\, s_1')s_2$ where $s_1 \to s_1'$.

By Lemma A.20,

$$t_1 \to^* \{s_2/x\}s_1' \tag{1}$$

By reduction rules,

$$t_2 \to \{s_2/x\}s_1' \tag{2}$$

Case: $t = \mathbf{bind}\ x\ =\ s_1\ \mathbf{in}\ s_2$
$\quad$ where $s_1 = \mathbf{bind}\ y\ =\ \mathbf{return}@[a]u_1\ \mathbf{in}\ u_2$

By assumption,

$$t_1 = \mathbf{bind}\ x\ =\ \{u_1/x\}u_2\ \mathbf{in}\ s_2 \tag{1}$$

$$t_2 = \mathbf{bind}\ y\ =\ \mathbf{return}@[a]u_1\ \mathbf{in}\ \mathbf{bind}\ x\ =\ u_2\ \mathbf{in}\ s_2$$
$$\text{and } y \notin fv(s_2) \tag{2}$$

By reduction R-BINDT,

$$t_2 \to \mathbf{bind}\ x\ =\ \{u_1/y\}u_2\ \mathbf{in}\ \{u_1/y\}s_2 \tag{3}$$

By (2), (3),

$$t_2 \to \mathbf{bind}\ x\ =\ \{u_1/y\}u_2\ \mathbf{in}\ s_2 \tag{4}$$

Case: $t = \mathbf{bind}\ x\ =\ s_1\ \mathbf{in}\ s_2$
$\quad$ where $s_1 = \mathbf{bind}\ y\ =\ u_1\ \mathbf{in}\ u_2$
$\quad$ and $u_1 = \mathbf{bind}\ z\ =\ w_1\ \mathbf{in}\ w_2$

By assumption,

$$t_1 = \mathbf{bind}\ x\ =\ s_1'\ \mathbf{in}\ s_2$$
$$\text{where } s_1' = \mathbf{bind}\ z\ =\ w_1\ \mathbf{in}\ \mathbf{bind}\ y\ =\ w_2\ \mathbf{in}\ u_2 \tag{1}$$
$$t_2 = \mathbf{bind}\ y\ =\ u_1\ \mathbf{in}\ \mathbf{bind}\ x\ =\ u_2\ \mathbf{in}\ s_2$$

By applying R-BINDC rule many times,

$$t_1 \to^* t_3 \tag{2}$$

$$t_2 \to^* t_3 \tag{3}$$

$$\text{where } t_3 = \mathbf{bind}\ z\ =\ w_1\ \mathbf{in}\ \mathbf{bind}\ y\ =\ w_2\ \mathbf{in}\ \mathbf{bind}\ x\ =\ u_2\ \mathbf{in}\ s_2$$

$\square$