

Effective Blame for Information-Flow Violations

Dave King[†] Trent Jaeger[†] Somesh Jha^{*} Sanjit A. Seshia[‡]

[†] The Pennsylvania State University
University Park, PA, 16802
{dhking,tjaeger}@cse.psu.edu

^{*} University of Wisconsin
Madison, WI, 53706
jha@cs.wisc.edu

[‡] University of California
Berkeley, CA, 94720
sseshia@eecs.berkeley.edu

ABSTRACT

Programs trusted with secure data should not release this data in ways contrary to system policy. However, when a program contains an illegal flow of information, current information-flow reporting techniques are inadequate for determining the cause of the error. Reasoning about information-flow errors can be difficult, as the flows involved can be quite subtle. We present a general model for *information-flow blame* that can explain the source of such security errors in code. This model is implemented by changing the information-flow verification procedure to: (1) generate supplementary information to reveal otherwise hidden program dependencies; (2) modify the constraint solver to construct a blame dependency graph; and (3) develop an explanation procedure that returns a *complete* and *minimal* error report. Our experiments show that information-flow errors can generally be explained and resolved by viewing only a small fraction of the total code.

1. INTRODUCTION

It is essential for programs, especially those that are entrusted with critical personal information, to enforce security goals [11]. Information-flow security is a key requirement for secure programs; if a program is information-flow secure, then secret information will not flow to public channels at runtime [9]. *Security-typed languages* can be used to build applications that are information-flow secure [10, 23, 24, 25]. Once the programmer has annotated the security requirements on variables and I/O channels, the compiler verifies that these annotations do not permit an illegal information flow.

Although secure systems have been built using security-typed languages, at present these can only be built from scratch [2, 6, 18]. To scale to real systems, we must be able to retrofit legacy code for security. However, such retrofitting is difficult for two reasons. Because the original code was not necessarily written with information-flow requirements in mind, it will contain a large number of security violations,

some of which may correspond to expected system behavior (releasing the result of a password check), and some of which may not (the unintended release of a patient's social security number). These errors can be quite complex, as any runtime program path may cause an error. Program paths can span procedures, and conditionals and exceptions create subtle dependencies between data. In even moderately-sized code, the number of possible program paths means that deducing the source of an error, given just the location of where an error manifested itself, can be daunting.

The *information-flow blame problem* is to identify the possible sources of an error in a program. An effective solution to the blame problem is an error report that is both complete (everything that caused the error is reported) and minimal (what is reported does not contain spurious information). The blame problem has been well-studied for runtime errors [5, 17], typing errors [22, 32], and incorrect value computation [19, 31, 33]. However, current methods for determining causes of information-flow errors [8, 16] are either incomplete, non-minimal, or do not scale to non-trivial programs.

In this paper, we present a model for information-flow blame and an algorithm for extracting the core reason for program inconsistency from a constraint solver. We implement our model by extending the Jif information-flow compiler [24] with an interprocedural constraint-based program analysis. We then demonstrate how our analysis performs on a variety of Java programs and briefly describe the kinds of information-flow errors that these programs permit. This paper makes the following contributions:

- We show how to build a *blame dependency graph* of constraints that enables queries to retrieve the slice of the graph that contributes to a violation.
- We show how the dependency graph can be used to generate a *complete* and *minimal* set of constraints that caused an error.
- We expand the Jif compiler to display error explanations generated by Java code.
- We apply our analysis to a variety of programs, showing that the errors reported accurately describe the cause of an illegal information flow. We find that our blame model reduces each error to viewing a small fraction of the program (less than 0.5% of the original source), making error resolution straightforward.

The remainder of the paper is organized as follows: Section 2 outlines the current blame procedure in security-typed languages and why it is presently inadequate. We then outline our model for determining blame for security errors in

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

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

programs. The formal background for a constraint solver commonly used for security-typed language extensions is detailed in Section 3. In Section 4, we describe the blame dependency graph, a data structure that allows tracing the reason why a program contains an illegal flow of information. In Section 5, we outline the modifications we needed to make to the Jif compiler for our analysis. In Section 6, we describe the result of our experiments on several Java programs. In Section 7, we describe some related work, and we give concluding remarks in Section 8.

2. PROBLEMS WITH CURRENT BLAME

We now outline some current problems in finding and resolving information-flow errors in program code. As an example of a typical information-flow analysis¹, we investigate the security-typed language Jif [24], an extension of JFlow [23]. We first give some background on how Jif verifies the information-flow security of its programs.

A program is *information-flow secure* if high security variables do not affect the values stored in low security variables [9]. Explicit information flows occur when high security data is directly written to low security data, while implicit information flows occur when low security data is otherwise affected by high security data. If h is a high security variable and l is a low security variable, then the assignment $l := h$ enables the *explicit* flow of information from h to l , while the conditional `if (h = 0) then l := 0` enables the *implicit* flow of information from h to l .

To determine if a program contains an illegal flow, Jif generates *label constraints* for the statements in the code. An assignment statement $v := e$ generates the constraint $L_2 \leq L_1$, where L_1 is the security level of v , L_2 is the security level of e , and $L_2 \leq L_1$ is read as “the security level of L_2 is less or equal to that of L_1 ”. This constraint requires the security level of the expression e to be lower than the security level of the variable v . If the security level of e is not known (as it may not be explicitly labeled), a variable β_e is introduced and used in place of L_2 .

To detect implicit information flows, the compiler keeps track of the security of the *program counter*; the security level of the program counter is the level of information released by executing a particular line in a program. Assignments that occur inside a conditional must be to variables of a level no lower than the program counter. For example, if program variables i , j , and k have security levels α_i , β_j , γ_k respectively, then conditional `if (i == 0) { j := k; }` generates the label constraint $\alpha_i \sqcup \gamma_k \leq \beta_j$: both i and k must not have a higher security level than j . Nested assignments, loops, and exceptions are treated in a similar way: any variable that affects the control-flow of the program taints the program counter.

Once the label constraints for representing the information flows permitted by the program have been generated, Jif runs a modified implementation of the Rehof-Mogensen constraint solver [26] (described in Section 3.1), generating an assignment of variables to security levels such that each constraint is satisfied. If no such assignment exists, the solver reports the first constraint that it could not satisfy, and then fails. The constraint generation and solving

¹Most implementations of information-flow checkers behave in the way described in this paper, though their underlying type systems may vary.

```

1 User[{Secret}] user; // information in User is Secret
2 OutputStream[{Public}] out; // output stream writes to Public
3
4 private void handleRetr() {
5     // fileIn is a secret file reader
6     BufferedReader fileIn = user.getReader(msgNum);
7     String currentLine = fileIn.readLine();
8     while (currentLine != null) {
9         this.write(currentLine);
10        currentLine = fileIn.readLine();
11    }
12 }
13
14 private void handleTop() {
15     // whether a message is deleted or not is secret information
16     if (user.getMessage(msgNum).isDeleted())
17         this.write(MESSAGE_NO_SUCH_MESSAGE);
18 }
19
20 private void write(String message) {
21     // writes a message out (Public sink)
22     out.print(message);
23 }

```

Figure 1: Fragment of code from the POP3 processor in the JES Email Server.

process is together called *label checking*.

2.1 Difficulties with Information-Flow Error Reporting

The primary difficulty with the above approach to information-flow blame is that blaming the first constraint at which an error was detected always blames a *sink* of secure information. Information-flow security annotations to programs fall into two different categories: *sources* or *sinks*. For example, if a field containing a PIN number is labeled as `Secret`, then that field is a source of `Secret` information. If a public output stream is labeled as requiring `Public` data, then it is a sink for `Public` information. A field or variable annotated by a security label is a source, while the variable of the left-hand-side of an assignment or a formal parameter of a method call are sinks. Sinks impose a security requirement on flows through the program: if a sink has security label L , then any source that flows to the sink must have a label L' such that L' is not more restrictive than L ($L' \leq L$).

Figure 1 contains code from the JES, an open source Java email server². The situation in the code is a common one: a secret source, here the user’s mailbox, flows to an public output, here a socket’s output stream, through more than one method. By examining the code, we can see that there are two information leaks: the leak of whether or not a message has been deleted (caused by line 16) and leaking the contents of a message (caused by line 9). We would prefer to treat each information flow separately: while information flows can be permitted by adding an explicit declassifier, it would not be good practice to automatically declassify any information leaving the system through the `write` method.

The approach that Jif takes to information-flow blame is to simply blame the first constraint that failed during label checking. Using that approach for analysis of a Java program, we would highlight the line `out.print(message)` and indicate that this line might send out secret information. We claim that this error message is unsatisfactory, as it does not highlight the *cause* of the error. In the situation where a client is sending a message `msg` out over a `Public` socket, it

²<http://www.ericdaugherty.com/java/mailserver/>

may be possible that `Secret` data has affected `msg`, but this error message does not display how.

We claim that for a solver’s error message to be satisfactory, it must be *complete*: everything that caused a constraint to fail must be available to the programmer. Moreover, the error message must be minimal: it should not show anything unrelated to a specific failure. Specifically, it should show how each line of the program’s code contributes to an illegal flow. In the previous example, we should report either that `message` was raised to contain `Secret` data because of the implicit flow from line 16 to the output, or because of the explicit flow from line 9.

Alternate approaches to finding information-flow errors in programs are also unsatisfactory. Past work has also used program slicing to find information-flow errors [29]. However, slices from information-flow errors are too large, as both the forward slice of code from the initialization of the `user` variable or the backwards slice of code from the output statement `out.println(message)` would include the entire example code. For example, in our analysis of the JES, we found 30 information-flows from `Secret` to `Public` in the POP3 Processor, a class containing 468 lines of code. We found that a program slicing tool was insufficient to find distinct information flows from `Secret` to `Public`, as it could not identify individual violating paths from source to sink. Additional analyses for information-flow errors do not track implicit flows [14], do not operate on a full programming language [8], or are too computationally expensive [16].

A final technical difficulty with using Jif for retrofitting existing Java code is that its analysis requires programmers to annotate, for all methods, the security value of each argument and the side effects the method may release. This quickly becomes difficult and leads to labeling conflicts rather than true security errors [6, 18]. This is not a common feature of other information-flow analyses [25, 30].

2.2 Our Approach

To find information-flow errors in existing Java codebases, we modified the Jif compiler to operate on Java programs; this required only minor changes to each aspect of the label checking procedure. Figure 2 summarizes our approach.

To use our tool, a programmer first gives security annotations on various program elements such as variables and fields. The blame algorithm will then output relevant parts of program paths that witness an information flow violation to the user. We find these program paths by instrumenting the solver contained in the Jif engine with a *blame dependency graph* that records information to determine why a label constraint became unsatisfiable. The constraint set is generated by an *interprocedural label analysis* that allows programmers to only label security-relevant code. The analysis infers the rest of the security annotations for the rest of the program. Our experiments show that our blame algorithm, provided with an interprocedural label analysis, can accurately identify illegal information flows in programs.

3. SOLVER BACKGROUND

The first step in building our comprehensive blame model is to modify the constraint solver to extract the constraints that contributed to a failure. In this section, we provide a formal description of the constraint solver used by Jif. This will be necessary to understand the motivation behind the *blame dependency graph*, given in the next section.

3.1 Rehof-Mogensen Constraint Solver

To verify the information-flow security of program code, Jif uses a variant of the Rehof-Mogensen solver, a linear-time constraint solver [26]. The Rehof-Mogensen solver operates in two phases: the first adjusts variables based on constraints with a variable on the right-hand side, while the second makes sure that constraints with a label from the security lattice on the right-hand side are still satisfied. We now give some theoretical background for this constraint solver; we will later revisit how the Jif compiler uses it.

Let P be a partially-ordered set (poset) and F be a finite set of monotone functions $f : P^{a_f} \rightarrow P$ where $a_f \geq 1$ is the arity of f . The pair $\Phi = (P, F)$ is called a *monotone function problem* or *MFP*. Given a MFP Φ , the set of Φ -terms (denoted by T_Φ) is given by the following grammar:

$$\tau = \alpha \mid L \mid f(\tau_1, \dots, \tau_{a_f}),$$

where L ranges over constants in P , α is a variable, $f \in F$, and $\tau, \tau_1, \dots, \tau_{a_f}$ are Φ -terms. Let the set of all variables, assumed to be a denumerably infinite set, be \mathcal{V} . The set of variables occurring in a term τ is denoted by $\text{Vars}(\tau)$. A constraint is of the form $\tau \leq \tau'$, where $\tau, \tau' \in T_\Phi$. A constraint set C is a finite set of constraints over Φ .

Let C be a constraint set. A *valuation* ρ is a function from \mathcal{V} to P . Given a fixed run of the solver, a partial valuation ρ_t refers to the t -th valuation. Given a valuation ρ , $\rho(\tau)$ is the value of the term τ under ρ . A valuation ρ satisfies the constraint $\tau \leq \tau'$ iff $\rho(\tau) \leq \rho(\tau')$; we write this $P, \rho \models \tau \leq \tau'$. A valuation $\rho : \mathcal{V} \rightarrow P$ satisfies C iff ρ satisfies every constraint in the set C . The set of all valuations that satisfy C , denoted by $\text{sol}(C)$, is called the set of all solutions to C . Given a MFP $\Phi = (P, F)$, the decision problem Φ -SAT is defined as follows:

Given a constraint set C over Φ , determine whether C is satisfiable.

For the rest of the paper, we assume that P is a lattice \mathcal{L} with bottom element \perp . The *join* and *meet* operators for \mathcal{L} are denoted by \sqcup and \sqcap , respectively. Let $\Phi = (\mathcal{L}, F)$ be a MFP. Let an *atom* be a variable (α or β) or a constant $L \in \mathcal{L}$. A constraint set C over Φ in which every inequality is of the form $\tau \leq A$, with an atom on the right hand side, is called *definite*. A definite set $C = \{\tau_i \leq A_i\}_{i \in I}$ can be written as $C = C_{var} \cup C_{cst}$ where C_{var} are constraints in C that have a variable on the right-hand-side (rhs) and C_{cst} are constraints in C that have a constant on the rhs. The algorithm for Φ -SAT where C is a definite set of constraints is shown in Figure 3. A detailed explanation for the algorithm, including proofs of its tractability and completeness, is given by Rehof and Mogensen [26].

In this paper we always assume that Φ -terms τ have the form $\tau \equiv \beta_0 \sqcup \dots \sqcup \beta_k \sqcup L_1 \sqcup \dots \sqcup L_j$. As information-flow is a monotonically increasing property, terms that include the meet (\sqcap) do not make sense for our analysis. This assumption simplifies the presentation of our blame algorithm in the upcoming section.

3.2 Constraint Solving Example

Let $C = \{L_1 \leq \beta_0, L_2 \sqcup \beta_0 \leq \beta_1, \beta_0 \sqcup \beta_1 \leq \beta_2\}$, where security labels L_1 and L_2 are incomparable. When run on C , the solver will construct a valuation ρ such that $\mathcal{L}, \rho \models \rho(C)$. Initially, $\rho(\beta) = \perp$ for all variables β . The solver first considers $L_1 \leq \beta_0$. Because β_0 is currently mapped to

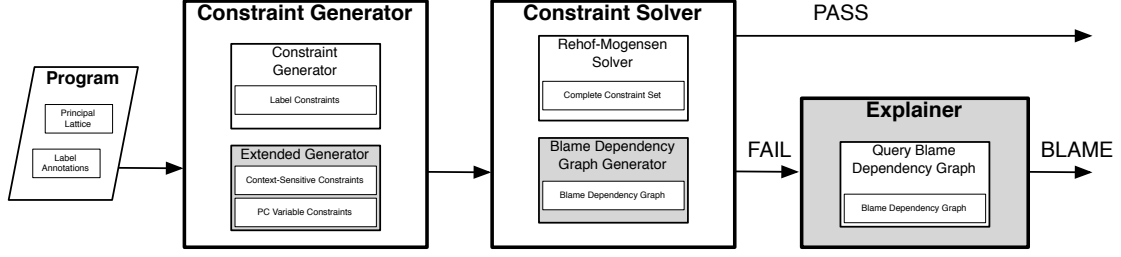


Figure 2: Modified Label Checking Process: (1) a partially-labeled program is input to a *constraint generator* that generates the label constraints from the code using an interprocedural analysis; (2) a *constraint solver* determines if there is a satisfying assignment for the label variables, building a *blame dependency graph* to determine the causes of a constraint failure; (3) if not, the *explainer* generates the security-relevant slice of code that witness security violations by repeatedly querying the blame dependency graph.

```

RMSOLVE( $C$ )
1   $\rho(\beta) \leftarrow \perp$  for all  $\beta \in \mathcal{V}$ 
2   $WL \leftarrow \{\tau \leq \beta \mid \tau \leq \beta \in C \text{ such that } \mathcal{L}, \rho \not\models \tau \leq \beta\}$ 
3  while  $WL$  is non-empty
4     $\tau \leq \beta \leftarrow \text{POP}(WL)$ 
5    if  $\mathcal{L}, \rho \not\models \tau \leq \beta$ 
6       $\rho(\beta) \leftarrow \rho(\beta) \sqcup \rho(\tau)$ 
7      for each  $\tau' \leq \alpha \in C$  with  $\beta \in \text{Vars}(\tau')$ 
8         $WL \leftarrow WL \cup \{\tau' \leq \alpha\}$ 
9    for each  $\tau \leq L \in C$ 
10     if  $\mathcal{L}, \rho \not\models \tau \leq L$ 
11       raise exception
12  return  $\rho$ 

```

Figure 3: The Rehof-Mogensen constraint solver

\perp under ρ , the solver raises $\rho(\beta_0)$ to L_1 . Next, the solver considers $L_2 \sqcup \beta_0 \leq \beta_1$. Because $\rho(\beta_1) = \perp$, the solver modifies $\rho(\beta_1) = L_2 \sqcup \rho(\beta_0) = L_1 \sqcup L_2$. Finally, the solver considers $\beta_0 \sqcup \beta_1 \leq \beta_2$; since $\rho(\beta_2) = \perp$, it sets $\rho(\beta_2) = \rho(\beta_0) \sqcup \rho(\beta_1) = L_1 \sqcup L_2$. With this ordering of constraints, the **for** loop in line 7 is not executed as its condition is always false. As there are no constraints of the form $\tau \leq L \in C$, the solver succeeds, returning ρ .

Suppose instead the constraint $\beta_0 \sqcup \beta_1 \leq \beta_2$ was considered first. Because initially all variables are mapped to \perp , the solver does not modify β_2 . However, after either β_0 or β_1 was modified, $\beta_0 \sqcup \beta_1 \leq \beta_2$ would be added to the worklist WL by line 8 and so β_2 would eventually be raised to $L_1 \sqcup L_2$.

For the constraint set $C' = C \cup \{\beta_2 \leq L_2\}$, the solver initially performs as described above. However, after each constraint of the form $\tau \leq \beta$ is considered, the solver attempts to check that $\beta_2 \leq L_2$ holds under ρ . However, as $\rho(\beta_2) = L_1 \sqcup L_2$ and $L_1 \sqcup L_2 \not\leq L_2$, the solver reports an error, specifically that it has failed on $\beta_2 \leq L_2$.

From the constraint set, it is not obvious why β_2 was raised above L_2 . We know that $\beta_0 \sqcup \beta_1 \leq \beta_2$ modified β_2 (as it is the only constraint with β_2 on the right-hand side), but knowing why β_2 was raised above L_2 involves knowing why either β_0 or β_1 were raised above L_2 . In larger constraint sets, determining which constraints caused an error can be even more difficult.

4. BLAME IN THE SOLVER

In this section, we present a structure for assessing blame at the solver level. This is a first step towards the blame

model described in Section 2. The *blame dependency graph* is a data structure that acts as a backend for failure explanation in the Rehof-Mogensen solver that can be queried to provide complete error explanations. We then present an algorithm that, for each error, reports a set of constraints X with the property that the constraints in X cause the error (the error reported is *complete*), and each constraint in X contributes to the error (the error reported is *minimal*).

4.1 The Blame Dependency Graph

As described in the previous section, the Rehof-Mogensen solver is used to determine if a Jif program has any illegal flows. In order to assign blame to information-flow errors in Java programs, we construct the blame dependency graph during the run of the Rehof-Mogensen solver. The intuition behind the dependency graph is to record the run of the solver by keeping track of which constraints $\tau \leq \beta$ raise the level of a variable β . If the constraint $\beta \leq L$ fails, meaning that β was raised too high, the solver can use this information and determine whether or not the particular constraint $\tau \leq \beta$ was responsible for this.

4.1.1 Definition

The blame dependency graph (hereafter referred to as the dependency graph) is a directed graph that records the history of the valuation produced by the solver. Let ρ_i be the valuation the i -th time that line 6 in Figure 5 is executed; we refer to this as *time i* . The dependency graph contains each of the ρ_i and information connecting each ρ_i with ρ_{i+1} .

Definition 4.1 (Blame Dependency Graph) *The blame dependency graph $\mathcal{B} = (V, E)$ is a graph that maintains the history of the valuations produced by a run of the solver. The dependency graph has the following properties:*

- If at time t , the valuation $\rho_t(\alpha) = L$, then the node $(\alpha, t, L) \in V$.
- For all t , if $(\alpha, t-1, L), (\alpha, t, L') \in V$, then the edge $((\alpha, t-1, L), (\alpha, t, L')) \in E$.
- If α is set to L' at time t because of a constraint $\tau \leq \alpha$, then for each $\beta \in \text{Vars}(\tau)$ and node $(\beta, t-1, L_\beta) \in V$, $((\beta, t-1, L), (\alpha, t, L')) \in E$, and this edge is labeled with the constraint $\tau \leq \alpha$.

A visualization of the dependency graph for the solver run over constraint set C from Section 3.2 is given in Figure 4. Initially each of the variables $\beta_0, \beta_1, \beta_2$ begin at \perp . The dependency graph indicates that the constraint $L_1 \leq \beta_0$ is




Figure 4: The dependency graph after a solver run

responsible for raising β_0 from \perp to L_1 . Next, it shows that β_1 was raised because of the constraint $L_2 \sqcup \beta_0 \leq \beta_1$. Finally, it identifies that the constraint $\beta_0 \sqcup \beta_1 \leq \beta_2$ was responsible for raising β_2 to $L_1 \sqcup L_2$.

4.1.2 Using the Blame Dependency Graph

As an abstract data type, the dependency graph supports two operations.

- $\text{RECORD}(t, \beta, L, \tau \leq \beta)$ records that the variable β has been modified at time t to the label L because of the constraint $\tau \leq \beta$.
- $\text{FAILURECAUSE}(\tau \leq L)$ returns the constraint $\tau' \leq \beta$ after which $\tau \leq L$ first became unsatisfiable under the valuation ρ constructed by the solver.

The operation $\text{RECORD}(t, \beta, L, \tau \leq \beta)$ adds the node (β, t, L) to the dependency graph, and adds an edge for each $\alpha \in \text{Vars}(\tau)$ from $(\alpha, t-1, \rho_{t-1}(\alpha))$ to (β, t, L) labeled with the constraint $\tau \leq \beta$. To implement $\text{FAILURECAUSE}(\tau \leq L)$, we trace through the dependency graph to find the time j at which $\rho_j(\tau) \leq L$, but $\rho_{j+1}(\tau) \not\leq L$. The variable β changed at time j was modified because of a constraint $\tau' \leq \beta$. FAILURECAUSE returns the constraint $\tau' \leq \beta$; at time j , $\rho_j(\tau) \leq L$ was true, while afterwards $\rho_{j+1}(\tau) \not\leq L$. The constraint $\tau \leq L$ fails because the variable β (occurring in τ) was raised to a security level above L by the constraint $\tau' \leq \beta$. This gives an explanation for the constraint's failure that can be determined from the blame dependency graph.

4.1.3 Generating the Blame Dependency Graph

To populate the dependency graph during the run of the RMSOLVE algorithm, after each execution of line 6 in the original solving algorithm, we record that at time t , β was modified from $\rho(\beta)$ to $\rho(\beta) \sqcup \rho(\tau)$ because of the equation $\tau \leq \beta$ and increment the current time. The modified solver algorithm is presented as $\text{RMSOLVE-DEPENDENCYGRAPH}$, shown in Figure 5.

4.2 Reporting the Cause of an Error

We described a procedure FAILURECAUSE that returns, for a single constraint $\tau \leq L$, the constraint $\tau' \leq \beta$ that caused it first to fail. Because of our assumption that labels contain only joins, this constraint $\tau' \leq \beta$ suffices to cause the failure of $\tau \leq L$. However, $\tau' \leq \beta$ may not provide every reason as to why the constraint $\tau \leq L$ failed. For example, τ' may contain variables, making it unclear why τ' itself was raised above L .

The algorithm RECURSIVEEXPLAIN , given in Figure 5, recursively explains why a constraint $\tau \leq L$ failed, returning

$\text{RMSOLVE-DEPENDENCYGRAPH}(C)$

```

1   $\rho(\beta) \leftarrow \perp$  for all  $\beta \in \mathcal{V}$ 
2   $WL \leftarrow \{\tau \leq \beta \mid \tau \leq \beta \in C \text{ such that } \mathcal{L}, \rho \not\leq \tau \leq \beta\}$ 
3   $t \leftarrow 0$ 
4  while  $WL$  is non-empty
5     $\tau \leq \beta \leftarrow \text{POP}(WL)$ 
6    if  $\mathcal{L}, \rho \not\leq \tau \leq \beta$ 
7       $\text{RECORD}(t, \beta, \rho(\beta) \sqcup \rho(\tau), \tau \leq \beta)$ 
8       $\rho(\beta) \leftarrow \rho(\beta) \sqcup \rho(\tau)$ 
9       $t \leftarrow t + 1$ 
10   for each  $\tau' \leq \alpha \in C$  with  $\beta \in \text{Vars}(\tau')$ 
11      $WL \leftarrow WL \cup \{\tau' \leq \alpha\}$ 
12 for each  $\tau \leq L \in C$ 
13   if  $\mathcal{L}, \rho \not\leq \tau \leq L$ 
14      $X \leftarrow \text{RECURSIVEEXPLAIN}(\rho(\tau) \leq L)$ 
15     raise exception “ $\tau \not\leq L$  because  $X$ ”
16 return  $\rho$ 

```

$\text{RECURSIVEEXPLAIN}(\tau \leq L)$

```

1  if  $\tau \leq L$  is unsatisfiable
2    return  $\emptyset$ 
3   $\tau' \leq \beta \leftarrow \text{FAILURECAUSE}(\tau \leq L)$ 
4  return  $\{\tau' \leq \beta\} \cup \text{RECURSIVEEXPLAIN}(\tau' \leq L)$ 

```

Figure 5: The Rehof-Mogensen Solver extended with error explanation

every constraint that contributed to its failure. When $\tau \leq L$ fails, RECURSIVEEXPLAIN consults the dependency graph to see why $\tau \leq L$ has failed, receiving the answer $\tau' \leq \beta$. Since $\tau \leq L$ failed because β was raised above $\rho(\tau')$, to determine why τ' was raised that high, the algorithm makes a recursive call to determine which constraints caused $\tau' \leq L$ to fail. This continues until the algorithm is called on an unsatisfiable constraint. In Section 4.3, we will see that this procedure can be optimized to run in $O(n)$ time, where n is the number of constraints given to the solver.

To concretely illustrate this procedure, we again revisit the example from Section 3.2 by examining the blame set returned by our explanation algorithm for the unsatisfiable set C' , which contains the extra constraint $\beta_2 \leq L$. When the solver is run on C' , the solver fails attempting to verify $\rho(\beta_2) \leq L_2$. Instead of reporting this constraint as the sole cause of the violation, RECURSIVEEXPLAIN searches for the constraint that first caused $\beta_2 \leq L_2$ to fail. The first point at which $\beta_2 \leq L_2$ fails is at time 3, when β_2 was raised to $L_1 \sqcup L_2$. The incoming edges to $(\beta_2, 3, L_1 \sqcup L_2)$ are marked with $\beta_0 \sqcup \beta_1 \leq \beta_2$, so we look for the time where $\beta_0 \sqcup \beta_1 \leq L_2$ first failed. This first failed at time 1, when $L_1 \leq \beta_0$ raised β_0 to L_1 . Finally, we attempt to explain $L_1 \leq \beta_0$, which is unsatisfiable. The algorithm then reports: “ $\beta_2 \not\leq L_2$ because $\{\beta_0 \sqcup \beta_1 \leq \beta_2, L_1 \leq \beta_1\}$ ”. These are the constraints that caused $\beta_2 \leq L_2$ to fail.

4.3 Error Reporting Properties

In this section, we describe some important properties of RECURSIVEEXPLAIN : it is guaranteed to terminate, and the error sets returned by RECURSIVEEXPLAIN are complete and minimal, i.e. they actually witness an error and do not contain any smaller explanations. All of these proofs are given in our companion technical report [21].

Our recursive explanation algorithm, when called, will return in a finite amount of time. The main proof idea is that

for each constraint $\tau \leq \beta$ is added to the returned error set, the time t at which constraint $\tau \leq L$ first became unsatisfiable is always decreasing. Together with the linearity of the Rehof-Mogensen solver, this implies RECURSIVEEXPLAIN can be implemented so that it takes time $O(n)$, where n is the size of the constraint set.

Lemma 4.2 RECURSIVEEXPLAIN *always terminates.*

Our main result is that RECURSIVEEXPLAIN accurately identifies a cause of an error without displaying useless constraints. We first define what it means for a set of constraints to be *unsatisfiable*.

Definition 4.3 A set C of constraints is *unsatisfiable* under security lattice \mathcal{L} if there does not exist a valuation ρ such that $\mathcal{L}, \rho \models C$.

An error set for $\tau \leq L$ is a set of constraints X that, together with $\tau \leq L$, is unsatisfiable. If X is an error set for $\tau \leq L$, then it is an explanation of $\tau \leq L$ that contains every constraint that caused an error: the error set X is complete.

Definition 4.4 (Completeness) Let C be a set of constraints and $X \subseteq C$. We say X is an error set for a constraint $\tau \leq L$ if $X \cup \{\tau \leq L\}$ is unsatisfiable.

We are interested in returning *minimal* error sets: these are error sets that do not contain any smaller error set. Minimal error sets are small witnesses to a specific information flow violation.

Definition 4.5 (Minimality) Let X be an error set for the constraint $\tau \leq L$. We say X is a *minimal error set* if there is no $X' \subset X$ such that X' is an error set for $\tau \leq L$.

The main property of our error reporting algorithm is its completeness and minimality: when our explanation algorithm is run on $\tau \leq L$, the returned set X will be a witness for the reason why $\tau \leq L$ failed, and it will not contain any smaller witness X' . Completeness follows from the fact that if $X \cup \{\tau \leq L\}$ is unsatisfiable, then $X \cup \tau \leq \beta, \beta \leq L$ is also unsatisfiable. Minimality follows by a specific structure of error sets: for all β , there is at most one τ such that $\tau \leq \beta \in X$.

Theorem 4.6 (Completeness and Minimality) Let $\text{RECURSIVEEXPLAIN}(\tau \leq L) = X$. Then X is a minimal error set for $\tau \leq L$.

PROOF. The proof is given in our companion technical report [21]. \square

The problem of finding the *minimum error set* X_{min} for a constraint set C is computationally difficult; the reduction follows from VERTEX-COVER.

Theorem 4.7 For a constraint set C , the problem of finding an $X \subseteq C$ such that X is an error set and for all other error sets X' , $|X| \leq |X'|$, is NP-Hard.

4.4 Error Traces From Code

Figure 6 shows the error traces returned by our blame algorithm for two of the errors as described in Figure 1. In them, a semicolon (;) represents the join of two labels, {Secret} and {Public} are explicit security labels, and everything else is a label variable. Variables beginning with pc

```
{message@callto:write:2; pc1} <= {Public}
1: {message@callto:write:2} == {inst(MESSAGE_NO_SUCH_MESSAGE);
    write_receiver:2; pc33}
2: {pc33} ==_{def} {isDeleted:value_returned; pc32}
3: {pc32} ==_{def} {getMessage:return_observed; pc31}
   {getMessage:return_observed} ==_{def} {Secret}
```

```
{message@callto:write:2; pc1} <= {Public}
1: {message@callto:write:2} == {currentLine; pc22}
2: {readLine:value_returned; pc18} <= {currentLine}
3: {readLine:value_returned} ==_{def} {L_var@1782}
   {L_var@1782} == {Secret}
```

Figure 6: Error sets returned by Recursive Explain for the code in Figure 1. Numbered constraints were generated by the correspondingly numbered lines of code in Figure 7.

represents the program counter at a specific program point: these are discussed more in Section 5.3. In both cases, the failed constraint is the same: however, the reason is different. Our tool first returned the topmost error trace, indicating the implicit flow from whether a specific user’s message was deleted. When we inserted a declassifier to allow this information flow, our analysis returned the second error trace, showing the explicit flow from the BufferedReader that operated on the message from the user. The variable `L_var@BufferedReader` is the variable representing the security of the BufferedReader returned by the `getReader` method.

As understanding what is happening only from constraints can be difficult, we present the information from the error set to the user by highlighting the expressions that created each constraint. Figure 7 shows two traces that our tool returns when run over the example code. In our figures, constraints that correspond with a line of code are indicated by the same number; so in the first error trace, the constraint 2: `{pc33} ==_{def} {isDeleted:value_returned; pc32}` was originally generated by the code expression `user.getMessage(msgNum).isDeleted()`.

By looking at each line of the error trace along with the error set, we can see the exact cause of the two reported information flow errors. For the first trace, the write out constraint fails because of an implicit flow from observing the return of the method `getMessage`, a method invocation performed on a `Secret` data structure. This implicit flow affects the value returned by `isDeleted`, which affects whether or not the `write` method is called. The second trace is caused by an explicit flow: the data written out is equal to `currentLine`, which is retrieved from data stored in the instance of `BufferedReader` (`L_var@1782`); this is explicitly set equal to `Secret` by the annotation on line 1.

5. BLAME IN PROGRAM CODE

The Rehof-Mogensen solver together with the dependency graph can give explanations for why constraints failed, independent of determining errors in Java code. We first show the type of error traces that our tool returns from Java code, and then detail modifications that we needed to make to the Jif constraint generator in order to effectively use our blame algorithm to find security errors in Java code, as described in Section 6.

5.1 Building the Blame Dependency Graph With The Jif Solver

The information-flow constraints generated by the Jif engine have a different form than Rehof-Mogensen constraints;

```

failed:
  {message@callto:write:2; pc1} <= {Public}
failure site:
  out.println(message);
1: if (user.getMessage(msgNum).isDeleted()) { ... }
   ~~~~~
2: user.getMessage(msgNum).isDeleted()
   ~~~~~
3: user.getMessage(msgNum)
   ~~~~~
why:
  {getMessage:return_observed} ==_{def} {Secret: }

```

```

failed:
  {message@callto:write:2; pc1} <= {}
failure site:
  out.println(message);
1: this.write(currentLine);
   ~~~~~
2: String currentLine = fileIn.readLine();
   ~~~~~
3: BufferedReader fileIn = user.getReader(msgNum);
   ~~~~~
why:
  {L_var@BufferedReader} == {Secret: }

```

Figure 7: Error traces from the code in Figure 1. Numbered lines correspond to the correspondingly numbered constraints from Figure 6.

a constraint c need not be definite, and so the right-hand side of its inequality may be a join of multiple labels [23]. The current implementation of the Jif solver solves constraints of this form through an implementation of the Rehof-Mogensen solver that uses backtracking when it needs to satisfy constraints with more than one component on the right-hand side. This backtracking is done only rarely: most of the constraints that Jif generates are definite. Only one of our code examples required backtracking; the analysis of the JES Email Server attempted to perform backtracking on 21 of its constraints (out of a total of 9539 constraints). None of the other examples required backtracking.

The key property of the Rehof-Mogensen solver that we use in our definitions of the blame dependency graph is its monotonicity: if a variable α is assigned to label L , then α will not later be assigned to a label L' such that $L \not\leq L'$. While the Jif compiler uses an implementation of an extended version of the Rehof-Mogensen solver, it is still monotonic; we can therefore use the blame dependency graph to find errors using the Jif solver. The only modification to the method described in Section 4 is that, when the Jif solver performs a backtracking step, a separate copy of the blame dependency graph is given to each recursive call to the backtracker. In the event that a backtracking call cannot satisfy a constraint without causing another constraint to fail, these copies can be used to provide an error explanation. Our experience has been that backtracking is performed very rarely, and so the overhead associated with this step will be minimal.

5.2 Interprocedural Label Checking

To focus on resolving security conflicts in Java code without annotating every formal argument and side effect bound to each method, we modified the Jif compiler to perform *interprocedural label checking*. Procedures whose arguments are tagged with security labels are checked normally. Otherwise, the constraints for a procedure are inferred using

method summaries, a standard technique in static analysis [28]. During label checking, we assign each procedure a list of constraints that must be satisfied by the arguments to the procedure for the code to be information-flow secure. Unlabeled methods are assigned *summary variables*, which represent security labels that are not known. If a procedure p has summary variables v_1, \dots, v_k and summary constraints C_p , then the call of p with actual label arguments a_1, \dots, a_k generates the summary constraints $C_p[a_1/v_1, \dots, a_k/v_k]$, where a_1/v_1 represents the substitution of a_1 for all instances of v_1 . We can only label-check a procedure if its constraints have no unbound summary variables.

For example, from the code from Figure 1, the function `void write(String message)` has following summary constraints:

```

{pc1} ==_{def} {sum(caller_pc@write)}
{sum(message); pc1} <= {Public}
{println:return_observed} ==_{def} {Public}
{println:value_returned} ==_{def} {Public}
{receiver:10} == {inst(out); this; pc1}
{arg1label:5} == {sum(message); pc1}

```

In the above, the variables `sum(message)` and `sum(caller_pc@write)` are the summary variables representing the security of the `message` argument and the side effects revealed by `write`. Observe that, as `Public` is the bottom of our security lattice, the constraints `{pc1} ==_{def} {sum(caller_pc@write)}` and `{sum(message); pc1} <= {Public}` force both of these security values to be `Public`. The special variable `inst(out)` represents the security value of the field `out`: this is inferred through whole-program analysis. The label variables `receiver:10` and `arg1label:5` represent the security value of the receiver (`out`) and argument (`message`) of the `println` call.

Because procedures may be recursive or mutually recursive, we generate summary constraints for each strongly-connected component of a program’s call-graph. Summary constraints for recursive procedures are generated by taking the fixed point of each strongly-connected component. This fixed point will always exist because the number of recursive calls within a strongly-connected component is finite and the join operator is idempotent.

In our experiments, we found that a context-insensitive interprocedural approach worked fairly well on a variety of Java programs. Our analysis reported several false positives arising from context-insensitive in cases where a program variable was treated with having two different security levels. For example, if a method is used to return both high-security and low-security values, our tool will report an error. We found only a few false positives when examining the program code: we found 2 false positives in JES, 8 in `tinySQL`, and none in the other three applications that we evaluated.

When we detected a false positive that occurred because of the context insensitivity of our analysis, we were able to resolve it by adding a *label parameter*, a feature of the underlying Jif language, to the enclosing class. A label parameter allows for variables in a class to be given a label based on an immutable label assigned to the enclosing instance of the class; this allows the class to be treated in a context-sensitive fashion. For example, the Java Card Purse (described in Section 6) used a special utility `Decimal` class to perform arithmetic operations. Some instances of

`Decimal` contained `Tainted` information (low-integrity), while some contained `Untainted` information (high-integrity). We parameterized `Decimal` with a label parameter, changing the definition to `Decimal[label L]`, and then annotated each field of `Decimal` as having security level `L`. We then annotated instances of `Decimal` that contained `Untainted` data as `Decimal[{Untainted}]`. Our analysis then automatically inferred the labels required for every other instance of `Decimal`. To aid programmers in the process of annotating programs, We are investigating ways to automatically determine which classes and methods may cause context-insensitive false positives.

5.3 Program Counter Variables

The dependency graph can trace the interactions between variables to explain why a label constraint of the form $\tau \leq L$ failed, assuming that such a constraint is satisfiable in the first place. However, under the standard implementation of the program counter, illegal implicit flows can generate unsatisfiable constraints that cannot be explained using the dependency graph.

For example, if the program counter is `{Bob}` before an assignment to an `{Alice}` variable, then the assignment statement will cause the constraint `{Bob} \sqcup $\beta_{data} \leq$ {Alice}` to be generated; here β_{data} is a label variable for expression on the right hand side of an assignment statement. This constraint cannot be satisfied as `{Bob}` and `{Alice}` security levels are incomparable, and so the value of β_{data} is irrelevant. Without a technique to explain why the program counter was set to `{Bob}`, we cannot explain error in a satisfactory way.

To better explain implicit flows with the dependency graph, we introduce temporary variables indicating when the label associated with the program counter has changed, called *program counter variables*. These variables are a layer of indirection that allow the dependency graph to trace errors across program counter changes. Whenever the program counter is changed from pc to L , a fresh variable α_{pc} is introduced, a definitional equality constraint $\alpha_{pc} == pc \sqcup L$ is created, and the program counter is set to α_{pc} . If a constraint generated by code checked under this program counter fails because α_{pc} is raised too high, there is then a clear path back to where the program counter was set.

This process is similar to converting a program to SSA form [7]. We found that using program counter variables greatly improved our error explanations at the cost of increasing their length.

6. EVALUATION

We evaluated our blame algorithm on several codebases, and show that it aids in the process of finding and resolving errors in Java programs. We compared the source-to-sink error traces returned from our tool to a backwards slice from that sink. We used the WALA libraries for program analysis [1] to create program slices. Our experiments demonstrate the following:

- The errors returned by our information-flow blame framework were much smaller than backwards slices on the violating sinks.
- Each of the error sets returned contained a reasonable target site to resolve the illegal information flow.

Our analysis ran on an Intel Core 2 Duo running at 2.20 GHz with 2 GB of memory. Though we ran our machine

Application	LOC	Constraint Set Size	Time (s)
Java Card Wallet	296	237	0.60
Mental Poker	1499	6021	5.12
JES Email Server	2595	9539	6.99
Java Card Purse	5581	20924	36.63
tinySQL	8240	23518	102.71

Table 1: A table summarizing the runtime behavior of our analysis. Column 1 contains the name of the application. Column 2 contains the size of the application in source-lines-of-code, (the lines of code without whitespace). Column 3 contains the total number of constraints generated by the program, while Column 4 contains the time for constraint generation.

on a multiprocessor system, our analysis is presently single threaded. Statistics about the total number of generated constraints and the running time of our analysis are given in Table 1.

The goal of our experiments was to demonstrate that the error sets returned by our tool contained an expression that caused of the error (completeness) and were small enough to easily find this expression. A complete catalogue of the source-to-sink violations found along with candidate resolutions is available at our website <http://www.cse.psu.edu/~dhking/jlift/>.

6.1 Java Card Wallet

The Java Card Wallet is a small example designed to teach Java programmers how to program in Java Card³. It represents a Java Card program that contains a balance that can be credited, debited, or retrieved. To protect the card from unauthorized tampering, the application stores the wallet’s PIN in a member field `OwnerPIN pin`. We labeled field `OwnerPIN pin` as having secret data and labeled the `APDU`, a Java Card data structure representing input and output, as having public data. We also labeled the security data revealed by termination of the main `process` method as `Public`.

We found eight total information flows from source to sink, some of which were duplicates. We found three violations: the `APDU` would be written to only if the PIN was validated (two instances), and failing to verify the PIN would throw an exception that was visible to the user. These errors were resolved by adding declassifiers around the calls to checking and verifying the PIN of the Wallet. In many of our other experiments, we often found that one declassifier would resolve multiple errors.

The backwards slices computed by Wallet were particularly small. The Wallet source code represents a best-case scenario for backwards slices: there were only three paths through the program, and each represented an information-flow violation.

6.2 Mental Poker

Mental Poker was one of the original implementations of a non-trivial application in a security-typed language [2]. We ran our analysis tool on a prototype Java implementation provided by the authors. Most classes in the implementation were designed to store either `Secret` data or `Public` data, making manual annotation of the field data in the classes simple. We found 8 separate resolution points in the pro-

³<http://developers.sun.com/mobility/javacard/articles/intro/>

gram, each corresponding with a cryptographic operation associated with the Mental Poker cryptographic protocol.

6.3 Java Card Purse

The PACAP Purse is a prototype designed to secure information flow in Java Card applications [4]. We added integrity labels to its source code: input from the user was marked as `Tainted`, while state kept in the Purse was marked as `Untainted`. Our analysis tool returned 110 errors, as each member of the `Purse` class was marked as being an `Untainted` container, meaning that each code location that modified a field of `Purse` was treated as a separate sink. However, we found that all of the errors could be resolved by adding six declassifiers at places in the code that corresponded to an existing card verification mechanism. These resolution sites were either a call to verify a PIN or a call to an access control table with one of five different arguments, corresponding to the operation that was about to be performed (initialize debit, initialize credit, initialize exchange, initialize PIN verification, or initialize administrative mode).

6.4 JES Email Server

As described in Section 2, JES is an open-source Java email server with full POP3 and SMTP functionality. Its functionality is a superset of JPMail [18], an information-flow secure email server that was manually developed in Jif. To determine how the data stored by the JES server interacted with the user's inputs and outputs, we labeled user data as `Secret` (confidential and high integrity) and labeled the input and output sockets as containing `Tainted` (public and low integrity) data. `Tainted` data should not be allowed to influence the user configuration file without being sanitized, and `Secret` data should not be released outside the system without being declassified.

We found that five distinct program points created errors in JES, and that these error sites were exclusively either integrity violations or confidentiality violations; no site caused both kinds of violations. In total, we found that we needed to insert 66 (51 confidentiality violations, 15 integrity violations) different resolutions to remove the information-flow errors in the mailbox. Most of the confidentiality violations corresponded to expected behavior of the mail server. For example, the POP implementation of the mail server sent information about the size of a user's mailbox in response to a STAT command. With respect to integrity violations, we found that the application already contained functions to parse and sanitize input; once these were marked appropriately, there were only a few other points in the program that caused an integrity violation. The process of manually inserting these fixes took a few hours.

6.5 tinySQL

tinySQL is a minimal implementation of an SQL client and server on DBF and text files⁴. We marked `System.out` as a `Public` output stream and marked the contents of a SQL table as `Secret`. We found that we needed to manually resolve 30 separate information-flow violations. The major difference between tinySQL and JES was that tinySQL did not separate logging output from normal system output: a debugging flag being enabled caused secret data to be sent to the screen, which we had labeled as `Public`. Using an automatic find/replace, we changed over 60 calls to `System.out` behind debugging conditionals to calls to a logger, and then

⁴<http://www.jepstone.net/tinySQL/>

manually resolved the remainder. We found that the errors corresponded to expected system behavior, with the exception of several unnecessary `System.out` calls in the code revealing data about unnecessary table structure during a query.

6.6 Evaluation

A summary of the error traces returned by our system is presented in Table 2. We computed backwards slices using the context insensitive slicer provided by the WALA program analysis libraries. The statistics given for error set size are the statistics that our tool gave during an initial run of our program, rather than during subsequent runs after fixes were applied. The largest error set sizes encountered during the error resolution process were: Wallet: 13, Purse: 18, JES: 57, Mental Poker: 14, tinySQL: 35.

We found that the presented error traces were much smaller than a backwards program slice. A one-to-one comparison between error trace size and backwards program slice size is difficult because a typical error trace only contains information corresponding to only one illegal program path, meaning that more than one error may have to be investigated for a given sink. The major benefits of using our blame analysis over a program slicers for resolving information-flow errors are: (1) the analysis does not report runtime program paths that do not cause an error, (2) the analysis allows a programmer to step through each information-flow error individually, rather than having to sort through the entire program slice to determine what caused an error, and (3) a type-based analysis for information flow will contain less statements than a value-based analysis such as program slicing.

In most cases there was a high degree of overlap between most error traces, meaning that often the whole error trace did not need to be inspected due to familiarity with past error traces. A common situation was that each error for a `Public` sink included a different reason (often within the same Java class) for `Secret` information tainting that sink, and then a common path back to a `Secret` sink. This corresponds with the expected use of program slices in practice [27].

Due to our recursive approach of blaming the *first* modification that caused a constraint to become unsatisfiable, our tool only reports one error for each constraint associated with a sink. Once that error is resolved, it may report another error for that sink, meaning that there was a different program path that caused a violation. Sometimes a resolution inserted to fix one error fixed many others; this corresponded to the situation where many sinks were hidden behind a common *authorization hook*, such as encryption (Mental Poker) or PIN/password verification (Java Card Purse). To free programmers from having to manually evaluate each error trace, We are working on a system for automatically suggesting candidate fixes.

6.7 Limitations

Our source code analysis has a few limitations; most are common to all static analysis techniques. We must have the source code to analyze a program. In particular, to analyze calls to a library, we must either have the source code for the library or make a simplifying assumption about the security behavior of a library. The Jif language handles library calls by relying on *class signatures*, which give an explicit security labeling for each call to an external method (these can be

Application	# Failed Constraints	# Fixes Required	Avg. Size of Error Set (# constraints)	Avg. Size of Error Set (# lines)	Avg. Backwards Slice Size (# bytecode instructions)
Java Card Wallet	8	3	10	9.88	31
Mental Poker	65	8	9.86	6.91	397.69
JES Email Server	5	66	19	13.40	356.20
Java Card Purse	110	6	16.36	15.49	1416.93
tinySQL	120	30	28.61	19.37	312.56

Table 2: A table containing information about how our error traces compare to other tools for determining information-flow errors. Column 1 contains the number of failed constraints in each program (the number that would be reported by Jif). Column 2 contains the number of resolutions, specific to each application, required for an application to be information-flow secure (performed by hand based on our analysis). Column 3 and 4 contains the average size of an error set returned by our tool on the initial run over each program, first in the number of constraints and second in the number of lines of program code. Column 5 contains the average size of a backwards program slice as computed by the WALA library.

thought of as a security annotation on header files). To automatically assign a conservative security policy to external library calls, we developed a tool for automatically generating Jif class signatures from Java source⁵. By default, external libraries are treated as containing security data at some immutable level L : all data that enters and leaves the library must be of security level L . This conservative labeling prevents programs from laundering data through mutable data structures in libraries. However, these files may still require manual annotation for libraries with system-specific security behavior (for example, all Sockets in an application might be treated as outputting public data). The number and size of the automatically generated signature files scales with the size of application being analyzed. The Wallet Java Card signature files consisted of 150 lines of code (disregarding comments), while the tinySQL signature files required 712 lines of code.

At present, our analysis framework does not handle dataflows as caused by threads or reflection. For the applications that we analyzed that used threads (most notably the JES email server), we replaced calls of the form `new Thread(ThreadClass.class).run()` with explicit calls to `new ThreadClass().run()`: for the purposes of our label analysis, these two calls have identical security behavior. The applications that we surveyed used Java’s reflection capacities in a very limited way; the most common method was using the Apache logger `log4j`, which is initialized with a `java.lang.Class` object.

While our tool represents a considerable step forward in finding and resolving information-flow errors in program code, these errors must still be resolved manually. We are currently implementing a system where resolutions are suggested automatically, allowing the programmer to accept or reject resolutions as necessary.

7. RELATED WORK

In this section, we describe some related work in more depth.

7.1 Determining Error Causes

There is a large amount of work towards explaining the cause of type inference errors in ML-like languages. Our work follows a common theme of adding information to the compiler in order to produce more helpful error messages [32]. Recent work explicitly views finding error messages as a search problem [22]. The major difference between

error messages in ML and our work is that a type error in ML may have a syntactic fix, while resolving an information-flow error requires a semantic fix.

Efficient generation of counterexamples and finding the root cause of a counterexample has also been studied by the model checking community. Explaining counterexamples by keeping auxiliary information during fixed-point computations has been investigated in model checking [12]. Finding root-causes of counterexample traces has been used in software verification for abstraction refinement [3]. Groce *et al.* [15] explore a semi-automatic method for determining whether an error trace from a model checker represents a real fault or a false positive from the imprecision of the static analysis.

7.2 Determining Causes for Information-Flow Errors

Hammer *et al.* [16] view detecting illegal information flow as a path traversal problem in the program dependency graph. Errors detected using this method will have clear error messages by observing the path from high information to low information. However, their analysis takes much longer to complete (up to 40 times longer when computing context-sensitive slices) because of a need to compute path-sensitive conditions for conditionals, while their computed relevant slices are much larger.

Deng and Smith have presented a method for security error explanations in a simple language featuring `while` loops and arrays [8]. Instead of generating constraints for information-flows, they use a customized solver algorithm that records the histories of which variables influenced the modification of another variable. When a program fails to type, the history of all of the variables that were involved in determining the type of the broken expression is recursively reported. While this gives good error messages in practice, such explanations may not be minimal. Additionally, as their approach is language-specific, providing error reports for different languages will require expanding and modifying their solver algorithm.

7.3 Type Qualifiers

Our tool has similar a goal to that of the CQual tool, which refines the C type system with qualifiers on types [13], allowing programmers to find inconsistencies between program points. CQual has been used to find bugs in the Linux kernel [20, 34]. A user of CQual annotates code with additional annotations, such as `kernel_ptr` or `user_ptr`, and applies a constraint-based analysis to the source code to de-

⁵<http://www.cse.psu.edu/~dhking/siggen>.

termine inconsistencies between the annotations, where, for example, a kernel pointer is inadvertently passed to user code. However, CQual does not handle implicit flows and so can only be used to find explicit information flows. In the event of an error, CQual can provide information as to the cause of a type error, which are similar to our error traces. Johnson and Wagner [20] provide heuristics for sorting and pruning error traces before displaying them to the user: we believe our error reporting system could benefit from this technique.

8. CONCLUSION

In this paper, we presented a blame model for finding the cause of information-flow errors in program code. The model extends the Jif compiler by adding an explanation procedure that returns a *complete* and *minimal* error set for a failing constraint. The explanation procedure uses a *blame dependency graph* that tracks relationships between constraints during the run of a solver. This is a general utility for explaining errors generated by the Rehof-Mogensen solver and can be extended to explain errors in other security-typed languages that also use a similar constraint solver.

9. REFERENCES

- [1] T.J. Watson Libraries for Analysis. <http://wala.sourceforge.net>.
- [2] ASKAROV, A., AND SABELFELD, A. Secure implementation of cryptographic protocols: A case study of mutual distrust. In *ESORICS '05*.
- [3] BALL, T., NAIK, M., AND RAJAMANI, S. K. From symptom to cause: localizing errors in counterexample traces. In *POPL '03*, pp. 97–105.
- [4] BIEBER, P., CAZIN, J., MAROUANI, A. E., GIRARD, P., LANET, J.-L., WIELS, V., AND ZANON, G. The PACAP prototype: A tool for detecting Java Card illegal flow. *Java Card Workshop (2000)*, 25–37.
- [5] BLANCHET, B., COUSOT, P., COUSOT, R., FERET, J., MAUBORGNE, L., MINÉ, A., MONNIAUX, D., AND RIVAL, X. A static analyzer for large safety-critical software. In *PLDI (2003)*.
- [6] CHONG, S., VIKRAM, K., AND MYERS, A. C. Sif: Enforcing confidentiality and integrity in web applications. In *USENIX Security (2007)*.
- [7] CYTRON, R., FERRANTE, J., ROSEN, B. K., WEGMAN, M. N., AND ZADECK, F. K. Efficiently computing static single assignment form and the control dependence graph. *TOPLAS 13*, 4 (Oct 1991), 451–490.
- [8] DENG, Z., AND SMITH, G. Type inference and informative error reporting for secure information flow. In *ACM-SE 44 (New York, NY, USA, 2006)*.
- [9] DENNING, D. E. A lattice model of secure information flow. *Commun. ACM 19*, 5 (1976), 236–243.
- [10] DENNING, D. E., AND DENNING, P. J. Certification of programs for secure information flow. *Commun. ACM 20*, 7 (1977), 504–513.
- [11] DEVANBU, P. T., AND STUBBLEBINE, S. G. Software engineering for security: a roadmap. In *ICSE - Future of SE Track (2000)*, pp. 227–239.
- [12] E.M. CLARKE, O. GRUMBERG, K.L. McMILLAN, AND X. ZHAO. Efficient Generation of Counterexamples and Witnesses in Symbolic Model Checking. In *DAC '95 (San Francisco, CA, USA, 1995)*, pp. 427–432.
- [13] FOSTER, J. S., FÄHNDRICH, M., AND AIKEN, A. A theory of type qualifiers. In *PLDI (1999)*.
- [14] GREENFIELDBOYCE, D., AND FOSTER, J. S. Type qualifier inference for Java. In *OOPSLA (2007)*.
- [15] GROCE, A., CHAKI, S., KROENING, D., AND STRICHMAN, O. Error explanation with distance metrics. *STTT 8*, 3 (2006), 229–247.
- [16] HAMMER, C., KRINKE, J., AND SNELTING, G. Information flow control for Java based on path conditions in dependence graphs. In *Proceedings of the IEEE International Symposium on Secure Software Engineering (2006)*.
- [17] HANGAL, S., AND LAM, M. S. Tracking down software bugs using automatic anomaly detection. In *ICSE (2002)*.
- [18] HICKS, B., AHMADIZADEH, K., AND MCDANIEL, P. In *ACSAC '06*, IEEE Computer Society.
- [19] HORWITZ, S., REPS, T., AND BINKLEY, D. Interprocedural slicing using dependence graphs. *ACM Trans. Program. Lang. Syst. 12*, 1 (1990), 26–60.
- [20] JOHNSON, R., AND WAGNER, D. Finding user/kernel pointer bugs with type inference. In *USENIX (2004)*.
- [21] KING, D., JAEGER, T., JHA, S., AND SESHIA, S. A. Effective blame for information flow violations. Tech. Rep. NAS-TR-0069-2007 (Updated March 2008), The Pennsylvania State University, 2008.
- [22] LERNER, B., FLOWER, M., GROSSMAN, D., AND CHAMBERS, C. Searching for type-error messages. In *PLDI '07*.
- [23] MYERS, A. C. JFlow: Practical mostly-static information flow control. In *POPL '99*, pp. 228–241.
- [24] MYERS, A. C., NYSTROM, N., ZHENG, L., AND ZDANCEWIC, S. Jif: Java + Information Flow. <http://www.cs.cornell.edu/jif>.
- [25] POTTIER, F., AND SIMONET, V. Information flow inference for ML. In *POPL '02*, pp. 319–330.
- [26] REHOF, J., AND MOGENSEN, T. A. Tractable constraints in finite semilattices. *Science of Computer Programming 35*, 2–3 (1999), 191–221.
- [27] RENIERIS, M., AND REISS, S. P. Fault localization with nearest neighbor queries. In *ASE (2003)*.
- [28] SHARIR, M., AND PNUELI, A. Two approaches to interprocedural dataflow analysis. In *Program Flow Analysis: Theory and Applications (1981)*, Prentice Hall, pp. 189–234.
- [29] SMITH, S. F., AND THOBER, M. Refactoring programs to secure information flows. In *PLAS (2006)*.
- [30] SUN, Q., BANERJEE, A., AND NAUMANN, D. Modular and constraint-based information flow inference for an object-oriented language. In *SAS (2004)*.
- [31] TIP, F. A survey of program slicing techniques. *Journal of programming languages 3 (1995)*, 121–189.
- [32] WAND, M. Finding the source of type errors. In *POPL '86 (1986)*, ACM Press, pp. 38–43.
- [33] WEISER, M. Program slicing. In *ICSE (1981)*.
- [34] ZHANG, X., EDWARDS, A., AND JAEGER, T. Using equal for static analysis of authorization hook placement. In *USENIX (2002)*.