



## Information Flow Type Systems — Limin Jia

Lecture 2 - June 24, 2025

### 1 Final thoughts from previous lecture

Recall the example of the Strava heatmap, which clearly is a security violation.

- An individual might not mind their anonymized data being shared with Strava. The problem only arises when the data is from a sensitive location.
- So, here, we have two classifications of sensitivity: the data itself, and the fact that the aggregated data reveals information about a sensitive area. Therefore, even aggregate data might not conform to noninterference.
- This is an example of our security properties not appropriately handling the declassification of information.
- From this, we see that information flow security is a *global* property and stronger than access control, which is just *local*.
- We want to establish local policies that guarantee non-interference.

### 2 Noninterference for concurrent and nondeterministic systems

Non-interference is a property of a system, and it can be defined as the absence of dependency between secret inputs and public outputs. In other words, a system satisfies non-interference if any set of inputs that differ *only* on secret values have indistinguishable public outputs.

Consider a nondeterministic system. We can formulate a new definition of non-interference for this system.

**Definition (Possible Noninterference)** Given two sets of inputs that differ only in secret inputs, any public outcome for the first inputs is a possible public outcome for the second inputs.

- This definition is okay for truly nondeterministic programs. However, true nondeterminism is impossible in practice. For one, any concurrent program's execution will depend on the scheduler. Most schedulers are deterministic, i.e. round-robin.

As an example of how this definition fails, consider the program in Figure 1 spread across two processes.

```
1 P1
2 if(secret == 0) long_computation;
3 print 0
4
5 P2
6 print 1
```

Figure 1: Example security flaw for concurrent programs

If both processes are running on a round-robin system, and we pass in 0 as the secret, it will almost always print 1 followed by 0.

### 3 Approaches

Suppose we are working with the following security labels: **S** (secret) and **P** (public). To ensure a valid information flow, we must consider the approach to adopt to prevent any invalid flows. One possibility is to enhance our type checker to assign security labels to terms in addition to just types. This is a static approach because the type checker catches errors during compile time. However, it is too *conservative* and may reject programs that are actually okay — false positives —. In other words, static analysis over-approximates security-related behaviors.

An alternative solution is to use a runtime monitor. This checks every step of the running program and either allows it to continue or terminates the program if the execution violates some security property. This approach is dynamic because infractions are caught during runtime, just before execution.

### 4 Information flow type system

We will work with the IMP language.

<b>Values</b>	$v ::= x \mid n$
<b>Terms</b>	$e ::= v \mid e_1 \text{ bop } e_2 \mid \text{uop } e$
<b>Commands</b>	$c ::= \text{skip} \mid \text{halt} \mid x := e \mid \text{if } e \text{ then } c_1 \text{ else } c_2 \mid \text{while } e \text{ do } c \mid c_1; c_2$
<b>Stores</b>	$\sigma ::= \bullet \mid \sigma, x \mapsto v$
<b>Security Labels</b>	$\ell ::= L \mid H$

In this terse language, values are variables ( $x$ ) or numerical constants ( $n$ ). Terms can be a value, or the binary and unary operators (*bop/uop*). Here,  $\sigma$  is a store that maps variables to values, representing our runtime state.  $\ell$  is the syntax category for secrecy labels, on top of which there exists a partial order  $\sqsubseteq$ , e.g.  $L \sqsubseteq H$ .

For our semantics, we imagine a program  $e$  in a store  $\sigma$  will result in a value  $v$ . Our programs evaluate like  $(\sigma, c) \rightarrow (\sigma', c')$ . Also, the meaning of an expression in a store is a value,  $\llbracket e \rrbracket_\sigma = v$ , since terms do not have effects.

### Contexts

$\boxed{\Gamma := \overline{x : \ell}}$  A context  $\Gamma$  is a list of secrecy labels (i.e.  $L, H$ ) assignments to variables. Normally, we would say something like  $x : \text{int } \ell$ . For the sake of simplicity, we will omit the concrete type part (e.g. *int*) and focus on labels ( $\ell$ ) from now on.

### Security typing

The judgment  $\boxed{\Gamma \vdash e : \ell}$  asserts that an expression  $e$  has label  $\ell$  under the context  $\Gamma$ . Take the typing rule for *bop* as an example:

$$\frac{\Gamma \vdash e_1 : \ell_1 \quad \Gamma \vdash e_2 : \ell_2}{\Gamma \vdash e_1 \text{ bop } e_2 : \ell_1 \sqcup \ell_2}$$

We can observe that the term has  $e_1 \text{ bop } e_2$ , which guarantees that the resulting label is the greatest upper bound among  $\ell_1$  and  $\ell_2$ . To materialize this, let us consider that  $\ell_1 = L$  and  $\ell_2 = H$ . An operation involving  $H$  should be treated as  $H$ .

### Typing commands

The judgment  $\boxed{\Gamma, pc \vdash c}$  checks if a command  $c$  respects the security context  $pc$  — a security label — in context  $\Gamma$ . The label  $pc$  represents the current security label of the systems. We start the execution of a program by assigning the lowest label ( $\perp$ ) to  $pc$ . Let us define the type checking rule for *if* statements:

$$\frac{\Gamma \vdash b : \ell \quad \Gamma, pc \sqcup \ell \vdash c_1 \quad \Gamma, pc \sqcup \ell \vdash c_2}{\Gamma, pc \vdash \text{if } b \text{ then } c_1 \text{ else } c_2}$$

In this rule,  $b$  has the security label  $\ell$ . Notice how  $c_1$  and  $c_2$  are type-checked under a combined label  $pc \sqcup \ell$ . This guarantees that we respect the higher label among  $pc$  and  $\ell$ . The purpose of this inference rule is that the type-checking for a branching statement will ensure that the label for either branch respects the combined label  $pc \sqcup \ell$ . This avoids leaks that could occur if  $b$  has a higher label than its branches.

Another example, for assignment:

$$\frac{\Gamma \vdash e : \ell \quad \Gamma(x) = \ell \quad pc \sqsubseteq \ell}{\Gamma, pc \vdash x := e}$$

This states that when we assign a term to a variable, the label of the variable and the term must be compatible and also respect the current  $pc$ .

Consider the following rule. We will examine how it violates non-interference:

$$\frac{\Gamma \vdash e : \ell_1 \quad \ell_1 \sqsubseteq \ell_2}{\Gamma \vdash e : \ell_2}$$

When combined with the previous rule, this rule violates non-interference. Namely, we are able to derive  $e$  as a high ( $\ell_2$ ) value from this rule, but using the previous rule (with  $x$  as  $e$ ,  $pc$  as  $\ell_1$ , and  $\ell$  as  $\ell_2$ ), we have that  $x$  is in a low context ( $pc$ , or  $\ell_1$ ).

### Attacker label

Let  $\ell_a$  denote the label of an attacker.

### $\ell_a$ Equivalence

We define  $v_1 \approx_{\ell_a} v_2 : \ell$  by

$$\frac{\ell \sqsubseteq \ell_a \quad v_1 = v_2}{v_1 \approx_{\ell_a} v_2 : \ell}$$

If the attacker is at a higher level than the values, the values must be *identical* in order for them to be *indistinguishable*.

$$\frac{\ell \not\sqsubseteq \ell_a}{v_1 \approx_{\ell_a} v_2 : \ell}$$

Otherwise, every value is indistinguishable to the attacker.

### $\ell_a$ Equivalence of stores

Similarly, we define  $\Gamma \vdash \sigma_1 \approx_{\ell_a} \sigma_2$  by

$$\Gamma \vdash \bullet \approx_{\ell_a} \bullet$$

$$\frac{\Gamma \vdash \sigma_1 \approx_{\ell_a} \sigma_2 \quad v_1 \approx_{\ell_a} v_2 : P(x)}{\Gamma \vdash \sigma_1, x \rightarrow v_1 \approx_{\ell_a} \sigma_2, x \rightarrow v_2}$$

**Lemma** (*Noninterference of terms*). If  $\Gamma \vdash e : \ell$ , and  $\Gamma \vdash \sigma_1 \approx_{\ell_a} \sigma_2$ , then  $\llbracket e \rrbracket_{\sigma_1} \approx_{\ell_a} \llbracket e \rrbracket_{\sigma_2} : \ell$ .

### Sequencing

$$\begin{array}{ll} \textbf{Evaluation Context} & E ::= \square \mid \square; c \\ \textbf{Stack} & s ::= \cdot \mid E \triangleright s \\ \textbf{State} & \Sigma ::= (\sigma, s, c) \end{array}$$

To actually run programs, we evaluate a sequence of commands. Think of the stack as containing continuations. The sequencing rule for pushing to the stack is

$$(\sigma, s, c_1; c_2) \rightarrow (\sigma, (\square; c_2) \triangleright s, c_1).$$

Likewise, for pop, when there is no command to execute immediately (*SKIP*) but there is some command pending in the stack  $((\square; c) \triangleright s)$ , we move it out to be the currently executing command.

$$(\sigma, (\square; c) \triangleright s, \text{SKIP}) \rightarrow (\sigma, s, c)$$

Here's an additional rule for *SKIP*, when we have no immediate command nor pending command in the stack, we can halt the machine.

$$(\sigma, \square, \text{SKIP}) \rightarrow (\sigma, \text{HALT})$$

**Lemma** (*Noninterference for Commands*). If  $\Gamma, \perp \vdash c$ , and  $\Gamma \vdash \sigma_1 \approx_{\ell_a} \sigma_2$ , and  $(\sigma_i, \square, c) \rightarrow^* (\sigma'_i, \text{HALT})$  (for  $i = 1, 2$ ) then  $\Gamma \vdash \sigma'_1 \approx_{\ell_a} \sigma'_2$ .

We can prove this by induction on the number of steps.