

# Type-Driven Repair for Information Flow Security

## Supplementary Material

**Well-Formedness**  $\boxed{\Gamma \vdash r} \boxed{\Gamma \vdash B} \boxed{\Gamma \vdash S}$

$$\text{WF-}\psi \frac{\Gamma \vdash \psi : \text{Bool}}{\Gamma \vdash \psi}$$

$$\text{WF-}\pi \frac{\Gamma \vdash \pi \bar{x} : \text{Bool} \quad \Gamma(\pi) \neq T[\ominus]}{\Gamma \vdash \pi \bar{x}}$$

$$\text{WF-}\alpha \frac{\Gamma(\alpha) \neq \ominus}{\Gamma \vdash \alpha} \quad \text{WF-SC} \frac{\Gamma \vdash B \quad \Gamma; \nu : B \vdash r}{\Gamma \vdash \{B \mid r\}}$$

$$\text{WF-FUN} \frac{\Gamma^- \vdash T_x \quad \Gamma; x : T_x \vdash T}{\Gamma \vdash T_x \rightarrow T}$$

$$\text{WF-D} \frac{\Gamma(D) = \overline{\forall_{\circ} \alpha_i. \forall_{\circ} \langle \pi_j : U_j \rangle. T} \quad |T_i| = |\alpha_i| \quad \Gamma \vdash p_j : U_j}{\Gamma \vdash D \overline{T_i} \langle \overline{p_j} \rangle}$$

$$\text{WF-FO} \frac{\Gamma^- \vdash \{B \mid r\} \quad \Gamma; x : \{B \mid r\} \vdash T}{\Gamma \vdash x : \{B \mid r\} \rightarrow T}$$

$$\text{WF-HO} \frac{T_x \text{ non-scalar} \quad \Gamma^- \vdash T_x \quad \Gamma \vdash T}{\Gamma \vdash T_x \rightarrow T}$$

$$\text{WF-}\forall\alpha \frac{\Gamma; \alpha : \circ \vdash S}{\Gamma \vdash \forall_{\circ} \alpha. S} \quad \text{WF-}\forall\pi \frac{\Gamma; \pi : T[\circ] \vdash S}{\Gamma \vdash \forall_{\circ} \langle \pi : T \rangle. S}$$

**Figure 1.** Well-formedness rules of  $\mathcal{BL}$ .

## 1. $\mathcal{BL}$ Static Semantics

Figures 1,2, and 3 give the full version of the static semantics of  $\mathcal{BL}$ .

### 1.1 Proving Non-Interference Using Tagged<sup>2</sup>

We now prove that executions involving the Tagged monad preserve *contextual noninterference*: if a sensitive value  $v$  may not flow to a given viewer, then any pair of executions

**Subtyping**  $\boxed{\Gamma \vdash T <: T'}$

$$\text{<:-REFL} \frac{}{\Gamma \vdash B <: B}$$

$$\text{<:-SC} \frac{\Gamma \vdash B <: B' \quad \text{Valid}(\llbracket \Gamma \rrbracket \wedge r \Rightarrow r')}{\Gamma \vdash \{B \mid r\} <: \{B' \mid r'\}}$$

$$\text{<:-FUN} \frac{\Gamma \vdash T_y <: T_x \quad \Gamma; y : T_y \vdash [y/x]T <: T'}{\Gamma \vdash x : T_x \rightarrow T <: y : T_y \rightarrow T'}$$

$$\text{<:-D} \frac{\Gamma(D) = \overline{\forall_{\circ} \alpha_i. \forall_{\circ} \langle \pi_j \rangle. T} \quad \Gamma \vdash T_i \sim_{\circ_i} T'_i \quad \Gamma \vdash p_j \sim_{\circ_j} p'_j}{\Gamma \vdash D \overline{T_i} \langle \overline{p_j} \rangle <: D \overline{T'_i} \langle \overline{p'_j} \rangle}$$

$$\frac{\Gamma \vdash T <: T' \quad \Gamma \vdash T' <: T}{\Gamma \vdash T \sim_{\oplus} T'} \quad \frac{\Gamma \vdash T <: T' \quad \Gamma \vdash T' <: T}{\Gamma \vdash T \sim_{\ominus} T'}$$

$$\frac{\Gamma \vdash T <: T' \quad \Gamma \vdash T' <: T}{\Gamma \vdash T \sim_{\circ} T'}$$

$$\frac{\Gamma; x : T \vdash p \sim_{\circ} p'}{\Gamma \vdash \lambda x : T. p \sim_{\circ} \lambda x : T. p'}$$

$$\frac{\Gamma \vdash \{() \mid r\} \sim_{\circ} \{() \mid r'\}}{\Gamma \vdash r \sim_{\circ} r'}$$

**Figure 2.** Subtyping rules of  $\mathcal{BL}$ .

involving different assignments to  $v$  should yield equivalent outputs.

Reasoning directly about noninterference is inconvenient because it requires talking about two executions. We simplify our noninterference proof using a technique similar to that of Pottier and Simonet [1]: we introduce auxiliary constructs that allow us to reason about two executions in one. Being able to encode security labels as a library makes the formalization particularly nice: the only auxiliary construct we need for the proof is an alternative definition of the Tagged monad. We introduce the Tagged<sup>2</sup> monad with new implementations of the four primitive operations, yielding the property that if a program type-checks with Tagged<sup>2</sup>, then it preserves contextual noninterference with Tagged.

**The Tagged<sup>2</sup> monad.** To simplify formalization of noninterference, we parameterize the semantics of  $\mathcal{BL}$  by the context, *i.e.* the principal who is observing the execution and the world at the time of output. More concretely, we assume that

## Type Checking $\boxed{\Gamma \vdash e :: S}$

$$\begin{array}{c}
\text{SUBT} \frac{\Gamma \vdash e :: T' \quad \Gamma \vdash T' <: T}{\Gamma \vdash e :: T} \\
\text{VAR-SC} \frac{\Gamma(x) = \{B \mid r\}}{\Gamma \vdash x :: \{B \mid \nu = x\}} \\
\text{VAR} \frac{\Gamma(x) = S \quad S \text{ non-scalar}}{\Gamma \vdash x :: S} \\
\text{ABS} \frac{\Gamma \vdash T_x \quad \Gamma; x: T_x \vdash e :: T}{\Gamma \vdash \lambda x: T_x. e :: (x: T_x \rightarrow T)} \\
\text{LET} \frac{\Gamma \vdash v_1 :: (y: T_y \rightarrow T') \quad \Gamma \vdash v_2 :: T_2 \quad \Gamma \vdash T_2 <: T_y \quad \Gamma; x: [v/y]T' \vdash e :: T}{\Gamma \vdash \text{let } x = v_1 \ v_2 \text{ in } e :: T} \\
\text{MATCH} \frac{\Gamma \vdash x :: \{D \bar{T}_x \langle \bar{p}_x \mid r_x \rangle \quad \Gamma(D) = \forall_o \bar{\alpha}. \forall_o \langle \bar{\pi} \rangle. \quad T_1 \rightarrow \dots \rightarrow T_n \rightarrow \{D \bar{\alpha} \langle \bar{\pi} \rangle \mid r\} \quad \Gamma; y_i: [\bar{T}_x / \bar{\alpha}] [\bar{p}_x \triangleright \bar{\pi}] T_i; \quad x: \{D \bar{T}_x \langle \bar{p}_x \rangle \mid r_x \wedge r\} \vdash e :: T}{\Gamma \vdash \text{match } x \text{ with } D \bar{x} \rightarrow e :: T} \\
\text{IF} \frac{\Gamma; [\top/\nu]r \vdash e_1 :: T \quad \Gamma; [\perp/\nu]r \vdash e_2 :: T}{\Gamma \vdash \text{if } x \text{ then } e_1 \text{ else } e_2 :: T} \\
\text{T-GEN} \frac{\Gamma; \alpha: o \vdash e :: S}{\Gamma \vdash e :: \forall_o \alpha. S} \\
\text{T-INST} \frac{\Gamma \vdash e :: \forall_o \alpha. S \quad \Gamma \vdash \{B \mid r\}}{\Gamma \vdash e :: \{\{B \mid r\} / \alpha\} S} \\
\text{P-GEN} \frac{\Gamma; \pi: T[o] \vdash e :: S \quad \Gamma \vdash T}{\Gamma \vdash e :: \forall_o \langle \pi: T \rangle. S} \\
\text{P-INST} \frac{\Gamma \vdash e :: \forall_o \langle \pi: T \rangle. S \quad \Gamma \vdash p: T}{\Gamma \vdash e :: [p \triangleright \pi] S}
\end{array}$$

Figure 3. Type-checking rules of  $\mathcal{B}\mathcal{L}$ .

the environment always contains two variables  $cw : W$  and  $cu : U$ ; when a program executes, it executes with all possible values of  $cw$  and  $cu$  “in parallel”, but in each of these parallel threads, **print** only performs the output when its arguments match  $cw$  and  $cu$ , so this parametric semantics has no effect on the output.

We first construct a *phantom encoding*: a new information flow monad,  $\text{Tagged}^2$ , that explicitly relates pairs of program executions. The intuition behind  $\text{Tagged}^2$  is as follows: it represents two versions of a sensitive value from two different executions of the program as seen by the current context. Mirroring what we want for our noninterference property, the two versions are only allowed to differ for those sensitive values that are *not visible* in the context. The  $\text{Tagged}^2$  constructor accepts two  $\alpha$  values,  $l$  and  $r$ , which we call *projections*. Its third argument *prop* serves as a proof of the property  $p \text{ cw } cu \Rightarrow l = r$ , that is, if the policy holds of the current context, the two projections must be equal.

```

module Tagged2 where

private cw: W, cu: U -- Current context

5 -- | Tagged data constructor
private Tagged2: ∀α. ∀⊖<p: W → U → Bool>.
  l: α → r: α → prop: ({() | p cw cu} → {() | l = r})
  → Tagged α <p>

10 return2: ∀α. ∀⊖<p: W → U → Bool>. α → Tagged α <p>
return2 = λx. Tagged2 x x (λz. ())

bind2: ∀α β. ∀⊖<p: W → U → Bool>. ∀<f: α → β →
Bool>.
  x: Tagged α <p> → (y: α → Tagged {β | f y ν} <p>)
  → Tagged {β | f (l x) ν} <p>
15 bind2 = λx. λg.
  match x with Tagged2 xl xr xp →
  match g xl with Tagged2 yl _ yp →
  match g xr with Tagged2 _ yr _ →
20 Tagged2 yl yr (λz. join (xp z) (yp z))

print2: ∀α. ∀⊖<p: W → U → Bool>.
  w: W → u: Tagged2 {U | p w ν}<p>
  → x: Tagged2 α <p> → W
25 print2 = λw. λu. λx.
  match u with Tagged2 ul ur up →
  if w ≠ cw ∨ (ul ≠ cu ∧ ur ≠ cu) then w
  else if ul ≠ ur then fail (up ())
  else match x with Tagged2 xl xr xp →
30 if xl ≠ xr
  then fail (xp ())
  else doPrint w xl

downgrade2: ∀⊖<p: W → U → Bool>. ∀<c: Bool>.
35 x: Tagged2 {Bool | ν ⇒ c} <λw u. p w u ∧ c>
  → Tagged2 {Bool | ν ⇒ c} <p>
downgrade2 = λx.
  match x with
  Tagged2 xl xr xp →
40 Tagged2 xl xr (λz. if xl ∨ xr then xp z else ())

```

Figure 4. The  $\text{Tagged}^2$  monad, which keeps track of two projections.

A  $\text{Tagged}^2$  value with different projections corresponds to Pottier and Simonet’s “bracket value” in [1], and the *prop* requirement corresponds to their rule that all bracket values are assigned high security labels. The main conceptual difference of our treatment is that the division between high and low security, as well as the notion of a leak, is context-specific.

We show the implementation of the  $\text{Tagged}^2$  in Fig. 4. The phantom encoding provides alternative implementations of the four primitive operations. The function **return**<sup>2</sup> gives the same value for both projections, while **bind**<sup>2</sup> applies the

function projection-wise. The  $\mathcal{BL}$  type checker can easily show both implementations type-safe.

The function `print`<sup>2</sup> is designed to *fail* when it detects interference: namely, whenever the target of the output is different in the two executions ( $u_l \neq u_r$ ) or because it outputs two different values ( $x_l \neq x_r$ ). We assume that `fail` has the type  $\{() \mid \text{False}\} \rightarrow a$ , so the only way to type-check `print`<sup>2</sup> is to prove that both failing branches are unreachable, which the  $\mathcal{BL}$  type checker successfully accomplishes. To understand why the first failing branch is unreachable, recall that from the type of  $u$  we know that  $p \ w \ u_l \wedge p \ w \ u_r$ ; we also know that  $w = cw$  and  $u_l = cu \vee u_r = cu$  from the path condition, thus  $p \ cw \ cu$  holds, which gives  $u_l = u_r$  guaranteed by the  $\text{Tagged}^2$  constructor.

The function `downgrade`<sup>2</sup> simply reconstructs its argument, but provides a proof that  $x_l = x_r$  under a weaker assumption. The proof can be understood as follows: if  $x_l \vee x_r$ , then  $c$  must hold, so we can invoke the proof  $x_p$  of  $x_l = x_r$  that we obtained from the argument; otherwise  $x_l = x_r = \text{False}$ .

**Contextual noninterference.** We now show that type-checking with  $\text{Tagged}^2$  implies contextual noninterference with  $\text{Tagged}$ . Because the  $\text{Tagged}^2$  functions type-check and because the type system of  $\mathcal{BL}$  is sound [2], we know that no type-correct program that manipulates  $\text{Tagged}^2$  values can go wrong, *i.e.* attempts to print the results of two executions that are different. Now we only have to formally connect computations with  $\text{Tagged}$  values and those with  $\text{Tagged}^2$  values, and show how type safety of the latter implies noninterference for the former.

We first show that replacing a  $\text{Tagged}^2$  value with its projection in  $\text{Tagged}$  at the beginning of an execution yields the same result as projecting at the end of an execution. A *projection* of an expression  $e$  (written  $\lfloor e \rfloor_j$ , for  $j = \{l, r\}$ ) is an expression where every occurrence of  $\text{Tagged}^2 \ x_l \ x_r$  in  $e$  is replaced by  $\text{Tagged} \ x_j$ .

**Lemma 1 (Projection).** If  $e \rightarrow^* e'$  then  $\lfloor e \rfloor_j \rightarrow^* \lfloor e' \rfloor_j$ , for  $j = \{l, r\}$ .

*Proof outline.* The only steps that are different in the evaluation of  $e$  and its projections are those resulting from the bodies of `bind` and `print`. By inspection of `bind`<sup>2</sup> it is easy to see that it applies the function projection-wise, and thus preserves the property of the lemma. In case of `print`<sup>2</sup>, since it does not fail, either it does not do any output, or the two pro-

jections are the same; in both cases, projections of its body will have the same behavior.  $\square$

**Theorem (Contextual Noninterference).** Let

$$\Gamma; x : \text{Tagged} \ \alpha \ \langle p \rangle \vdash e :: W$$

and  $\neg(p \ cw \ cu)$ . Let for  $j \in \{l, r\}$ ,  $\Gamma \vdash v_j :: \alpha$  and  $\llbracket (\text{Tagged} \ v_j) / x \rrbracket e \rightarrow^* w_j$ . Then  $w_l = w_r$ .

*Proof outline.* Since  $\neg(p \ cw \ cu)$ , we know

$$\Gamma \vdash \text{Tagged}^2 \ v_l \ v_r \ \text{id} :: \text{Tagged} \ \alpha \ \langle p \rangle$$

for any  $v_l, v_r$ . Let  $e^2$  be  $\llbracket (\text{Tagged}^2 \ v_l \ v_r \ \text{id}) / x \rrbracket e$ ; note that  $\lfloor e^2 \rfloor_j = \llbracket (\text{Tagged} \ v_j) / x \rrbracket e$ . By inspection of typing rules of  $\mathcal{BL}$ , substitution of a subterm with the same type does not change the type of the term, so  $\Gamma \vdash e^2 :: W$ . By soundness of the type system,  $e^2$  either diverges or reduces to a value  $w$  of type  $W$ . Note that the execution of  $e^2$  differs from the executions of either  $\llbracket (\text{Tagged} \ v_j) / x \rrbracket e$  only in the bodies of `bind` and `print` functions; since none of them introduces divergence,  $e^2$  cannot diverge either. By Lemma 1,  $\lfloor e^2 \rfloor_j \rightarrow^* \lfloor w \rfloor_j$ , that is  $w_j = \lfloor w \rfloor_j$ , but  $\lfloor w \rfloor_l = \lfloor w \rfloor_r$  since  $w$  is a value of type  $W$ , which is different from  $\text{Tagged}$ .  $\square$

**A note on the proof technique.** Being able to express tagged values as a data type with a phantom predicate parameter is not only simpler, but also allows us to prove non-interference over pairs of traces simply by grounding phantom predicates. In the information flow monad  $\text{Tagged}$ , policies are phantom predicates that do not appear in the arguments of data constructors. In  $\text{Tagged}^2$ , the predicates are no longer phantom, but appear negatively in the type of `prop`, consistent with its variance annotation. Using these predicates for explicitly relating multiple program executions helps simplify the formalization and proof of non-interference.

## References

- [1] F. Pottier and V. Simonet. Information flow inference for ML. *ACM Transactions on Programming Languages and Systems*, 25(1), Jan. 2003.
- [2] N. Vazou, P. M. Rondon, and R. Jhala. Abstract refinement types. In *ESOP*, 2013.