We are Family:
Relating Information-Flow Trackers

Musard Balliu, Daniel Schoepe, and Andrei Sabelfeld

Chalmers University of Technology, Gothenburg, Sweden

Abstract. While information-flow security is a well-established area, there is an unsettling gap between heavyweight *information-flow control*, with formal guarantees yet limited practical impact, and lightweight *tainting* techniques, useful for bug finding yet lacking formal assurance. This paper proposes a framework for exploring the middle ground in the range of enforcement from tainting (tracking data flows only) to fully-fledged information-flow control (tracking both data and control flows). We formally illustrate the trade-offs between the soundness and permissiveness that the framework allows to achieve. The framework is deployed in a staged fashion, statically embedding a dynamic monitor, being parametric in security policies, as they do not need to be fixed until the final deployment. This flexibility facilitates a secure app store architecture, where the static stage of verification is performed by the app store and the dynamic stage is deployed on the client. To illustrate the practicality of the framework, we implement our approach for a core of Java and evaluate it on a use case with enforcing privacy policies in the Android setting. We also show how a state-of-the-art dynamic monitor for JavaScript can be easily adapted to implement our approach.

Keywords: Language-based security · Information-flow control · Taint Tracking

1 Introduction

Motivation The sheer bulk of sensitive information that software manipulates makes security a major concern. A recent report shows that several of the top 10 most popular flashlight apps on the Google Play store may send sensitive information such as pictures and video, users' location, and the list of contacts, to untrusted servers [49]. Unfortunately, trusted code also incurs serious security flaws, as proven by the Heartbleed bug [51] found in the OpenSSL library.

Information-flow control [44] offers an appealing approach to security assurance by design. It helps tracking the flow of information from confidential/untrusted sources to public/trusted sinks, ensuring, for *confidentiality*, that confidential inputs are not leaked to public outputs, and, for *integrity*, that untrusted inputs do not affect trusted outputs.

Background Applications can leak information through programming-language constructs, giving rise to two basic types of information flows: *explicit* and *implicit* flows [21]. Consider a setting with variables *secret* and *public* for storing confidential (or *high*) and public (or *low*) information, respectively. Explicit flows

occur whenever sensitive information is passed explicitly by an assignment, e.g., as in $public := secret$. Implicit flows arise via control-flow structures of programs, e.g. conditionals and loops, as in **if** $secret$ **then** $public := 0$ **else** $public := 1$. The final value of $public$ depends on the initial value of $secret$ because of a *low assignment*, i.e., assignment to a low variable, made in a *high context*, i.e., branch of a conditional with a secret guard.

Information-flow control is typically categorized as static and dynamic: (1) *Static* techniques mainly impose Dennings' approach [21] by assigning security labels to input data, e.g. variables, APIs, and ensuring separation between secret and public computation, essentially by maintaining the invariant that no low assignment [56,44,32] occurs in a high context. Other static techniques include program logics [10,13], model checking [8,23], abstract interpretations [27] and theorem proving [20,40]. However, static techniques face *precision* (high false-positive rate) challenges, rejecting many secure programs. These challenges include *dynamic code evaluation* and *aliasing*, as illustrated by the snippet $x.f := 0$; $y.f := secret$; $\mathbf{out}(\mathbf{L}, x.f)$. A non-trivial static analysis would have to approximate whether object references x and y are aliases. Moreover, the fact that security policies are to be known at verification time makes them less suitable in dynamic contexts. (2) *Dynamic* techniques use program runtime information to track information flows [26,5,43]. The execution of the analyzed program is monitored for security violations. Broadly, the monitor enforces the invariant that no assignment from high to low variables occurs either explicitly or implicitly. Dynamic techniques are particularly useful in highly dynamic contexts and policies, where the code is often unknown until runtime. However, since the underlying semantic condition, *noninterference* [28], is not a trace property [38], dynamic techniques face challenges with branches not taken by the current execution. Consider the secure program that manipulates *location* information: **if** $(MIN \leq loc)$ && $(loc \leq MAX)$ **then** $tmp := loc$ **else** **skip**. If the user's (secret) location loc is within an area bound by constants MIN and MAX, the program stores the exact location in a temporary variable tmp, without ever sending it to a public observer. A dynamic analysis, e.g. No-Sensitive Upgrade [58,5], incorrectly rejects the program (due to a security label upgrade in a high context), although neither loc nor tmp are ever sent to an attacker. Permissive Upgrade [6] increases precision, however, it will incorrectly rule out any secure program that subsequently branches on variable tmp.

Combining dynamic and static analysis, *hybrid* approaches have recently received increased attention [36,18,37,39,31]. While providing strong formal guarantees, to date the practical impact of all these approaches has been limited, largely due to low precision (or *permissiveness*). Moreover, static, dynamic, and hybrid information-flow analysis require knowledge of the control-flow graph to properly propagate the *program counter* security label that keeps track of the sensitivity of the context. This label is difficult to recover whenever code has undergone heavyweight optimization and obfuscation, e.g. to protect its intellectual property, or in presence of reflection.

In contrast, *taint tracking* is a practical success story in computer security, with many applications at all levels of hardware and software stack [47,45]. Taint tracking is a pure data dependency analysis that only tracks explicit flows. It

is successful thanks to its lightweight nature, ignoring any control-flow dependencies that would be otherwise required for fully-fledged information-flow control. On the downside, taint tracking is mainly used as a bug finding technique, providing, with a few exceptions [57,45,46], no formal guarantees. Importantly, implicit flows may occur not only in malicious code [42,33], but also in trusted programs (written by a trusted programmer) [35,50,34,11].

These considerations point to an unsettling gap between heavyweight techniques for information-flow control, with formal guarantees yet limited practical impact, and lightweight tainting techniques that are useful for bug finding yet lacking formal assurance.

Approach By considering the trade-offs between *soundness* and *permissiveness*, this paper explores the middle ground, by a framework for a range of enforcement mechanisms from tainting to fully-fledged information-flow control. We address *trusted* and *malicious* code. However, we make a key distinction between two kinds of implicit flows: *observable* implicit flows and *hidden* implicit flows, borrowing the terminology of Staicu and Pradel [50]. Observable implicit flows arise whenever a variable is updated under a high security context and later output to an attacker. Not all implicit flows are, however, observable, since also the absence of a variable update can leak information (cf. Example 3); we call these hidden implicit flows. Tracking explicit flows and observable implicit flows raises the security bar for trusted code [50]. It allows for permissive, lightweight and purely dynamic enforcement in the spirit of taint tracking, yet providing higher security assurance. To evaluate soundness and permissiveness of the technique, we propose *observable secrecy*, a novel security condition that captures the essence of observable implicit flows. It helps us answer the question: "what is the security price we pay for having fewer false positives for useful programs"? We remark that the distinction between observable and hidden implicit flows is purely driven by ease of enforcement and permissiveness. Moreover, we leverage existing techniques and extend the framework to account for hidden implicit flows, thus addressing malicious code. We then present a family of flow-sensitive dynamic monitors that enforce a range of security policies by adapting a standard information-flow monitor from the literature [5,43].

The framework is deployed in a staged fashion. We statically embed dynamic monitors for (observable and/or hidden) implicit flows into the program code by lightweight program transformation, and leverage a dynamic taint tracker to enforce stronger policies. For malicious code, we use the *cross-copying* technique, originally proposed by Vachharajani et al. [53] for systems code, to transform hidden implicit flows into observable implicit flows. The transformations and soundness proofs for theorems can be found in the full version of the paper [14].

Secure App Store The flexibility of the approach on the policy and enforcement side facilitates a secure app store architecture, depicted in Figure 1. Developers deliver the code to the App Store, which computes sources and sinks, and leverages the control-flow graph to convert implicit flows into explicit flows. For trusted (non-malicious) apps, a lightweight transformation converting observable implicit flows into explicit may be sufficient, otherwise cross copying is needed. Subsequently, the App Store can perform code optimizations and obfuscations, and publish the resulting APK file (together with sources and sinks) on behalf of

Fig. 1. Secure App Store Architecture

the developer. Finally, end users can download the app, define their own security policies and run the app on a dynamic taint tracker, remarkably, with no need of the program's control-flow graph. Alternatively, end users can leverage static taint trackers [1,29] to verify their policies against the code.

We implement the transformations for a core of Java and evaluate them on the use case of a Pedometer app. We run the transformed app on TaintDroid [24] and check it against user-defined policies. We also show how JSFlow [30], a dynamic monitor for JavaScript, can provide higher precision by changing the security condition to observable secrecy.

Structure and Contributions In summary, the paper makes the following contributions: (*i*) observable secrecy, a security condition for validating soundness and precision wrt. observable implicit and explicit flows (Sect. 2); (*ii*) a framework that allows expressing a range of enforcement mechanisms from tainting to information-flow control (Sect. 3); (*iii*) lightweight transformations that leverage dynamic taint tracking for higher security assurance (Sect. 4); (*iv*) a flexible app store architecture and a prototype implementation for Android apps (Sect. 5).

2 Security Framework

We employ knowledge-based definitions [4,9,10] to introduce security conditions ranging from weak/explicit secrecy [57,45] to noninterference [28].

2.1 Language

Consider a simple imperative language with I/O primitives, SIMPL. The language expressions consist of variables $x \in Var$, built-in values $n \in Val$ such as integers and booleans, binary operators $\oplus$ and unary operators $\ominus$. We write **tt** for boolean value *true* and **ff** for boolean value *false*. The language constructs contain skip, assignment, conditional, loops, input and output. The full grammar of SIMPL can be found in Figure 2.

$$e ::= x \mid n \mid e \oplus e \mid \ominus e$$
$$P ::= \mathbf{skip} \mid P; P \mid x := e \mid x \leftarrow \mathbf{in}(\ell) \mid \mathbf{out}(\ell, e)$$
$$\mid \mathbf{if}\ e\ \mathbf{then}\ P\ \mathbf{else}\ P \mid \mathbf{while}\ e\ \mathbf{do}\ P$$

Fig. 2. SIMPL Language Grammar

We use input and output channels to model communication of the program with the external world. We label input and output channels with *security levels* ℓ (defined below) that indicate the confidentiality level of the information

transmitted on the corresponding channel. We denote the set of SIMPL programs by $\mathcal{P}$. We write $\bar{x}$ for a set of variables $\{x_1, \cdots, x_n\}$ such that for all $1 \le i \le n, x_i \in Var$, and $Vars(e)$ for the set of free variables of expression e.

We assume a bounded lattice of security levels $(\mathcal{L}, \sqsubseteq, \sqcup, \sqcap)$. A level $\ell \in \mathcal{L}$ represents the confidentiality of a piece of data present on a given channel or program variable. We assume that there is one channel for each security level $\ell \in \mathcal{L}$. As usual, $\sqsubseteq$ denotes the ordering relation between security levels and, $\sqcup$ and $\sqcap$ denote the *join* and *meet* lattice operators, respectively. We write $\top$ and $\bot$ to denote the top and the bottom element of the lattice. In the examples, we use a two-level security lattice $\mathcal{L} = \{\mathbf{L}, \mathbf{H}\}$ consisting of level $\mathbf{H}$ (high) for variables/channels containing confidential information and level $\mathbf{L}$ (low) for variables/channels containing public information, and $\mathbf{L} \sqsubseteq \mathbf{H}$. We focus on confidentiality, noting that integrity is similar through dualization [16].

We model input by environments $\mathcal{E} \in Env$ mapping channels to streams of input values. For simplicity, we consider one stream for each level $\ell \in \mathcal{L}$. An environment $\mathcal{E} : \mathcal{L} \to \mathbb{N} \to Val$ maps levels to infinite sequences of values. Two environments $\mathcal{E}_1$ and $\mathcal{E}_2$ are ℓ-equivalent, written $\mathcal{E}_1 \approx_\ell \mathcal{E}_2$, iff $\forall \ell'. \ell' \sqsubseteq \ell \Rightarrow \mathcal{E}_1(\ell') = \mathcal{E}_2(\ell')$. Another source of input are the initial values of program variables. We model memory as a mapping $m : Var \to Val$ from variables to values. We use $m, m_0, m_1, \dots$ to range over memories. We write $m[x \mapsto n]$ to denote a memory m with variable x assigned the value n. We write $m(e)$ for the value of expression e in memory m. A *security environment* $\Gamma : Var \mapsto \mathcal{L}$ is a mapping from program variables to lattice elements. The security environment assigns security levels to the memory through program variables. We use the terms security *level* and security *label* as synonyms. Two memories m_1 and m_2 are ℓ-equivalent, written $m_1 \approx_\ell m_2$, iff $\forall x \in Var. \Gamma(x) \sqsubseteq \ell \Rightarrow m_1(x) = m_2(x)$.

An *observation* $\alpha \in Obs$ is a pair of a security level and a value, i.e. $Obs = \mathcal{L} \times Val$, or the empty observation ϵ. A *trace* τ is a finite sequence of observations. We write $\tau.\tau'$ for concatenation of traces τ and τ', and $|\tau|$ for the length of a trace τ. We denote by $\tau \restriction_\ell$ the projection of trace τ at security level ℓ. Formally, we have $\epsilon \restriction_\ell = \epsilon$ and $(\ell', n).\tau' \restriction_\ell = (\ell', n).(\tau' \restriction_\ell)$ if $\ell' \sqsubseteq \ell$; otherwise $(\ell', n).\tau' \restriction_\ell = \tau' \restriction_\ell$. Two traces τ_1, τ_2 are ℓ-equivalent, written $\tau_1 \approx_\ell \tau_2$, iff $\tau_1 \restriction_\ell = \tau_2 \restriction_\ell$.

2.2 Semantics

The operational semantics of SIMPL is standard and it is reported in the full version [14]. A state $(\mathcal{E}, m)$ is a pair of an environment $\mathcal{E} \in Env$ and a memory $m \in Mem$. A configuration $\mathcal{E} \vdash \langle P, m \rangle$ consists of an environment $\mathcal{E}$, a program P and a memory m. We write $\mathcal{E} \vdash \langle P, m \rangle \xrightarrow{\alpha} \mathcal{E}' \vdash \langle P', m' \rangle$ to denote that a configuration $\mathcal{E} \vdash \langle P, m \rangle$ evaluates in one step to configuration $\mathcal{E}' \vdash \langle P', m' \rangle$, producing an observation $\alpha \in Obs$. We write $\to^*$ or $\xrightarrow{\tau}^*$ to denote the reflexive and transitive closure of $\to$. We write $\mathcal{E} \vdash \langle P, m \rangle \xrightarrow{\tau'}$ whenever the configuration is unimportant. We use ε to denote program termination.

2.3 Defining Secrecy

The goal of this subsection is to provide an attacker-centric definition of secrecy. The condition requires that the knowledge acquired by observing program outputs does not enable the attacker to learn sensitive information about the initial

program state (inputs and memories). We assume the attacker knows the program code and has *perfect recall* of all the past observations. We first illustrate the security condition by an example, and then provide the formal definition.

Example 1. Let $P = \textbf{if } h \textbf{ then out}(\textbf{L}, 1) \textbf{ else out}(\textbf{L}, 2)$ be a SIMPL program and h a secret variable, i.e. $\Gamma(h) = \textbf{H}$. Depending on the initial value of h, the program outputs either $\textbf{out}(\textbf{L}, 1)$ or $\textbf{out}(\textbf{L}, 2)$ on a channel of security level $\textbf{L}$.

An attacker at security level $\textbf{L}$ can reason about the initial value of h as follows: (i) Before seeing any output, the attacker considers any boolean value as possible for h, therefore the knowledge is $h \in \{\textbf{tt}, \textbf{ff}\}$. (ii) If the statement $\textbf{out}(\textbf{L}, 1)$ is executed, the attacker can refine the knowledge to $h \in \{\textbf{tt}\}$ and thus learn the initial value of h. (iii) Similarly, if the statement $\textbf{out}(\textbf{L}, 2)$ is executed, the attacker learns that h was initially false. Hence, the program is insecure.

We now define the knowledge that an attacker at level ℓ acquires from observing a trace of a program P. We capture this by considering the set of initial states that the attacker considers possible based on their observations. Concretely, for a given initial state $(\mathcal{E}_0, m_0)$ and a program P, an initial state $(\mathcal{E}, m)$ is considered possible if $\mathcal{E} \approx_\ell \mathcal{E}_0$, $m \approx_\ell m_0$, and it matches the trace produced by $\mathcal{E}_0 \vdash \langle P, m_0 \rangle$. We define the attacker's knowledge in the standard way [4]:

Definition 1 (Knowledge). *The* knowledge set *for program P, initial state $(\mathcal{E}_0, m_0)$, security level ℓ and trace τ is given by* $k(P, \mathcal{E}_0, m_0, \tau) = \{(\mathcal{E}, m) \mid \mathcal{E} \approx_\ell \mathcal{E}_0 \wedge m \approx_\ell m_0 \wedge (\exists P', \mathcal{E}', m', \tau'. \mathcal{E} \vdash \langle P, m \rangle \xrightarrow{\tau}^* \mathcal{E}' \vdash \langle P', m' \rangle \wedge \tau \approx_\ell \tau')\}$.

We focus on *progress-insensitive* security, which ignores information leaks through the observation of computation progress, e.g. program divergence [3]. To this end, we relax the requirement that the attacker learns nothing at each execution step, by allowing leaks that arise from observing the progress of computation. Concretely, we define progress knowledge as the set of initial states that the attacker considers possible based on the fact that *some* output event has occurred, independently of what the exact output value was.

Definition 2 (Progress Knowledge). *The* progress knowledge set *for program P, initial state $(\mathcal{E}_0, m_0)$, level ℓ, and trace τ is given by* $k_P(P, \mathcal{E}_0, m_0, \tau) = \{(\mathcal{E}, m) \mid \mathcal{E} \approx_\ell \mathcal{E}_0 \wedge m \approx_\ell m_0 \wedge (\exists P', \mathcal{E}', m', \tau', \alpha \neq \epsilon. \mathcal{E} \vdash \langle P, m \rangle \xrightarrow{\tau}^* \mathcal{E}' \vdash \langle P', m' \rangle \xrightarrow{\alpha}^* \wedge \alpha \upharpoonright_\ell = \alpha \wedge \tau \approx_\ell \tau')\}$.

We can now define a *progress-insensitive* secrecy by requiring that progress knowledge after observing a trace τ is the same as the knowledge obtained after observing the trace $\tau.\alpha$. Consequently, what the attacker learns from observing the exact output value is the same as what they learn from observing the computation progress, i.e. that some output event has occurred.

Definition 3 (Progress-insensitive Secrecy). *A program P satisfies* Progress-insensitive Secrecy *at level ℓ, written $Sec(\ell) \vDash P$, iff whenever $\mathcal{E} \vdash \langle P, m \rangle \xrightarrow{\tau.\alpha}^* \mathcal{E}' \vdash \langle P', m' \rangle \wedge \alpha \upharpoonright_\ell = \alpha \wedge \alpha \neq \epsilon$, we have $k_P(P, \mathcal{E}, m, \tau) = k(P, \mathcal{E}, m, \tau.\alpha)$. P satisfies* Progress-insensitive Secrecy, *written $Sec \models P$ iff $Sec(\ell) \vDash P$, for all ℓ.*

We can see that the program in Ex. 1 does not satisfy progress-insensitive secrecy at security level $\mathbf{L}$, as the progress knowledge of observing *some* output, i.e. either $\mathbf{out}(\mathbf{L}, 1)$ or $\mathbf{out}(\mathbf{L}, 2)$, is $h \in \{\mathbf{tt}, \mathbf{ff}\}$, while the knowledge of observing the *exact* output, e.g. $\mathbf{out}(\mathbf{L}, 1)$, is $h \in \{\mathbf{tt}\}$.

2.4 Security Conditions

Information-flow monitors can enforce progress-insensitive secrecy, thus preventing both implicit and explicit flows. Taint tracking, on the other hand, is an enforcement mechanism that only prevents explicit flows, otherwise ignores any control-flow dependencies [21]. In contrast to noninterference, security conditions for taint tracking [57,45] serve more as semantic criteria for evaluating soundness and precision of the underlying enforcement mechanism rather than providing an intuitive meaning of security. Driven by the same motivation, we propose a family of security conditions that allows exploring the space of enforcement mechanisms from taint tracking to information-flow control.

Our security conditions rely on the observational power of an attacker over the program code and executions. We model attackers with respect to their per-run view of the program code and extract the program *slice* that an attacker considers possible for any concrete execution. This allows to re-use the same condition as in Def. 3 for the program *slice* that the attacker can observe.

Concretely, a security condition for taint tracking can be modelled as secrecy with respect to an attacker that only observes explicit statements (input, output and assignment) extracted from any concrete execution of a program P. Similarly,

$$l_1 := \mathbf{tt} ; l_2 := \mathbf{tt} ; \qquad (1)$$
$$\mathbf{if}\ h\ \mathbf{then}\ l_1 := \mathbf{ff}\ \mathbf{else}\ \mathbf{skip} \qquad (2)$$
$$\mathbf{if}\ l_1\ \mathbf{then}\ l_2 := \mathbf{ff}\ \mathbf{else}\ \mathbf{skip} \qquad (3)$$
$$\mathbf{out}(\mathbf{L}, l_2) \qquad (4)$$

Fig. 3. Leaking through label upgrades

(termination-insensitive) noninterference [3] corresponds secrecy for an attacker that has a whole view of P.

We will use the example in Figure 3 to illustrate the security conditions. Consider the program P with boolean variable h of level $\mathbf{H}$ and boolean variables l_1, l_2 of level $\mathbf{L}$. It can be seen that P outputs the initial value of variable h to an observer at security level $\mathbf{L}$ through a sequence of control flow decisions. In fact, the program does not satisfy the condition in Def. 3.

We introduce *extraction contexts* C as a *gadget* to model the observational power of an attacker over the program code. Extraction contexts provide a mechanism to leverage the operational semantics of the language and extract the program slice that an attacker observes for any given concrete execution.

$$C ::= [] \mid \mathbf{skip} \mid x := e \mid x \leftarrow \mathbf{in}(\ell) \mid \mathbf{out}(\ell, e) \mid C; C \mid \mathbf{if}\ e\ \mathbf{then}\ C\ \mathbf{else}\ C$$

Syntactically, extraction contexts are programs that may contain *holes* []. For our purposes, contexts will contain at most one hole that represents a placeholder for the program statements that are yet to be evaluated by the program execution at hand. We extend the operational semantics to transform contexts in order to extract programs for *weak secrecy* and *observable secrecy*.

Weak Secrecy Weak secrecy [57], a security condition for taint tracking, states that every sequence of explicit statements executed by any program run

must be secure. We formalize weak secrecy as secrecy (cf. Def. 3) for the program, i.e. the sequence of explicit statements, extracted from any (possibly incomplete) execution of the original program. We achieve this by extending the configurations with extraction contexts. Here we discuss a few interesting rules as reported in Figure 4. The complete set of rules can be found in [14].

W-ASSIGN
$$\frac{m(e) = n}{\mathcal{E} \vdash \langle x := e, m, C \rangle \to \mathcal{E} \vdash \langle \varepsilon, m[x \mapsto n], C[x := e] \rangle}$$

W-OUT
$$\frac{m(e) = n}{\mathcal{E} \vdash \langle \mathbf{out}(\ell, e), m, C \rangle \xrightarrow{[(\ell, n)]} \mathcal{E} \vdash \langle \varepsilon, m, C[\mathbf{out}(\ell, e)] \rangle}$$

W-IFTRUE
$$\frac{m(e) = \mathbf{tt}}{\mathcal{E} \vdash \langle \mathbf{if}\ e\ \mathbf{then}\ P_1\ \mathbf{else}\ P_2, m, C \rangle \to \mathcal{E} \vdash \langle P_1, m, C \rangle}$$

W-SEQ
$$\frac{\mathcal{E} \vdash \langle P_1, m, C \rangle \xrightarrow{\alpha} \mathcal{E}' \vdash \langle P_1', m', C' \rangle}{\mathcal{E} \vdash \langle P_1\ ;\ P_2, m, C \rangle \xrightarrow{\alpha} \mathcal{E}' \vdash \langle P_1'\ ;\ P_2, m', C' \rangle}$$

Fig. 4. Excerpt of Extraction Rules for Weak Secrecy

Each program execution starts with the empty context []. To extract explicit statements, we propagate assignment and output commands into the context, while conditionals are simply ignored (cf. the context remains unchanged). Sequential composition ensures that the sequence of explicit statements is propagated correctly. It can be shown that complete (terminated) executions contain no holes and incomplete executions contain exactly one hole.

We define weak secrecy in terms of secrecy for explicit statements extracted from any program execution. We write $C[\mathbf{skip}]$ to denote the result of replacing the hole with command $\mathbf{skip}$ in a context C. Otherwise, if the context contains no hole, we have $C[\mathbf{skip}] = C$. This is needed because the security condition is defined for any execution, including complete and incomplete executions.

Definition 4 (Weak secrecy). *A program P satisfies* weak secrecy *for initial state $(\mathcal{E}, m)$, written $WS \vDash_{\mathcal{E}, m} P$, iff whenever $\mathcal{E} \vdash \langle P, m, [] \rangle \xrightarrow{\tau}{}^{*} \mathcal{E}' \vdash \langle P', m', C \rangle$, we have $Sec \vDash C[\mathbf{skip}]$. A program P satisfies* weak secrecy*, written $WS \vDash P$, iff $WS \vDash_{\mathcal{E}, m} P$ for all states $(\mathcal{E}, m)$.*

Consider the program from Ex. 3 and an initial state $(\mathcal{E}_0, m_0)$. Depending on whether $m_0(h) = \mathbf{tt}$ and $m_0(h) = \mathbf{ff}$, we extract program (5) or program (6), respectively, shown in Figure 5.

We can see that none of the programs contains variable h, hence they both satisfy secrecy (Def. 3).

$$l_1 := \mathbf{tt}\ ;\ l_2 := \mathbf{tt}\ ;\ l_1 := \mathbf{ff}\ ;\ \mathbf{skip}\ ;\ \mathbf{out}(\mathbf{L}, l_2) \qquad (5)$$
$$l_1 := \mathbf{tt}\ ;\ l_2 := \mathbf{tt}\ ;\ \mathbf{skip}\ ;\ l_2 := \mathbf{ff}\ ;\ \mathbf{out}(\mathbf{L}, l_2) \qquad (6)$$

Fig. 5. Extracted programs

As a result, the original program P satisfies weak secrecy.

Observable Secrecy We now present a novel security condition, dubbed *observable secrecy*, that captures the intuition of *observable* implicit flows. Observable implicit flows are implicit flows that arise whenever a variable is modified in the high branch that is currently executed by the program, and later it is output to the attacker. Preventing observable implicit flows is of interest for purely

O-SKIP

$$\mathcal{E} \vdash \langle \mathbf{skip}, m, C \rangle \to \mathcal{E} \vdash \langle \varepsilon, m, C[\mathbf{skip}] \rangle$$

O-IN

$$\frac{\mathcal{E}' = \mathcal{E}[\ell \mapsto n \mapsto \mathcal{E}(\ell)(n+1)] \qquad m' = m[x \mapsto \mathcal{E}(\ell)(0)]}{\mathcal{E} \vdash \langle x \leftarrow \mathbf{in}(\ell), m, C \rangle \to \mathcal{E}' \vdash \langle \varepsilon, m', C[x \leftarrow \mathbf{in}(\ell)] \rangle}$$

O-ASSIGN

$$\frac{m(e) = n}{\mathcal{E} \vdash \langle x := e, m, C \rangle \to \mathcal{E} \vdash \langle \varepsilon, m[x \mapsto n], C[x := e] \rangle}$$

O-SEQ

$$\frac{\mathcal{E} \vdash \langle P_1, m, C \rangle \xrightarrow{\alpha} \mathcal{E}' \vdash \langle P_1', m', C' \rangle}{\mathcal{E} \vdash \langle P_1 \; ; P_2, m, C \rangle \xrightarrow{\alpha} \mathcal{E}' \vdash \langle P_1' \; ; P_2, m', C' \rangle}$$

O-OUT

$$\frac{m(e) = n}{\mathcal{E} \vdash \langle \mathbf{out}(\ell, e), m, C \rangle \xrightarrow{[(\ell, n)]} \mathcal{E} \vdash \langle \varepsilon, m, C[\mathbf{out}(\ell, e)] \rangle}$$

O-WHILEFALSE

$$\frac{m(e) = \mathbf{ff}}{\mathcal{E} \vdash \langle \mathbf{while} \; e \; \mathbf{do} \; P, m, C \rangle \to \mathcal{E} \vdash \langle \varepsilon, m, C[\mathbf{skip}] \rangle}$$

O-WHILETRUE

$$\frac{m(e) = \mathbf{tt}}{\mathcal{E} \vdash \langle \mathbf{while} \; e \; \mathbf{do} \; P, m, C \rangle \to \mathcal{E} \vdash \langle P \; ; \mathbf{while} \; e \; \mathbf{do} \; P, m, C[\mathbf{if} \; e \; \mathbf{then} \; [] \; \mathbf{else} \; \mathbf{skip}] \rangle}$$

O-IFTRUE

$$\frac{m(e) = \mathbf{tt}}{\mathcal{E} \vdash \langle \mathbf{if} \; e \; \mathbf{then} \; P_1 \; \mathbf{else} \; P_2, m, C \rangle \to \mathcal{E} \vdash \langle P_1, m, C[\mathbf{if} \; e \; \mathbf{then} \; [] \; \mathbf{else} \; \mathbf{skip}] \rangle}$$

O-SEQEMPTY

$$\mathcal{E} \vdash \langle \varepsilon \; ; P_2, m, C \rangle \to \mathcal{E} \vdash \langle P_2, m, C \; ; [] \rangle$$

O-IFFALSE

$$\frac{m(e) = \mathbf{ff}}{\mathcal{E} \vdash \langle \mathbf{if} \; e \; \mathbf{then} \; P_1 \; \mathbf{else} \; P_2, m, C \rangle \to \mathcal{E} \vdash \langle P_2, m, C[\mathbf{if} \; e \; \mathbf{then} \; \mathbf{skip} \; \mathbf{else} \; []] \rangle}$$

O-IFTRUE

$$\frac{m(e) = \mathbf{tt}}{\mathcal{E} \vdash \langle \mathbf{if} \; e \; \mathbf{then} \; P_1 \; \mathbf{else} \; P_2, m, C \rangle \to \mathcal{E} \vdash \langle P_1, m, C[\mathbf{if} \; e \; \mathbf{then} \; [] \; \mathbf{else} \; \mathbf{skip}] \rangle}$$

O-SEQEMPTY

$$\mathcal{E} \vdash \langle \varepsilon \; ; P_2, m, C \rangle \to \mathcal{E} \vdash \langle P_2, m, C \; ; [] \rangle$$

Fig. 6. Extraction Rules for Observable Secrecy

dynamic mechanisms as it provides higher security compared to weak secrecy, yet allowing for dynamic monitors that are more permissive than monitors for noninterference. Permissiveness, however, comes at the price of ignoring hidden implicit flows. The following program, where h has security level $\mathbf{H}$, contains an observable implicit flow whenever $m_0(h) = \mathbf{tt}$, otherwise the flow is hidden.

$$l := \mathbf{ff} \; ; \mathbf{if} \; h \; \mathbf{then} \; \{l := \mathbf{tt}\} \; \mathbf{else} \; \{\mathbf{skip}\} \; ; \mathbf{out}(\mathbf{L}, l)$$

The security condition considers an attacker that only observes the instructions (both control-flow and explicit statements) executed by the concrete program execution, otherwise it ignores (i.e. replaces with **skip**) any instruction occurring in the untaken branches. To capture these flows, we extend the small-step operational semantics to extract the program code observable by this attacker, as shown in Figure 6.

The rules for assignment, input, output and sequential composition are the same as for weak secrecy. Rules for conditionals propagate the *observable* conditional into the context C to keep track of the executed branch and replace the untaken branch with **skip**. The new hole [] ensures that the commands under the executed branch are properly modified by the new context. We unfold loop statements into conditionals and handle them similarly. Sequential composition ensures that the sequence of observable statements is propagated correctly. When rule O-SEQEMPTY is applied, the context C does not contain any holes, hence a new hole is introduced to properly handle the remaining command P_2.

Definition 5 (Observable secrecy). *A program P satisfies* observable secrecy *for initial state $(\mathcal{E}, m)$, written $OS \vDash_{\mathcal{E},m} P$, iff whenever $\mathcal{E} \vdash \langle P, m, [] \rangle \xrightarrow{\tau}^* \mathcal{E}' \vdash \langle P', m', C \rangle$, we have $Sec \models C[\textbf{skip}]$. A program P satisfies* observable secrecy, *written $OS \vDash P$, iff $OS \vDash_{\mathcal{E},m} P$ for all states $(\mathcal{E}, m)$.*

For the above example, the operational semantics rules for observable secrecy yield the programs:

> $l := \textbf{ff}$; **if** h **then** $\{l := \textbf{tt}\}$ **else** $\{\textbf{skip}\}$; $\textbf{out}(\textbf{L}, l)$
> $l := \textbf{ff}$; **if** h **then** $\{\textbf{skip}\}$ **else** $\{\textbf{skip}\}$; $\textbf{out}(\textbf{L}, l)$

The first program does not satisfy secrecy (Def. 3), while the second program does. Therefore the original program does not satisfy observable secrecy.

Full Secrecy Full secrecy is a security condition that models secrecy with respect to an attacker that has a complete knowledge of program code and therefore can learn information through explicit and (observable or hidden) implicit flows. This corresponds to progress-insensitive noninterference (Def. 3).

Definition 6 (Full secrecy). *A program P satisfies* full secrecy *for initial state $(\mathcal{E}, m)$, written $FS \vDash_{\mathcal{E},m} P$, iff whenever $\mathcal{E} \vdash \langle P, m \rangle \xrightarrow{\tau}^* \mathcal{E}' \vdash \langle P', m' \rangle$, we have $Sec \models P$. A program P satisfies* full secrecy, *written $FS \vDash P$, iff $FS \vDash_{\mathcal{E},m} P$ for all states $(\mathcal{E}, m)$.*

3 Enforcement Framework

We employ variants of flow-sensitive dynamic monitors (trackers) to enforce the security conditions presented in the last section. Compared to existing work (cf. Sect. 6), we use semantic security conditions, weak secrecy and observable secrecy, to justify soundness of *weak* tracking and *observable* tracking mechanisms.

Figure 7 presents the instrumented semantics which is parametric on the security labels, transfer functions and constraints. By instantiating each of the parameters (Table. 1), we show how the semantics implements sound dynamic trackers for weak secrecy (Thrm. 1), observable secrecy (Thrm. 2) and full secrecy (Thrm. 3). All proofs are reported in the full version [14].

The instrumented semantics assumes a bounded lattice $(\mathcal{L}, \sqsubseteq, \sqcup, \sqcap)$ and an initial security environment Γ, as defined in Sect. 2.1. We use a *program counter* stack of security levels pc to keep track of the security context, i.e. the security level of conditional and loop expressions, at a given execution point. We write $\ell :: pc$ to denote a stack of labels, where the label ℓ is its top element. Abusing notation, we also write pc to represent the upper bound on the security levels of the stack elements. The monitored semantics introduces the special instruction **end** to remember the join points in the control flow and update the pc stack accordingly. Instrumented configurations $\Gamma, pc, \mathcal{E} \vdash \langle P, m \rangle$ extend original configurations with the security environment Γ and security context stack pc. We write $\Gamma, pc, \mathcal{E} \vdash \langle c, m \rangle \xrightarrow{\alpha} \Gamma', pc', \mathcal{E}' \vdash \langle c', m' \rangle$ to denote that an instrumented configuration $\Gamma, pc, \mathcal{E} \vdash \langle c, m \rangle$ evaluates in one step to instrumented configuration $\Gamma', pc', \mathcal{E}' \vdash \langle c', m' \rangle$, producing observations $\alpha \in Obs$. We write $\twoheadrightarrow^*$ or

S-Skip

$$\frac{}{\Gamma, pc, \mathcal{E} \vdash \langle \mathbf{skip}, m \rangle \twoheadrightarrow \Gamma, pc, \mathcal{E} \vdash \langle \varepsilon, m \rangle}$$

S-In-F

$$\frac{\phi_{inF}}{\Gamma, pc, \mathcal{E} \vdash \langle x \leftarrow \mathbf{in}(\ell), m \rangle \twoheadrightarrow \mathbf{\xi}}$$

S-Out

$$\frac{m(e) = n \qquad \phi_{outT}}{\Gamma, pc, \mathcal{E} \vdash \langle \mathbf{out}(\ell, e), m \rangle \xrightarrow{[(\ell, n)]} \Gamma, pc, \mathcal{E} \vdash \langle \varepsilon, m \rangle}$$

S-WhileFalse

$$\frac{m(e) = \mathbf{tt} \qquad \phi_{wh}}{\Gamma, pc, \mathcal{E} \vdash \langle \mathbf{while}\ e\ \mathbf{do}\ P, m \rangle \twoheadrightarrow \Gamma, pc', \mathcal{E} \vdash \langle \mathbf{end}, m \rangle}$$

S-In

$$\frac{\mathcal{E}' = \mathcal{E}[\ell \mapsto n \mapsto \mathcal{E}(\ell)(n+1)] \qquad m' = m[x \mapsto \mathcal{E}(\ell)(0)] \qquad \Gamma' = \Gamma[x \mapsto \ell \sqcup pc] \qquad \phi_{inT}}{\Gamma, pc, \mathcal{E} \vdash \langle x \leftarrow \mathbf{in}(\ell), m \rangle \twoheadrightarrow \Gamma', pc, \mathcal{E}' \vdash \langle \varepsilon, m' \rangle}$$

S-IfFalse

$$\frac{m(e) = \mathbf{ff} \qquad \phi_{if}}{\Gamma, pc, \mathcal{E} \vdash \langle \mathbf{if}\ e\ \mathbf{then}\ P_1\ \mathbf{else}\ P_2, m \rangle \twoheadrightarrow \Gamma, pc', \mathcal{E} \vdash \langle P_2\ ;\ \mathbf{end}, m \rangle}$$

S-Assign-F

$$\frac{\phi_{asgF}}{\Gamma, pc, \mathcal{E} \vdash \langle x := e, m \rangle \twoheadrightarrow \mathbf{\xi}}$$

S-IfTrue

$$\frac{m(e) = \mathbf{tt} \qquad \phi_{if}}{\Gamma, pc, \mathcal{E} \vdash \langle \mathbf{if}\ e\ \mathbf{then}\ P_1\ \mathbf{else}\ P_2, m \rangle \twoheadrightarrow \Gamma, pc', \mathcal{E} \vdash \langle P_1\ ;\ \mathbf{end}, m \rangle}$$

S-Out-F

$$\frac{\phi_{outF}}{\Gamma, pc, \mathcal{E} \vdash \langle \mathbf{out}(\ell, e), m \rangle \twoheadrightarrow \mathbf{\xi}}$$

S-SeqEmpty

$$\frac{}{\Gamma, pc, \mathcal{E} \vdash \langle \varepsilon\ ;\ P_2, m \rangle \twoheadrightarrow \Gamma, pc, \mathcal{E} \vdash \langle P_2, m \rangle}$$

S-End

$$\frac{\phi_{End}}{\Gamma, pc, \mathcal{E} \vdash \langle \mathbf{end}, m \rangle \twoheadrightarrow \Gamma, pc', \mathcal{E} \vdash \langle \varepsilon, m \rangle}$$

S-WhileTrue

$$\frac{m(e) = \mathbf{tt} \qquad \phi_{wh}}{\Gamma, pc, \mathcal{E} \vdash \langle \mathbf{while}\ e\ \mathbf{do}\ P, m \rangle \twoheadrightarrow \Gamma, pc', \mathcal{E} \vdash \langle P\ ;\ \mathbf{end}\ ;\ \mathbf{while}\ e\ \mathbf{do}\ P, m \rangle}$$

S-Assign

$$\frac{m(e) = n \qquad \Gamma' = \Gamma[x \mapsto pc \sqcup \Gamma(e)] \qquad \phi_{asgT}}{\Gamma, pc, \mathcal{E} \vdash \langle x := e, m \rangle \twoheadrightarrow \Gamma', pc, \mathcal{E} \vdash \langle \varepsilon, m[x \mapsto n] \rangle}$$

S-Seq

$$\frac{\Gamma, pc, \mathcal{E} \vdash \langle P_1, m \rangle \xrightarrow{\alpha} \Gamma', pc', \mathcal{E}' \vdash \langle P_1', m' \rangle}{\Gamma, pc, \mathcal{E} \vdash \langle P_1\ ;\ P_2, m \rangle \xrightarrow{\alpha} \Gamma', pc', \mathcal{E}' \vdash \langle P_1'\ ;\ P_2, m' \rangle}$$

Fig. 7. Instrumented Semantics

$\xrightarrow{\tau}{}^*$ to denote the reflexive and transitive closure of $\xrightarrow{\alpha}$. We write $\Gamma(e)$ for $\sqcup_{x \in Vars(e)} \Gamma(x)$ and $\mathbf{\xi}$ for abnormal termination.

In what follows, we use the constraints in Table 1 to instantiate the rules in Figure 7, and present a family of dynamic monitors for weak tracking (known as taint tracking), observable tracking, and full tracking (known as No-Sensitive Upgrade [5]). The monitors implement the *failstop* strategy and terminate the program abnormally (cf. rules for $\mathbf{\xi}$) whenever a potentially insecure statement is executed. Note that abnormal termination does not produce any observable event and it is treated as a progress channel, similarly to nontermination. We write $\mathcal{I} \vdash_{\mathcal{E}, m} P$ for an execution of a monitored program P from initial state $(\mathcal{E}, m)$, initial security environment Γ and initial stack $\bot$, where $\mathcal{I} \in \{WS, OS, FS\}$.

Monitored executions may change the semantics of the original program by collapsing insecure executions into abnormal termination. To account for the monitored semantics, we instantiate the security conditions from Section 2.4 with the semantics of instrumented executions and, abusing notation, write $\mathcal{I} \models_{\mathcal{E}, m} P$ to refer to an execution of P under the instrumented semantics. We then show that any program executed under an instrumented execution, i.e., $\mathcal{I} \vdash_{\mathcal{E}, m} P$, satisfies the security condition, i.e., $\mathcal{I} \models_{\mathcal{E}, m} P$.

Weak Tracking Weak tracking is a dynamic mechanism that prevents explicit flows from sources of higher security levels to sinks of lower security levels. Weak tracking allows leaks through implicit flows. The second column in Table 1 gives the set of constraints that a typical taint analysis would implement for our language.

Since the analysis ignores all implicit flows, the pc stack is redundant and we never update it during the monitor execution. For the same reason, we apply no side conditions to the rules for conditionals and loops. Rule S-ASSIGN propagates the security level of the expression on the right-hand side to the variable on the left-hand side to track potential explicit flows, while rule S-ASSIGN-F never applies. Rule S-OUT ensures that only direct flows from lower levels affect a given output level. If the constraint is not satisfied, the program terminates abnormally (cf. S-OUT-F).

RULE	WEAK	OBSERVABLE	FULL
ϕ_{asgT}	tt	tt	$pc \sqsubseteq \Gamma(x)$
ϕ_{asgF}	ff	ff	$pc \not\sqsubseteq \Gamma(x)$
ϕ_{outT}	$\Gamma(e) \sqsubseteq \ell$	$\Gamma(e) \sqsubseteq pc \sqcup \ell$	$\Gamma(e) \sqsubseteq pc \sqcup \ell$
ϕ_{outF}	$\Gamma(e) \not\sqsubseteq \ell$	$\Gamma(e) \not\sqsubseteq pc \sqcup \ell$	$\Gamma(e) \not\sqsubseteq pc \sqcup \ell$
ϕ_{inT}	tt	$pc \sqsubseteq \ell$	$pc \sqsubseteq \ell$
ϕ_{inF}	ff	$pc \not\sqsubseteq \ell$	$pc \not\sqsubseteq \ell$
ϕ_{end}	tt	$pc = \ell :: pc'$	$pc = \ell :: pc'$
ϕ_{if}/ϕ_{wh}	tt	$\ell' = pc \sqcup \Gamma(e)$ $pc' = \ell' :: pc$	$\ell' = pc \sqcup \Gamma(e)$ $pc' = \ell' :: pc$

Table 1. Constraints for Monitors in Fig. 7

To illustrate the weak tracking monitor, consider the program from Ex. 3. Initially, the security environment Γ assigns the label **L** to variables l_1 and l_2, and the label **H** to variable h. After the execution of line (1), the security environment Γ' does not change since $pc = \mathbf{L}$ and, $\Gamma(n) = \mathbf{L}$ for all $n \in Val$, therefore $\Gamma'(l_1) = \Gamma'(l_2) = \mathbf{L} \sqcup \Gamma(\mathbf{ff}) = \mathbf{L}$ (cf. rule S-ASSIGN). Moreover, the lines (2) and (3) do not modify Γ' (cf. rules S-IFTRUE and S-IFFALSE). Finally, the output in line (4) is allowed since $\Gamma(l_2) = \mathbf{L} \sqsubseteq \mathbf{L}$ (cf. rule S-OUT). In fact, the program satisfies weak secrecy (Def. 4), and it is accepted by weak tracking.

We show that any program that is executed under the weak tracking monitor, i.e. $\mathcal{I} = WS$, satisfies weak secrecy.

Theorem 1. $WS \vdash_{\mathcal{E},m} P \Rightarrow WS \vDash_{\mathcal{E},m} P$

Observable Tracking Observable tracking is a dynamic security mechanism that accounts for explicit flows and observable implicit flows. Observable implicit flows occur whenever a low security variable that is updated in a high security context is later output to a low security channel. The condition justifies the security of a program with respect to an attacker that only knows the control-flow path of the current execution. Observable tracking has the appealing property of only propagating the security label of variables in a concrete program execution, without analyzing variables modified in the untaken branches. This is remarkable as it sidesteps the need for convoluted static analysis otherwise required for languages with dynamic features such as reflection. Moreover, as we discuss later, observable tracking is more permissive than existing enforcement mechanisms such as NSU [5] or Permissive Upgrade [6]. Permissiveness is achieved at the expense of enforcing a different security condition, i.e. observable secrecy, instead of full secrecy. For trusted code, observable secrecy might be sufficient to determine unintentional security bugs. Otherwise, for malicious code, we present a transformation (Sect. 4) that enables observable tracking to enforce full secrecy, yet being more permissive than full tracking.

The instrumented semantics for observable tracking (cf. third column in Table 1) strengthens the constraints for weak tracking by: (i) introducing the pc stack to properly track changes of security labels for variables updated in a high context; (ii) disallowing input from low security channels in a high context; (iii) and constraining the output on a low channel by disallowing low expressions that depend on a high context.

Consider again the program in Ex. 3 under the instrumented semantics for observable tracking. After executing the assignments in (1), the variables l_1 and l_2 have security level **L**. If h is **tt**, the variable l_1 has security level **H** after the first conditional in (2) (cf. S-IFTRUE rule). As a result, the guard of the second conditional in (3) is false, and we execute the **else** branch. The security level of the variable l_2 remains **L**, therefore the output on the **L** channel in (4) is allowed (cf. S-OUT rule). Otherwise, if h is **ff**, then the **else** branch is executed and l_1 has security level **L**. The second conditional does not change the security level of l_2, although the **then** branch is executed. In fact, the guard only depends on **L** variables, i.e. l_1, hence security level of l_2 remains **L** and the subsequent output is allowed. The program, in fact, satisfies observable secrecy.

We prove that any program that is executed under the observable tracking monitor, i.e. $\mathcal{I} = OS$, satisfies observable secrecy.

Theorem 2. $OS \vdash_{\mathcal{E},m} P \Rightarrow OS \vDash_{\mathcal{E},m} P$

Full Tracking Full tracking, best known as No-Sensitive Upgrade [58,5], prevents both explicit and (observable or hidden) implicit flows from sources of higher security levels to sinks of lower security levels. This is achieved by disallowing changes of variables' security labels in high contexts (as opposed to the strategy followed by observable tracking). While sound for full secrecy, this strategy incorrectly terminates any program that updates a low security variable in a high security context, even if that variable is never output to low channel. This is unfortunate as it rejects secure programs that only use sensitive data for internal computations without ever sending them on low channels.

The semantics for full tracking adds additional constraints to the rules for observable tracking (cf. fourth column in Table 1). In particular, rule S-ASSIGN only allows low assignments in low security contexts, i.e. whenever $pc \sqsubseteq \Gamma(x)$.

Consider again the program in Ex. 3 and the semantics for full tracking. As before, initially $\Gamma(l_1) = \Gamma(l_2) = \mathbf{L}$, and $\Gamma(h) = \mathbf{H}$. If the value of h is true, the **then** branch of the first conditional is executed, and the program is stopped because of a low assignment in a high context. This is a sound behavior of full tracking as the original program does not satisfy full secrecy. Unfortunately, full tracking will also stop any secure programs that contain the conditional statement in (2). For example, if we replace the output statement in (4) with **out**($\mathbf{L}, 1$) or **out**($\mathbf{H}, l_2$), the resulting program clearly satisfies full secrecy. However, whenever h is true, full tracking will incorrectly stop the program.

We show that any program that is executed under the full tracking monitor, i.e. $\mathcal{I} = FS$, satisfies full secrecy.

Theorem 3. $FS \vdash_{\mathcal{E},m} P \Rightarrow FS \vDash_{\mathcal{E},m} P$

PROGRAM $\quad \Gamma(h) = \mathbf{H}$, $\Gamma(l) = \Gamma(k) = \mathbf{L}$ and $h = \mathbf{tt}$	WEAK	FULL	PU	OT
P_0 $l := \mathbf{tt}$; **if** h **then** $\{l := h\}$; $\mathbf{out}(\mathbf{L}, l)$	−	−	−	−
P_1 **if** h **then** $l := \mathbf{tt}$	+	−	+	+
P_2 **if** h **then** $l := \mathbf{tt}$; **if** l **then skip**	+	−	−	+
P_3 $l := \mathbf{tt}$; $k := \mathbf{tt}$; **if** h **then** $\{l := \mathbf{ff}\}$; **if** l **then** $\{k := \mathbf{ff}\}$; $\mathbf{out}(\mathbf{L}, 1)$	+	−	−	+
P_4 **if** h **then** $\mathbf{out}(\mathbf{L}, 1)$ **else** $\mathbf{out}(\mathbf{L}, 1)$	+	−	−	−
P_5 $l := \mathbf{tt}$; $k := \mathbf{tt}$; **if** h **then** $\{l := \mathbf{ff}\}$; **if** l **then** $\{k := \mathbf{ff}\}$; $\mathbf{out}(\mathbf{L}, k)$	+	✗	✗	+

Table 2. Permissiveness

4 Staged Information-Flow Control

Two main factors hinder the adoption of dynamic information-flow control in practice: *challenging implementation* and *permissiveness*. To properly update the program counter stack at runtime, observable and full tracking require the knowledge of the program's control-flow graph. This requirement is unrealistic for unstructured, heavily optimized or obfuscated code, such as the code delivered to end users (cf. Sect. 1). In contrast, weak tracking disregards the control-flow graph and only considers explicit statements. As a result, the enforcement is more permissive and easier to implement.

In the full version [14], we present a staged analysis that first applies lightweight program transformations to convert implicit flows into explicit flows, thus delegating the task of enforcing observable and full secrecy to a weak tracker. Concretely, we inline the program counter stack into the source code in a semantics-preserving manner by introducing fake dependencies that cause a weak tracker to capture potential observable and/or hidden implicit flows. The transformation is completely transparent to the underlying security policy, which makes it suitable for the scenarios envisioned in Sect. 1.

Soundness vs Permissiveness We use the examples in Table 2 to illustrate soundness and permissiveness for existing dynamic trackers. Except for the program P_5, all programs are secure for full secrecy. We summarize the relations between the security conditions (solid ovals) and enforcement mechanisms (dashed ovals) in Figure 4. The security conditions are incomparable, as shown by the programs P_0, P_4 and P_5 from Table 2. Moreover, there is a strict inclusion between the set of secure programs accepted by the trackers (cf. Table 2).


Fig. 8. Soundness vs Permissiveness

Theorem 4. $FT \vdash_{\mathcal{E}, m} P \Rightarrow OT \vdash_{\mathcal{E}, m} P \Rightarrow WT \vdash_{\mathcal{E}, m} P$

Table 2 illustrates permissiveness for the state-of-the-art purely dynamic trackers. All trackers account for explicit flows, however, as illustrated by program P_0, they can be imprecise (cf. "−") due to approximation. P_1 will be rejected by full tracking, i.e NSU [5], while program P_2 will be rejected by Permissive

Upgrade [6], although none of them performs any outputs. P_3 encodes the value of the high boolean variable h into the final value of variable k through hidden implicit flows, however, k is never output. Observable tracking (column 6 and 7) correctly accepts the program, thus decreasing the number of false positives that the other trackers would otherwise report. P_0 and P_4 will be rejected by most trackers due to over-approximation. Arguably, program patterns like P_0 and P_4 are unlikely to be used, and, for trusted code, they can be fixed, e.g. by code transformations.

These considerations make a good case for using observable tracking as a permissive purely dynamic mechanism for security testing. However, programs may still leak through hidden implicit flows. The insecure program P_5 will be correctly rejected by NSU and Permissive Upgrade (cf. "✗") and, it will be correctly accepted by observable tracking.

5 Implementation and Evaluation

Implementation Our tool is a prototype built on top of the *Soot* framework [54] and it uses an intermediate bytecode language, *Jimple* [54], to implement the static transformations presented in Sect. 3. We provide a description of Jimple and discuss advanced language features in the full version [14].We implemented the code transformation for Android applications. The instrumented applications are then run using TaintDroid [24]. The code of the implementation is available online [14]. Overall, the implementation of static transformations proved to be straight-forward, due to the use of Jimple as an intermediate language and the modularity of the transformations. This indicates that this approach is indeed lightweight compared to elaborate information-flow trackers.

Use Case: Pedometer To evaluate our approach, we apply the presented implementation to an open-source step counting application [41] from the popular F-Droid repository. By default, the application performs no network output. To check if illegal flows are properly detected, we add network communication in a number of scenarios. We give condensed forms of these examples in this section to abstract from Android-specific issues regarding sensor queries; we refer the reader to the implementation's source code for the full examples [14].

Usage statistics: The step counting application may want to report usage information to the developer. However, a user may not want the actual step count to be reported to the developer. By tracking observable implicit flows, reporting usage information in a low context does not generate a false positive. However, disclosing the actual step count or reporting that the app was used on certain day in a high context will yield an error.

Declassifying average pace: The application may additionally send the average pace to a server to provide comparisons with other users. However, the actual step count should still not be disclosed.

if $(stepSensor.newStep() == $ **true**$)$
    then $steps := steps + 1$ **else skip**
out$(L, $"App used on " $+ ($**new** $Date()))$

Fig. 9. Step counter example

We implement a where-style declassification policy as described in [14].

Location information: To show the user more detailed information, we also extended the application with rudimentary location tracking to allow for displaying information such as the number of steps per city. As location information is

sensitive, our transformation ensures that nothing about the user's coordinates is leaked through explicit or observable implicit flows. We then modified the program to leak location information through hidden implicit flows as in Ex. 3. Again, our cross-copying transformation ensured that such leaks are prevented.

Use Case: JSFlow Existing information-flow tools, such as JSFlow [30], can be easily modified to enforce observable secrecy instead of noninterference. For the latest release of JSFlow, version 1.1, it was sufficient to comment out as few as 4 lines of code to change to enforcing observable secrecy.

Work on value sensitivity in the context of JSFlow [31] points out precision issues due to the No-Sensitive Upgrade policy, as in examples like ($x := 1$; **if** h **then** $x := 2$ **else skip** ; **out**($\mathbf{L}, 1$)). A standard information-flow monitor such as JSFlow would stop this program to avoid upgrading the label of x in a secret context, even though x is never output later in the program. Modifying JSFlow to enforce observable secrecy however accepts the program.

6 Related Work

Referring to the surveys on language-based information-flow security [44] and taint tracking [47], we only discuss the most closely related work.

Information-Flow Policies Contrasting noninterference [28], Volpano [57] introduces weak secrecy, a security condition for taint tracking. Schoepe et al. generalize weak secrecy by explicit secrecy [45] and enforce it by faceted values [46]. Our work explores observable secrecy as the middle ground. Similarly to weak secrecy and noninterference, observable secrecy is not a trace property.

Several authors study knowledge-based conditions [4,3,9,10]. We explore the attacker's view of program code to discriminate polices, relating in particular to the *forgetful* attackers by Askarov and Chong [2], though the exact relation is subject to further investigation. While implicit flows in the wild are important [42,33], they can also appear in trusted code [35,34]. By tracking explicit and observable implicit flows, we raise the security bar wrt. taint tracking.

Staged Analysis Our work takes inspiration from Beringer [15], who provides formal arguments of using taint tracking to enforce noninterference policies. Beringer also leverages the cross copying technique to consider hidden implicit flows. By contrast, we justify soundness of the enforcement mechanism in terms of semantic conditions like weak secrecy with respect to *uninstrumented* semantics. On the other hand, Beringer introduces a notion of *path tracking* to account for termination-sensitive noninterference, and supports the theory (for an imperative language *without* I/O) by a formalization in Coq. Our work distinguishes between malicious and trusted code, providing security conditions and enforcement mechanisms for both settings (including a prototype implementation).

Rifle [53] treats implicit flows by cross-copying program instrumentation and taint tracking, with separate taint registers for explicit and implicit flows. The focus is on efficiency, as soundness is only justified informally. Like Beringer's, our work gives formal and practical evidence for the usefulness of Rifle's ideas.

Other works leverage the cross-copying technique to enforce noninterference policies. Le Guernic [36] uses cross-copying in a hybrid monitor for noninterference, and refers to observable and hidden implicit flows as implicit and explicit indirect flows, respectively. Chugh et al. [19] present a hybrid approach to

handling JavaScript code. Their approach first computes statically a dynamic residual, which is checked at runtime in a second stage. For trusted code, Kang et al. [34] study targeted (called *culprit*) implicit flows. Bao et at. [11] identify *strict* control dependences and evaluate their effectiveness for taint tracking empirically. These works illuminate the benefits of observable implicit flows.

Dynamic Enforcement and Inlining Fenton [26] studies purely dynamic information-flow monitors. Austin and Flanagan [5] leverage No-Sensitive Upgrade [58] to enforce noninterference for JavaScript and propose Permissive Upgrade [6] to improve precision. We show that NSU can be too restrictive, and propose solutions to improve precision for malicious and trusted code. Chudnov and Naumann [18] and Magazinius et al. [37] propose information-flow monitor inlining, integrating the NSU strategy into program's code. Bielova and Rezk [17] survey recent work in (information-flow) monitor inlining. Our transformations can be seen as lightweight inlining of dynamic monitors, for (observable and/or hidden) implicit flows. Russo and Sabelfeld [43] discuss trade-offs between static and dynamic flow-sensitive analysis. We leverage their flow-sensitive monitor.

Secure multi-execution [22] and faceted values [7] enforce noninterference: programs are executed as many times as there are security levels, with outputs at each level computed by the respective runs. Barthe et al. [12] study program transformations to implement secure multi-execution. These techniques are secure by construction and provide high precision. However, they require synchronization between computations at different security levels, and face challenges for languages with side-effects and I/O. Also, they may modify the semantics and introduce crashes, thus making it difficult to detect attacks. By contrast, we focus on failstop monitoring, trading full permissiveness to avoids such pitfalls.

Static and Hybrid Enforcement Volpano et al. [56] formalize the soundness of Dennings' static analysis [21] with respect to noninterference by a security type system, extended by further work with advanced features [44]. Hunt and Sands [32] present flow-sensitive security types. Our work leverages dynamic analysis to enforce similar policies. Other analysis for information flow include program logics [10,13], model checking [8,23], abstract interpretations [27] and theorem proving [20,40]. While more precise than security type systems, these approaches may face several challenges with scalability.

Hybrid enforcement combines static and dynamic analysis. Le Guernic [36] proposes hybrid flow-sensitive mechanisms supporting for sequential and concurrent languages. Venkatakrishnan et al [55] present a hybrid monitor for a language with procedures and show that it enforces noninterference. Shroff et al. [48] present a monitor with dynamic dependency analysis for a language with heap. Tripp et al. [52] study hybrid security for JavaScript code by combining static analysis and dynamic partial evaluation. Moore and Chong [39] propose two optimizations of hybrid monitors for efficiency: selective tracking of variable security levels and memory abstractions for languages with dynamic memory. Hybrid approaches use static analysis to approximate computational effects for program paths that are not visited by a given execution. This can be challenging for languages with complex features, e.g. reflection, and unstructured control flow. We strike the balance by performing static analysis for implicit flows (ba-

sically boolean expressions) and delegating the resolution of complex features to a dynamic taint tracker.

Mobile App Security There exists a large body of works on information-flow analysis in the mobile app domain. The majority of these analysis only accounts for explicit flows. This is due to the presence of complex language features and highly dynamic lifecycles, however, for potentially malicious and trusted code, implicit flows are important to address. Our proposal in Figure 1 enables existing work to provide stronger guarantees in a flexible manner. TaintDroid [24] is a dynamic taint tracker developed to capture privacy violations in Android apps. We use TaintDroid as dynamic component in our implementation. Most static analysis works certify security with respect to weak secrecy [1,29]. Despite the great progress in improving precision, the false positive rate remains high [29].

Ernst et al. [25] propose collaborative verification of information-flow requirements for a high-integrity app store. Developers and the app store collaborate to reduce the overall verification cost. Concretely, developers provide the source code with information-flow specifications (security types), while the app store verifies their correctness. Our model is complementary and, by contrast, user-centric, allowing for more flexible policies and reducing the developers' burden.

7   Conclusion

We have presented a framework of information-flow trackers, allowing us to relate a range of enforcement from taint tracking to information-flow control. We have explored the middle ground by distinguishing malicious and trusted code and considering trade-offs between soundness and permissiveness. We have deployed the framework in a staged fashion by combining lightweight static analysis with dynamic taint tracking, enabling us to envision a secure app store architecture. We have experimented with the approach by a prototype implementation.

Future work includes dynamic security policies and case studies from the F-Droid repository. While the current framework allows for parametric policies on users' side, we conjecture that the static transformations, being transparent to the underlying policy, can be extended to handle rich dynamic policies.

References

1. Arzt, S., Rasthofer, S., Fritz, C., Bodden, E., Bartel, A., Klein, J., Traon, Y.L., Octeau, D., McDaniel, P.: Flowdroid: precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android apps. In: PLDI (2014)
2. Askarov, A., Chong, S.: Learning is change in knowledge: Knowledge-based security for dynamic policies. In: CSF (2012)
3. Askarov, A., Hunt, S., Sabelfeld, A., Sands, D.: Termination-insensitive noninterference leaks more than just a bit. In: ESORICS (2008)
4. Askarov, A., Sabelfeld, A.: Gradual release: Unifying declassification, encryption and key release policies. In: S&P (2007)
5. Austin, T.H., Flanagan, C.: Efficient purely-dynamic information flow analysis. SIGPLAN Not. (2009)
6. Austin, T.H., Flanagan, C.: Permissive dynamic information flow analysis. In: PLAS (2010)

7. Austin, T.H., Yang, J., Flanagan, C., Solar-Lezama, A.: Faceted execution of policy-agnostic programs. In: PLAS (2013)
8. Balliu, M., Dam, M., Guernic, G.L.: ENCoVer: Symbolic Exploration for Information Flow Security. In: CSF (2012)
9. Balliu, M., Dam, M., Le Guernic, G.: Epistemic Temporal Logic for Information Flow Security. In: PLAS (2011)
10. Banerjee, A., Naumann, D.A., Rosenberg, S.: Expressive Declassification Policies and Modular Static Enforcement. In: S & P (2008)
11. Bao, T., Zheng, Y., Lin, Z., Zhang, X., Xu, D.: Strict control dependence and its effect on dynamic information flow analyses. In: ISSTA (2010)
12. Barthe, G., Crespo, J.M., Devriese, D., Piessens, F., Rivas, E.: Secure multi-execution through static program transformation. In: IFIP (2012)
13. Barthe, G., D'Argenio, P.R., Rezk, T.: Secure information flow by self-composition. MSCS (2011)
14. We are family: Relating information flow trackers (Extended Version). `http://www.cse.chalmers.se/research/group/security/family`
15. Beringer, L.: End-to-end multilevel hybrid information flow control. In: Jhala, R., Igarashi, A. (eds.) APLAS (2012)
16. Biba, K.J.: Integrity considerations for secure computer systems. Tech. rep., MITRE Corp. (1977)
17. Bielova, N., Rezk, T.: A taxonomy of information flow monitors. In: POST (2016)
18. Chudnov, A., Naumann, D.A.: Information flow monitor inlining. In: CSF (2010)
19. Chugh, R., Meister, J.A., Jhala, R., Lerner, S.: Staged information flow for javascript. In: PLDI (2009)
20. Darvas, A., Hähnle, R., Sands, D.: A theorem proving approach to analysis of secure information flow. In: SPC (2005)
21. Denning, D.E., Denning, P.J.: Certification of programs for secure information flow. Commun. ACM (1977)
22. Devriese, D., Piessens, F.: Noninterference through secure multi-execution. In: S&P 2010 (2010)
23. Dimitrova, R., Finkbeiner, B., Kovács, M., Rabe, M.N., Seidl, H.: Model checking information flow in reactive systems. In: VMCAI (2012)
24. Enck, W., Gilbert, P., Han, S., Tendulkar, V., Chun, B.G., Cox, L.P., Jung, J., McDaniel, P., Sheth, A.N.: Taintdroid: An information-flow tracking system for realtime privacy monitoring on smartphones. ACM Trans. Comput. Syst. (2014)
25. Ernst, M.D., Just, R., Millstein, S., Dietl, W., Pernsteiner, S., Roesner, F., Koscher, K., Barros, P.B., Bhoraskar, R., Han, S., Vines, P., Wu, E.X.: Collaborative verification of information flow for a high-assurance app store. In: CCS (2014)
26. Fenton, J.S.: Memoryless subsystems. Comput. J. 17(2), 143–147 (1974)
27. Giacobazzi, R., Mastroeni, I.: Abstract Non-Interference: Parameterizing Non-Interference by Abstract Interpretation. In: POPL (2004)
28. Goguen, J.A., Meseguer, J.: Security policies and security models. In: S&P (1982)
29. Gordon, M.I., Kim, D., Perkins, J.H., Gilham, L., Nguyen, N., Rinard, M.C.: Information flow analysis of android applications in droidsafe. In: NDSS (2015)
30. Hedin, D., Birgisson, A., Bello, L., Sabelfeld, A.: JSFlow: Tracking information flow in JavaScript and its APIs. In: SAC (2014)
31. Hedin, D., Bello, L., Sabelfeld, A.: Value-sensitive hybrid information flow control for a javascript-like language. In: CSF (2015)
32. Hunt, S., Sands, D.: On flow-sensitive security types. In: POPL. pp. 79–90 (2006)
33. Jang, D., Jhala, R., Lerner, S., Shacham, H.: An empirical study of privacy-violating information flows in JavaScript web applications. In: CCS (2010)
34. Kang, M.G., McCamant, S., Poosankam, P., Song, D.: DTA++: dynamic taint analysis with targeted control-flow propagation. In: NDSS (2011)

35. King, D., Hicks, B., Hicks, M., Jaeger, T.: Implicit flows: Can't live with 'em, can't live without 'em. In: ICISS (2008)
36. Le Guernic, G.: Confidentiality Enforcement Using Dynamic Information Flow Analyses. Ph.D. thesis, Kansas State University (2007)
37. Magazinius, J., Russo, A., Sabelfeld, A.: On-the-fly inlining of dynamic security monitors. In: Computers & Security (2010)
38. McLean, J.: A general theory of composition for trace sets closed under selective interleaving functions. In: S&P (1994)
39. Moore, S., Chong, S.: Static analysis for efficient hybrid information-flow control. In: CSF (2011)
40. Nanevski, A., Banerjee, A., Garg, D.: Dependent type theory for verification of information flow and access control policies. ACM Trans. Program. Lang. (2013)
41. `https://f-droid.org/repository/browse/?fdid=name.bagi.levente.pedometer`
42. Russo, A., Sabelfeld, A., Li, K.: Implicit flows in malicious and nonmalicious code. Marktoberdorf Summer School (IOS Press) (2009)
43. Russo, A., Sabelfeld, A.: Dynamic vs. static flow-sensitive security analysis. In: CSF (2010)
44. Sabelfeld, A., Myers, A.C.: Language-based information-flow security. JSAC (2003)
45. Schoepe, D., Balliu, M., Pierce, B.C., Sabelfeld, A.: Explicit secrecy: A policy for taint tracking. In: EuroS&P (2016)
46. Schoepe, D., Balliu, M., Piessens, F., Sabelfeld, A.: Let's face it: Faceted values for taint tracking. In: ESORICS (2016)
47. Schwartz, E.J., Avgerinos, T., Brumley, D.: All you ever wanted to know about dynamic taint analysis and forward symbolic execution (but might have been afraid to ask). In: S&P 2010 (2010)
48. Shroff, P., Smith, S., Thober, M.: Dynamic dependency monitoring to secure information flow. In: CSF (2007)
49. SnoopWall: Flashlight Apps Threat Assessment Report. `https://www.snoopwall.com/reports` (2014)
50. Staicu, C., Pradel, M.: An empirical study of implicit information flow (2015), poster at PLDI. `https://www.informatik.tu-darmstadt.de/fileadmin/user_upload/Group_SOLA/Papers/poster-pldi2015-src.pdf`
51. `http://www.heartbleed.com` (2015)
52. Tripp, O., Ferrara, P., Pistoia, M.: Hybrid security analysis of web javascript code via dynamic partial evaluation. In: ISSTA (2014)
53. Vachharajani, N., Bridges, M.J., Chang, J., Rangan, R., Ottoni, G., Blome, J.A., Reis, G.A., Vachharajani, M., August, D.I.: RIFLE: An Architectural Framework for User-Centric Information-Flow Security. In: MICRO (2004)
54. Vallée-Rai, R., Co, P., Gagnon, E., Hendren, L.J., Lam, P., Sundaresan, V.: Soot - a java bytecode optimization framework. In: CASCR (1999)
55. Venkatakrishnan, V.N., Xu, W., DuVarney, D.C., Sekar, R.: Provably correct runtime enforcement of non-interference properties. In: ICICS (2006)
56. Volpano, D., Smith, G., Irvine, C.: A sound type system for secure flow analysis. JCS (1996)
57. Volpano, D.M.: Safety versus secrecy. In: SAS (1999)
58. Zdancewic, S.A.: Programming Languages for Information Security. Ph.D. thesis, Cornell University, Ithaca, NY, USA (2002)