# Foundations for Auditing Assurance

Sepehr Amir-Mohammadian
University of Vermont
samirmoh@uvm.edu

Stephen Chong
Harvard University
chong@seas.harvard.edu

Christian Skalka
University of Vermont
ceskalka@uvm.edu

## ABSTRACT

Retrospective security is an important element of layered security systems. Auditing is central to the theory and practice of retrospective security, however, in systems where auditing is used, programs are typically instrumented to generate audit logs using manual, ad-hoc strategies. This is a potential source of error even if log auditing techniques are formal, since the relation of the log itself to program execution is unclear. This paper focuses on provably correct program rewriting algorithms for instrumenting formal logging specifications. Correctness guarantees that execution of an instrumented program produces sound and complete audit logs, properties defined by an information containment relation between logs and the program's logging semantics. As an application example, we consider auditing for break the glass policies, wherein authorization is replaced by auditing in emergency conditions.

## 1. INTRODUCTION

*Retrospective security* is the enforcement of security, or detection of security violations, after program execution [31, 35, 39]. Many real-world systems use retrospective security. For example, the financial industry corrects errors and fraudulent transactions not by prospectively preventing suspicious transactions, but by retrospectively correcting or undoing these problematic translations. Another example is a hospital whose employees are trusted to access confidential patient records, but who might (rarely) violate this trust [15]. Upon detection of such violations, security is enforced retrospectively by holding responsible employees accountable [40].

Retrospective security is often used in combination with prospective security methods such as access control [33, 23]. These approaches coexist since retrospective security cannot be achieved entirely by prospective computer security mechanisms. Reasons include that detection of violations may be external to the computer system (such as consumer reports of fraudulent transactions, or confidential patient information appearing in news media), the high cost of access denial (e.g., preventing emergency-room physicians from accessing medical records) coupled with high trust of systems users (e.g., users are trusted employees that rarely violate this trust) [41]. In addition, remediation actions to address violations may also be external to the computer system, such as reprimanding employees, prosecuting law suits, or otherwise holding users accountable for their actions [40].

*Auditing* underlies retrospective security frameworks and has become increasingly important to the theory and practice of cyber security. By recording appropriate aspects of a computer system's execution an audit log (and subsequent examination of the audit log) can enable detection of violations, and provide sufficient evidence to hold users accountable for their actions and support other remediation actions. For example, an audit log can be used to determine *post facto* which users performed dangerous operations, and can provide evidence for use in litigation.

However, despite the importance of auditing to real-world security, relatively little work has focused on the formal foundations of auditing, particularly with respect to defining and ensuring the correctness of audit log generation. Indeed, correct and efficient audit log generation poses at least two significant challenges. First, it is necessary to record sufficient and correct information in the audit log. If a program is manually instrumented, it is possible for developers to fail to record relevant events. Recent work showed that major health informatics systems do not log sufficient information to determine compliance with HIPAA policies [28]. Second, an audit log should ideally not contain more information than needed. While it is straightforward to collect sufficient information by recording essentially *all* events in a computer system, this can cause performance issues, both slowing down the system due to generating massive audit logs, and requiring the handling of extremely large audit logs. Excessive data collection is a key challenge for auditing [22, 12, 27], and is a critical factor in the design of tools that generate and employ audit logs (e.g., spam filters [13]).

A main goal of this paper is to establish a formal foundation for audit logging, especially to establish general correctness conditions for audit logs. Our broader goal is to eventually reason about and implement assured systems combining both prospective and retrospective security measures [23]. We define a general semantics of audit logs using the theory of *information algebra* [29]. We interpret both program execution traces and audit logs as information elements in an information algebra. A *logging specification* defines the intended relation between the information in traces and in audit logs. An audit log is correct if it satisfies this relation. A benefit of this formulation is that it separates logging specifications from programs, rather than burying them in code and implementation details.

Separating logging specifications from programs supports clearer definitions and more direct reasoning. Additionally, it enables algorithms for implementing general classes of logging specifications. Our formal theory establishes conditions that guarantee enforcement of logging specifications by such algorithms. As we will

show, correct instrumentation of logging specifications is a safety property, hence enforceable by security automata [37]. Inspired by related approaches to security automata implementation [20], we focus on program rewriting to automatically enforce correct audit instrumentation. Program rewriting has a number of practical benefits versus, for example, program monitors, such as lower OS process management overhead.

This approach would allow system administrators to define logging specifications which are automatically instrumented in code, including legacy code. Implementation details and matters such as optimization can be handled by the general program rewriting algorithm, not the logging specification. Furthermore, establishing correctness of a program rewriting algorithm provides an important security guarantee. Such an algorithm ensures that logging specifications will be implemented correctly, even if the rewritten source code contains malicious code or programmer errors.

## 1.1 A Motivating Example: Break the Glass

Although audit logs contain information *about* program execution, they are not just a straightforward selection of program events. Illustrative examples from practice include so-called "break the glass policies" used in electronic medical record systems [33]. These policies use access control to disallow care providers from performing sensitive operations such as viewing patient records, however care providers can "break the glass" in an emergency situation to temporarily raise their authority and access patient records, *with the understanding that subsequent sensitive operations will be logged and potentially audited*. This is a clear example from practice of the interaction between prospective and retrospective security methods. One potential accountability goal is the following:

> *In the event that a patient's sensitive information is inappropriately leaked, determine who accessed a given patient's files due to "breaking the glass."*

Since it cannot be predicted a priori whose information may leak, this goal can be supported by using an audit log that records all reads of sensitive files following glass breaking. To generate correct audit logs, programs must be instrumented for logging appropriately, i.e., to implement the following *logging specification* that we call $LS_H$:

> $LS_H$ : *Record in the log all patient information file reads following a break the glass event, along with the identity of the user that broke the glass.*

If at some point in time in the future it is determined that a specific patient **P**'s information was leaked, logs thus generated can be analyzed with the following query that we call $LQ_H$:

> $LQ_H$ : *Retrieve the identity of all users that read **P**'s information files.*

The specification $LS_H$ and the query $LQ_H$ together constitute an auditing policy that directly supports the above-stated accountability goal. Their separation is useful since at the time of execution the information leak is unknown, hence **P** is not known. Thus while it is possible to implement $LS_H$ as part of program execution, $LQ_H$ must be implemented retrospectively.

It is crucial to the enforcement of the above accountability goal that $LS_H$ is implemented correctly. If logging is incomplete then some potential recipients may be missed. If logging is overzealous then bloat is possible and audit logs become "write only". These types of errors are common in practice [28]. To establish formal correctness of instrumentation for audit logs, it is necessary to define a formal language of logging specifications, and establish techniques to guarantee that instrumented programs satisfy logging specifications. That is the focus of this paper. Other work has focused on formalisms for querying logs [38, 16], however these works presuppose correctness of audit logs for true accountability.

## 1.2 Threat Model

With respect to program rewriting (i.e., automatic techniques to instrument existing programs to satisfy a logging specification), we regard the program undergoing instrumentation as untrusted. That is, the program source code may have been written to avoid, confuse, or subvert the automatic instrumentation techniques. We do, however, assume that the source code is well-formed (valid syntax, well-typed, etc.). Moreover, we trust the compiler, the program rewriting algorithm, and the runtime environment in which the instrumented program will ultimately be executed. Non-malleability of generated audit logs, while important, is beyond the scope of this paper.

# 2. A SEMANTICS OF AUDIT LOGGING

Our goal in this Section is to formally characterize logging specifications and correctness conditions for audit logs. To obtain a general model, we leverage ideas from the theory of *information algebra* [30, 29], which is an abstract mathematical framework for information systems. In short, we interpret program traces as information, and logging specifications as functions from traces to information. This separates logging specifications from their implementation in code, and defines exactly the information that should be in an audit log. This in turn establishes correctness conditions for audit logging implementations.

Following [37], an *execution trace* $\tau = \kappa_0\kappa_1\kappa_2\ldots$ is a possibly infinite sequence of configurations $\kappa$ that describe the state of an executing program. We deliberately leave configurations abstract, but examples abound and we explore a specific instantiation for a $\lambda$-calculus in Section 4. Note that an execution trace $\tau$ may represent the partial execution of a program, i.e. the trace $\tau$ may be extended with additional configurations as the program continues execution. We use metavariables $\tau$ and $\sigma$ to range over traces.

An *information algebra* contains information elements $X$ (e.g. a set of logical assertions) taken from a set $\Phi$ (the algebra). A partial ordering is induced on $\Phi$ by the so-called *information ordering* relation $\leq$, where intuitively for $X, Y \in \Phi$ we have $X \leq Y$ iff $Y$ contains at least as much information as $X$, though its precise meaning depends on the particular algebra. We assume given a function $\lfloor\cdot\rfloor$ that is an injective mapping from traces to $\Phi$. This mapping *interprets a given trace as information*, where the injective requirement ensures that information is not lost in the interpretation. For example, if $\sigma$ is a proper prefix of $\tau$ and thus contains strictly less information, then formally $\lfloor\sigma\rfloor \leq \lfloor\tau\rfloor$. We intentionally leave both $\Phi$ and $\lfloor\cdot\rfloor$ underspecified for generality, though application of our formalism to a particular logging implementation requires instantiation of them. We discuss an example in Section 3.

We let $LS$ range over *logging specifications*, which are functions from traces to $\Phi$. As for $\Phi$ and $\lfloor\cdot\rfloor$, we intentionally leave the language of specifications abstract, but consider a particular instantiation in Section 3. Intuitively, $LS(\tau)$ denotes the information that should be recorded in an audit log during the execution of $\tau$ given specification $LS$, regardless of whether $\tau$ actually records any log information, correctly or incorrectly. We call this the semantics of the logging specification $LS$.

We assume that auditing is implementable, requiring at least that all conditions for logging any piece of information must be met in a finite amount of time. As we will show, this restriction implies that correct logging instrumentation is a safety property [37].

*Definition 1.* We require of any logging specification $LS$ that for all traces $\tau$ and information $X \leq LS(\tau)$, there exists a finite prefix $\sigma$ of $\tau$ such that $X \leq LS(\sigma)$.

It is crucial to observe that some logging specifications may *add* information not contained in traces to the auditing process. Security information not relevant to program execution (such as ACLs), interpretation of event data (statistical or otherwise), etc., may be added by the logging specification. As an example consider the OpenMRS system [34], in which logging of sensitive operations includes a human-understandable "type" designation, not used by any other code. Thus, given a trace $\tau$ and logging specification $LS$, it is *not* necessarily the case that $LS(\tau) \leq \lfloor \tau \rfloor$. Audit logging is not just a filtering of program events.

## 2.1 Correctness Conditions for Audit Logs

A logging specification defines what information should be contained in an audit log. In this section we develop formal notions of *soundness* and *completeness* as audit log correctness conditions. We use metavariable $\mathbb{L}$ to range over audit logs. Again, we intentionally leave the language of audit logs unspecified, but assume that the function $\lfloor \cdot \rfloor$ is extended to audit logs, i.e. $\lfloor \cdot \rfloor$ is an injective mapping from audit logs to $\Phi$. Intuitively, $\lfloor \mathbb{L} \rfloor$ denotes the information in $\mathbb{L}$, interpreted as an element of $\Phi$.

An audit log $\mathbb{L}$ is sound with respect to a logging specification $LS$ and trace $\tau$ if the log information is contained in $LS(\tau)$. Similarly, an audit log is complete with respect to a logging specification if it contains all of the information in the logging specification's semantics. Crucially, both definitions are independent of the implementation details that generate $\mathbb{L}$.

*Definition 2.* Audit log $\mathbb{L}$ is *sound with respect to logging specification $LS$ and execution trace $\tau$* iff $\lfloor \mathbb{L} \rfloor \leq LS(\tau)$.

*Definition 3.* Audit log $\mathbb{L}$ is *complete with respect to logging specification $LS$ and execution trace $\tau$* iff $LS(\tau) \leq \lfloor \mathbb{L} \rfloor$.

*The relation to log queries..*

As discussed in Section 1.1, we make a distinction between logging specifications such as $LS_H$ which define how to record logs, and log queries such as $LQ_H$ which ask questions of logs, and our notions of soundness and completeness apply strictly to logging specifications. However, any logging query must assume a logging specification semantics, hence a log that is demonstrably sound and complete provides the same answers on a given query that an "ideal" log would. This is an important property that is discussed in previous work, e.g. as "sufficiency" in [4].

## 2.2 Correct Logging Instrumentation is a Safety Property

In case program executions generate audit logs, we write $\tau \rightsquigarrow \mathbb{L}$ to mean that trace $\tau$ generates $\mathbb{L}$, i.e. $\tau = \kappa_0 \ldots \kappa_n$ and $logof(\kappa_n) = \mathbb{L}$ where $logof(\kappa)$ denotes the audit log in configuration $\kappa$, i.e. the residual log after execution of the full trace. Ideally, information that *should* be added to an audit log, *is* added to an audit log, immediately as it becomes available. This ideal is formalized as follows.

*Definition 4.* For all logging specifications $LS$, the trace $\tau$ is *ideally instrumented for $LS$* iff for all finite prefixes $\sigma$ of $\tau$ we have $\sigma \rightsquigarrow \mathbb{L}$ where $\mathbb{L}$ is sound and complete with respect to $LS$ and $\sigma$.

We observe that the restriction imposed on logging specifications by Definition 1, implies that ideal instrumentation of any logging specification is a safety property in the sense defined by Schneider [37][1].

THEOREM 1. *For all logging specifications $LS$, the set of ideally instrumented traces is a safety property.*

This result implies that e.g. edit automata can be used to enforce instrumentation of logging specifications. However, theory related to safety properties and their enforcement by execution monitors [37, 2] do not provide an adequate semantic foundation for audit log generation, nor an account of soundness and completeness of audit logs.

## 2.3 Implementing Logging Specifications with Program Rewriting

The above-defined correctness conditions for audit logs provide a foundation on which to establish correctness of logging implementations. Here we consider program rewriting approaches. Since rewriting concerns specific languages, we introduce an abstract notion of programs $\mathfrak{p}$ with an operational semantics that can produce a trace. We write $\mathfrak{p} \Downarrow \tau$ iff program $\mathfrak{p}$ can produce execution trace $\tau$, either deterministically or non-deterministically.

A rewriting algorithm $\mathcal{R}$ is a (partial) function that takes a program $\mathfrak{p}$ in a source language and a logging specification $LS$ and produces a new program, $\mathcal{R}(\mathfrak{p}, LS)$, in a target language.[2] The intent is that the target program is the result of instrumenting $\mathfrak{p}$ to produce an audit log appropriate for the logging specification $LS$. A rewriting algorithm may be partial, in particular because it may only be intended to work for a specific set of logging specifications.

Ideally, a rewriting algorithm should preserve the semantics of the program it instruments. That is, $\mathcal{R}$ is semantics-preserving if the rewritten program simulates the semantics of the source code, modulo logging steps. We assume given a correspondence relation $:\approx$ on execution traces. A coherent definition of correspondence should be similar to a bisimulation, but it is not necessarily symmetric nor a bisimulation, since the instrumented target program may be in a different language than the source program. We deliberately leave the correspondence relation underspecified, as its definition will depend on the instantiation of the model. Possible definitions are that traces produce the same final value, or that traces when restricted to a set of memory locations are equivalent up to stuttering. We provide an explicit definition of correspondence for $\lambda$-calculus source and target languages in Section 4.

*Definition 5.* Rewriting algorithm $\mathcal{R}$ is *semantics preserving* iff for all programs $\mathfrak{p}$ and logging specifications $LS$ such that $\mathcal{R}(\mathfrak{p}, LS)$ is defined, all of the following hold:

1. For all traces $\tau$ such that $\mathfrak{p} \Downarrow \tau$ there exists $\tau'$ with $\tau :\approx \tau'$ and $\mathcal{R}(\mathfrak{p}, LS) \Downarrow \tau'$.
2. For all traces $\tau$ such that $\mathcal{R}(\mathfrak{p}, LS) \Downarrow \tau$ there exists a trace $\tau'$ such that $\tau' :\approx \tau$ and $\mathfrak{p} \Downarrow \tau'$.

In addition to preserving program semantics, a correctly rewritten program constructs a log in accordance with the given logging specification. More precisely, if $LS$ is a given logging specification and a trace $\tau$ describes execution of a source program, rewriting should produce a program with a trace $\tau'$ that corresponds to $\tau$ (i.e., $\tau :\approx \tau'$), where the log $\mathbb{L}$ generated by $\tau'$ contains the same

[1]The proofs of Theorems 1-5 in this text are omitted for brevity, but are available from the authors upon request.

[2]We use metavariable $\mathfrak{p}$ to range over programs in either the source or target language; it will be clear from context which language is used.

information as $LS(\tau)$, or at least a sound approximation. Some definitions of $:\approx$ may allow several target-language traces to correspond to source-language traces (as for example in Section 4, Definition 12). In any case, we expect that at least one simulation exists. Hence we write $simlogs(\mathfrak{p}, \tau)$ to denote a nonempty set of logs $\mathbb{L}$ such that, given source language trace $\tau$ and target program $\mathfrak{p}$, there exists some trace $\tau'$ where $\mathfrak{p} \Downarrow \tau'$ and $\tau :\approx \tau'$ and $\tau' \rightsquigarrow \mathbb{L}$. The name *simlogs* evokes the relation to logs resulting from simulating executions in the target language.

The following definitions then establish correctness conditions for rewriting algorithms. Note that satisfaction of either of these conditions only implies condition (i) of Definition 5, not condition (ii), so semantics preservation is an independent condition.

*Definition 6.* Rewriting algorithm $\mathcal{R}$ is *sound* iff for all programs $\mathfrak{p}$, logging specifications $LS$, and finite traces $\tau$ where $\mathfrak{p} \Downarrow \tau$, for all $\mathbb{L} \in simlogs(\mathcal{R}(\mathfrak{p}, LS), \tau)$ it is the case that $\mathbb{L}$ is sound with respect to $LS$ and $\tau$.

*Definition 7.* Rewriting algorithm $\mathcal{R}$ is *complete* iff for all programs $\mathfrak{p}$, logging specifications $LS$, and finite traces $\tau$ where $\mathfrak{p} \Downarrow \tau$, for all $\mathbb{L} \in simlogs(\mathcal{R}(\mathfrak{p}, LS), \tau)$ it is the case that $\mathbb{L}$ is complete with respect to $LS$ and $\tau$.

# 3. LANGUAGES FOR LOGGING SPECIFICATIONS

Now we go into more detail about information algebra and why it is a good foundation for logging specifications and semantics. For a detailed account of information algebra, the reader is referred to a definitive survey paper [30]– available space disallows a detailed account here. In short, in addition to a definition of the elements of $\Phi$, any information algebra $\Phi$ includes two basic operators:

- Combination: The operation $X \otimes Y$ *combines* the information in elements $X, Y \in \Phi$.
- Focusing: The operation $X^{\Rightarrow S}$ isolates the elements of $X \in \Phi$ that are relevant to a *sublanguage* $S$, i.e. the subpart of $X$ specified by $S$.

Focusing and combination must additionally satisfy certain properties. The definitions of elements $X \in \Phi$, sublanguages $S$, combination, and focusing constitute the definition of the algebra. In all cases, the relation $X \leq Y$ holds iff $X \otimes Y = Y$. Proving that $\otimes$ has been correctly defined for an algebra implies that $\leq$ is a partial order [30].

## 3.1 Support for Various Approaches

Various approaches are taken to audit log generation and representation, including logical [16], database [1], and probabilistic approaches [42]. Information algebra is sufficiently general to contain relevant systems as instances, so our notions of soundness and completeness can apply broadly. Here we discuss logical and database approaches.

### *First Order Logic (FOL).*

Logics have been used in several well-developed auditing systems [24, 8], for the encoding of both audit logs and queries. FOL in particular is attractive due to readily available implementation support, e.g. Datalog and Prolog.

Let Greek letters $\phi$ and $\psi$ range over FOL formulas and let capital letters $X, Y, Z$ range over sets of formulas. We posit a sound and complete proof theory supporting judgements of the form $X \vdash \phi$. In this text we assume without loss of generality a natural deduction proof theory.

Elements of our algebra are sets of formulas closed under logical entailment. Intuitively, given a set of formulas $X$, the closure of $X$ is the set of formulas that are logically entailed by $X$, and thus represents all the information contained in $X$. In spirit, we follow the treatment of sentential logic as an information algebra explored in related foundational work [29], however our definition of closure is syntactic, not semantic.

*Definition 8.* We define a closure operation $C$, and a set $\Phi_{FOL}$ of closed sets of formulas:

$$C(X) = \{\phi \mid X \vdash \phi\} \qquad \Phi_{FOL} = \{X \mid C(X) = X\}$$

Note in particular that $C(\varnothing)$ is the set of logical tautologies.

Let $Preds$ be the set of all predicate symbols, and let $S \subseteq Preds$ be a set of predicate symbols. We define *sublanguage* $L_S$ to be the set of well-formed formulas over predicate symbols in $S$ (and including boolean atoms $T$ and $F$, and closed under the usual first-order connectives and binders). We will use sublanguages to define refinement operations in our information algebra. Subset containment induces a lattice structure, denoted $\mathcal{S}$, on the set of all sublanguages, with $\mathcal{F} = L_{Preds}$ as the top element.

Now we can define the focus and combination operators, which are the fundamental operators of an information algebra. Focusing isolates the component of a closed set of formulas that is in a given sublanguage. Combination closes the union of closed sets of formulas. Intuitively, the focus of a closed set of formulas $X$ to sublanguage $L$ is the refinement of the information in $X$ to the formulas in $L$. The combination of closed sets of formulas $X$ and $Y$ combines the information of each set.

*Definition 9.* Define:

1. Focusing: $X^{\Rightarrow S} = C(X \cap L_S)$ where $X \in \Phi_{FOL}$, $S \in Preds$
2. Combination: $X \otimes Y = C(X \cup Y)$ where $X, Y \subseteq \Phi_{FOL}$

These definitions of focusing and combination enjoy a number of properties within the algebra, as stated in the following Theorem, establishing that the construction is an information algebra. FOL has been treated as an information algebra before, but our definitions of combination and focusing and hence the result are novel.

THEOREM 2. *Structure $(\Phi_{FOL}, \mathcal{S})$ with focus operation $X^{\Rightarrow S}$ and combination operation $X \otimes Y$ forms a* domain-free information algebra.

In addition, to interpret traces and logs as elements of this algebra, i.e. to define the function $\lfloor \cdot \rfloor$, we assume existence of a function $toFOL(\cdot)$ that injectively maps traces and logs to sets of FOL formulas, and then take $\lfloor \cdot \rfloor = C(toFOL(\cdot))$. To define the range of $toFOL(\cdot)$, that is, to specify how trace information will be represented in FOL, we assume the existence of *configuration description predicates* $P$ which are each at least unary. Each configuration description predicate fully describes some element of a configuration $\kappa$, and the first argument is always a natural number $t$, indicating the time at which the configuration occurred. A set of configuration description predicates with the same timestamp describes a configuration, and traces are described by the union of sets describing each configuration in the trace. In particular, the configuration description predicates include predicate $\mathrm{Call}(t, f, x)$, which indicates that function $f$ is called at time $t$ with argument $x$. We will fully define $toFOL(\cdot)$ when we discuss particular source and target languages for program rewriting.

EXAMPLE 1. *We return to the example described in Section 1.1 to show how FOL can express break the glass logging specifications. Adapting a logic programming style, the trace of a program can be viewed as a fact base, and the logging specification $LS_H$ performs resolution of a* LoggedCall *predicate, defined via the following Horn clause we call $\psi_H$:*

$$\forall t, d, s, u.(\mathrm{Call}(t, \mathbf{read}, d) \wedge \mathrm{Call}(s, \mathbf{breakGlass}, u) \wedge s < t \\ \wedge \mathrm{PatientInfo}(d)) \implies \mathrm{LoggedCall}(t, \mathbf{read}, u, d)$$

*Here we imagine that* **breakGlass** *is a break the glass function where $u$ identifies the current user and* PatientInfo *is a predicate specifying which files contain patient information. The log contains only valid instances of* LoggedCall *given a particular trace, which specify the user and sensitive information accessed following glass breaking, which otherwise would be disallowed by a separate access control policy.*

Formally, we define logging specifications in a logic programming style by using combination and focusing. Any logging specification is parameterized by a sublanguage $S$ that identifies the predicate(s) to be resolved and Horn clauses $X$ that define it/them, hence we define a functional *spec* from pairs $(X, S)$ to specifications $LS$, where we use $\lambda$ as a binder for function definitions in the usual manner:

*Definition 10.* The function *spec* is given a pair $(X, S)$ and returns a *FOL logging specification*, i.e. a function from traces to elements of $\Phi_{FOL}$:

$$spec(X, S) = \lambda\tau.(\lfloor\tau\rfloor \otimes C(X))^{\Rightarrow S}.$$

In any logging specification $spec(X, S)$, we call $X$ the *guidelines*.

The above example $LS_H$ would then be formally defined as $spec(\psi_H, \{\mathrm{LoggedCall}\})$.

### *Relational Database.*

Relational algebra is a canonical example of an information algebra. Following standard practice [29], we define relations $X$ as sets of *tuples* $f$. We write $((a_1 : x_1), ..., (a_n : x_1))$ to denote an $n$-ary tuple with attributes (aka label) $a_i$ associated with values $x_i$. Relations are elements of the information algebra, and sublanguages $S$ are sets of attributes. Focusing is projection, i.e. we take $X^{\Rightarrow S} = \pi_S(X)$ in standard notation, and combination is natural join $\bowtie$. Hence, letting $\leq_{RA}$ denote the relational algebra information ordering, $X \leq_{RA} Y$ iff $X \bowtie Y = Y$. We refer to this algebra as $\Phi_{RA}$. In this context, a trace can be interpreted as a relation, logging specifications can be defined using selects, as has been done for the related problem of trace-based pointcuts for AspectJ [1]. Relational databases are also heavily used for storing and querying audit logs.

## 3.2 Transforming and Combining Audit Logs

Multiple audit logs from different sources are often combined in practice. Also, logging information is often transformed for storage and communication. For example, log data may be generated in common event format (CEF), which is parsed and stored in relational database tables, and subsequently exported and communicated via JSON In all cases, it is crucial to characterize the effect of transformation (if any) on log information, and relate queries on various representations to the logging specification semantics. Otherwise, it is unclear what is the relation of log queries to log-generating programs.

To address this, information algebra provides another useful concept called *monotone mapping*. Given two information algebras $\Psi_1$ and $\Psi_2$ with ordering relations $\leq_1$ and $\leq_2$ respectively, a mapping $\mu$ from elements $X, Y$ of $\Psi_1$ to elements $\mu(X), \mu(Y)$ of $\Psi_2$ is monotone iff $X \leq_1 Y$ implies $\mu(X) \leq_2 \mu(Y)$. For example, assuming that $\Psi_1$ is our FOL information algebra while $\Psi_2$ is relational algebra, we can define a monotone mapping using a *least Herbrand interpretation* [9], denoted $\mathfrak{H}$, and by positing a function $attrs$ from $n$-ary predicate symbols to functions mapping numbers $1, ..., n$ to labels. That is, $attrs(\mathrm{P})(n)$ is the label associated with the $n$th argument of predicate P. We require that if $\mathrm{P} \neq \mathrm{Q}$ then $attrs(\mathrm{P})(j) \neq attrs(\mathrm{Q})(k)$ for all $j, k$. To map predicates to tuples we have:

$$tuple(\mathrm{P}(x_1, \ldots, x_n)) = \\ ((attrs(\mathrm{P})(1) : x_1), \ldots, (attrs(\mathrm{P})(n) : x_n))$$

Then to obtain a relation from all valid instances of a particular predicate P given formulas $X$ we define:

$$R_{\mathrm{P}}(X) = \{tuple(\mathrm{P}(x_1, \ldots, x_n)) \mid \mathrm{P}(x_1, \ldots, x_n) \in \mathfrak{H}(X)\}$$

Now we define the function $rel$ which is the cartesian product[3] $\times$ of all the relations obtained from $X$, where $\mathrm{P}_1, ..., \mathrm{P}_n$ are the predicate symbols occurring in $X$:

$$rel(X) = R_{\mathrm{P}_1}(X) \times \cdots \times R_{\mathrm{P}_n}(X)$$

Since $\bowtie$ is equivalent to cartesian product over relations with disjoint attributes, the following result then holds by definition of $attrs$, $\Phi_{FOL}$ and $\Phi_{RA}$, and properties of least Herbrand models.

THEOREM 3. *$rel$ is a monotone mapping.*

Thus, if we wish to generate an audit log $\mathbb{L}$ as a set of FOL formulas, but ultimately store the data in a relational database, we are still able to maintain a formal relation between stored logs and the semantics of a given trace $\tau$ and specification $LS$. E.g., if a log $\mathbb{L}$ is sound with respect to $\tau$ and $LS$, then $rel(\lfloor\mathbb{L}\rfloor) \leq_{RA} rel(LS(\tau))$. While the data in $rel(\lfloor\mathbb{L}\rfloor)$ may very well be broken up into multiple relations $\mathbf{R}$ in practice, e.g. to compress data and/or for query optimization, the formalism also establishes correctness conditions for the transformation that relate resulting information to the logging semantics $LS(\tau)$ by way of the mapping.

# 4. REWRITING PROGRAMS WITH LOGGING SPECIFICATIONS

Since correct logging instrumentation is a safety property (2.2), there are various implementation strategies. For example, one could define an edit automata that enforces the property that could be implemented either as a separate program monitor or using IRM techniques [20]. But since we are interested in program rewriting for a particular class of logging specifications, the approach we discuss here is more simply stated and proven correct than a general IRM methodology.

We specify a class of logging specifications of interest, along with a program rewriting algorithm that is sound and complete for it. We consider a basic $\lambda$-calculus that serves as a prototypical case study. The supported class of logging specifications is predicated on temporal properties of function calls and characteristics of their arguments. This class has practical potential since security-sensitive operations are often packaged as functions or methods (e.g. in medical records software [36]), and the supported class allows complex policies such as break the glass to be expressed. The language of logging specifications is FOL, and we use $\Phi_{FOL}$ to

[3]Recall that $\times$ is associative and commutative in relational algebra.

define the semantics of logging and prove correctness of the algorithm.

## 4.1 Source Language

We first define a source language $\Lambda_{\text{call}}$, including the definitions of configurations, execution traces, and function $toFOL(\cdot)$ that shows how we concretely model execution traces in FOL.

Language $\Lambda_{\text{call}}$ is a simple call-by-value $\lambda$-calculus with named functions. A $\Lambda_{\text{call}}$ program is a pair $(e, \mathcal{C})$ where $e$ is an expression, and $\mathcal{C}$ is a *codebase* which maps function names to function definitions. A $\Lambda_{\text{call}}$ configuration is a triple $(e, n, \mathcal{C})$, where $e$ is the expression remaining to be evaluated, $n$ is a timestamp (a natural number) that indicates how many steps have been taken since program execution began, and $\mathcal{C}$ is a codebase. The codebase does not change during program execution.

The syntax of $\Lambda_{\text{call}}$ is as follows.

$$
\begin{array}{rr}
v ::= x \mid \mathbf{f} \mid \lambda x.\ e & \textit{values} \\
e ::= e\ e \mid v & \textit{expressions} \\
E ::= [\,] \mid E\ e \mid v\ E & \textit{evaluation contexts} \\
\kappa ::= (e, n, \mathcal{C}) & \textit{configurations} \\
\mathfrak{p} ::= (e, \mathcal{C}) & \textit{programs}
\end{array}
$$

The small-step semantics of $\Lambda_{\text{call}}$ is defined as follows.

$$\beta \quad \frac{}{((\lambda x.\ e)\ v, n, \mathcal{C}) \to (e[v/x], n+1, \mathcal{C})}$$

$$\beta_{\text{Call}} \quad \frac{\mathcal{C}(\mathbf{f}) = \lambda x.\ e}{(\mathbf{f}\ v, n, \mathcal{C}) \to (e[v/x], n+1, \mathcal{C})}$$

$$\text{Context} \quad \frac{(e, n, \mathcal{C}) \to (e', n', \mathcal{C})}{(E[e], n, \mathcal{C}) \to (E[e'], n', \mathcal{C})}$$

An execution trace $\tau$ is a sequence of configurations, and for a program $\mathfrak{p} = (e, \mathcal{C})$ and execution trace $\tau = \kappa_0 \ldots \kappa_n$ we define $\mathfrak{p} \Downarrow \tau$ if and only if $\kappa_0 = (e, 0, \mathcal{C})$ and for all $i \in 1..n$ we have $\kappa_{i-1} \to \kappa_i$.

We now show how to model a configuration as a set of ground instances of predicates, and then use this to model execution traces. We posit predicates Call, App, Value, Context, and Codebase to logically denote run time entities. For $\kappa = (e, n, \mathcal{C})$, we define $toFOL(\kappa)$ by cases, where

$$\langle \mathcal{C} \rangle_n = \bigcup_{\mathbf{f} \in \text{dom}(\mathcal{C})} \{\text{Codebase}(n, \mathbf{f}, \mathcal{C}(\mathbf{f}))\}^4.$$

$$
\begin{aligned}
&toFOL(v, n, \mathcal{C}) = \{\text{Value}(n, v)\} \cup \langle \mathcal{C} \rangle_n \\
&toFOL(E[\mathbf{f}\ v], n, \mathcal{C}) = \{\text{Call}(n, \mathbf{f}, v), \text{Context}(n, E)\} \cup \langle \mathcal{C} \rangle_n \\
&toFOL(E[(\lambda x.\ e)\ v)], n, \mathcal{C}) = \\
&\qquad \{\text{App}(n, (\lambda x.e), v), \text{Context}(n, E)\} \cup \langle \mathcal{C} \rangle_n
\end{aligned}
$$

We define $toFOL(\tau)$ for execution trace $\tau = \kappa_0 \ldots \kappa_n$ as follows.

$$toFOL(\kappa_0 \ldots \kappa_n) = toFOL(\kappa_0) \cup \cdots \cup toFOL(\kappa_n)$$

[4]While $\Lambda_{\text{call}}$ expressions and evaluation contexts appear as predicate arguments, their syntax can be written as string literals to conform to typical Datalog or Prolog syntax.

Function $toFOL(\cdot)$ is injective up to $\alpha$-equivalence since $toFOL(\tau)$ fully and uniquely describes the execution trace $\tau$.

## 4.2 Specifications Based on Function Call Properties

We define a class **Calls** of logging specifications that capture temporal properties of function calls, such as those reflected in break the glass policies. We restrict specification definitions to safe Horn clauses to ensure applicability of well-known results and total algorithms such as Datalog [9]. Specifications in **Calls** support logging of calls to a specific function $\mathbf{f}$ that happen after functions $\mathbf{g}_1, \ldots, \mathbf{g}_n$ are called. Conditions on all function arguments, and times of their invocation, can be defined via a predicate $\phi$. Hence more precise requirements can be imposed, e.g. a linear ordering on function calls, particular values of functions arguments, etc.

*Definition 11.* **Calls** is the set of all logging specifications, denoted by $spec(X, \{\text{LoggedCall}\})$, where $X$ contains a safe Horn clause of the following form:

$$
\begin{aligned}
&\forall t_0, \ldots, t_n, x_0, \ldots, x_n\ . \\
&\qquad \text{Call}(t_0, \mathbf{f}, x_0) \bigwedge_{i=1}^{n} (\text{Call}(t_i, \mathbf{g}_i, x_i) \wedge t_i < t_0) \wedge \\
&\qquad \phi((x_0, t_0), \ldots, (x_n, t_n)) \implies \text{LoggedCall}(t_0, \mathbf{f}, x_0).
\end{aligned}
$$

While set $X$ may contain other safe Horn clauses, in particular definitions of predicates occurring in $\phi$, no other Horn clause in $X$ uses the predicate symbols LoggedCall, Value, Context, Call, App, or Codebase. For convenience in the following, we define $Logevent(LS) = \mathbf{f}$ and $Triggers(LS) = \{\mathbf{g}_1, ..., \mathbf{g}_n\}$.

We note that specifications in **Calls** clearly satisfy Definition 1, since preconditions for logging a particular call to $\mathbf{f}$ must be satisfied at the time of that call.

## 4.3 Target Language

The syntax of target language $\Lambda_{\text{log}}$ extends $\Lambda_{\text{call}}$ syntax with a command to track relevant function calls ($callEvent(\mathbf{f}, v)$) and a command to emit log entries ($emit(\mathbf{f}, v)$). Configurations are extended to include a set $X$ of relevant function calls, and an audit log $\mathbb{L}$.

$$
\begin{array}{rr}
e ::= \ldots \mid callEvent(\mathbf{f}, v); e \mid emit(\mathbf{f}, v); e & \textit{expressions} \\
\kappa ::= (e, X, n, \mathbb{L}, \mathcal{C}) & \textit{configurations}
\end{array}
$$

The semantics of $\Lambda_{\text{log}}$ extends the semantics of $\Lambda_{\text{call}}$ with new rules for commands $callEvent(\mathbf{f}, v)$ and $emit(\mathbf{f}, v)$, which update the set of relevant function calls and audit log respectively. An instrumented program uses the set of relevant function calls to determine when it should emit events to the audit log. The semantics is parameterized by a guideline $X_{Guidelines}$, typically taken from a logging specification. Given the definition of **Calls**, these seman-

tics would be easy to implement using e.g. a Datalog proof engine.

RelevantCall

$$\frac{}{(callEvent(\mathbf{f},v);e,X,n,\mathbb{L},\mathcal{C}) \rightarrow (e,X \cup \{\text{Call}(n-1,\mathbf{f},v)\},n,\mathbb{L},\mathcal{C})}$$

Log

$$\frac{X \cup X_{Guidelines} \vdash \text{LoggedCall}(n-1,\mathbf{f},v)}{(emit(\mathbf{f},v);e,X,n,\mathbb{L},\mathcal{C}) \rightarrow (e,X,n,\mathbb{L} \cup \{\text{LoggedCall}(n-1,\mathbf{f},v)\},\mathcal{C})}$$

NoLog

$$\frac{X \cup X_{Guidelines} \nvdash \text{LoggedCall}(n-1,\mathbf{f},v)}{(emit(\mathbf{f},v);e,X,n,\mathbb{L},\mathcal{C}) \rightarrow (e,X,n,\mathbb{L},\mathcal{C})}$$

Note that to ensure that these instrumentation commands do not change execution behavior, the configuration's time is not incremented when $callEvent(\mathbf{f},v)$ and $emit(\mathbf{f},v)$ are evaluated. That is, the configuration time counts the number of source language computation steps.

The rules Log and NoLog rely on checking whether $X_{Guidelines}$ and relevant function calls $X$ entail $\text{LoggedCall}(n-1,\mathbf{f},v)$. This can be accomplished using off-the-shelf theorem provers for Horn clause logics, such as Datalog or Prolog.

For a target language program $\mathfrak{p} = (e,\mathcal{C})$ and execution trace $\tau = \kappa_0 \ldots \kappa_n$ we define $\mathfrak{p} \Downarrow \tau$ if and only if $\kappa_0 = (e,\emptyset,0,\emptyset,\mathcal{C})$ and for all $i \in 1..n$ we have $\kappa_{i-1} \rightarrow \kappa_i$.

To establish correctness of program rewriting, we need to define a correspondence relation $:\approx$. Source language execution traces and target language execution traces correspond if they represent the same expression evaluated to the same point. We make special cases for when the source execution is about to perform a function application that the target execution will track or log via a $callEvent(\mathbf{f},v)$ or $emit(\mathbf{f},v)$ command. In these cases, the target execution may be ahead by one or two steps, allowing time for addition of information to the log.

*Definition 12.* Given the source language execution trace $\tau = \kappa_0 \ldots \kappa_m$ and the target language execution trace $\tau' = \kappa'_0 \ldots \kappa'_n$, where $\kappa_i = (e_i,t_i,\mathcal{C}_i)$ and $\kappa'_i = (e'_i,X_i,t'_i,\mathbb{L}_i,\mathcal{C}'_i)$, $\tau :\approx \tau'$ iff $e_0 = e'_0$ and either

1. $e_m = e'_n$ (taking $=$ to mean syntactic equivalence); or
2. $e_m = e'_{n-1}$ and $e'_n = callEvent(\mathbf{f},v);e'$ for some expressions $\mathbf{f}$, $v$, and $e'$; or
3. $e_m = e'_{n-2}$ and $e'_n = emit(\mathbf{f},v);e'$ for some expressions $\mathbf{f}$, $v$, and $e'$.

Finally, we need to define $toFOL(\mathbb{L})$ for audit logs $\mathbb{L}$ produced by an instrumented program. Since our audit logs are just sets of formulas of the form $\text{LoggedCall}(t,\mathbf{f},v)$, we define $toFOL(\mathbb{L}) = \mathbb{L}$.

## 4.4 Program Rewriting Algorithm

Our program rewriting algorithm $\mathcal{R}_{\Lambda_{\text{call}}}$ takes a $\Lambda_{\text{call}}$ program $\mathfrak{p} = (e,\mathcal{C})$, a logging specification $LS \in \mathbf{Calls}$ defined as $LS = spec(X_{Guidelines},\{\text{LoggedCall}\})$, and produces a $\Lambda_{\log}$ program $\mathfrak{p}' = (e',\mathcal{C}')$ such that $e$ and $e'$ are identical, and $\mathcal{C}'$ is identical to $\mathcal{C}$ except for the addition of $callEvent(\mathbf{h},v)$ and $emit(\mathbf{h},v)$ commands. The algorithm is straightforward: we modify the codebase to add $callEvent(\mathbf{h},v)$ to the definition of any function $\mathbf{h} \in Triggers(LS) \cup \{Logevent(LS)\}$ and add $emit(\mathbf{f},v)$ to the definition of function $\mathbf{f} = Logevent(LS)$.

*Definition 13.* For $\Lambda_{\text{call}}$ program $\mathfrak{p} = (e,\mathcal{C})$ and logging specifications $LS \in \mathbf{Calls}$, define:

$$\mathcal{R}_{\Lambda_{\text{call}}}((e,\mathcal{C}),LS) = (e,\mathcal{C}')$$

where $\mathcal{C}'(\mathbf{f}) =$

$$\begin{cases} \lambda x.callEvent(\mathbf{f},x);emit(\mathbf{f},x);e_{\mathbf{f}} & \text{if } \mathbf{f} = Logevent(LS) \text{ and } \mathcal{C}(\mathbf{f}) = \lambda x.e_{\mathbf{f}} \\ \lambda x.callEvent(\mathbf{f},x);e_{\mathbf{f}} & \text{if } \mathbf{f} \in Triggers(LS) \text{ and } \mathcal{C}(\mathbf{f}) = \lambda x.e_{\mathbf{f}} \\ \mathcal{C}(\mathbf{f}) & \text{otherwise} \end{cases}$$

This algorithm obeys the required properties, i.e. it is both semantics preserving and sound and complete for a given logging specification.

THEOREM 4. *Program rewriting algorithm $\mathcal{R}_{\Lambda_{\text{call}}}$ is semantics preserving (Definition 5).*

THEOREM 5 (SOUNDNESS AND COMPLETENESS). *Program rewriting algorithm $\mathcal{R}_{\Lambda_{\text{call}}}$ is sound and complete (Definitions 6 and 7).*

# 5. RELATED WORK

Previous work by DeYoung et al. has studied audit policy specification for medical (HIPAA) and business (GLBA) processes [18, 19, 17]. This work illustrates the effectiveness and generality of a temporal logic foundation for audit policy specification, which is well-founded in a general theory of privacy [16]. Their auditing system has also been implemented in a tool similar to an interactive theorem prover [24]. Their specification language inspired our approach to logging specification semantics. However, this previous work assumes that audit logs are given, and does not consider the correctness of logs. Some work does consider trustworthiness of logs [5], but only in terms of tampering (malleability). In contrast, our work provides formal foundations for the correctness of audit logs, and considers algorithms to automatically instrument programs to generate correct logs.

Other work applies formal methods (including predicate logics [14, 8], process calculi and game theory [26]) to model, specify, and enforce auditing and accountability requirements in distributed systems. In that work, audit logs serve as evidence of resource access rights, an idea also explored in Aura [38] and the APPLE system [21]. In Aura, audit logs record machine-checkable proofs of compliance in the Aura policy language. APPLE proposes a framework based on trust management and audit logic with log generation functionality for a limited set of operations, in order to check user compliance.

In contrast, we provide a formal foundation to support a broad class of logging specifications and relevant correctness conditions. In this respect our proposed system is closely related to PQL [32], which supports program rewriting with instrumentation to answer queries about program execution. From a technical perspective, our approach is also related to trace matching in AspectJ [1], especially in the use of logic to specify trace patterns. However, the concern in that work is aspect pointcut specification, not logging correctness, and their method call patterns are restricted to be regular expressions with no conditions on arguments, whereas the latter is needed for the specifications in **Calls**.

Logging specifications are related to safety properties [37] and are enforceable by security automata, as we have shown. Hence

IRM rewriting techniques could be used to implement them [20]. However, the theory of safety properties does not address correctness of audit logs as we do, and our approach can be viewed as a logging-specific IRM strategy.

Guts et al. [25] develop a static technique to guarantee that programs are properly instrumented to generate audit logs with sufficient evidence for auditing purposes. As in our research, this is accomplished by first defining a formal semantics of auditing. However, they are interested in evidence-based auditing for specific distributed protocols.

Other recent work [22] has proposed log filters as a required improvement to the current logging practices in the industry due to costly resource consumption and the loss of necessary log information among the collected redundant data. This work is purely empirical, not foundational, but provides practical evidence of the relevance of our efforts since logging filters could be defined as logging specifications.

Audit logs can be considered a form of *provenance*: the history of computation and data. Several recent works have considered formal semantics of provenance [7, 6]. Cheney [10] presents a framework for provenance, built on a notion of system traces. Recently, W3C has proposed a data model for provenance, called PROV [3], which enjoys a formal description of its specified constraints and inferences in first-order logic, [11], however the given semantics does not cover the relationship between the provenance record and the actual system behavior.

# 6. CONCLUSION

In this paper we have addressed the problem of audit log correctness. In particular, we have considered how to separate logging specifications from implementations, and how to formally establish that an implementation satisfies a specification. This separation allows security administrators to clearly define logging goals independently from programs, and inspires program rewriting tools that support correct, automatic instrumentation of logging specifications in legacy code.

By leveraging the theory of information algebra, we have defined a semantics of logging specifications as functions from program traces to information. By interpreting audit logs as information, we are then able to establish correctness conditions for audit logs via an information containment relation between log information and logging specification semantics. These conditions allow proof of correctness of program rewriting algorithms that automatically instrument general classes of logging specifications. To demonstrate, we define a prototype rewriting algorithm for a functional calculus that instruments a class of logging specifications defined in first order logic, and prove the algorithm correct.

# References


[1] Allan, C., Avgustinov, P., Christensen, A.S., Hendren, L.J., Kuzins, S., Lhoták, O., de Moor, O., Sereni, D., Sittampalam, G., Tibble, J.: Adding trace matching with free variables to AspectJ. In: Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2005, October 16-20, 2005, San Diego, CA, USA. pp. 345–364 (2005)

[2] Bauer, L., Ligatti, J., Walker, D.: More enforceable security policies. Tech. Rep. TR-649-02, Princeton University (2002)

[3] Belhajjame, K., B'Far, R., Cheney, J., Coppens, S., Cresswell, S., Gil, Y., Groth, P., Klyne, G., Lebo, T., McCusker, J., Miles, S., Myers, J., Sahoo, S., Tilmes, C.: PROV-DM: The PROV data model. http://www.w3.org/TR/2013/REC-prov-dm-20130430 (2013), accessed: 2015-02-07

[4] Biswas, D., Niemi, V.: Transforming privacy policies to auditing specifications. In: 13th IEEE International Symposium on High-Assurance Systems Engineering, HASE 2011, Boca Raton, FL, USA, November 10-12, 2011. pp. 368–375 (2011)

[5] Böck, B., Huemer, D., Tjoa, A.M.: Towards more trustable log files for digital forensics by means of "trusted computing". In: Proceedings of the 24th IEEE International Conference on Advanced Information Networking and Applications. pp. 1020–1027. AINA '10, IEEE Computer Society, Washington, DC, USA (2010)

[6] Buneman, P., Chapman, A., Cheney, J.: Provenance management in curated databases. In: Proceedings of the ACM SIGMOD International Conference on Management of Data. pp. 539 – 550 (2006)

[7] Buneman, P., Khanna, S., Tan, W.C.: Why and where: A characterization of data provenance. Lecture Notes in Mathematics - Springer Verlag pp. 316–330 (2000)

[8] Cederquist, J.G., Corin, R., Dekker, M.A.C., Etalle, S., den Hartog, J.I., Lenzini, G.: Audit-based compliance control. International Journal of Information Security 6(2-3), 133–151 (2007)

[9] Ceri, S., Gottlob, G., Tanca, L.: What you always wanted to know about Datalog (And never dared to ask). IEEE Transactions on Knowledge and Data Engineering 1(1), 146–166 (1989)

[10] Cheney, J.: A formal framework for provenance security. In: Proceedings of the 24th IEEE Computer Security Foundations Symposium. pp. 281–293 (2011)

[11] Cheney, J.: Semantics of the PROV data model. http://www.w3.org/TR/2013/NOTE-prov-sem-20130430 (2013), accessed: 2015-02-07

[12] Chuvakin, A.: Beautiful log handling. In: Oram, A., Viega, J. (eds.) Beautiful security: Leading security experts explain how they think. O'Reilly Media Inc. (2009)

[13] Cook, D., Hartnett, J., Manderson, K., Scanlan, J.: Catching spam before it arrives: Domain specific dynamic blacklists. In: Proceedings of the Fourth Australasian Symposium on Grid Computing and e-Research (AusGrid 2006) and the Fourth Australasian Information Security Workshop. pp. 193–202. Australian Computer Society, Inc. (2006)

[14] Corin, R., Etalle, S., den Hartog, J.I., Lenzini, G., Staicu, I.: A logic for auditing accountability in decentralized systems. In: Formal Aspects in Security and Trust. pp. 187–201 (2004)

[15] CPMC Press Release: Audit finds employee access to patient files without apparent business or treatment purpose. http://www.cpmc.org/about/press/News2015/phi.html (2015), accessed: 2015-01-30

[16] Datta, A., Blocki, J., Christin, N., DeYoung, H., Garg, D., Jia, L., Kaynar, D.K., Sinha, A.: Understanding and protecting privacy: Formal semantics and principled audit mechanisms. In: Proceedings of the 7th International Conference on Information Systems Security. pp. 1–27 (2011)

[17] DeYoung, H., Garg, D., Jia, L., Kaynar, D., Datta, A.: Privacy policy specification and audit in a fixed-point logic: How to enforce HIPAA, GLBA, and all that. Tech. Rep. CMU-CyLab-10-008, Carnegie Mellon University (April 2010)

[18] DeYoung, H., Garg, D., Jia, L., Kaynar, D.K., Datta, A.: Experiences in the logical specification of the HIPAA and GLBA privacy laws. In: Proceedings of the 2010 ACM Workshop on Privacy in the Electronic Society. pp. 73–82 (2010)

[19] DeYoung, H., Garg, D., Kaynar, D., Datta, A.: Logical specification of the GLBA and HIPAA privacy laws. Tech. Rep. CMU-CyLab-10-007, Carnegie Mellon University (April 2010)

[20] Erlingsson, Ú.: The inlined reference monitor approach to security policy enforcement. Ph.D. thesis, Cornell University (2003)

[21] Etalle, S., Winsborough, W.H.: A posteriori compliance control. In: Proceedings of the 12th ACM Symposium on Access Control Models and Technologies. pp. 11–20 (2007)

[22] Fu, Q., Zhu, J., Hu, W., Lou, J., Ding, R., Lin, Q., Zhang, D., Xie, T.: Where do developers log? An empirical study on logging practices in industry. In: Proceedings of the 36th International Conference on Software Engineering. pp. 24–33 (2014)

[23] Ganapathy, V., Jaeger, T., Skalka, C., Tan, G.: Assurance for defense in depth via retrofitting. In: Layered Assurance Workshop (2014)

[24] Garg, D., Jia, L., Datta, A.: Policy auditing over incomplete logs: Theory, implementation and applications. In: Proceedings of the 18th ACM Conference on Computer and Communications Security. pp. 151–162 (2011)

[25] Guts, N., Fournet, C., Nardelli, F.Z.: Reliable evidence: Auditability by typing. In: Proceedings of the 14th European Conference on Research in Computer Security. pp. 168–183. ESORICS'09, Springer-Verlag, Berlin, Heidelberg (2009)

[26] Jagadeesan, R., Jeffrey, A., Pitcher, C., Riely, J.: Towards a theory of accountability and audit. In: Proceedings of the 14th European Symposium on Research in Computer Security. pp. 152–167 (2009)

[27] Kemmerer, R.A., Vigna, G.: Intrusion detection: A brief history and overview. Computer 35(4), 27–30 (2002)

[28] King, J.T., Smith, B., Williams, L.: Modifying without a trace: General audit guidelines are inadequate for open-source electronic health record audit mechanisms. In: Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium. pp. 305–314. ACM (2012)

[29] Kohlas, J.: Information Algebras: Generic Structures For Inference. Discrete mathematics and theoretical computer science, Springer (2003)

[30] Kohlas, J., Schmid, J.: An algebraic theory of information: An introduction and survey. Information 5(2), 219–254 (2014)

[31] Lampson, B.W.: Computer security in the real world. IEEE Computer 37(6), 37–46 (2004)

[32] Martin, M., Livshits, B., Lam, M.S.: Finding application errors and security flaws using PQL: A program query language. In: Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications. pp. 365–383. ACM (2005)

[33] Matthews, P., Gaebel, H.: Break the glass. In: HIE Topic Series. Healthcare Information and Management Systems Society (2009), http://www.himss.org/files/himssorg/content/files/090909breaktheglass.pdf

[34] OpenMRS Wiki: Usage statistics module. https://wiki.openmrs.org/display/docs/Usage+Statistics+Module (2010), accessed: 2015-06-15

[35] Povey, D.: Optimistic security: A new access control paradigm. In: Proceedings of the 1999 Workshop on New Security Paradigms. pp. 40–45 (1999)

[36] Rizvi, S.Z., Fong, P.W.L., Crampton, J., Sellwood, J.: Relationship-based access control for an open-source medical records system. In: ACM Symposium on Access Control Models and Technologies (2015)

[37] Schneider, F.B.: Enforceable security policies. ACM Transactions on Information and System Security 3(1), 30–50 (2000)

[38] Vaughan, J.A., Jia, L., Mazurak, K., Zdancewic, S.: Evidence-based audit. In: Proceedings of the 21st IEEE Computer Security Foundations Symposium. pp. 177–191 (2008)

[39] Weitzner, D.J.: Beyond secrecy: New privacy protection strategies for open information spaces. IEEE Internet Computing 11(5), 94–96 (2007)

[40] Weitzner, D.J., Abelson, H., Berners-Lee, T., Feigenbaum, J., Hendler, J.A., Sussman, G.J.: Information accountability. Communications of the ACM 51(6), 82–87 (2008)

[41] Zhang, W., Chen, Y., Cybulski, T., Fabbri, D., Gunter, C.A., Lawlor, P., Liebovitz, D.M., Malin, B.: Decide now or decide later? Quantifying the tradeoff between prospective and retrospective access decisions. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security. pp. 1182–1192 (2014)

[42] Zheng, A.X., Jordan, M.I., Liblit, B., Naik, M., Aiken, A.: Statistical debugging: Simultaneous identification of multiple bugs. In: Proceedings of the 23rd International Conference on Machine Learning. pp. 1105–1112. ICML '06, ACM, New York, NY, USA (2006), http://doi.acm.org/10.1145/1143844.1143983