June 2005

Practical Information-flow Control in Web-based Information Systems

Peng Li

University of Pennsylvania

Stephan A. Zdancewic

University of Pennsylvania, stevez@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_papers

Recommended Citation

Peng Li and Stephan A. Zdancewic, "Practical Information-flow Control in Web-based Information Systems", . June 2005.

Copyright 2005 IEEE. Reprinted from *Proceedings of the 18th IEEE Workshop on Computer Security Foundations 2005 (CSFW-18 2005)*, pages 2-15.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_papers/221
For more information, please contact repository@pobox.upenn.edu.

Practical Information-flow Control in Web-based Information Systems

Abstract

This paper presents a practical application of language-based information-flow control, namely, a domain-specific web scripting language designed for interfacing with databases. The primary goal is to provide strong enforcement of confidentiality and integrity policies: confidential data can be released only in permitted ways and trustworthy data must result from expected computations or conform to expected patterns. Such security policies are specified in the database layer and statically enforced for the rest of the system in an end-to-end fashion.

In contrast with existing web-scripting languages, which provide only ad hoc mechanisms for information security, the scripting language described here uses principles based on the well-studied techniques in information-flow type systems. However, because web scripts often need to downgrade confidential data and manipulated untrusted user input, they require practical and convenient ways of downgrading secure data. To achieve this goal, the language allows safe downgrading according to *downgrading policies* specified by the programmer. This novel, pattern-based approach provides a practical instance of recent work on *delimited release* and *relaxed noninterference* and extends that work by accounting for integrity policies.

Comments

Copyright 2005 IEEE. Reprinted from *Proceedings of the 18th IEEE Workshop on Computer Security Foundations 2005 (CSFW-18 2005)*, pages 2-15.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Practical Information-flow Control in Web-based Information Systems

Peng Li
University of Pennsylvania
lipeng@cis.upenn.edu

Steve Zdancewic
University of Pennsylvania
stevez@cis.upenn.edu

Abstract

This paper presents a practical application of language-based information-flow control, namely, a domain-specific web scripting language designed for interfacing with databases. The primary goal is to provide strong enforcement of confidentiality and integrity policies: confidential data can be released only in permitted ways and trustworthy data must result from expected computations or conform to expected patterns. Such security policies are specified in the database layer and statically enforced for the rest of the system in an end-to-end fashion.

In contrast with existing web-scripting languages, which provide only ad hoc mechanisms for information security, the scripting language described here uses principles based on the well-studied techniques in information-flow type systems. However, because web scripts often need to downgrade confidential data and manipulated untrusted user input, they require practical and convenient ways of downgrading secure data. To achieve this goal, the language allows safe downgrading according to downgrading policies specified by the programmer. This novel, pattern-based approach provides a practical instance of recent work on delimited release and relaxed noninterference and extends that work by accounting for integrity policies.

1. Introduction

This paper presents a language-based approach to enforcing confidentiality and integrity of data in typical web-based information systems. Such systems are usually implemented using a layered design in which data is stored in the database and accessed by using a web browser. The database management system (DBMS) provides a data query language (for example, SQL) to store, modify and extract information from databases. Application software connects to the DBMS via some programming language interface and submits *queries* as requests for information from the DBMS. The application software then processes the

data and send the computation results to the end-users, typically bundled as HTML. This paper focuses on a simple yet widely used design, where relational databases such as MySQL are used as the DBMS and web scripting languages such as PHP are used for developing application software.

In practice, there are many security concerns for such systems. For example, an unsafe PHP script could use strings from untrusted inputs to compose SQL queries and then have the DBMS execute the query, which potentially allows an attacker to insert arbitrary commands in the SQL query. PHP scripts that access confidential data in the database must also release them only in permitted ways. For example, one might require that a password can be compared against user inputs but cannot be printed to the web page verbatim, or that only the last several digits of a social security number or a credit card number should be displayed in the HTML output.

As these web applications become complex, the security of the system becomes hard to manage. In the worst case, the programmer must walk through all the code and check every line to make sure that there are no security violations. In addition, the queries performed by the script must comply with the desired policies on the data stored in the database. Ensuring that all of the security requirements are met is difficult to do manually, tedious, and error-prone.

Currently, there exist only ad hoc ways to (partially) enforce such security policies. To prevent accidental use of untrusted inputs as parameters to safety-critical operations—thereby preventing format string attacks and malicious DBMS queries—some scripting languages provide mechanisms to track the uses of untrusted inputs, dynamically checking that they are not used inappropriately. Perl, for instance, can be run in a “taint checking” mode in which user input strings are considered tainted until they are matched against a programmer-supplied pattern, which establishes that the untrusted input actually conforms to an expected form.

In an effort to enforce the confidentiality policies on the database, experienced software developers implement most of the security-sensitive operations in the DBMS as procedure calls in the query language, exposing explicit, restricted interfaces. Programmers should use only these in-

terfaces to perform queries in the web scripting language. This limited interface means that the database query engine can itself enforce the desired policies. Such encapsulation means that the programmer needs to check at least two things: (1) that the sensitive operations are correctly implemented in the DBMS and, (2) that the query interfaces are properly used by the web scripts. These checks must be performed manually.

The above ad hoc approaches have several drawbacks. First, they encourage that a significant portion of the business logic be implemented in the DBMS, which makes the interfaces less modular and less reusable. For example, we may require that only the last 4 digits of the credit card number can be displayed to the user. If this limitation is implemented in a procedure call in the DBMS, the web script can only get the last 4 digits of the number using that call. Security is guaranteed, but the web script cannot use the full number to perform further queries to the database. Instead, a dedicated procedure for this web script must be written in the database that connect these queries together. As the business logic becomes complex, the procedures in the DBMS become hard to manage.

Second, doing the dynamic checks may be inefficient. For example, the user inputs in web forms usually have constraints on them that restrict their data ranges and data types. Such constraints are often dynamically checked multiple times by the web scripts, the DBMS procedures, and the triggers in the databases. The reason is that the programming interface between the DBMS and the application software is either untyped or dynamically typed, and many constraints cannot be easily expressed as data types.

Third, the intended security policy of the system is not apparent from its implementation. There is no explicit description of what data is considered to be confidential or what the requirements are for checking the validity of untrusted inputs. This makes the software more fragile (local changes to the system can be inconsistent with the desired global security policy) and much harder to maintain over time.

1.1. Contributions

This paper proposes a language-based solution to the above problems. Instead of implementing all the access control mechanisms in the DBMS procedures or dynamic checks, we allow the programmer to specify security policies on the application programming interfaces of the DBMS. Such interfaces are strongly typed. The security policies are statically enforced in the scripting language using an information-flow type system.

Following the ideas of recent language prototypes such as Jif [7], which extends Java, and FlowCaml [14, 9, 10], which extends Caml, we design a security-typed lan-

guage suitable for writing web scripts. Web scripting languages differ from general-purpose programming languages in many aspects, several of which simplify our information-flow analyses. For example, most web scripts contain little or no state, and very limited looping constructs; they involve little computation and are intended to terminate quickly. These domain-specific features allow us to deal with covert channels (such as the timing channels) and some side effects more easily than in general purpose languages.

Importantly, when we reason about the information flows in such web scripts, downgrading is very common. Many web scripts read sensitive data from databases and release them to the end user (i.e. declassification of confidential data). Conversely, web scripts also take untrusted user inputs and use them to synthesize database queries which can alter trusted data (i.e. endorsement of low integrity data). To make information-flow control effective, downgrading must be controlled in a safe manner.

Our contributions include several important extensions to prior research [5, 12]. First, we present a simple yet general architecture for building secure web systems; this paper emphasizes the scripting-language component of that architecture. Second, our language addresses the problem of downgrading in information-flow systems by providing a practical instance of our recent theoretical framework on *downgrading policies* [5]. Third, we extend the framework to include integrity policies and conditional downgrading policies that use a novel pattern-matching sublanguage to express information-flow constraints. Conditional downgrading policies are extremely useful for specifying requirements that involve run-time identity tests.

The rest of the paper is organized as the following. Section 2 discusses a layered architecture for information flow control in online information systems. Section 3 presents a language of security policies, defines the security levels and formalize downgrading. Section 4 presents the abstract syntax of the web scripting language, shows a program example and introduces its type system. Section 5 discusses issues with untrusted code and related work. Section 6 concludes.

2. Mandatory access control in online information systems

2.1. System architecture

Typical web-based systems consist of a database, a database query language and various components such as scripts and web servers that work together in a multi-tiered fashion. These systems can have massive amounts of confidential and trusted information, with quite complex security policies. For example, recent laws in the United States man-

Figure 1: Information-flow control in web-based information systems

date complicated procedures for releasing medical information about patients; such policies must be enforced in the information management systems used by health care offices. Given the ubiquity and familiarity of the web, providing a web interface to such information would be valuable, but enforcing the appropriate security policies is a challenging task.

We propose a simplified, yet general architecture for such web-based information systems, as depicted in Figure 1. The key elements are:

- The storage layer: a database management system that stores and manages confidential and trustworthy data using database schemas, tables, views and a database query language.
- The access control layer: a query interface language that controls the release and update of information in the database. A query interface provides means to access the database, but also specifies language-level information-flow policies on the interface.
- The presentation layer: a web scripting language that executes the queries, manipulate the results and generate web pages for end users. The information flow policies from the query interfaces are enforced in the web scripting language.

It is desirable to use security-typed languages to program these systems, because language-based information flow control provides strong end-to-end security guarantees [13]: information-flow policies are specified at the endpoints of the program (on variables and module interfaces) and enforced in the whole program. However, in these practical systems, confidential data lives outside the programs that manipulate them. The source of such data is not program variables or function call interfaces, but rather the databases themselves. To enforce information flow policies across the whole system, it is necessary to specify the policies directly at the very end of the system, namely, at the database level, and enforce these information flow policies across the boundaries of different components.

The following sections outline how the pieces of this architecture fit together; subsequent sections explain the scripting language in more detail.

2.2. Policy specification in the DBMS interface

We require that all database queries are accessed via strongly typed programming interfaces provided by the DBMS. This is shown in Figure 1 as the access control layer, which, for this example contains two queries `GetID` and `FetchRecords`. The application programmer can use only the declared interfaces to access the database but cannot execute arbitrary SQL queries. As described in the introduction, this is a standard engineering practice, as it man-

dates that all the security-related queries be implemented as procedure calls in the DBMS layer. However, because our scripting language types include security policy information, this layer does not have to *enforce* all the security policies in the DBMS. Instead, it can *specify* security policies as types on the output of the queries and let the application software enforce these policies using language-based information-flow control. For example, the programmer can declare the output of a query to have security level “secret”, so the web scripts using this query will not leak its information to public places.

One issue is how to propagate the typing information from databases queries to the scripting language. In relational databases, a query returns a table that has a fixed number of columns. The type of each column is statically known. The script reads the query results iteratively in a row-by-row fashion, so the schema of a table can be translated to a tuple type that represents the data types for each column. As shown below, our scripting language provides built-in constructs for accessing rows of query results.

In our architecture, information-flow policies are specified on the query interface types. For confidentiality, we can annotate the output as *secret* so that it is not allowed to be leaked to public places. For integrity, we can annotate the argument of a query as *untainted* so that its value is not affected by tainted inputs. Besides these built-in levels, programmers can specify *downgrading policies* as security levels that describe how secret values can become public and how untrusted values become trusted. Section 3 presents a generalized framework of these information flow policies.

2.3. Policy enforcement in the web scripting language

This remainder of this paper presents the design of a web scripting language similar to PHP, which is used for the presentation layer shown in Figure 1. The major difference with PHP is that the queries to databases are strongly typed. The arguments and results of queries are annotated with information-flow policies. The interface to each query is explicitly defined in the program, which makes auditing the software for security and maintenance purposes easier for humans and provides the information needed by our type checker to determine whether the web script satisfies the policies. At run-time, when the script connects to the DBMS, the type signature of each query interface is matched against the interface provided in the DBMS.

At a high level, a web script takes some input data provided by interaction with the user and the DBMS and produces HTML data as output. The input data coming from the client web browser can be treated as having the security labels *public* and *tainted*. Results obtained from database queries have security policy specifications defined

Labels	$l ::= cl \mid il$
Confidentiality Labels	$cl ::= \{cp_1, \dots, cp_k\} \mid \text{secret} \mid \text{public}$
Integrity Labels	$il ::= \{ip_1, \dots, ip_k\} \mid \text{tainted} \mid \text{untainted}$
Confidentiality Policies	$cp ::= c \mid s \mid \star \mid \oplus cp_1, \dots, cp_k \mid \text{if } cp \text{ } cp \text{ } cp \mid \text{this}$
Integrity Policies	$ip ::= c \mid s \mid \star \mid \oplus ip_1, \dots, ip_k \mid \text{if } ip \text{ } ip \text{ } ip$
Trusted Variables	$s ::= s_i$
Constants	$c ::= c_i$
Operators	$\oplus ::= + \mid - \mid \text{hash} \mid \text{built-in operators...}$
Downgrading Actions	$a ::= c \mid s \mid \star \mid \oplus a_1, \dots, a_k \mid \text{if } a \text{ } a \text{ } a \mid \text{this}$

Figure 2: The pattern language for security labels

by the query interfaces. The web script performs some computation using these inputs from different security levels and produce strings that constitute the resulting web page, which has security level public. Additionally, the query arguments sent along with the query are also outputs and they may have required security levels such as untainted. In our model, the security policies on the outputs of the script are enforced by an information-flow type system, which provides end-to-end security guarantees on the input–output relationship of the program. The following sections present our scripting language features in detail.

3. Security levels for confidentiality and integrity

In many conventional information-flow type systems, the security levels considered are limited to the simple lattices $\text{public} \sqsubseteq \text{secret}$ for confidentiality and $\text{untainted} \sqsubseteq \text{tainted}$ for integrity. Our scripting language interprets these security levels in a more general framework, where *downgrading policies* [5] are used to express the security levels of data. A security level is simply a non-empty set of downgrading policies, where each policy describes the necessary computation that must be performed in order to downgrade the data at this security level. Such policies are expressed in a small pattern-matching language. Confidentiality policies and integrity policies are mostly symmetric in our language. The following subsections present the syntax and the semantics of these security levels.

3.1. Input variables and the pattern language

Each program takes some input data and produce some outputs. We classify the program inputs into two kinds:

- **Untrusted input:** inputs with security level tainted and public. Untrusted inputs come from the user inputs via web forms.

- **Trusted input:** inputs with security level untainted and secret, or other predefined security levels. Trusted inputs come from database query interfaces and run-time API calls. In our scripting language, such inputs are represented as *trusted variables*. For example, the root password can be modeled as a secret and untainted input to the login process. We use the metavariable s to represent such trusted input variables.

The syntax of security labels is defined in Figure 2. Each label has its confidentiality part and integrity part. Each confidentiality (or integrity) label is a non-empty set of policies. Special names such as *secret*, *public*, *untainted* and *tainted* can also be used as labels; they can be interpreted as sets of policies as shown in Figure 3. The next two subsections explain the semantics of these labels.

<i>secret</i>	$\equiv$	$\{c_0\}$	<i>public</i>	$\equiv$	$\{\text{this}\}$
<i>tainted</i>	$\equiv$	$\{\star\}$	<i>untainted</i>	$\equiv$	$\{ip \mid \star \notin ip\}$

Figure 3: Special security Labels

3.2. Confidentiality labels

Syntactically, a **confidentiality policy** is an expression embedded with *this* and $\star$. The meaning of such a policy is some computation required to declassify (or downgrade) the secret. When *this* is replaced by the annotated secret value and every $\star$ is replaced by a value at security level public, the result of the expression is at public. For example, if the password p has a policy “ $\text{this} = \star$ ” and x is a public value, then the expression $p=x$ is a public value. This policy allows the password be compared to untrusted values, but does not allow the password be leaked by other means.

Policy equivalence: We write $cp_1 \equiv cp_2$ to denote that the two policies have the same semantic meaning. For example, $x + y \equiv y + x$. To keep things simple in this paper, we use only syntactical equivalence.

A **confidentiality label** is a non-empty set of confidentiality policies. Such a label describes a security level for confidentiality. If data x has a confidentiality label $cl = \{cp_1, \dots, cp_k\}$, it means x can only be downgraded to public by using one of the policies cp_i in that label. In other words, a label specifies possible ways (policies) to downgrade the data it annotates. For example, the label $\{\text{this} = \star, \text{hash}(\text{this})\}$ on a password value also allows the password be leaked by computing its hash. We define the **interpretation** of a label cl to be an infinite set of policies:

$$\mathbb{S}(cl) \triangleq \{cp_1 \mid cp_1 \equiv [cp_2/\text{this}]cp_3, cp_2 \in cl, s_i \notin cp_3\}$$

Intuitively, it means that if cp_2 is a valid downgrading policy, i.e. cp_2 represents a public expression, we can put cp_2 in another context cp_3 to yield another downgrading policy cp_1 , as long as the context cp_3 does not contain trusted secret variables (which we denote as $s_i \notin cp_3$). The interpretation of a label is the set of all policies that can be derived from the policies in that label. For example, if “hash(this)” is a valid downgrading policy, then it implies the policy “hash(this)%16” is also valid.

In this framework, both secret and public can be expressed using confidentiality policies: secret can be represented by a constant policy $\{c_0\}$, which trivially hides the secret after substituting this with the secret. The interpretation of secret is all the pattern expressions that does not contain this and secret variables. The security level public can be represented by the expression $\{\text{this}\}$, which says the annotated value can be straightforwardly treated as public data. The interpretation of public includes all the pattern expressions that does not contain trusted secret variables. The **ordering** on labels is then defined using set inclusion:

$$cl_1 \sqsubseteq cl_2 \iff \mathbb{S}(cl_2) \subseteq \mathbb{S}(cl_1)$$

As an example, it is easy to verify that $\text{public} \sqsubseteq \{\text{this} = \star, \text{hash}(\text{this})\} \sqsubseteq \{\text{this} = \star\} \sqsubseteq \text{secret}$. Intuitively, higher security levels contains fewer downgrading policies. The **join** of labels is interpreted as the join of the label interpretations, which can be conservatively approximated by taking the intersection of two labels.

These security labels allow us to formally define **downgrading** in programs. We use the metavariable a to represent **actions**, which specifies a fragment of computation on a value. In an action expression, this denotes the value of interest and all the $\star$ represent values of public level. A label cl can be downgraded to another label using an action a :

$$\Downarrow(cl, a) = \{cp_1 \mid [a/\text{this}]cp_1 \equiv cp_2, cp_2 \in cl\}$$

Here, $\Downarrow(cl, a)$ is the label of data obtained by taking data with label cl and performing action a on it. For example, suppose we define the following labels and actions:

$$\begin{aligned} cl_1 &\triangleq \{(\text{hash}(\text{this})\%4)=\star\} & a_1 &\triangleq \text{hash}(\text{this}) \\ cl_2 &\triangleq \{(\text{this}\%4)=\star\} & a_2 &\triangleq \text{this}\%4 \\ cl_3 &\triangleq \{\text{this}=\star\} & a_3 &\triangleq \{\text{this}=\star\} \end{aligned}$$

We have $\Downarrow(cl_1, a_1) = cl_2$, $\Downarrow(cl_2, a_2) = cl_3$, $\Downarrow(cl_3, a_3) = \text{public}$. If a variable x has security level cl_1 and y is public, then the expression $\text{hash}(x)$ has level cl_2 , $\text{hash}(x)\%4$ has level cl_3 , $(\text{hash}(x)\%4) = y$ has level public. Intuitively, an action is a pattern that matches the computation in the program: this matches a secret value to be downgraded and $\star$ matches any expressions at level public. Each confidentiality label corresponds to a state machine that models downgrading, where states are labels and transitions are downgrading actions.

3.3. Integrity labels

Integrity labels are largely the dual of confidentiality labels with several subtle differences. Confidentiality policies specify what can be done with the data in the future, integrity policy specify what has been done to the data in the past. An **integrity policy** is an expression embedded with $\star$, meaning an expression that has computed the annotated value as a result, where each $\star$ represents an untrusted, tainted input to the expression. For example, if the variable x is tainted, then the expression $x\%4$ has an integrity policy “ $\star\%4$ ”, which states an integrity constraint on the result. An **integrity label** is a non-empty set of integrity policies, where each policy describes an expression that could have computed the value as a result. If the value x has an integrity label $il = \{ip_1, \dots, ip_k\}$, then it must have been computed using one of the expression $ip_j \in il$. Most interesting integrity labels have only one policy in them, because adding policies to a label only *weakens* the integrity guarantee. Similar to confidentiality labels, the **interpretation** of integrity labels is defined as:

$$\mathbb{S}(il) \triangleq \{ip_1 \mid ip_1 \equiv [ip_2/\star_j]ip_3, ip_3 \in il\}$$

The **ordering** on integrity labels is defined as:

$$il_1 \sqsubseteq il_2 \iff \mathbb{S}(il_1) \subseteq \mathbb{S}(il_2)$$

In this framework, tainted can simply be represented as $\{\star\}$, as this is dual case of public in confidentiality. The security label untainted corresponds to an infinite label $\{ip \mid \star \notin ip\}$, which includes all expressions that do not use tainted inputs. Although we have untainted $\sqsubseteq$ tainted, many interesting policies do not sit between these two security levels. For example, $\{\min(\star, 10)\}$ has no direct ordering with untainted. This explains why traditional definitions of noninterference gives a weak form of integrity

Variables	$x ::= x_i$
Expressions	$e ::= c \mid s \mid x \mid \oplus e_1, \dots, e_k \mid \text{declassify}(e, L : a) \mid \text{endorse}(e, L : a)$
Commands	$p ::= \epsilon \mid p; p \mid x := e \mid [L :] \text{if } e \text{ p } p \mid \text{while } e \text{ p } p \mid \text{print } e$ $\mid x := \text{query name } (e_1, \dots, e_k) \mid (s_1, \dots, s_k) := \text{readrow } x$
Types	$\tau ::= \text{string } l \mid \text{query name}$
Programs	$Prog ::= Inputs \ Queries \ Vars \ Body$
Input	$Input ::= field \Rightarrow s$
Query	$Query ::= name(x_1 : il_1, \dots, x_j : il_j) \Rightarrow (s_1 : cl_1, \dots, s_k : cl_k)$
Variables	$Var ::= x : cl \ ! \ il$
Body	$Body ::= c \ Body \mid p \ Body \mid \epsilon$

Figure 4: Abstract syntax

guarantee: there are many interesting integrity policies besides untainted. In fact, untainted is a very coarse security level and it can be further strengthened. For example, the integrity label $il_1 \triangleq \{s_2 + s_3\}$ satisfies $il_1 \sqsubseteq \text{untainted}$. It says a very strong integrity guarantee: values at this integrity level must be equal to the sum of two trusted program inputs. An even stronger integrity label $\{0\}$ works like a singleton type where the only inhabitant is zero. **Downgrading** for integrity labels is formalized as:

$$\Downarrow (il, a) \triangleq \{ip_1 \mid [ip_2/\text{this}]a \equiv ip_1, ip_2 \in il\}$$

For example, suppose

$$\begin{aligned} a_1 &\triangleq \text{ToInt}(\text{this}) & il_1 &\triangleq \{\text{ToInt}(\star)\} \\ a_2 &\triangleq \min(\text{this}, 10) & il_2 &\triangleq \{\min(\text{ToInt}(\star), 10)\} \\ a_3 &\triangleq \max(\text{this}, 5) & il_3 &\triangleq \{\max(\min(\text{ToInt}(\star), 10), 5)\} \end{aligned}$$

We have $\Downarrow (\text{tainted}, a_1) = il_1$, $\Downarrow (il_1, a_2) = il_2$, $\Downarrow (il_2, a_3) = il_3$. If a variable x has security level il_2 , then the expression $\max(x, 5)$ has level il_3 .

4. The web scripting language

4.1. Language syntax and semantics

The web scripting language provides a programming model similar to PHP. The abstract syntax is shown in Figure 4. Program fragments are inserted to the web page using special tags like `<?ssp . . . !ssp>`. At the top level, a web script consists of a header and several program fragments. A header includes the mapping from the form inputs to variable names and the definition of query interfaces. A query interface includes the name of the query, the query arguments, the result variables and the security levels of these variables. Each program fragment is a command. Commands can be sequential composition of commands, assignments, branches, loops, print statements, and query

operations. For simplicity, function calls are not presented in this paper, although they are not fundamentally difficult to include. One other difference from PHP and other information-flow languages is that a branch statement can have a tag L on it; such tags can be used in declassify and endorse expressions, which will be explained later. When the script is executed by the web server, each program fragment is substituted by its output using the print statement. Except query handles, all values are simply strings in this language.

The type system of the web scripting language is presented in Figures 5, 7 and 8. The type system statically controls the information flow in the programs. Confidentiality labels and integrity labels are tracked separately. Like Jif and FlowCaml, the type system generally disallows information flow from high security levels to low security levels, where the ordering of security labels is defined as in Section 3. For confidentiality, implicit information flows are also tracked by adding a program counter label pc to the typing context. The pc label is only permitted to be either public or secret, because the downgrading policies written by the programmer apply only to expressions, not to control flow.

A typing judgment for an expression has the form $\Gamma, \Phi \vdash e : \tau$, where Γ is the typing context for variables and Φ is the context of conditional expressions, explained below. The typing judgments for commands are of the form $\Gamma, \Phi, pc \vdash p$ meaning that the command p is well-typed under the contexts Γ, Φ and the program counter pc . For example, the C-ASSIGN rule only allows information flow from low security levels to high security levels. It also take pc label into account to track implicit flows.

The program writes its output using the `print` statement. Output data is publicly visible to the user, so the type system must ensure that the confidentiality label of the output is public in the C-PRINT rule. Furthermore, the pc

$\frac{\Gamma, \Phi, pc \vdash p_1 \quad \Gamma, \Phi, pc \vdash p_2}{\Gamma, \Phi, pc \vdash p_1; p_2}$	C-COMPOSITION
$\frac{\Gamma(x) = \text{string } cl_1 ! il_1 \quad \Gamma, \Phi \vdash e : \text{string } cl_2 ! il_2 \quad cl_2 \sqsubseteq cl_1 \sqcup pc \quad il_2 \sqsubseteq il_1}{\Gamma, \Phi, pc \vdash x := e}$	C-ASSIGN
$\frac{\Gamma, \Phi \vdash e : \text{string } cl ! il \quad cl \equiv \text{public} \quad pc \equiv \text{public}}{\Gamma, \Phi, pc \vdash \text{print } e}$	C-PRINT
$\frac{\Gamma, \Phi \vdash e : \text{string } cl ! il \quad cl \equiv \text{public} \quad \Gamma, \Phi, \text{public} \vdash p \quad pc \equiv \text{public}}{\Gamma, \Phi, pc \vdash \text{while } e \text{ } p}$	C-WHILE
$\frac{\Gamma, \Phi \vdash e : \text{string } cl ! il \quad cl \sqcup pc \sqsubseteq \text{public} \quad \Gamma, \Phi \cup (L^+ : e), \text{public} \vdash p_1 \quad \Gamma, \Phi \cup (L^- : e), \text{public} \vdash p_2}{\Gamma, \Phi, pc \vdash L : \text{if } e \text{ } p_1 \text{ } p_2}$	C-IF-PUB
$\frac{\Gamma, \Phi \vdash e : \text{string } cl ! il \quad \Gamma, \Phi, \text{secret} \vdash p_1 \quad \Gamma, \Phi, \text{secret} \vdash p_2}{\Gamma, \Phi, pc \vdash L : \text{if } e \text{ } p_1 \text{ } p_2}$	C-IF-SEC
$\frac{Q(\text{name}) = (x_1 : il_1, \dots, x_k : il_k) \Rightarrow (\dots) \quad \Gamma, \Phi \vdash e_i : \text{string } cl_i ! il'_i \quad cl_i \equiv \text{public} \quad il'_i \sqsubseteq il_i \quad pc \equiv \text{public} \quad \Gamma \cup (x : \text{query name}), \Phi, pc \vdash p}{\Gamma, \Phi, pc \vdash x := \text{query name } (e_1, \dots, e_k); p}$	C-QUERY
$\frac{\Gamma(x) = \text{query name} \quad Q(\text{name}) = (\dots) \Rightarrow (S_1 : cl_1, \dots, S_k : cl_k) \quad cl'_i \triangleq pc \sqcup [s_1/S_1] \dots [s_k/S_k] cl_i \quad pc \equiv \text{public} \quad \Gamma \cup (s_i : \text{string } cl'_i ! \{s_i\}), \Phi, pc \vdash p}{\Gamma, \Phi, pc \vdash (s_1, \dots, s_k) := \text{readrow } x; p}$	C-READROW

Figure 5: Command typing rules: $\boxed{\Gamma, \Phi, pc \vdash p}$

```

01 <?ssp_header
02   FormInputs ( "UserName" => u, "Password" => p, "QueryYear" => y );
03
04   Query GetID ( username: !tainted ) => (
05     PASSWORD : {this=*},
06     ID       : {if (PASSWORD=*) this 0}
07   );
08
09   Query FetchRecords( index: !untainted, year: !{Integer(*)} ) => (
10     ORDERID  : public,
11     AMOUNT   : public,
12     CCNUM    : {tailstr(this,4)}
13   );
14
15   Variables ( pub_id: public!untainted );
16
17 !ssp_header>
18 <html><head><title>....</title>
19 <meta http-equiv="content-type" content="text/html; charset=UTF-8">
20 </head><body>
21 .....
22 <?ssp
23   q1 := query GetID(u);
24   if ( empty(q1) ) {
25     print 'Unknown username';
26   } else {
27     (pwd, id) := readrow(q1);
28     L_AUTH: if (pwd=p)
29     {
30       print 'Username = '; print u;
31       pub_id := declassify(id, L_AUTH:(pwd=*));
32       print 'School ID ='; print pub_id;
33       q2 := query FetchRecords( pub_id, Integer(y) );
34       while (!empty(q2)) {
35         (orderid, amount, ccnum) := readrow(q2);
36         print 'Order ID = '; print orderid;
37         print 'Amount = '; print amount;
38         print 'Credit Card = XXXX-XXXX-XXXX-';
39         print tailstr(ccnum, 4);
40       }
41     } else {
42       print 'Wrong password';
43     }
44   }
45 !ssp>
46 .....

```

Figure 6: A web script example

$\frac{}{\Gamma, \Phi \vdash c : \text{string public ! } \{c\}}$	E-CONST
$\frac{\Gamma(s) = \text{string } cl ! il}{\Gamma, \Phi \vdash s : \text{string } cl ! il}$	E-TRUSTVAR
$\frac{\Gamma(x) = \text{string } cl ! il}{\Gamma, \Phi \vdash x : \text{string } cl ! il}$	E-VAR
$\frac{\Gamma, \Phi \vdash x : \text{query name}}{\Gamma, \Phi \vdash \text{empty}(x) : \text{string public ! tainted}}$	E-EMPTYTEST
$\frac{\Gamma, \Phi \vdash e_1 : \text{string } cl_1 ! il_1 \quad \Gamma, \Phi \vdash e_2 : \text{string } cl_2 ! il_2 \quad CL(\oplus, cl_1, cl_2, il_1, il_2) = cl_3 \quad IL(\oplus, cl_1, cl_2, il_1, il_2) = il_3}{\Gamma, \Phi \vdash e_1 \oplus e_2 : \text{string } cl_3 ! il_3}$	E-OP
$\frac{\Gamma, \Phi \vdash e_1 : \text{string } cl_1 ! il \quad \Phi(L^+) = e_2 \quad \text{match}(\Gamma, e_2, a) \quad cl_2 \stackrel{\Delta}{\Downarrow} (cl, \text{if } a \text{ this } 0)}{\Gamma, \Phi \vdash \text{declassify}(e_1, L : a) : \text{string } cl_2 ! il}$	E-DECLASSIFY+
$\frac{\Gamma, \Phi \vdash e_1 : \text{string } cl_1 ! il \quad \Phi(L^-) = e_2 \quad \text{match}(\Gamma, e_2, a) \quad cl_2 \stackrel{\Delta}{\Downarrow} (cl, \text{if } a \text{ } 0 \text{ this})}{\Gamma, \Phi \vdash \text{declassify}(e_1, L : a) : \text{string } cl_2 ! il}$	E-DECLASSIFY-

Figure 7: Expression typing rules: $\boxed{\Gamma, \Phi \vdash e : \tau}$

label in the typing context of the `print` statement must also be public to prevent implicit information flow such as “if secret then print 1 else print 0”. Secret data must be downgraded to public data before they can be printed to web pages.

Figure 6 shows an actual web script. It has two query interfaces to the database. The script takes the user input, performs queries to the database, reads the results from the query and generates a web page for the end user. The following subsections walk through this example step-by-step to present the language features.

4.2. Query interfaces and type declarations

Query interfaces and variables are declared in the header section of a script. In Figure 6, the script header from line 1 to line 17 specifies all the input/output data types. Line 2 specifies the inputs fields submitted from the web forms in a HTTP request and maps them to variables `u`, `p` and `y` in the scripting language. All of these variables are considered to be public, untrusted data; they have security level “public ! tainted” in the type system.

Two query interfaces are defined in lines 4-13. At runtime, they must match their specifications in the DBMS interfaces. The first query, `GetID`, looks up a user name in the database and returns the user’s password and identity number; both are secrets and cannot be directly released to the public. The second query `FetchRecords` uses the user’s identity number to look up the user’s transaction history in the database.

Query arguments: Integrity labels are specified for query arguments in the interface; their confidentiality labels are public by default as required by the C-QUERY type-checking rule. For example, the query argument `index` on line 9 has an integrity label untainted, which makes it impossible to pass an untrusted tainted value to `index`. The argument `year` on line 9 requires a mandatory conversion from an arbitrary string to a string that contain only an integer, which forbids certain SQL string attacks.

Query results: Query results are modeled as trusted variables in the type system. All the trusted variables are treated as read only in the type system: they cannot be updated using direct assignments. Confidentiality labels are

$\frac{}{CL(\oplus, cl_1, cl_2, il_1, il_2) = \text{secret}}$	CL-SECRET
$\frac{cl_1 \sqsubseteq \text{public} \quad cl_2 \sqsubseteq \text{public}}{CL(\oplus, cl_1, cl_2, il_1, il_2) = \text{public}}$	CL-PUBLIC
$\frac{il_2 \sqsubseteq \{ip\} \quad a_1 \stackrel{\Delta}{=} \text{this} \oplus ip \quad \text{match}(a_1, a_2) \quad cl_3 \stackrel{\Delta}{=} \Downarrow (cl_1, a_2)}{CL(\oplus, cl_1, \text{public}, il_1, il_2) = cl_3}$	CL-DOWNGRADE(L)
$\frac{}{IL(\oplus, il_1, il_2) = \text{tainted}}$	IL-TAINTED
$\frac{il_1 \sqsubseteq \text{untainted} \quad il_2 \sqsubseteq \text{untainted}}{IL(\oplus, il_1, il_2) = \text{untainted}}$	IL-UNTAINTED
$\frac{}{IL(\oplus, \{ip_1\}, \{ip_2\}) = \{ip_1 \oplus ip_2\}}$	IL-COMPOSE

Figure 8: Downgrading rules

specified for query results, and their integrity labels are implicitly defined by the C-READROW rule: a variable s has an integrity label $\{s\}$. On line 5, the confidentiality label for PASSWORD states that the only possible way to leak information about the password is to compare it with a value at level public.

A row of the query result is read together using the readrow command. The confidentiality policy of a query result variable can mention names of other variables in the same query. For example, the variable ID on line 6 has a policy that mentions the variable PASSWORD, saying that the ID string can only be disclosed if a publicly known string matches PASSWORD on the same row of the query result. This policy specifies a run-time test of identity information.

4.3. Downgrading

Downgrading is the key feature of this type system. The downgrading policies specified by the query interfaces control how confidential data is released and how trustworthy information is updated. There are two downgrading mechanisms in the type system.

Implicit downgrading: Downgrading happens implicitly in each step of computation that uses the built-in operators. Without loss of generality, we present only the typing rules for binary operators in this paper. In traditional security type systems, if x has security level l_1 and y has level l_2 , the result of $x \oplus y$ has security level $l_1 \sqcup l_2$, which is an upper-bound of l_1 and l_2 . The E-OP rule in Figure 7 is backward-

compatible with those type systems. However, E-OP examines the labels of the operands more carefully using rules in Figure 8. The CL-SECRET, CL-PUBLIC, IL-TAINTED and IL-UNTAINTED rules are standard rules—they give the label of the result by approximating the join of the arguments. The IL-COMPOSE rule, however, attempts to compute the integrity label for the result when both operands have good integrity guarantees. The CL-DOWNGRADE(L) rule declassifies the left operand using an action that corresponds to the use of the $\oplus$ operator (there is a symmetric version that operates on the right side of $\oplus$). The integrity label of the other operand is also used to describe this action.

It is possible that more than one action a_2 can be chosen in the CL-DOWNGRADE rule. For example, the action $\text{this} = c_0$ can match the policy $\text{this} = \star$. The predicate match determines whether the action matches a downgrading policy pattern; we omit the straightforward definition of when such patterns unify. We make this downgrading implicit because most useful downgrading policies are simple and it is easy to search (in the implementation of match) for an usable downgrading action. To avoid searching, the language could be extended with an optional construct that specifies the downgrading action as annotations on the operator so that the type checker knows which action to use.

Conditional downgrading: The conditional expressions in the policy language allow us to specify downgrading patterns with branches. This is achieved in the type system by tracking the “active conditions” on the current execution path. The if statements can have an optional tag L

in its syntax. In the C-IF-PUB rule, the context Φ is used to keep track of all the conditional expressions tags on the current execution path. These tags are also annotated with either $+$, indicating the “true” branch, or $-$, indicating the “false” branch. Programs can use declassify and endorse statements to downgrade the label of a value by specifying the tag of the conditional expression that corresponds to the if statement, and an action that *matches* the conditional expression. Here again, a $\text{match}(\Gamma, e, a)$ predicate is needed to determine whether the expression e can instantiate action a in context Γ . Intuitively, the tag L mentioned in the downgrading operation provides the justification for control flow that validates the use of the downgrading action. This is a novel extension of our earlier type system [5] without breaking the *relaxed noninterference* result: the relaxed noninterference of conditional patterns can be justified by an equivalence rule “if $e_1 E[e_2] e_3 \equiv \text{if } e_1 E[\text{if } e_1 e_2 c] e_3$ ” where E is an evaluation context, with some side conditions on the typing contexts and variable bindings.

4.4. Information flow control in the example

In Figure 6, lines 22-44 show a program fragment in the web script. It uses the query `GetID` to authenticate the user and uses the query `FetchRecords` to look up the user’s history of transactions.

Reading query results: Line 23 submits a query to the database and returns a handle `q1` to the instance of this query. Line 27 reads a row from the query results. When the variables (pwd, id) are added to the context, their types are added to the typing context according to the database interfaces. The variable `pwd` has security level $\{\text{this}=\star\}!\{\text{pwd}\}$ and `id` has security level $\{\text{if } (\text{pwd}=\star) \text{ this } 0\}!\{\text{id}\}$. The variable names in the query interfaces are substituted by the actual instances in the C-READROW rule.

Implicit downgrading: Line 28 performs an implicit downgrading in the conditional expression. The expression `pwd=p` corresponds to an action “`this = *`” for the variable `pwd`, and we have $\Downarrow (\{\text{this} = \star\}, \{\text{this} = \star\}) = \text{public}$. Therefore, the expression `pwd=p` has security level `public ! tainted` by using the E-OP rule, the CL-DOWNGRADE(L) rule and the IL-TAINTED rule. The information leak on this expression is permitted by the policy and the *pc* label inside the branch will be public. In contrast, in languages like Jif, downgrading must be performed by using its `declassify` expression. Implicit downgrading also happens on line 33 for the variable `y` and on line 39 for the variable `ccnum`. The type system provides substantial guarantees about downgrading—all the downgrading must follow permissible downgrading paths specified by the policies.

Conditional downgrading: When the `if` statement is typechecked on line 28, the tag `L_AUTH` and the conditional expression $(\text{pwd}=p)$ are stored in the context Φ . This information is used to justify further downgrading inside the body of the branches. Line 31 shows such an example. The declassify operation is a no-op at run-time. It merely serves as a hint to the typechecker that a conditional test is in the current execution path and the expression `id` can be downgraded using an action corresponding to the conditional test on line 28. First, the E-DECLASSIFY+ rule verifies that the action $(\text{pwd}=\star)$ matches the conditional expression $(\text{pwd}=p)$ that correspond to the tag `L_AUTH` in the typing context. Then, the security level of `id` is downgraded using the action $(\text{if } (\text{pwd}=\star) \text{ this } 0)$ where 0 is an arbitrary constant. Thus, the resulting confidentiality label for `pub_id` is public. If the programmer does not perform the required identity test (specified on line 6), the type system will not permit the program to output the ID. This conditional downgrading policy effectively enforces a run-time identity test [15].

Writing query arguments: The `declassify` statement only affects the confidentiality label. The integrity label for `pub_id` is still $\{\text{id}\}$. Line 33 calls another query `FetchRecords` using the value `pub_id`. According to the C-QUERY rule, the interface of `FetchRecords` demands that the first argument has an untainted integrity level. This is satisfied because $\{\text{id}\} \sqsubseteq \text{untainted}$ in our framework. If the `FetchRecords` query used a tainted value—perhaps obtained from user input—the type system will detect such an error.

5. Discussion

5.1. Untrusted code, timing channels and side effects

Our web scripting language is primarily intended as a tool to help web-systems builders create more secure systems. As such, the main focus of this paper has been on *trusted code* which is written without malicious intent. Dealing with *untrusted code* is a much more difficult problem. However, the downgrading and trust model described in this paper differs from previous work in a couple significant ways.

In Jif, downgrading is controlled using the decentralized label model, where each principal can only downgrade its own policies. The DLM uses the notion of *authority* as justification for privileged operations; but *authority* is not connected to the program semantics. As a result, untrusted code (code without the authority of a principal P) cannot downgrade data owned by P . Our language provides a complementary ability to specify downgrading policies based on *required computation* rather than using *code*

privileges. This makes it possible to allow untrusted code to perform downgrading in a safe manner as long as the downgrading policies are correctly specified (and the untrusted code passes the typechecker). Of course, for untrusted code, the downgrading policies must be specified conservatively with possible attacks in mind. For example, the `FetchRecords` query is not safe to use in an untrusted setting because the attacker can enumerate possible identity numbers and steal information from the database. This problem can be solved by posing stronger policies on the query interface, for example, using run-time identity tests like the policy on line 6.

For confidentiality, untrusted code can also leak information through covert channels such as timing channels and side effects. This problem is solved by requiring the *pc* label be public at all places where observable side effects are possible to happen: loops, reading rows from queries, etc. In the type system, the `C-PRINT`, `C-WHILE`, `C-QUERY` and `C-READROW` all require the *pc* label to be public. This solution is impractical for languages like Jif and FlowCaml, because the explicit declassification needed make it too clumsy to allow useful programs be written. However, our language makes it more practical, because many secret data can be implicitly downgraded to public data before they affect the *pc* label. The *pc* label is indeed public everywhere in the example in Figure 6. It is also worth pointing out that the control flow in a web script is often simpler than other programs. Web scripts naturally use the continuation-passing programming style and many scripts execute for a very short time. By limiting the confidentiality label of the loop condition in the `C-WHILE` rule, timing channel leaks are largely eliminated.

For trusted code, the requirement on the *pc* label can be relaxed. Instead of rejecting a program, the typechecker can raise appropriate warnings in the `C-PRINT`, `C-WHILE`, `C-QUERY` and `C-READROW` rules, where confidential information can be leaked through covert channels.

The web scripting language presented in this paper is intended to be a practical instantiation of a more theoretical language in our earlier paper [5]. However, despite confidence in the type system presented here (as justified by that previous work), we have not proved any formal security guarantees about it. The security goal is harder to formalize than the *relaxed noninterference* result proved previously, because this language includes side effects and state. A promising future direction is to formalize the security guarantee for this language, perhaps by using a functional variant of this language with monadic effects.

5.2. Related work

Language-based information-flow control has been studied for some time [13]. Recent language prototypes such as

Jif [7], which extends Java, and FlowCaml [14], which extends Caml, provide nonstandard type systems that enforce information-flow policies. The security guarantee of such type systems is usually formalized as *noninterference* [4, 6], an end-to-end extensional guarantee that requires that no information propagates from high security levels to low security levels. However, there have been very few practical applications that demonstrate the use of these security-typed languages. In this paper, we have proposed to apply these techniques to web scripting languages, for which security concerns are increasingly important. The current state-of-the-art in web scripting protects data confidentiality and (perhaps more importantly) integrity in ad hoc ways. Here we aim to do better, yet still provide a practical enforcement mechanism.

One important challenge in making such an approach practical is the problem of *downgrading* [8, 16, 3, 11, 2, 5]. Noninterference alone is too strong for practical use. For confidentiality, it is usually necessary that secret data can be leaked to public places, but only in controlled ways. One approach, the decentralized label model (DLM) [8] can control downgrading by using privileges associated with the code. The DLM allows us to specify policies about *who* can downgrade the data, but does not specify *how* the data should be downgraded and *what* what is downgraded. As a result, the end-to-end noninterference guarantee no longer holds for code with downgrading.

Another problem is information *integrity policies* [1]. Although confidentiality and integrity are usually considered as duals in information-flow systems, the resulting integrity guarantee is weak. Noninterference guarantees only that tainted data does not affect the values of the trusted, untainted data, but it does not say anything about how the trusted data are manipulated in the system. There is absolutely no integrity guarantee for data coming from untrusted code in the Jif language, because a malicious program can manipulate the trusted data in arbitrary ways without using tainted data. As a result, the two-dimensional DLM degenerates to a one-dimensional *trust model* or *writers model* for integrity policies.

Recent advances in the research on *downgrading* extend the notion of noninterference by specifying *downgrading policies* as security levels and studies *how* the data are downgraded. *Relaxed noninterference* [5] and *delimited release* [12] provide end-to-end security guarantees on downgrading. In the *relaxed noninterference* framework [5], an information-flow type system is used to control downgrading in a fine-grained manner according to the *downgrading policies* specified by the programmer. A secure program can be proved to be equivalent to a special form where all the downgrading are explicit and external to the body of the program. The security guarantee can then be interpreted in the model of *delimited release* [12] proposed

by Sabelfeld and Myers, which states a weakened and backward-compatible version of noninterference.

This paper integrates these theoretical frameworks together in a practical, domain-specific programming language. The *relaxed noninterference* framework expresses the downgrading policies as lambda calculus terms; manipulating the policies requires higher-order unification and extensive proof searching. In this paper, we simplify the policy language by using patterns to represent policies and use straightforward pattern matching in type-checking to avoid extensive searching. Instead of using an effect type system to enforce delimited release [12], we simply require that all the confidential input variables are read-only variables. Furthermore, the conditional downgrading policies can be used to enforce run-time identity tests and achieve similar goals with *run-time principals* [15].

6. Conclusion

This paper presents an architecture for obtaining strong, end-to-end security in web-based online information systems and motivates the use of language-based information-flow control in a web scripting language. In this approach, information-flow policies are specified in the database query interfaces and enforced in the web scripting language by a static type checker.

Based on prior research, this paper presents a framework of downgrading policies using a simple and tractable pattern language that connects implicit downgrading to computations in the script. Integrity policies and confidentiality policies are treated symmetrically, leading to a clean and intuitive way for programmers to describe their policies. Moreover, this paper presents a novel downgrading mechanism that works by tracking the conditional expressions in the typing context and using them to enforce policies on run-time conditions such as identity tests.

References

- [1] K. J. Biba. Integrity considerations for secure computer systems. Technical Report ESD-TR-76-372, USAF Electronic Systems Division, Bedford, MA, Apr. 1977.
- [2] S. Chong and A. Myers. Security policies for downgrading. In *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS)*, 2004.
- [3] R. Giacobazzi and I. Mastroeni. Abstract non-interference: Parameterizing non-interference by abstract interpretation. In *Proc. 31st ACM Symp. on Principles of Programming Languages (POPL)*, pages 186–197, Venice, Italy, Jan. 2004.
- [4] J. A. Goguen and J. Meseguer. Security policies and security models. In *Proc. IEEE Symposium on Security and Privacy*, pages 11–20. IEEE Computer Society Press, Apr. 1982.
- [5] P. Li and S. Zdancewic. Downgrading policies and relaxed noninterference. In *Proc. 32th ACM Symp. on Principles of Programming Languages (POPL)*, 2005.
- [6] J. McLean. Security models and information flow. In *Proc. IEEE Symposium on Security and Privacy*, pages 180–187. IEEE Computer Society Press, 1990.
- [7] A. C. Myers. JFlow: Practical mostly-static information flow control. In *Proc. 26th ACM Symp. on Principles of Programming Languages (POPL)*, pages 228–241, San Antonio, TX, Jan. 1999.
- [8] A. C. Myers and B. Liskov. Protecting privacy using the decentralized label model. *ACM Transactions on Software Engineering and Methodology*, 9(4):410–442, 2000.
- [9] F. Pottier and S. Conchon. Information flow inference for free. In *Proc. 5th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 46–57, Sept. 2000.
- [10] F. Pottier and V. Simonet. Information flow inference for ML. In *Proc. 29th ACM Symp. on Principles of Programming Languages (POPL)*, Portland, Oregon, Jan. 2002.
- [11] A. W. Roscoe and M. H. Goldsmith. What is intransitive noninterference? In *Proc. of the 12th IEEE Computer Security Foundations Workshop*, 1999.
- [12] A. Sabelfeld and A. Myers. A model for delimited information release. In *Proceedings of the International Symposium on Software Security (ISSS'03)*, 2004.
- [13] A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, Jan. 2003.
- [14] V. Simonet. Flow Caml in a nutshell. In G. Hutton, editor, *Proceedings of the first APPSEM-II workshop*, pages 152–165, Nottingham, United Kingdom, Mar. 2003.
- [15] S. Tse and S. Zdancewic. Run-time Principals in Information-flow Type Systems. In *Proc. IEEE Symposium on Security and Privacy*, 2004.
- [16] D. Volpano and G. Smith. Verifying secrets and relative secrecy. In *Proc. 27th ACM Symp. on Principles of Programming Languages (POPL)*, pages 268–276. ACM Press, Jan. 2000.