

Chapter 10

Information Flow Control: Advanced

There are various reasons that a λ -observer could have access to the variables in $V_{\in \lambda}$ while a program S is executing.

- Variables in $V_{\in \lambda}$ might be used to model the channels that a λ -observer reads while S is executing. Outputs produced during execution would be implemented as updates to those variables.
- Malware concurrently executing on the same computer as S might be able to monitor the state as S executes. If the program that was co-opted by the malware has label λ then the malware would become a λ -observer.

Therefore, we now consider noninterference policies for settings where the Batch and Asynchronous assumptions (page 253) that we had adopted for TINI are replaced by:

Interactive. For a program S with variables V , a λ -observer can read variables in $V_{\in \lambda}$ initially, during execution of S , and after terminating executions of S .

Termination Detection. A λ -observer can detect the termination of S .

Termination Detection implies that a λ -observer seeing no change to the variables in $V_{\in \lambda}$ can determine whether the apparent lack of activity is because execution of S has terminated or because all of the variables that S is changing are from $V_{\notin \lambda}$.

10.1 Progress Sensitive Noninterference

Progress sensitive noninterference (PSNI) specifies that the initial values of variables in $V_{\notin \lambda}$ do not affect whether S terminates and do not affect the values

of variables $V_{\subseteq\lambda}$ in the sequence of intermediate states produced by a terminating or a non-terminating execution of S . So PSNI specifies limits on what could be revealed about the initial values of variables in $V_{\not\subseteq\lambda}$ to a λ -observer that can detect if S has terminated and that can monitor the values of variables $V_{\subseteq\lambda}$ during an execution.¹

The following program illustrates how information about the initial values of variables from $V_{\not\subseteq\lambda}$ might be revealed through the sequence of values assigned to variables in $V_{\subseteq\lambda}$ during an execution—even though nothing about the initial values of variables from $V_{\not\subseteq\lambda}$ is revealed in the final values of variables in $V_{\subseteq\lambda}$.

$$\begin{aligned}
 &\ell_1: y := 0; \ell_2: z := 0; \\
 &\ell_3: \text{if } x = 0 \text{ then } \ell_4: y := 1; \ell_5: z := 2 \\
 &\quad \quad \quad \text{else } \ell_6: x := 4; \ell_7: z := 2; \ell_8: y := 1 \\
 &\text{fi } \ell_9:
 \end{aligned} \tag{10.1}$$

The initial value of x determines the order of updates to y and z in the **if** statement. So if $\Gamma(x) = \text{H}$, $\Gamma(y) = \text{L}$, and $\Gamma(z) = \text{L}$ hold, then the sequence of values assigned to variables $y, z \in V_{\subseteq\text{L}}$ in intermediate states reveals information about the initial value of variable $x \in V_{\not\subseteq\text{L}}$. Yet program (10.1) does satisfy TINI, because the final values of y and z are the same for all initial values of x .

For formally specifying PSNI, we define a predicate $\mathcal{V} \xrightarrow[\text{ps}]{S} \mathcal{W}$ that holds if and only if executions of S from initial states that differ only in the values of variables in \mathcal{V} produce indistinguishable sequences of values for the variables in \mathcal{W} . As an example, $\{x\} \xrightarrow[\text{ps}]{S} \{y, z\}$ does not hold if S is program (10.1), because different initial values for x can cause executions where the sequences of values assigned to variables y and z during one execution can differ from the sequences assigned during another. Figure 10.1 illustrates by giving two *execution traces*, each specifying an initial state followed by the intermediate states produced during an execution of program (10.1). We write “ $x \mapsto \text{val}$ ” to indicate that a state maps variable name x to a value val , where ? represents an unknown value. So execution trace σ describes the execution from an initial state satisfying $x = 0$, and execution trace σ' describes execution from an initial state satisfying $x = 1$.

For $\{x\} \xrightarrow[\text{ps}]{S} \{y, z\}$ to hold, the sequences of updates to the values for y and z in σ and in σ' is required to be indistinguishable. We can check this requirement by constructing *projected execution traces* $\langle \sigma|_{\{y, z\}} \rangle$ and $\langle \sigma'|_{\{y, z\}} \rangle$,

¹*Progress insensitive noninterference* (PINI) has also been studied. It drops the Termination Detection assumption. PINI thus requires that the initial values of variables in $V_{\not\subseteq\lambda}$ not affect the values of variables $V_{\subseteq\lambda}$ in any prefix of the sequence of intermediate states produced by executing S . PINI, however, does allow the initial values of variables in $V_{\not\subseteq\lambda}$ to affect when or whether S terminates.

$$\begin{array}{l}
\sigma: \begin{bmatrix} x \mapsto 0 \\ y \mapsto ? \\ z \mapsto ? \\ pc \mapsto \ell_1 \end{bmatrix} \begin{bmatrix} x \mapsto 0 \\ y \mapsto 0 \\ z \mapsto ? \\ pc \mapsto \ell_2 \end{bmatrix} \begin{bmatrix} x \mapsto 0 \\ y \mapsto 0 \\ z \mapsto 0 \\ pc \mapsto \ell_3 \end{bmatrix} \begin{bmatrix} x \mapsto 0 \\ y \mapsto 0 \\ z \mapsto 0 \\ pc \mapsto \ell_4 \end{bmatrix} \begin{bmatrix} x \mapsto 0 \\ y \mapsto 1 \\ z \mapsto 0 \\ pc \mapsto \ell_5 \end{bmatrix} \begin{bmatrix} x \mapsto 0 \\ y \mapsto 1 \\ z \mapsto 2 \\ pc \mapsto \ell_9 \end{bmatrix} \\
\sigma': \begin{bmatrix} x \mapsto 1 \\ y \mapsto ? \\ z \mapsto ? \\ pc \mapsto \ell_1 \end{bmatrix} \begin{bmatrix} x \mapsto 1 \\ y \mapsto 0 \\ z \mapsto ? \\ pc \mapsto \ell_2 \end{bmatrix} \begin{bmatrix} x \mapsto 1 \\ y \mapsto 0 \\ z \mapsto 0 \\ pc \mapsto \ell_3 \end{bmatrix} \begin{bmatrix} x \mapsto 1 \\ y \mapsto 0 \\ z \mapsto 0 \\ pc \mapsto \ell_6 \end{bmatrix} \begin{bmatrix} x \mapsto 4 \\ y \mapsto 0 \\ z \mapsto 0 \\ pc \mapsto \ell_7 \end{bmatrix} \begin{bmatrix} x \mapsto 4 \\ y \mapsto 0 \\ z \mapsto 2 \\ pc \mapsto \ell_8 \end{bmatrix} \begin{bmatrix} x \mapsto 0 \\ y \mapsto 1 \\ z \mapsto 2 \\ pc \mapsto \ell_9 \end{bmatrix}
\end{array}$$

Figure 10.1: Sequences σ and σ' of intermediate states for executions of (10.1)

which give the sequences of changed values for y and z .

$$\langle \sigma|_{\{y,z\}} \rangle: \begin{bmatrix} y \mapsto ? \\ z \mapsto ? \end{bmatrix} \begin{bmatrix} y \mapsto 0 \\ z \mapsto ? \end{bmatrix} \begin{bmatrix} y \mapsto 0 \\ z \mapsto 0 \end{bmatrix} \begin{bmatrix} y \mapsto 1 \\ z \mapsto 0 \end{bmatrix} \begin{bmatrix} y \mapsto 1 \\ z \mapsto 2 \end{bmatrix} \quad (10.2)$$

$$\langle \sigma'|_{\{y,z\}} \rangle: \begin{bmatrix} y \mapsto ? \\ z \mapsto ? \end{bmatrix} \begin{bmatrix} y \mapsto 0 \\ z \mapsto ? \end{bmatrix} \begin{bmatrix} y \mapsto 0 \\ z \mapsto 0 \end{bmatrix} \begin{bmatrix} y \mapsto 0 \\ z \mapsto 2 \end{bmatrix} \begin{bmatrix} y \mapsto 1 \\ z \mapsto 2 \end{bmatrix} \quad (10.3)$$

The fourth states of (10.2) and (10.3) are different, so $\{x\} \not\stackrel{S}{\rightarrow}_{\text{ps}} \{y, z\}$ does not hold—information about the initial value of x would be revealed to an observer that is monitoring y and z .

Construction of a projected execution trace $\langle \sigma|_{\mathcal{V}} \rangle$ involves two steps: (i) variables not in \mathcal{V} are deleted from each state of σ and (ii) runs of identical states in the result are collapsed into a single state.² And predicate $\sigma =_{\mathcal{V}} \sigma'$ is defined to hold if de-stuttered trace projections $\langle \sigma|_{\mathcal{V}} \rangle$ and $\langle \sigma'|_{\mathcal{V}} \rangle$ form trace projections having the same length and having identical i^{th} states, for each i :

$$\sigma =_{\mathcal{V}} \sigma': \quad \langle \sigma|_{\mathcal{V}} \rangle = \langle \sigma'|_{\mathcal{V}} \rangle$$

We collapse runs of identical states when forming a projected execution trace $\langle \sigma|_{\mathcal{V}} \rangle$ in order to obtain the sequence of states that would be seen by an observer monitoring \mathcal{V} , since state transitions are detectable by such an observer only

²We formalize these steps, as follows. For a sequence σ states, define $\sigma[i]$ to be its i^{th} state, $\sigma[i..]$ to be the suffix starting with $\sigma[i]$, and *trace projection* $\sigma|_{\mathcal{V}}$ to be the sequence of state projections for the states in σ :

$$\sigma|_{\mathcal{V}}: \sigma[1]|_{\mathcal{V}} \sigma[2]|_{\mathcal{V}} \dots$$

The *de-stutter* operation $\langle \sigma \rangle$ collapses runs of identical states in σ is defined formally by:

$$\langle \sigma \rangle: \begin{cases} \sigma & \text{if } \text{len}(\sigma) = 1 \vee (\forall 1 < i \leq \text{len}(\sigma): \sigma[1] = \sigma[i]) \\ \sigma[1] \langle \sigma[2..] \rangle & \text{if } \text{len}(\sigma) > 1 \wedge \sigma[1] \neq \sigma[2] \\ \langle \sigma[2..] \rangle & \text{otherwise} \end{cases}$$

if the value of some variable in \mathcal{V} changes. Notice, a projected execution trace $\langle \sigma|_{\mathcal{V}} \rangle$ might have finite length even though σ has infinite length—this case arises with a non-terminating loop in which all assignment statements in the loop body update variables that are not visible to the observer.

The formal definition for $\mathcal{V} \xrightarrow[\text{ps}]{S} \mathcal{W}$ employs a function $\llbracket S \rrbracket^{\text{tr}}(s)$ that, for a deterministic program S , maps an initial state s to an execution trace that begins with s and is followed by the (possibly infinite length) sequence of intermediate states that would be produced by executing S . For defining PSNI, we are concerned with pairs of projected execution traces that start in initial states s and s' satisfying $s =_{\overline{\mathcal{V}}} s'$.

$$\mathcal{V} \xrightarrow[\text{ps}]{S} \mathcal{W}: (\forall s, s' \in \text{Init}_S: s =_{\overline{\mathcal{V}}} s' \Rightarrow \llbracket S \rrbracket^{\text{tr}}(s) =_{\mathcal{W}} \llbracket S \rrbracket^{\text{tr}}(s')) \quad (10.4)$$

So $\mathcal{V} \xrightarrow[\text{ps}]{S} \mathcal{W}$ holds if different initial values for variables in \mathcal{V} result in indistinguishable projected execution sequences for variables in \mathcal{W} .

To illustrate the use of definition (10.4), we check whether $\{x\} \xrightarrow[\text{ps}]{S} \{y, z\}$ holds when S is program (10.1). $\overline{\mathcal{V}}$ is $\{y, z\}$ (since \mathcal{V} is $\{x\}$), and \mathcal{W} is $\{y, z\}$, resulting in:

$$(\forall s, s' \in \text{Init}_S: s =_{\{y, z\}} s' \Rightarrow \llbracket S \rrbracket^{\text{tr}}(s) =_{\{y, z\}} \llbracket S \rrbracket^{\text{tr}}(s')) \quad (10.5)$$

The earlier claim that $\{x\} \xrightarrow[\text{ps}]{S} \{y, z\}$ does not hold would be confirmed by showing that (10.5) does not hold. The initial states of σ and σ' in Figure 10.1 do satisfy antecedent $s =_{\{y, z\}} s'$, but consequent $\llbracket S \rrbracket^{\text{tr}}(s) =_{\{y, z\}} \llbracket S \rrbracket^{\text{tr}}(s')$ of (10.5) does not hold, since (10.2) and (10.3) are different. As expected, we have shown that $\{x\} \xrightarrow[\text{ps}]{S} \{y, z\}$ does not hold.

We now have the building blocks needed for specifying that the values of variables $V_{\neq \lambda}$ in initial states are not allowed to affect the termination of a program S or affect the values of variables that a λ -observer can monitor.

Progress Sensitive Noninterference (PSNI). For a deterministic program S where the variables have labels from a set Λ with a partial order \sqsubseteq :

$$(\forall \lambda \in \Lambda: V_{\neq \lambda} \xrightarrow[\text{ps}]{S} V_{\sqsubseteq \lambda}) \quad \square$$

10.1.1 PSNI Enforcement

The possibility of non-termination by a **while** statement can cause an implicit flow that violates PSNI. For example, whether **while** statement ℓ_W terminates in IMP program

$$\begin{aligned} S: & x := 0; \\ & \ell_{\text{if}}: \text{if } z = 0 \text{ then skip} \\ & \quad \text{else } \ell_W: \text{while } y \neq 0 \text{ do } y := y + 1 \text{ end} \\ & \text{fi} \\ \ell: & x := 23 \end{aligned} \quad (10.6)$$

depends on which of $y > 0$ or $y \leq 0$ holds initially. Because final assignment statement ℓ is neither in the body of **while** statement ℓ_W nor in the body of **if** statement ℓ_{if} , guards $z = 0$ and $y \neq 0$ are not in $\Theta_S(\ell)$. These guards nevertheless can affect whether assignment statement ℓ is reached: for ℓ to be reached, $z = 0$ must hold initially or $y = 0$ must hold eventually so that **while** statement ℓ_W terminates. Thus, there is an implicit flow from the variables in those guards to target x of assignment statement ℓ . If $\Gamma_{\mathcal{E}}(z = 0) \sqcup \Gamma_{\mathcal{E}}(y \neq 0) \not\sqsubseteq \Gamma(x)$ holds then this implicit flow also is an illicit flow—even though ℓ complies with Θ_S -Safe Assignment Statements condition (9.26). We conclude that restrictions beyond those required for enforcing TINI are needed for enforcing PSNI.

In general, there will be an implicit flow to any statement executed after a **while** statement. This implicit flow conveys information about the values of the variables in the **while** statement guard as well as the variables in the other guards that could affect whether that **while** statement is reached. To characterize this implicit flow to a statement ℓ in a program S , we form the set $\Delta_S(\ell)$ containing those guards that affect whether ℓ can be reached in some execution of S . So the elements of $\Delta_S(\ell)$ come from (i) control-flow statements in S having a body that contains ℓ (i.e., $\Theta_S(\ell)$) and (ii) the subset \mathcal{W}_ℓ of the **while** statements in S that, by not terminating, would prevent ℓ from being reached, where G_W is the guard on a **while** statement ℓ_W :

$$\Delta_S(\ell) : \Theta_S(\ell) \cup \bigcup_{\ell_W \in \mathcal{W}_\ell} \{G_W\} \cup \Delta_S(\ell_W) \quad (10.7)$$

Beware that for a statement ℓ in the body of a **while** statement, subset \mathcal{W}_ℓ used for (10.7) might include **while** statements appearing as alternatives to ℓ or appearing after ℓ in the program text. We see this for statement ℓ : S_1 in the program fragment:

```

while  $G_1$  do if  $G_2$  then  $\ell$ :  $S_1$ 
                else while  $G_3$  do  $S_2$  end
fi;
    while  $G_4$  do  $S_3$  end
end

```

Certain choices for guards G_i and statements S_i in this program fragment could result in G_3 (a guard appearing in an alternative to ℓ) and G_4 (a guard appearing after ℓ) being members of $\Delta_S(\ell)$. Both cases arise if ℓ can be executed after the first iteration of outer-most **while** statement.

We account for illicit flows caused by guards affecting a **while** statement that does not have assignment statement ℓ in its body by replacing $\Theta_S(\ell)$ with $\Delta_S(\ell)$ in Θ_S -Safe Assignment Statements condition (9.26).

Δ_S -Safe Assignment Statements. Ensure that

$$\left(\Gamma_{\mathcal{E}}(expr) \sqcup \left(\bigsqcup_{G \in \Delta_S(\ell)} \Gamma_{\mathcal{E}}(G) \right) \right) \sqsubseteq \Gamma(w) \quad (10.8)$$

holds for each assignment statement ℓ : $w := expr$ that S executes. \square

Δ_S -Safe Assignment Statements condition (10.8) is more-stringent than Θ_S -Safe Assignment Statements condition (9.26), since $\Theta_S(\ell) \subseteq \Delta_S(\ell)$ holds according to definition (10.7) of $\Delta_S(\ell)$. This should not be surprising. TINI ignores non-terminating executions; PSNI does not. The additional guards in $\Delta_S(\ell)$ are those guards that affect whether ℓ cannot be reached due to non-termination, causing illicit flows ignored by TINI but not by PSNI.

Calculation of $\Delta_S(\ell)$ is undecidable because $\Delta_S(\ell)$ is defined in terms of statement reachability—specifically, which **while** statements can prevent ℓ from being reached. However, the calculation of $\Delta_S(\ell)$ becomes straightforward for programs S that comply with some easily-checked restrictions.

Guard Restrictions for PSNI. If for every **while** statement ℓ_W with guard G_W in a program S

- (i) $\Gamma_{\mathcal{E}}(G_W) = \perp_{\Lambda}$ holds, and
- (ii) $\Gamma_{\mathcal{E}}(G) = \perp_{\Lambda}$ holds for every guard $G \in \Theta_S(\ell_W)$.

then for every statement ℓ in S :³ $\bigsqcup_{G \in \Delta_S(\ell)} \Gamma_{\mathcal{E}}(G) = \bigsqcup_{G \in \Theta_S(\ell)} \Gamma_{\mathcal{E}}(G)$ □

So in programs satisfying restrictions (i) and (ii), Δ_S -Safe Assignment Statements condition (10.8) can be discharged by checking Θ_S -Safe Assignment Statements condition (9.26). The former is undecidable but the latter can be discharged using a straightforward static analysis of S .

Extending IMP to Relax Guard Restrictions. Guard Restrictions for PSNI requires that all iteration be controlled by variables with label \perp_{Λ} .⁴ Some relaxation of the restrictions are possible, though, for looping that is guaranteed to terminate. Only loops that sometimes terminate can cause implicit flows.

³Here is that derivation. From definition (10.7) for $\Delta_S(\ell)$ we get:

$$\bigsqcup_{G \in \Delta_S(\ell)} \Gamma_{\mathcal{E}}(G) = \bigsqcup_{G \in \Theta_S(\ell)} \Gamma_{\mathcal{E}}(G) \sqcup \bigsqcup_{\substack{\ell_W \in \mathcal{W}_{\ell} \\ G \in (\Delta_S(\ell_W) \cup \{G_W\})}} \Gamma_{\mathcal{E}}(G) \quad (10.9)$$

Restrictions (i) and (ii) in Guard Restrictions for PSNI imply that for every subset \mathcal{W} of the **while** statements in a program S we have:

$$\bigsqcup_{\substack{\ell_W \in \mathcal{W} \\ G \in (\Delta_S(\ell_W) \cup \{G_W\})}} \Gamma_{\mathcal{E}}(G) = \perp_{\Lambda}$$

So substituting into (10.9) (with \mathcal{W} instantiated by \mathcal{W}_{ℓ}) we get

$$\bigsqcup_{G \in \Delta_S(\ell)} \Gamma_{\mathcal{E}}(G) = \bigsqcup_{G \in \Theta_S(\ell)} \Gamma_{\mathcal{E}}(G) \sqcup \perp_{\Lambda}$$

⁴Compliance with this restriction on iteration also is necessary if we want to ensure that delays caused by execution of **while** statements reveal information only about the values of public variables. So if we are concerned with preventing side-channel attacks (as discussed in §12.2) then compliance with this restriction is going to be required, anyway.

A static analysis cannot determine whether a **while** statement will always terminate, due to the undecidability of the halting problem. However, a static analysis can identify iteration that always terminates if the statements for defining loops in the programming language are restricted. The **for**-loop statement

$$\text{for } v := \text{expr} \text{ to } \text{expr}' \text{ do } T \text{ end} \quad (10.10)$$

would be an example of such a statement provided it is considered syntactically correct only if: (i) v is an integer variable and considered an assignment statement target; (ii) expr and expr' are integer-valued expressions; and (iii) body T does not contain any assignment statements or **for**-loop statements with v as a target or with any variable in $\text{Vars}(\text{expr}')$ as a target. In executions of **for**-loop statement (10.10), body T is executed between zero and some bounded number of times. The guard for **for**-loop statement (10.10) is defined to be $\text{expr} \leq v \leq \text{expr}'$, since this predicate holds each time body T starts executing. So for every statement ℓ in body T of a **for**-loop statement appearing in a program S , $\text{expr} \leq v \leq \text{expr}' \in \Theta_S(\ell)$ will hold.

10.1.2 Enforcing PSNI with Typing Rules

To use type-correctness for enforcing PSNI, we employ judgements $\Gamma, \gamma \vdash_{\text{ps}} S$, where typing context Γ gives a label $\Gamma(v)$ to each variable $v \in \text{Vars}(S)$, control context γ is a label, and S is an IMP program. Validity for judgements $\Gamma, \gamma \vdash_{\text{ps}} S$ is defined as follows.

Valid Judgements for PSNI. Judgement $\Gamma, \gamma \vdash_{\text{ps}} S$ for a deterministic program S is *valid* if and only if

- (i) $(\forall \lambda \in \Lambda: V_{\neq \lambda} \xrightarrow[S]{\text{ps}} V_{\in \lambda})$.
- (ii) $\gamma \sqsubseteq \Gamma(w)$ holds for target w of every assignment statement in S . \square

A program S is defined to be type-correct if judgement $\Gamma, \perp_{\Lambda} \vdash_{\text{ps}} S$ can be derived using the typing rules in Figure 10.2.⁵ Each of these typing rules derives a judgement that is valid whenever all of the rule's hypotheses are valid. So if a program S is type-correct then $\Gamma, \perp_{\Lambda} \vdash_{\text{ps}} S$ is valid and, due to (i) in Valid Judgements for PSNI, the program will satisfy PSNI.

The typing rules in Figure 10.2 resemble the TINI typing rules in Figure 9.8, but with rule **WHILE** modified to ensure that every **while** statement ℓ_W in a type-correct program satisfies the requirements of Guard Restrictions for PSNI (page 280). In particular, the use of \perp_{Λ} for the control context in the hypothesis

⁵Here is a typing rule in case IMP is being extended with **for**-loop statements.

$$\text{FOR: } \frac{\begin{array}{l} v \notin \text{tgts}(S), \text{tgts}(S) \cap \text{Vars}(\text{expr}') = \emptyset, \\ \Gamma_{\mathcal{E}}(\text{expr}) = \lambda, \Gamma_{\mathcal{E}}(\text{expr}') = \lambda', \Gamma(v) = \lambda'', \\ \gamma \sqcup \lambda \sqcup \lambda' \sqsubseteq \lambda'', \Gamma, \gamma \sqcup \lambda \sqcup \lambda' \sqcup \lambda'' \vdash_{\text{ps}} S \end{array}}{\Gamma, \gamma \vdash_{\text{ps}} \text{for } v := \text{expr} \text{ to } \text{expr}' \text{ do } S \text{ end}}$$

where $\text{tgts}(S)$ denotes the set of variables that are the targets of assignment statements in S .

$$\begin{array}{c}
\text{SKIP:} \frac{}{\Gamma, \gamma \vdash_{\text{ps}} \text{skip}} \qquad \text{ASSIGN:} \frac{\gamma \sqcup \Gamma_{\mathcal{E}}(expr) \sqsubseteq \Gamma(v)}{\Gamma, \gamma \vdash_{\text{ps}} v := expr} \\
\\
\text{IF:} \frac{\Gamma_{\mathcal{E}}(expr) = \lambda, \quad \Gamma, \gamma \sqcup \lambda \vdash_{\text{ps}} S, \quad \Gamma, \gamma \sqcup \lambda \vdash_{\text{ps}} S'}{\Gamma, \gamma \vdash_{\text{ps}} \text{if } expr \text{ then } S \text{ else } S' \text{ fi}} \\
\\
\text{WHILE:} \frac{\Gamma_{\mathcal{E}}(expr) = \perp_{\Lambda}, \quad \Gamma, \perp_{\Lambda} \vdash_{\text{ps}} S}{\Gamma, \perp_{\Lambda} \vdash_{\text{ps}} \text{while } expr \text{ do } S \text{ end}} \qquad \text{SEQ:} \frac{\Gamma, \gamma \vdash_{\text{ps}} S, \quad \Gamma, \gamma \vdash_{\text{ps}} S'}{\Gamma, \gamma \vdash_{\text{ps}} S; S'}
\end{array}$$

Figure 10.2: Typing rules for PSNI compliance

and conclusion of rule **WHILE** forces $\Gamma_{\mathcal{E}}(G) = \perp_{\Lambda}$ to hold for all $G \in \Theta_S(\ell_W)$. Since every type-correct **while** statement ℓ_W is being restricted in this way, we can conclude that $\Gamma_{\mathcal{E}}(G) = \perp_{\Lambda}$ holds for all $G \in \Delta_S(\ell_W)$. Rule **ASSIGN**, which ensures that Θ_S -Safe Assignment Statements condition (9.26) holds for an assignment statement, then also ensures that Δ_S -Safe Assignment Statements condition (10.8) holds for that assignment statement.

10.1.3 Dynamic Enforcement of PSNI

By monitoring state changes as execution proceeds, a λ -observer might be able to detect that a dynamic enforcement mechanism has blocked an execution. To account for this, we add

Blocking Detection. A λ -observer can detect when execution of a program becomes blocked by a dynamic enforcement mechanism.

to the Interactive assumption and the Termination Detection assumption.

Detecting that execution is blocked can result in an illicit implicit flow, as the following IMP program illustrates if $0 \leq x_H$, $\Gamma(i_L) = L$, $\Gamma(x_H) = H$, and $H \notin L$ hold.

$$\begin{array}{l}
S: \quad i_L := 0; \\
\quad \text{while } i_L \leq N \text{ do} \\
\quad \quad \text{if } x_H = i_L \text{ then } \ell: x_L := i_L \text{ else skip fi} \\
\quad \quad i_L := i_L + 1 \\
\quad \text{end}
\end{array} \tag{10.11}$$

Assignment statement ℓ does not comply with Δ_S -Safe Assignment Statements condition (10.8), because ℓ appears in the body of an **if** statement with guard $x_H = i_L$ and $\Gamma_{\mathcal{E}}(x_H = i_L) \not\sqsubseteq \Gamma(x_L)$ holds. Moreover, a dynamic enforcement mechanism that blocks execution upon reaching ℓ would not prevent an L-observer from deducing the initial value of x_H by monitoring i_L —the last value of i_L that

the L-observer reads before detecting that execution has blocked is the initial value of x_H .

We conclude that a dynamic enforcement mechanism for PSNI must not only prevent illicit explicit and implicit flows caused by assignment statements. It must also prevent illicit *detection flows*, wherein detecting that an execution has blocked allows a λ -observer to learn information about the initial value of a variable in $V_{\neq\lambda}$ based on the values assigned to variables in $V_{\in\lambda}$ during the execution. An illicit detection flow occurs whenever (i) a λ -observer is able to determine that execution became blocked upon reaching some statement ℓ , and (ii) $\Gamma(v) \notin \lambda$ holds for some variable $v \in \text{Vars}(G)$ and some guard $G \in \Delta_S(\ell)$ that was evaluated prior to execution becoming blocked. For example, program (10.11) exhibits an illicit detection flow because reaching assignment statement ℓ is the only reason that execution could become blocked, guard $x_H = i_L$ is evaluated before execution becomes blocked, and $\Gamma(x_H) \notin L$ holds.

The following requirements, then, would be sufficient to ensure that a reference monitor enforces PSNI without causing illicit detection flows.

- (i) To avoid illicit flows from an assignment statement ℓ : $w := \text{expr}$, execution of the monitored program must be blocked before reaching ℓ if ℓ does not satisfy Δ_S -Safe Assignment Statements condition (10.8).
- (ii) To avoid illicit flows due to non-termination of a **while** statement ℓ_W with guard G_W , execution of the monitored program must be blocked before reaching ℓ_W if $\Gamma_{\mathcal{E}}(G_W) \neq \perp_{\Lambda}$ holds or if $\Gamma_{\mathcal{E}}(G) \neq \perp_{\Lambda}$ holds for some $G \in \Delta_S(\ell_W)$ that has been evaluated.
- (iii) To avoid illicit detection flows, execution must not be blocked upon reaching a statement ℓ unless $\Gamma_{\mathcal{E}}(G) = \perp_{\Lambda}$ holds, for all guards $G \in \Delta_S(\ell)$.

Notice, requirements (i) and (ii) can be satisfied by blocking execution of the monitored program at a control point that was reached before a problematic statement is executed. Even given this flexibility to block execution early, requirement (iii) presents an implementation challenge, since a reference monitor must operate ignorant of code that has not yet been executed. So a reference monitor that is invoked when a control-flow statement ℓ is reached cannot determine if the body of ℓ contains a statement that should cause execution to be blocked in anticipation of the need to satisfy requirements (i) or (ii).

One solution to the implementation challenge for requirement (iii) is to employ a reference monitor \mathcal{R}_{PS} that is conservative and, therefore, proceeds as if the body of every control-flow statement contains a statement that would cause execution to be blocked for satisfying requirements (i) and (ii). Such a reference monitor would block any execution upon reaching a control-flow statement having a guard G where $\Gamma_{\mathcal{E}}(G) \neq \perp_{\Lambda}$ holds. Given this (admittedly draconian) restriction, $\Gamma_{\mathcal{E}}(G) = \perp_{\Lambda}$ holds for all guards that have been evaluated. Compliance with requirements (ii) and (iii) is thus guaranteed. Also, because no guard G will be evaluated where $\Gamma_{\mathcal{E}}(G) \neq \perp_{\Lambda}$ holds, all guards will have label \perp_{Λ} in composition $\mathcal{R}_{PS} \triangleright S$ of the reference monitor \mathcal{R}_{PS} and S . So program

upon S reaching •	action to be performed
• $w := expr$	require ($\Gamma_{\mathcal{E}}(expr) \sqsubseteq \Gamma(w)$)
• if G then ...	require ($\Gamma_{\mathcal{E}}(G) = \perp_{\Lambda}$)
• while G do ...	require ($\Gamma_{\mathcal{E}}(G) = \perp_{\Lambda}$)

Figure 10.3: Reference monitor \mathcal{R}_{PS} for PSNI

$\mathcal{R}_{PS} \triangleright S$ satisfies Guard Restrictions for PSNI (page 280) and, therefore Δ_S -Safe Assignment Statements condition (10.8) is equivalent to $\Gamma_{\mathcal{E}}(expr) \sqsubseteq \Gamma(w)$, which discharges requirement (i). The actions to implement reference monitor \mathcal{R}_{PS} are given in Figure 10.3.

By design, any program that reference monitor \mathcal{R}_{PS} executes without blocking is type-correct according to the the typing rules in Figure 10.2. A type-correct program might be blocked by \mathcal{R}_{PS} , though. As an illustration, if $\Gamma(x_H) \neq \perp_{\Lambda}$ holds then \mathcal{R}_{PS} would block execution by

$$\text{if } x_H = 0 \text{ then } x_H := x_H + 1 \text{ else skip fi} \quad (10.12)$$

even though this program is type-correct. This establishes that \mathcal{R}_{PS} is less permissive than type-checking is for PSNI enforcement.

Moreover, no reference monitor can be more permissive than \mathcal{R}_{PS} , because (by definition) the statements in the body of a control flow statement cannot be known by any reference monitor when that control flow statement is reached. So, for example, if $x_H = 0$ holds in initial states, then (by definition) a reference monitor cannot distinguish between an execution of (10.12), a program that satisfies PSNI, and an execution of

```

if  $x_H = 0$  then  $x_H := x_H + 1$ 
    else while true do skip end
fi

```

a program that violates PSNI. For enforcing PSNI, the improved permissiveness that is the goal of dynamic enforcement is not achieved by using a reference monitor.

Hybrid Enforcement for PSNI. With a *hybrid enforcement mechanism*, a reference monitor invokes a static analyzer before and/or during executions of the monitored program. Results from the static analyzer enable reference monitor actions to take into account code that could be executed in the future and code that was an alternative to code had been executed.

Our implementation of a PSNI hybrid enforcement mechanism employs a reference monitor \mathcal{R}_{HPS} that invokes a static analyzer $\mathcal{T}_S(\cdot)$, where $\mathcal{T}_S(T)$ returns *true* if fragment T of program S is guaranteed to terminate because T

contains (i) no **while** statements and (ii) no assignment statement that would be blocked by \mathcal{R}_{HPS} .

Figure 10.4 details the actions comprising \mathcal{R}_{HPS} . Execution of control-flow statements are blocked by \mathcal{R}_{HPS} , as follows.

- A **while** statement with guard G_W is blocked if $\Gamma_{\mathcal{E}}(G_W) \neq \perp_{\Lambda}$ holds.
- An **if** statement with guard G_{IF} and body T is blocked if $\Gamma_{\mathcal{E}}(G_{IF}) \neq \perp_{\Lambda}$ holds and T contains a **while** statement that \mathcal{R}_{HPS} would block or T contains an assignment statement that \mathcal{R}_{HPS} would block.

This blocking prevents illicit detection flows, because a blocked execution reveals no information about results obtained by evaluating any guard G that satisfies $\Gamma_{\mathcal{E}}(G) \neq \perp_{\Lambda}$. In addition, \mathcal{R}_{HPS} ensures compliance with Guard Restrictions for PSNI, which allows Δ_S -Safe Assignment Statements condition (10.8) to be checked for an assignment statement ℓ by using a stack where $\text{top}()$ satisfies (9.30). \mathcal{R}_{HPS} implements such a stack with the actions it performs when an **if** or a **fi** is reached. (No change to the stack is needed when a **while** statement is entered, because the guard label will be \perp_{Λ} , so (9.30) continues to hold.)

\mathcal{R}_{HPS} and \mathcal{R}_{PS} both block the same **while** statements. Some **if** statements that \mathcal{R}_{PS} blocks are not blocked by \mathcal{R}_{HPS} . For example, \mathcal{R}_{HPS} does not block execution of (10.12) which, as we observed earlier, \mathcal{R}_{PS} blocks. So \mathcal{R}_{HPS} is more permissive than \mathcal{R}_{PS} . Moreover, since any static analysis can be formulated as a type system, hybrid enforcement can be as permissive as any type system or other form of static analysis. To be more permissive than this, a hybrid enforcement mechanism would have to identify and ignore program fragments that would violate PSNI but cannot be reached. No static analyzer can perform such an analysis, since it requires solving an undecidable problem.

10.2 Flow-Sensitive Noninterference

We now consider programs in which some variables have labels that vary. A *fixed variable* has the same fixed label throughout all executions; a *flexible variable*,

upon S reaching •	action to be performed
• if G then S' else S'' fi	$\text{require}(\Gamma_{\mathcal{E}}(G) = \perp_{\Lambda} \vee (\mathcal{T}_S(S') \wedge \mathcal{T}_S(S''))$ $\text{push}(sps, \text{top}(sps) \sqcup \Gamma_{\mathcal{E}}(G))$
... fi •	$\text{pop}(sps)$
• while G do S end	$\text{require}(\Gamma_{\mathcal{E}}(G) = \perp_{\Lambda})$
• $w := \text{expr}$	$\text{require}(\text{top}(sps) \sqcup \Gamma_{\mathcal{E}}(\text{expr}) \sqsubseteq \Gamma(w))$

Figure 10.4: Hybrid enforcement \mathcal{R}_{HPS} for PSNI

which we distinguish with tilde-topped names (e.g., \tilde{y}), has a label that can be changed during an execution. Changing the label of a flexible variable allows that variable to be used for different purposes during an execution.

To specify label assignments for both fixed and flexible variables, program states come with a function-valued variable $\widehat{\Gamma}(\cdot)$. The label that a program state s gives to a fixed variable v is $s.\widehat{\Gamma}(v)$, the label s gives a flexible variable \tilde{v} is $s.\widehat{\Gamma}(\tilde{v})$, and the label s gives to an expression E is $s.\widehat{\Gamma}_{\mathcal{E}}(E)$ defined by

$$s.\widehat{\Gamma}_{\mathcal{E}}(E): \begin{cases} \perp_{\Lambda} & \text{if } E \text{ is a constant } c \\ s.\widehat{\Gamma}(v) & \text{if } E \text{ is a variable } v \\ \bigsqcup_{1 \leq i \leq n} s.\widehat{\Gamma}_{\mathcal{E}}(E_i) & \text{if } E \text{ is } f(E_1, E_2, \dots, E_n) \end{cases}$$

We write just $\widehat{\Gamma}(v)$ or $\widehat{\Gamma}_{\mathcal{E}}(E)$ when the state is clear from the context or when discussing fixed variables (since they have the same label in all states).

For defining flow-sensitive noninterference (F-PSNI), we generalize state projections $s|_{\mathcal{W}}$ and then use $\mathcal{V} \xrightarrow[\text{ps}]{S} \mathcal{W}$. Definition (10.4) for $\mathcal{V} \xrightarrow[\text{ps}]{S} \mathcal{W}$, employs state projections $s|_{\mathcal{W}}$ for comparing values of a fixed set \mathcal{W} of variables in corresponding program states in certain pairs of executions. For F-PSNI, the variables that are compared depends on their label assignments in the states being compared. We materialize that dependence by using state projections $s|_{F(\cdot)}$, where $F(\cdot)$ is a function giving the set $F(s)$ of variables to use for constructing the state projection. The value of an anchored or flexible variable v in a state projection $s|_{F(\cdot)}$ is thus defined by:

$$s|_{F(\cdot)}.v: \begin{cases} s.v & \text{if } v \in F(s) \\ ? & \text{otherwise} \end{cases}$$

Then, to define F-PSNI, instead of using sets $V_{\varepsilon\lambda}$ and $V_{\neq\lambda}$ for the state projections in PSNI formal definition (10.4), we use the following functions $\widehat{V}_{\varepsilon\lambda}(\cdot)$ and $\widehat{V}_{\neq\lambda}(\cdot)$ from program states to sets of fixed and flexible variables:

$$\widehat{V}_{\varepsilon\lambda}(s): \{v \in \text{Vars}(S) \mid s.\widehat{\Gamma}(v) \varepsilon \lambda\} \quad \widehat{V}_{\neq\lambda}(s): \{v \in \text{Vars}(S) \mid s.\widehat{\Gamma}(v) \neq \lambda\}$$

We get a definition for F-PSNI by combining these elements:

F-PSNI. For a deterministic program S having states that assign values to fixed variables, to flexible variables, and to the function $\widehat{\Gamma}$ mapping variable names to labels from a set Λ with a partial order ε :

$$(\forall \lambda \in \Lambda: \widehat{V}_{\neq\lambda}(\cdot) \xrightarrow[\text{ps}]{S} \widehat{V}_{\varepsilon\lambda}(\cdot)) \quad \square$$


10.2.1 Flexible Variables and Enforcement

We might expect that an update $\ell: \tilde{y} := \text{expr}$ (say) to a flexible variable \tilde{y} would not violate F-PSNI if executing ℓ also gives \tilde{y} a new label:

$$\widehat{\Gamma}(\tilde{y}) := \widehat{\Gamma}_{\mathcal{E}}(\text{expr}) \sqcup \bigsqcup_{G \in \Delta_S(\ell)} \widehat{\Gamma}_{\mathcal{E}}(G). \quad (10.13)$$

However, as we shall see, such a label update can cause an F-PSNI violation.

The problem can be seen in a simple IMP program, where the set of labels is Λ_{LH} and $\widehat{\Gamma}(x_{\text{H}}) = \text{H}$ holds. Two executions of such a program are depicted in the following diagram—one execution where $x_{\text{H}} = 0$ holds initially and the other where it doesn't.



Both executions terminate with the same value for \tilde{y} but different labels. Letting s and s' denote those two final states, $s|_{\tilde{V}_{\text{EL}}(\cdot)} = s'|_{\tilde{V}_{\text{EL}}(\cdot)}$ does not hold, since only one of $s|_{\tilde{V}_{\text{EL}}(\cdot)}$ and $s'|_{\tilde{V}_{\text{EL}}(\cdot)}$ gives a value to \tilde{y} . So, by definition, F-PSNI does not hold. The value of guard $x_{\text{H}} = 0$ is being leaked through the different label assignments.

A straightforward way to avoid such leaks is to associate a single label assignment Γ_{ℓ} with each control point ℓ by adjusting the exit control points for control-flow statements, as follows.

Preventing Variation in Label Assignments. For each flexible variable \tilde{x} , give \tilde{x} a more restrictive label if needed to eliminate variation in label $\widehat{\Gamma}(\tilde{x})$ when the exit control point ℓ for a control-flow statement is reached. \square

In (10.14), for example, variation in the label for \tilde{y} at the exit control point for the **if** statement would be eliminated by changing that label from **L** to **H** before or after the **skip**. Executions from any initial values of x_{H} then would terminate satisfying $\widehat{\Gamma}(\tilde{y}) = \text{H}$ (without also granting **L**-observers additional access).

Leaks also can be avoided if there are different label assignments when the exit control point for a control-flow statement is reached provided the following holds.⁶


Matching Outcomes to Guards. If the body of a control-flow statement S with guard G_S contains an assignment statement to a flexible variable \tilde{x} then make label $\widehat{\Gamma}(\tilde{x})$ more restrictive, if necessary, to ensure that $\widehat{\Gamma}_{\mathcal{E}}(G_S) \sqsubseteq \widehat{\Gamma}(\tilde{x})$ holds when the exit control point for S is reached. \square

⁶To see that Matching Outcomes to Guards stops a label $\widehat{\Gamma}(\tilde{y})$ from leaking the value of guard G_S to a λ -observer, we analyze scenarios in which $\widehat{\Gamma}_{\mathcal{E}}(G_S) \not\sqsubseteq \lambda$ holds and, therefore, learning G_S is a leak. Consider two executions of a control-flow statement S , where guard G_S was evaluated to different values and the executions terminated in states giving different labels λ_1 and λ_2 to \tilde{y} . For a λ -observer to detect a difference in these labels then either $\lambda_1 \sqsubseteq \lambda$ and $\lambda_2 \not\sqsubseteq \lambda$ must hold or $\lambda_2 \sqsubseteq \lambda$ and $\lambda_1 \not\sqsubseteq \lambda$ must hold. Without loss of generality, we consider the former and show that it is not possible because it results in a contradiction. Matching Outcomes to Guards requires that $\widehat{\Gamma}_{\mathcal{E}}(G_S) \sqsubseteq \lambda_1$ holds. So, by transitivity, we conclude $\widehat{\Gamma}_{\mathcal{E}}(G_S) \sqsubseteq \lambda$, but that contradicts the earlier assumption that $\widehat{\Gamma}_{\mathcal{E}}(G_S) \not\sqsubseteq \lambda$ holds.

Note that in implementing Matching Outcomes to Guards, an execution path that itself has no assignment statement to a flexible variable \tilde{y} (say) might nevertheless require an update to label $\widehat{\Gamma}(\tilde{y})$ if alternative execution paths by that control-flow statement do have assignment statements where \tilde{y} is the target. For IMP programs, this implies:

- An **if** statement's **then** or **else** alternative might have to update the label of a flexible variable because that variable is a target in the other alternative of that **if** statement.
- A **while** statement S that does not execute its body (because its guard was initially *false*) might still have to update the label of a flexible variable that is a target in the body of S .

To illustrate an application of Matching Outcomes to Guards, consider a program having a set $\Lambda_{\text{LMH}} = \{\text{L}, \text{M}, \text{H}\}$ of labels, where $\text{L} \sqsubset \text{M} \sqsubset \text{H}$ holds. Fixed variables x_{M} and x_{H} in this program satisfy $\widehat{\Gamma}(x_{\text{M}}) = \text{M}$ and $\widehat{\Gamma}(x_{\text{H}}) = \text{H}$. Two executions are depicted.



The updates to label $\widehat{\Gamma}(\tilde{y})$ inserted after each assignment statement in (10.15) are to comply with (10.13); the update after the **skip** statement is to comply with Matching Outcomes to Guards. To see that the update after the **skip** is necessary, observe that guard $x_{\text{M}} = 0$ affects label $\widehat{\Gamma}(\tilde{y})$ by affecting whether assignment statement ℓ is executed. So $\widehat{\Gamma}_{\mathcal{E}}(x_{\text{M}} = 0) \sqsubseteq \widehat{\Gamma}(\tilde{y})$ needs to hold when the **if** statement terminates—whether the **then** alternative or the the **else** alternative has been taken.

Another way to justify the update to $\widehat{\Gamma}(\tilde{y})$ after the **skip** in (10.15) is by analyzing what a λ -observer can learn when the **fi** is reached. The table to the right lists that $\widehat{\Gamma}(\tilde{y})$ is either M or H at that control point. Consequently (as listed in the L column of the table), an L-observer

state at if	$\widehat{\Gamma}(\tilde{y})$ at fi	λ -observer read \tilde{y} ?		
		L	M	H
$x_{\text{M}} = 0$	H	no	no	yes
$x_{\text{M}} \neq 0$	M	no	yes	yes

cannot read \tilde{y} no matter what value guard $x_{\text{M}} = 0$ had. We conclude that an L-observer can learn nothing about the guard evaluation from the different labels for \tilde{y} when execution reaches **fi**. But an M-observer can detect whether $x_{\text{M}} = 0$ evaluated to *true* by attempting to read \tilde{y} , since $\widehat{\Gamma}(\tilde{y}) = \text{M}$ holds if guard $x_{\text{M}} = 0$ evaluated to *false* but $\widehat{\Gamma}(\tilde{y}) = \text{H}$ holds if the guard evaluated to *true*. This difference does not cause a leak, though, because $\widehat{\Gamma}_{\mathcal{E}}(x_{\text{M}} = 0) = \text{M}$ holds, so an

M-observer is allowed to learn the value of guard $x_M = 0$. Finally, an H-observer is able to read \tilde{y} whether the **then** or the **else** alternative was executed, and the H-observer is also allowed to read guard $x_M = 0$.

10.2.2 Enforcing F-PSNI with Typing Rules

Typing rules can ensure compliance with Preventing Variation in Label Assignments. In that case, label assignments for the control points in a program S could be conveyed by a *label annotated program*—the program text for S along with a label assignment Γ_ℓ for each control point ℓ in S .⁷

One way to represent a label annotated program is by using the notation $S/\widehat{\Gamma}^*$ where S identifies a program and $\widehat{\Gamma}^*$ is a function that maps each control point ℓ in S (and perhaps other control points) to a label assignment. So upon reaching control point ℓ , function-valued variable $\widehat{\Gamma}(\cdot)$ would be set to $\widehat{\Gamma}^*(\ell)$.

Another representation is to insert a *label annotation* at each control point in the program text. The label annotation “ $\{\Gamma_\ell\}$ ” inserted at a control point ℓ indicates that function-valued variable $\widehat{\Gamma}(\cdot)$ be set to Γ_ℓ whenever control point ℓ is reached during an execution. IMP programs have a control point before and after each statement, so a label annotated IMP program can be represented by inserting a label annotation before and after each statement in the program text. Here is an example of a label annotated IMP program, assuming that the S_i are IMP assignment or **skip** statements and the Γ_i are label assignments.⁸

$$\begin{array}{l} \{\Gamma_1\} \ell_1: S_1 \{\Gamma_2\} \\ \ell_2: \text{if } G \text{ then } \{\Gamma_3\} \ell_3: S_2; \{\Gamma_4\} \\ \quad \text{else } \{\Gamma_5\} \ell_4: S_3 \{\Gamma_6\} \\ \text{fi } \{\Gamma_7\} \end{array}$$

It sometimes is helpful to combine the two kinds of representations for label annotated programs, as follows.

$\{\Gamma\} S/\widehat{\Gamma}^*$: label assignment Γ is associated with the entry control point for S , and the label assignments given by $\widehat{\Gamma}^*$ are used for all other control points in S .

$S/\widehat{\Gamma}^* \{\Gamma'\}$: label assignment Γ' is associated with the exit control point for S , and the label assignments given by $\widehat{\Gamma}^*$ are used for all other control points in S .

⁷Flexible variables might have to be given more restrictive labels than required by Matching Outcomes to Guards. So there will be programs that satisfy F-PSNI but are not type-correct. Program (10.15) is an example. The label update $\widehat{\Gamma}(\tilde{y}) := M$ performed after the **skip** statement is less restrictive from the label update $\widehat{\Gamma}(\tilde{y}) := H$ that would be required by Preventing Variation in Label Assignments.

⁸The placement of label annotation “ $\{\Gamma_2\}$ ” immediately after S_1 and before **if** statement ℓ_2 illustrates that IMP sequential compositions $S; S'$ have a single control point serving both as the exit control point of S and as the entry control point of S' . Also notice that the exit control point for an **if** statement is distinct from the exit control points for its **then** alternative and its **else** alternative. Therefore, the exit control point for an **if** statement always will have two direct predecessor control points.

$\{\Gamma\} S / \widehat{\Gamma}^* \{\Gamma'\}$: label assignment Γ is associated with the entry control point for S , label assignment Γ' is associated with the exit control point for S , and the label assignments given by $\widehat{\Gamma}^*$ are used for all other control points in S .

Judgements and Typing Rules for F-PSNI. To use type-checking for enforcing F-PSNI, we employ typing rules that derive judgements $\gamma \vdash_{\widehat{\Gamma}^*} S / \widehat{\Gamma}^*$ for label annotated programs $S / \widehat{\Gamma}^*$. Validity of these judgements is defined as follows.

Valid Judgements for FSNI. Judgement $\gamma \vdash_{\widehat{\Gamma}^*} S / \widehat{\Gamma}^*$, where S is a deterministic program, is *valid* if and only if

- (i) $(\forall \lambda \in \Lambda: \widehat{V}_{\neq \lambda}(\cdot) \xrightarrow[S]{\text{ps}} \widehat{V}_{\in \lambda}(\cdot)).$
- (ii) $\gamma \sqsubseteq \Gamma(w)$ holds for target w of every assignment statement in S . \square

A label annotated program $S / \widehat{\Gamma}^*$ is considered type-correct if the typing rules in Figure 10.5 can derive the judgement $\perp_{\Lambda} \vdash_{\widehat{\Gamma}^*} S / \widehat{\Gamma}^*$. By design, those typing rules only derive valid judgements from valid hypotheses. So a type-correct label annotated program $S / \widehat{\Gamma}^*$ will satisfy F-PSNI, due to (i) in the definition validity for $\perp_{\Lambda} \vdash_{\widehat{\Gamma}^*} S / \widehat{\Gamma}^*$.

The typing rules for control-flow statements are the key to having a label annotated program be the conclusion of each typing rule, thereby ensuring that different executions reaching the same control point give the same labels to the flexible variables. Different label assignments arise when a given control point is reached only if different execution paths are taken. That requires executing a control-flow statement. But rule IF mandates the same label assignment for the exit control point of the **then** alternative as for the **else** alternative, and rule WHILE mandates the same label assignment before each evaluation of the **while** statement guard. So in type-correct programs, different label assignments cannot arise when a given control point is reached.

To enforce F-PSNI, the typing rules also must reject programs that would allow λ -observers to distinguish executions starting from initial values that differ only in the values of anchored or flexible variables v satisfying $\widehat{\Gamma}(v) \notin \lambda$. The PSNI typing rules in Figure 10.2 achieve this for programs that only have fixed variables. The F-PSNI typing rules in Figure 10.5 can be seen as generalizing these PSNI typing rules in order to accomodate flexible variables. Specifically, judgements in the F-PSNI typing rules have label annotated programs in place of the programs appearing in the judgements of the PSNI typing rules, the F-PSNI typing rules include label annotations giving the mandates discussed above regarding label updates due to control-flow statements, and rule F-ASSIGN has been added for characterizing the label assignment produced by assignment statements that have flexible variables as targets.

Figure 10.6 illustrates the use of the F-PSNI typing rules, giving a derivation to show that program (10.14) is type-correct and, therefore, satisfies F-PSNI.⁹

⁹The set of labels used by this program is $\Lambda_{\text{LH}} = \{\text{L}, \text{H}\}$. Therefore $\perp_{\Lambda_{\text{LH}}}$ is L, so type-correctness requires derivation of judgement $\text{L} \vdash_{\widehat{\Gamma}^*} S / \widehat{\Gamma}^*$.

$$\begin{array}{c}
\text{SKIP:} \frac{}{\gamma \vdash_{\text{ps}} \{\Gamma\} \text{ skip } \{\Gamma\}} \quad \text{SEQ:} \frac{\gamma \vdash_{\text{ps}} S_1 / \widehat{\Gamma}_1^* \{\Gamma\}, \gamma \vdash_{\text{ps}} \{\Gamma\} S_2 / \widehat{\Gamma}_2^*}{\gamma \vdash_{\text{ps}} S_1 / \widehat{\Gamma}_1^*; \{\Gamma\} S_2 / \widehat{\Gamma}_2^*} \\
\\
\text{A-ASSIGN:} \frac{\gamma \sqcup \Gamma_{\mathcal{E}}(expr) \sqsubseteq \Gamma(w)}{\gamma \vdash_{\text{ps}} \{\Gamma\} x := expr \{\Gamma\}} \quad \text{F-ASSIGN:} \frac{}{\gamma \vdash_{\text{ps}} \{\Gamma\} \tilde{v} := expr \{\Gamma[\tilde{v} \leftarrow \gamma \sqcup \Gamma_{\mathcal{E}}(expr)]\}} \\
\\
\text{IF:} \frac{\Gamma_{\mathcal{E}}(expr) = \lambda, \gamma \sqcup \lambda \vdash_{\text{ps}} \{\Gamma\} S_1 / \widehat{\Gamma}_1^* \{\Gamma'\}, \gamma \sqcup \lambda \vdash_{\text{ps}} \{\Gamma\} S_2 / \widehat{\Gamma}_2^* \{\Gamma'\}}{\gamma \vdash_{\text{ps}} \{\Gamma\} \text{ if } expr \text{ then } \{\Gamma\} S_1 / \widehat{\Gamma}_1^* \{\Gamma'\} \text{ else } \{\Gamma\} S_2 / \widehat{\Gamma}_2^* \{\Gamma'\} \text{ fi } \{\Gamma'\}} \\
\\
\text{WHILE:} \frac{\Gamma_{\mathcal{E}}(expr) = \perp_{\Lambda}, \perp_{\Lambda} \vdash_{\text{ps}} \{\Gamma\} S / \widehat{\Gamma}^* \{\Gamma\}}{\perp_{\Lambda} \vdash_{\text{ps}} \{\Gamma\} \text{ while } expr \text{ do } \{\Gamma\} S / \widehat{\Gamma}^* \{\Gamma\} \text{ end } \{\Gamma\}} \\
\\
\text{RELAB:} \frac{\gamma' \sqsubseteq \gamma, \Gamma'_1 \sqsubseteq \Gamma_1, \Gamma_2 \sqsubseteq \Gamma'_2, \gamma \vdash_{\text{ps}} \{\Gamma_1\} S / \widehat{\Gamma}^* \{\Gamma_2\}}{\gamma' \vdash_{\text{ps}} \{\Gamma'_1\} S / \widehat{\Gamma}^* \{\Gamma'_2\}}
\end{array}$$

Notation:

$$\begin{array}{l}
\Gamma[\tilde{x} \leftarrow \Gamma_{\mathcal{E}}(expr)]: \Gamma, \text{ except that } \tilde{x} \text{ has label } \Gamma_{\mathcal{E}}(expr) \\
\Gamma \sqsubseteq \Gamma': (\forall v, \tilde{v} \in \text{Vars}(S): \Gamma(v) = \Gamma'(v) \wedge \Gamma(\tilde{v}) \sqsubseteq \Gamma'(\tilde{v}))
\end{array}$$

Figure 10.5: Typing rules for F-PSNI compliance

Steps 4 and 5 derive judgements for the **then** and **else** alternatives. These judgements have exit control points that mandate different labels for flexible variable \tilde{y} , necessitating step 7 to change label $\widehat{\Gamma}(\tilde{y})$ at the exit control point of the **else** alternative (step 5) to match the label assignment for \tilde{y} at the exit control point of the **then** alternative (step 4). The hypotheses for rule IF then can be discharged using the judgements in steps 4 and 7.

When using the F-PSNI typing rules to derive a label annotated program for an **if** statement, it is not unusual to start by deriving judgements for the **then** and **else** alternatives, but with different label assignments mandated for their exit control points:

$$\gamma \vdash_{\text{ps}} \{\Gamma\} S_1 / \widehat{\Gamma}_1^* \{\Gamma_1\} \quad \gamma \vdash_{\text{ps}} \{\Gamma\} S_2 / \widehat{\Gamma}_2^* \{\Gamma_2\}$$

Because labels are elements of a lattice, there will always exist a label assignment Γ_3 where $\Gamma_3(\tilde{x})$ is $\Gamma_1(\tilde{x}) \sqcup \Gamma_2(\tilde{x})$ for each flexible variable \tilde{x} (and $\Gamma_3(x)$ is $\Gamma_1(x)$ for each fixed variable x). By construction, Γ_3 satisfies $\Gamma_1 \sqsubseteq \Gamma_3$ and $\Gamma_2 \sqsubseteq \Gamma_3$. Rule RELAB then can be used to derive

$$\gamma \vdash_{\text{ps}} \{\Gamma\} S_1 / \widehat{\Gamma}_1^* \{\Gamma_3\} \quad \gamma \vdash_{\text{ps}} \{\Gamma\} S_2 / \widehat{\Gamma}_2^* \{\Gamma_3\}$$

1. $\mathsf{L} \vdash_{\text{ps}} \{\Gamma\} \tilde{y} := 0; \{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\}$
... instance of rule F-ASSIGN.
2. $\mathsf{L} \sqcup \Gamma_{\mathcal{E}}(x_{\mathsf{H}} = 0) \vdash_{\text{ps}} \{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\}$
 $\tilde{y} := 0$
 $\{(\Gamma[\tilde{y} \leftarrow \mathsf{L}])[\tilde{y} \leftarrow \mathsf{L} \sqcup \Gamma_{\mathcal{E}}(x_{\mathsf{H}} = 0) \sqcup \mathsf{L}]\}$
 ... instance of rule F-ASSIGN.
3. $(\Gamma[\tilde{y} \leftarrow \mathsf{L}])[\tilde{y} \leftarrow \mathsf{L} \sqcup \Gamma_{\mathcal{E}}(x_{\mathsf{H}} = 0) \sqcup \mathsf{L}] = \Gamma[\tilde{y} \leftarrow \mathsf{H}]$
... substitution and simplification.
4. $\mathsf{L} \sqcup \Gamma_{\mathcal{E}}(x_{\mathsf{H}} = 0) \vdash_{\text{ps}} \{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\} \tilde{y} := 0 \{(\Gamma[\tilde{y} \leftarrow \mathsf{H}])\}$
... instance of rule RELAB with 2, 3.
5. $\mathsf{L} \sqcup \Gamma_{\mathcal{E}}(x_{\mathsf{H}} = 0) \vdash_{\text{ps}} \{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\} \text{skip} \{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\}$
... instance of rule SKIP.
6. $\Gamma[\tilde{y} \leftarrow \mathsf{L}] \sqsubseteq \Gamma[\tilde{y} \leftarrow \mathsf{H}]$
... definition of \sqsubseteq .
7. $\mathsf{L} \sqcup \Gamma_{\mathcal{E}}(x_{\mathsf{H}} = 0) \vdash_{\text{ps}} \{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\} \text{skip} \{\Gamma[\tilde{y} \leftarrow \mathsf{H}]\}$
... instance of rule RELAB with 5, 6.
8. $\mathsf{L} \vdash_{\text{ps}} \{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\}$
 if $x_{\mathsf{H}} = 0$ then $\{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\} \tilde{y} := 0 \{\Gamma[\tilde{y} \leftarrow \mathsf{H}]\}$
 else $\{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\} \text{skip} \{\Gamma[\tilde{y} \leftarrow \mathsf{H}]\}$
 fi
 $\{\Gamma[\tilde{y} \leftarrow \mathsf{H}]\}$
 ... instance of rule IF with 4, 7.
9. $\mathsf{L} \vdash_{\text{ps}} \{\Gamma\} \tilde{y} := 0; \{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\}$
 if $x_{\mathsf{H}} = 0$ then $\{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\} \tilde{y} := 0 \{\Gamma[\tilde{y} \leftarrow \mathsf{H}]\}$
 else $\{\Gamma[\tilde{y} \leftarrow \mathsf{L}]\} \text{skip} \{\Gamma[\tilde{y} \leftarrow \mathsf{H}]\}$
 fi
 $\{\Gamma[\tilde{y} \leftarrow \mathsf{H}]\}$
 ... instance of rule SEQ with 1, 8.

Figure 10.6: Example F-PSNI derivation

which are the hypotheses needed for rule IF.

10.2.3 Dynamic Enforcement of F-PSNI

A reference monitor alone cannot ensure compliance with Preventing Variation in Label Assignments or with Matching Outcomes to Guards. This is because reference monitor actions based on untaken execution paths would be necessary for that compliance. For example, with the programs in (10.14) and (10.15), an **if** statment executes either the **then** alternative or the **else** alternative but the label updates required when that **if** statment terminates depend on the assignment statement targets in both the **then** and the **else** alternatives.

With hybrid enforcement mechanisms, a reference monitor is augmented with the capability to analyze code that has not yet been executed. Seeking greater permissiveness than allowed by Preventing Variation in Label Assignments (required for type-correctness), we develop a hybrid enforcement mechanism for compliance with Matching Outcomes to Guards. It uses an analyzer $ftgts(S)$ that provides the set of flexible variables that are targets of assignment statements in a code fragment S . Compliance with Matching Outcomes to Guards is then maintained by executing

forall $\tilde{x} \in ftgts(S)$ **do** $\widehat{\Gamma}(\tilde{x}) := \widehat{\Gamma}(\tilde{x}) \sqcup \widehat{\Gamma}_{\mathcal{E}}(G_S)$ **end**

whenever execution reaches the end of a control-flow statement S with guard G_S .

A dynamic enforcement mechanism to implement these label adjustments is given in Figure 10.7. It employs a stack sf containing pairs $\langle V, \lambda_G \rangle$, where V is a set of flexible variables with labels to be adjusted and λ_G is the label to use for the adjustment. To enforce F-PSNI, this mechanism would be combined with \mathcal{R}_{HPS} in Figure 10.4. Or, it could be combined with Reference monitor \mathcal{R}_{TI} in Figure 9.12 to enforce TINI in programs that have fixed and flexible variables.

10.3 *Other Noninterference Policies

Noninterference policies have the form

$$(\forall s, s' \in Init_S: s \sim_{\eta} s' \Rightarrow \llbracket S \rrbracket^{\eta}(s) \approx_{\eta} \llbracket S \rrbracket^{\eta}(s')) \quad (10.16)$$

where predicate $s \sim_{\eta} s'$ is satisfied by initial states s and s' that are indistinguishable to an attacker, function $\llbracket S \rrbracket^{\eta}(s)$ evaluates to a description of the execution effects produced when S is started in state s , and predicate $\xi \approx_{\eta} \xi'$ is satisfied if ξ and ξ' are descriptions of execution effects that would be indistinguishable to attackers. Different choices for predicate \sim_{η} , function $\llbracket S \rrbracket^{\eta}(\cdot)$, and predicate \approx_{η} result in the various different noninterference policies discussed in this chapter as well as many others.

The execution effects that $\llbracket S \rrbracket^{\eta}(\cdot)$ captures reflect the assumptions we are making about what can be monitored by attackers. For TINI, we assumed

upon S reaching •	action to be performed
• if G then S' else S'' fi	$\text{push}(sf, \langle \text{ftgts}(S') \cup \text{ftgts}(S''), \widehat{\Gamma}_{\mathcal{E}}(G) \rangle)$
... fi •	$\langle \widetilde{V}, \lambda_G \rangle := \text{top}(sf)$ forall $\tilde{x} \in \widetilde{V}$ do $\widehat{\Gamma}(\tilde{x}) := \widehat{\Gamma}(\tilde{x}) \sqcup \lambda_G$ end ; $\text{pop}(sf)$
• while G do S end	$\text{push}(sf, \langle \text{ftgts}(S), \widehat{\Gamma}_{\mathcal{E}}(G) \rangle)$
while G do S end •	$\langle \widetilde{V}, \lambda_G \rangle := \text{top}(sf)$ forall $\tilde{x} \in \widetilde{V}$ do $\widehat{\Gamma}(\tilde{x}) := \widehat{\Gamma}(\tilde{x}) \sqcup \lambda_G$ end ; $\text{pop}(sf)$

Figure 10.7: Compliance for matching outcomes to guards

that attackers can monitor only the final states of terminating executions, so $\llbracket S \rrbracket^n(\cdot)$ is a function (*viz.* $\llbracket S \rrbracket(\cdot)$) from initial states to final states; for PSNI, we assumed that an attacker can monitor intermediate states of all executions, so $\llbracket S \rrbracket^n(\cdot)$ is a function (*viz.* $\llbracket S \rrbracket^{\text{tr}}(\cdot)$) from initial states to execution traces.

But thus far we have considered only deterministic programs. Nondeterministic programs and concurrent programs can be handled, too. We use a function $\llbracket S \rrbracket^n(\cdot)$ that evaluates to sets of execution effects.

- With a nondeterministic program S , each element in set $\llbracket S \rrbracket^n(s)$ would represent the execution effects resulting from one of the possible sequences of outcomes for the nondeterministic choices made during an execution that starts in initial state s .
- With a concurrent program S , each element in set $\llbracket S \rrbracket^n(s)$ would represent the execution effects produced by one of the possible interleavings of atomic actions by the processes that comprise S , starting from initial state s . Which interleavings are considered possible would depend on the scheduler, capacity bounds on resources, and the semantics of any synchronization mechanisms that processes use.

The straightforward definition for relation $\mathcal{X} \approx_{\eta} \mathcal{X}'$ in (10.16) is $\mathcal{X} = \mathcal{X}'$. However, in situations where an attacker can account for differences due to nondeterminacy or scheduling choices, a weaker definition is more suitable:

$$\mathcal{X} \approx_{\eta} \mathcal{X}': (\forall \xi \in \mathcal{X}: (\exists \xi' \in \mathcal{X}': \xi \approx_{\eta} \xi'))$$

With this weaker definition, $\llbracket S \rrbracket^n(s) \approx_{\eta} \llbracket S \rrbracket^n(s')$ holds if an attacker observing the execution effects from an initial state s cannot rule out having witnessed the execution effects from a different initial state s' . Noninterference policies where \approx_{η} involves an existential quantifier are known as *possibilistic noninterference*.

With nondeterministic and with concurrent programs, we might also want to impose limits on what an observer might learn by running multiple experiments. Such a set of experiments could allow the attacker to approximate the likelihood of an initial state given how often certain execution effects are produced.¹⁰ For those settings, we would want to use a version of noninterference where $\llbracket S \rrbracket^\eta(\cdot)$ produces a probability distribution for possible execution effects and $\llbracket S \rrbracket^\eta(s) \approx_\eta \llbracket S \rrbracket^\eta(s')$ holds if the probability distributions $\llbracket S \rrbracket^\eta(s)$ and $\llbracket S \rrbracket^\eta(s')$ are indistinguishable to attackers.

Finally, we should acknowledge that an attacker must monitor some system interface to learn about execution effects. The noninterference policies discussed in this chapter avoided the details of those interfaces by using states to describe the information that is available to observers. The domain and range for those states reflect assumptions about what execution effects an attacker is and is not able to monitor. For example, TINI and PSNI are defined in terms of states that map variables to values and labels, because these policies assume that attackers are limited to reading (certain) variables.

Changes to variables are not the only observable execution effects, though. The passage of time is an execution effect, and measuring the time that elapses between changes to a variable while a given program executes can reveal information about the values of other variables. One reason is that the time required for performing certain operations can depend on the values of the operands. Another reason is that access to a variable may depend on whether that variable was recently accessed and, therefore, resides in a cache. So from information about execution timing, an attacker sometimes can make inferences about the values of variables that the attacker cannot read directly. A noninterference policy could be formulated that prohibits such leaks—but its states would have to provide information about execution times, cache contents, and any other hardware or operating system resources that affect execution timings.

Notes and Reading for Chapter 10

Success with enforcing TINI in sequential programs prompted researchers to investigate defending against more-capable attackers, developing enforcement mechanisms that would be more permissive, and supporting programming languages that had nondeterminacy and concurrency.

The first step was new typing rules in Volpano and Smith [18] for enforcing (what later became known as) termination-sensitive noninterference (TSNI). With TSNI, attackers are assumed able to distinguish between terminating and

¹⁰As an example, consider the nondeterministic program

S : if $x_H > 0$ then $[x_L := 1 \ \text{■}_{.99} \ x_L := 2]$ else $[x_L := 1 \ \text{■}_{.5} \ x_L := 2]$ fi

where $[S \ \text{■}_p \ S']$ is the syntax for a statement that executes S or S' , choosing S with probability $1 - p$ and choosing S' with probability p . An L-observer who instigates multiple executions of S should be able to predict whether $x_H > 0$ holds based on the distribution of the final values observed for x_L in those executions—if $x_L = 1$ holds often then $x_H > 0$ holds with high probability.

nontermination executions. Therefore, TSNI defends against more-capable attackers than TINI, since TINI ignores nonterminating executions.

Another significant step towards a more realistic model for attackers was the advent of an interactive model of computation in Askarov et al. [2] to replace the batch model used for TINI and TSNI. In this interactive model, attackers could observe outputs and/or intermediate states during an execution. The additional leaks that became possible were captured by new variants of noninterference, later named *progress insensitive noninterference* (PINI) and *progress sensitive noninterference* (PSNI) by Askarov and Myers [3]. Askarov et al. [2] also shows that the certification conditions in Denning [5] to enforce TINI in the batch model of computing suffice for enforcing PINI in this interactive model of computing, which explains why the typing rules in Figure 10.2 resemble the typing rules in Figure 9.8. Finally, Askarov et al. [2] debunks a widely-held belief that only a single secret bit would be leaked when observers can detect that an execution has been terminated by a dynamic enforcement mechanism.

The design of more-permissive enforcement mechanisms has also attracted considerable attention. Most type systems had employed flow-insensitive analyses. So the program $x_L := x_H; x_L := 63$, which satisfies TINI, would be rejected because it contains a subprogram (*viz.* $x_L := x_H$) that—in isolation—violates TINI. Hunt and Sands [7] gives typing rules that implement a flow-sensitive analysis, thereby avoiding such specious rejections.¹¹ Those typing rules (the basis for the typing rules in Figure 10.5) associate assertions with control points; each assertion gives a label assignment. Labels assigned to variables are allowed to change as execution proceeds, a flexibility Denning [5] allows but the subsequent type system formalization in Volpano, Smith, and Irvine [17] did not. Moreover, Hunt and Sands [7] observes that the label assignments generated by their typing rules is an abstract characterization of dependencies, leading to the surprising result: a flow-sensitive analysis of a program having flexible variables can be replaced by a flow-insensitive analysis of an equivalent program where all variables have a fixed label assignment if assignments statements to fresh variables can be added.

Dynamic enforcement mechanisms—reference monitors and hybrid enforcement mechanisms—were seen as another promising avenue for achieving increased permissiveness. Hedin and Sabelfeld [6, §4.3] surveys dynamic enforcement mechanisms found in the literature prior to 2012. These mechanisms differ in the information flow policy they enforce, the events they intercept, the actions they take to prevent a policy violation, the analyses they use, and whether they deliver increased permissiveness over a type system. Space does not permit giving a detailed enumeration here.

We might have expected that a dynamic enforcement mechanism would have to analyze code that will not be executed, since implicit flows can be caused

¹¹Hunt and Sands [7] is not the first flow-sensitive analysis for checking noninterference. It is preceded by Amtoft and Banerjee [1], which gives a Hoare-style logic that implements a flow-sensitive analysis for verifying that certain variables are independent. The first Hoare-style logic for reasoning about information flow policies of sequential and concurrent programs appears in Reitman and Andrews [11], but that analysis was not flow-sensitive.

by assignment statements in the untaken alternative of an `if` statement or an unexecuted body of a `while` statement. However, Sabelfeld and Russo [14] shows that a reference monitor like \mathcal{R}_{TI} can enforce TINI without analyzing untaken code. The relationship between the permissiveness of type systems (which typically analyze all statements in a program) and reference monitors (which only analyze the statements that execute) is complex. For flow-sensitive labels and programs that produce output, Russo and Sabelfeld [12] shows that a reference monitor to enforce PINI will not accept all executions of programs that are considered type-correct by the Hunt and Sands [7] flow-sensitive type system. So neither enforcement mechanism is more permissive than the other. However, Russo and Sabelfeld [12] also shows that a hybrid enforcement mechanism can be more permissive than the typing rules in Hunt and Sands [7].

To learn more about information flow control, good starting points are the Sabelfeld and Myers survey [13] of language-based approaches and the Hedin and Sabelfeld tutorial [6] on how various forms of noninterference can be enforced. See Kozyri et al. [8] for an in-depth exploration of the various kinds of information flow policies. Sabelfeld and Sands [15] is the authoritative treatment of declassification. Also, consider experimenting with a programming language like Jif [10, 9] or Flow Caml [16], where types enforce information flow control. Jif extends Java and has been used to build non-trivial applications. One such application that is well documented is the Civitas [4] coercion-resistant, universally and voter verifiable electronic voting system. Flow Caml is a prototype that extends the Caml language.

Bibliography

- [1] Torben Amtoft and Anindya Banerjee. Information flow analysis in logical form. In Roberto Giacobazzi, editor, *Proceedings 11th International Symposium on Static Analysis (SAS)*, volume 3148 of *Lecture Notes in Computer Science*, pages 100–115, Berlin, Heidelberg, August 2004. Springer-Verlag.
- [2] Aslan Askarov, Sebastian Hunt, Andrei Sabelfeld, and David Sands. Termination-insensitive noninterference leaks more than just a bit. In *Proceedings of the 13th European Symposium on Research in Computer Security: Computer Security, ESORICS '08*, pages 333–348, Berlin, Heidelberg, October 2008. Springer-Verlag.
- [3] Aslan Askarov and Andrew Myers. A semantic framework for declassification and endorsement. In Andrew D. Gordon, editor, *Programming Languages and Systems, ESOP'10*, pages 64–84, Berlin, Heidelberg, March 2010. Springer-Verlag.
- [4] Michael R. Clarkson, Stephen Chong, and Andrew C. Myers. Civitas: Toward a secure voting system. In *2008 IEEE Symposium on Security and Privacy*, pages 354–368. IEEE Computer Society, May 2008.

- [5] Dorothy E. Denning. *Secure Information Flow in Computer Ssystems*. PhD thesis, Purdue University, USA, 1975.
- [6] Daniel Hedin and Andrei Sabelfeld. A perspective on information-flow control. In Tobias Nipkow, Orna Grumberg, and Benedikt Hauptmann, editors, *Software Safety and Security – Tools for Analysis and Verification*, volume 33 of *NATO Science for Peace and Security Series – D: Information and Communication Security*, pages 319–347. IOS Press, 2012.
- [7] Sebastian Hunt and David Sands. On flow-sensitive security types. In J. Gregory Morrisett and Simon L. Peyton Jones, editors, *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '06, pages 79–90. ACM, January 2006.
- [8] Elisavet Kozyri, Stephen Chong, and Andrew C. Myers. Expressing information flow properties. *Foundations and Trends Privacy and Security*, 3(1):1–102, 2022.
- [9] A.C. Myers, L. Zheng, S. Zdancewic, S. Chong, and N. Nystrom. Jif: Java information flow (software release), July 2001. <http://www.cs.cornell.edu/jif>.
- [10] Andrew C. Myers. Jflow: Practical mostly-static information flow control. In Andrew W. Appel and Alex Aiken, editors, *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '99, pages 228–241. ACM, January 1999.
- [11] Richard P. Reitman and Gregory R. Andrews. Certifying information flow properties of programs: An axiomatic approach. In Alfred V. Aho, Stephen N. Zilles, and Barry K. Rosen, editors, *Conference Record of the Sixth Annual ACM Symposium on Principles of Programming Languages*, POPL '79, pages 283–290. ACM, January 1979.
- [12] Alejandro Russo and Andrei Sabelfeld. Dynamic vs. static flow-sensitive security analysis. In *Proceedings of the 23rd IEEE Computer Security Foundations Symposium*, CSF '10, pages 186–199. IEEE Computer Society, July 2010.
- [13] Andrei Sabelfeld and Andrew C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communication*, 21(1):5–19, 2003.
- [14] Andrei Sabelfeld and Alejandro Russo. From dynamic to static and back: Riding the roller coaster of information-flow control research. In Amir Pnueli, Irina B. Virbitskaite, and Andrei Voronkov, editors, *Perspectives of Systems Informatics, 7th International Andrei Ershov Memorial Conference (PSI)*, volume 5947 of *Lecture Notes in Computer Science*, pages 352–365. Springer, June 2009.

- [15] Andrei Sabelfeld and David Sands. Declassification: Dimensions and principles. *Journal of Computer Security*, 17(5):517–548, 2009.
- [16] Vincent Simonet. The Flow Caml system (version 1.00): Documentation and user’s manual. <http://crystal.inria.fr/~simonet/soft/flowcaml/manual/index.html>, July 2003.
- [17] Dennis M. Volpano, Cynthia E. Irvine, and Geoffrey Smith. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(2/3):167–188, 1996.
- [18] Dennis M. Volpano and Geoffrey Smith. Eliminating covert flows with minimum typings. In *10th Computer Security Foundations Workshop (CSFW ’97)*, pages 156–169. IEEE Computer Society, June 1997.