

Information-Flow Control and Effects

NWPT2021

Carlos Tomé Cortiñas¹ Fabian Ruch

Chalmers University of Technology

November 6, 2021

Information-flow control (IFC) aims to protect *confidentiality* of data by controlling flows of information in a system.

Information-flow control (IFC) aims to protect *confidentiality* of data by controlling flows of information in a system.

IFC ensures that **secret** data cannot influence the **publicly** observable behaviour of programs.

Information-flow control (IFC) aims to protect *confidentiality* of data by controlling flows of information in a system.

IFC ensures that **secret** data cannot influence the **publicly** observable behaviour of programs.

We are interested in IFC for programs with effects.

Consider a security specification

$$\text{Bool}_{\text{public}} \times \text{Bool}_{\text{secret}} \Rightarrow \text{Bool}_{\text{public}}$$

wrt the security policy $\text{public} \not\sqsupseteq \text{secret}$.

Consider a security specification

$$\text{Bool}_{\text{public}} \times \text{Bool}_{\text{secret}} \Rightarrow \text{Bool}_{\text{public}}$$

wrt the security policy $\text{public} \not\sqsupseteq \text{secret}$.

Consider the programs

$$f = \lambda x, y. \text{not}(x)$$

$$g = \lambda x, y. \text{and}(x, y)$$

Consider a security specification

$$\text{Bool}_{\text{public}} \times \text{Bool}_{\text{secret}} \Rightarrow \text{Bool}_{\text{public}}$$

wrt the security policy $\text{public} \not\sqsupseteq \text{secret}$.

Consider the programs

$$f = \lambda x, y. \text{not}(x)$$

$$g = \lambda x, y. \text{and}(x, y)$$

Noninterference: the public outputs of a program do not depend on its secret inputs, for all programs.

Enforcing Noninterference: Dependency Core Calculus

The Dependency Core Calculus (DCC) [1] is a simply-typed lambda calculus enhanced with a family of type constructors R_l for $l \in \{\text{public}, \text{secret}\}$ (pronounced “redaction”).

Enforcing Noninterference: Dependency Core Calculus

The Dependency Core Calculus (DCC) [1] is a simply-typed lambda calculus enhanced with a family of type constructors R_l for $l \in \{\text{public}, \text{secret}\}$ (pronounced “redaction”).

R is a graded monad over the lattice $\text{public} \sqsubseteq \text{secret}$.

Enforcing Noninterference: Dependency Core Calculus

The Dependency Core Calculus (DCC) [1] is a simply-typed lambda calculus enhanced with a family of type constructors R_l for $l \in \{\text{public}, \text{secret}\}$ (pronounced “redaction”).

R is a graded monad over the lattice $\text{public} \sqsubseteq \text{secret}$.

In DCC all programs are secure.

Consider Moggi's monadic metalanguage [4] as a prog. language for printing to a publicly-observable channel:

Consider Moggi's monadic metalanguage [4] as a prog. language for printing to a publicly-observable channel:

the type of computations is explicit $T_{\text{public}} A$.

Consider Moggi's monadic metalanguage [4] as a prog. language for printing to a **publicly**-observable channel:

the type of computations is explicit $T_{\text{public}} A$.

there is a do-nothing computation and sequencing of computations.

$$\frac{\Gamma \vdash t : A}{\Gamma \vdash \text{return}(t) : T_{\text{public}} A} \quad \frac{\Gamma \vdash t : T_{\text{public}} A \quad \Gamma, x : A \vdash u : T_{\text{public}} B}{\Gamma \vdash \text{let}(t, x. u)}$$

Consider Moggi's monadic metalanguage [4] as a prog. language for printing to a `publicly`-observable channel:

the type of computations is explicit $T_{\text{public}} A$.

there is a do-nothing computation and sequencing of computations.

$$\frac{\Gamma \vdash t : A}{\Gamma \vdash \text{return}(t) : T_{\text{public}} A} \quad \frac{\Gamma \vdash t : T_{\text{public}} A \quad \Gamma, x : A \vdash u : T_{\text{public}} B}{\Gamma \vdash \text{let}(t, x. u)}$$

there is a primitive that prints. $\frac{\Gamma \vdash b : \text{Bool}}{\Gamma \vdash \text{print}(b) : T_{\text{public}} \text{Unit}}$

Consider Moggi's monadic metalanguage [4] as a prog. language for printing to a **publicly**-observable channel:

the type of computations is explicit $T_{\text{public}} A$.

there is a do-nothing computation and sequencing of computations.

$$\frac{\Gamma \vdash t : A}{\Gamma \vdash \text{return}(t) : T_{\text{public}} A} \quad \frac{\Gamma \vdash t : T_{\text{public}} A \quad \Gamma, x : A \vdash u : T_{\text{public}} B}{\Gamma \vdash \text{let}(t, x. u)}$$

there is a primitive that prints. $\frac{\Gamma \vdash b : \text{Bool}}{\Gamma \vdash \text{print}(b) : T_{\text{public}} \text{Unit}}$

A prog. language for IFC with printing to **public** = DCC + Moggi^{print}

Consider the type

$$R_{\text{public}} \text{ Bool} \times R_{\text{secret}} \text{ Bool} \Rightarrow T_{\text{public}} (R_{\text{public}} \text{ Unit})$$

Consider the type

$$R_{\text{public}} \text{ Bool} \times R_{\text{secret}} \text{ Bool} \Rightarrow T_{\text{public}} (R_{\text{public}} \text{ Unit})$$

where printing to `public` is part of the `public` outputs of a program.

Consider the type

$$R_{\text{public}} \text{ Bool} \times R_{\text{secret}} \text{ Bool} \Rightarrow T_{\text{public}} (R_{\text{public}} \text{ Unit})$$

where printing to `public` is part of the `public` outputs of a program.

We can construct a `public` computation from a `public` Boolean:

$$R_{\text{public}}(\text{print}) : R_{\text{public}} \text{ Bool} \Rightarrow R_{\text{public}} (T_{\text{public}} \text{ Unit})$$

Consider the type

$$R_{\text{public}} \text{ Bool} \times R_{\text{secret}} \text{ Bool} \Rightarrow T_{\text{public}} (R_{\text{public}} \text{ Unit})$$

where printing to `public` is part of the `public` outputs of a program.

We can construct a `public` computation from a `public` Boolean:

$$R_{\text{public}}(\text{print}) : R_{\text{public}} \text{ Bool} \Rightarrow R_{\text{public}} (T_{\text{public}} \text{ Unit})$$

however, we cannot “run” it since the types `Rpublic` and `Tpublic` *do not* interact.

We could add a new primitive

$$\text{print}' : R_{\text{public}} \text{ Bool} \Rightarrow T_{\text{public}} \text{ Unit}$$

We could add a new primitive

$$\text{print}' : R_{\text{public}} \text{ Bool} \Rightarrow T_{\text{public}} \text{ Unit}$$

and prove again that our language satisfies noninterference.

We could add a new primitive

$$\text{print}' : R_{\text{public}} \text{ Bool} \Rightarrow T_{\text{public}} \text{ Unit}$$

and prove again that our language satisfies noninterference.

For combinations of Moggi^{print} and other calculi for “effect-free” IFC, e.g. Sealing Calculus (SC) [5] or Moggi’s metalanguage (again), we would have to redo all the work.

We could add a new primitive

$$\text{print}' : R_{\text{public}} \text{ Bool} \Rightarrow T_{\text{public}} \text{ Unit}$$

and prove again that our language satisfies noninterference.

For combinations of Moggi^{print} and other calculi for “effect-free” IFC, e.g. Sealing Calculus (SC) [5] or Moggi’s metalanguage (again), we would have to redo all the work.

And what about other effects?

We could add a new primitive

$$\text{print}' : R_{\text{public}} \text{ Bool} \Rightarrow T_{\text{public}} \text{ Unit}$$

and prove again that our language satisfies noninterference.

For combinations of Moggi^{print} and other calculi for “effect-free” IFC, e.g. Sealing Calculus (SC) [5] or Moggi’s metalanguage (again), we would have to redo all the work.

And what about other effects?

This approach is not modular.

To achieve (some) modularity, we abstract over concrete choices of IFC calculi and study effects in the *classified sets* model of Abadi et al. [1] and Kavvos [3].

To achieve (some) modularity, we abstract over concrete choices of IFC calculi and study effects in the *classified sets* model of Abadi et al. [1] and Kavvos [3].

The model captures dependency (independency) by interpreting types as sets *endowed* with a family of relations indexed by security levels that programs (functions) need to preserve. For instance,

To achieve (some) modularity, we abstract over concrete choices of IFC calculi and study effects in the *classified sets* model of Abadi et al. [1] and Kavvos [3].

The model captures dependency (independency) by interpreting types as sets *endowed* with a family of relations indexed by security levels that programs (functions) need to preserve. For instance,

$$R_{\text{public}} \text{ Bool} = (\{tt, ff\}, R_{\text{public}} = \begin{matrix} \text{ff} \\ \circlearrowleft \\ \text{tt} \end{matrix}, R_{\text{secret}} = \begin{matrix} \text{ff} \\ \circlearrowleft \\ \text{tt} \end{matrix})$$

To achieve (some) modularity, we abstract over concrete choices of IFC calculi and study effects in the *classified sets* model of Abadi et al. [1] and Kavvos [3].

The model captures dependency (independency) by interpreting types as sets *endowed* with a family of relations indexed by security levels that programs (functions) need to preserve. For instance,

$$\begin{aligned}
 R_{\text{public}} \text{ Bool} &= (\{tt, ff\}, R_{\text{public}} = \begin{array}{c} \text{ff} \\ \circlearrowleft \\ \text{tt} \end{array}, R_{\text{secret}} = \begin{array}{c} \text{ff} \\ \circlearrowleft \\ \text{tt} \end{array}) \\
 R_{\text{secret}} \text{ Bool} &= (\{tt, ff\}, R_{\text{public}} = \begin{array}{c} \text{ff} \\ \text{tt} \end{array}, R_{\text{secret}} = \begin{array}{c} \text{ff} \\ \circlearrowleft \\ \text{tt} \end{array})
 \end{aligned}$$

To achieve (some) modularity, we abstract over concrete choices of IFC calculi and study effects in the *classified sets* model of Abadi et al. [1] and Kavvos [3].

The model captures dependency (independency) by interpreting types as sets *endowed* with a family of relations indexed by security levels that programs (functions) need to preserve. For instance,

$$\begin{aligned}
 R_{\text{public}} \text{ Bool} &= (\{tt, ff\}, R_{\text{public}} = \begin{array}{c} \text{ff} \\ \circlearrowleft \\ \text{tt} \end{array}, R_{\text{secret}} = \begin{array}{c} \text{ff} \\ \circlearrowleft \\ \text{tt} \end{array}) \\
 R_{\text{secret}} \text{ Bool} &= (\{tt, ff\}, R_{\text{public}} = \begin{array}{c} \text{ff} \\ \text{tt} \end{array}, R_{\text{secret}} = \begin{array}{c} \text{ff} \\ \circlearrowleft \\ \text{tt} \end{array})
 \end{aligned}$$

A relation at level l captures what observers at l can distinguish, and programs preserving relations means they are forbidden to map indistinguishable inputs to distinguishable outputs.

Let the security policy be a join semilattice $(\mathcal{L}, \sqsubseteq, \perp, \vee)$.

Let the security policy be a join semilattice $(\mathcal{L}, \sqsubseteq, \perp, \vee)$.

The redaction monad is given by:

$$\begin{aligned} U(R_l(A)) &:= U(A) \\ R_{l'}(R_l(A))(a, b) &\Leftrightarrow \begin{cases} \top & l \not\sqsubseteq l' \\ R_{l'}(A)(a, b) & l \sqsubseteq l' \end{cases} \end{aligned}$$

Let the security policy be a join semilattice $(\mathcal{L}, \sqsubseteq, \perp, \vee)$.

The redaction monad is given by:

$$\begin{aligned} U(R_l(A)) &:= U(A) \\ R_{l'}(R_l(A))(a, b) &\Leftrightarrow \begin{cases} \top & l \not\sqsubseteq l' \\ R_{l'}(A)(a, b) & l \sqsubseteq l' \end{cases} \end{aligned}$$

In words: it forces the relation of a classified set A to be the everywhere true relation at any level l' such that $l \not\sqsubseteq l'$.

Printing Effects in Classified Sets

We generalize the previous example to *printing on multiple channels* $c \in \mathcal{C}$ where a function $\text{label} : \mathcal{L} \rightarrow \mathcal{C}$ specifies the security level of channels.

We generalize the previous example to *printing on multiple channels* $c \in \mathcal{C}$ where a function $\text{label} : \mathcal{L} \rightarrow \mathcal{C}$ specifies the security level of channels.

For a subset $C \subseteq \mathcal{C}$, the monoid $(\text{Out}_C, \epsilon_C, \cdot_C)$ is defined by:

$$\begin{aligned} \mathbf{U}(\text{Out}_C) &:= C \rightarrow \text{List}(\mathbb{2}) \\ \mathbf{R}_l(\text{Out}_C)(o_1, o_2) &:\Leftrightarrow \forall c \in C. \text{label}(c) \sqsubseteq l \Rightarrow o_1(c) = o_2(c) \end{aligned}$$

We generalize the previous example to *printing on multiple channels* $c \in \mathcal{C}$ where a function $\text{label} : \mathcal{L} \rightarrow \mathcal{C}$ specifies the security level of channels.

For a subset $C \subseteq \mathcal{C}$, the monoid $(\text{Out}_C, \epsilon_C, \cdot_C)$ is defined by:

$$\begin{aligned} \mathbf{U}(\text{Out}_C) &:= C \rightarrow \text{List}(\mathbb{2}) \\ \mathbf{R}_l(\text{Out}_C)(o_1, o_2) &:\Leftrightarrow \forall c \in C. \text{label}(c) \sqsubseteq l \Rightarrow o_1(c) = o_2(c) \end{aligned}$$

The graded monad $(W, \eta, \mu, \text{up})$ is given by $W_C A := A \times \text{Out}_C$.

Consider programs of type $R_l(W_C A)$

Printing Effects in Classified Sets (C'ed)

Consider programs of type $R_l(W_C A)$, when are these secure to “run”?

Printing Effects in Classified Sets (C'ed)

Consider programs of type $R_l(W_C A)$, when are these secure to “run”?

Intuitively, when the information at level l is allowed to flow to every channel $c \in C$.

Printing Effects in Classified Sets (C'ed)

Consider programs of type $R_l(W_C A)$, when are these secure to “run”?

Intuitively, when the information at level l is allowed to flow to every channel $c \in C$.

Indeed, in classified sets there is a map

$$\delta_{l,C,A} : R_l(W_C A) \rightarrow W_C(R_l A)$$

exactly when $l \sqsubseteq \text{label}(c)$ for all $c \in C$.

In the example of $\text{DCC} + \text{Moggi}^{\text{print}}$, we add the primitive

$$\frac{\Gamma \vdash t : R_{\text{public}}(T_{\text{public}} A)}{\Gamma \vdash \text{distr}(t) : T_{\text{public}}(R_{\text{public}} A)}$$

to the language.

In the example of DCC + Moggi^{print}, we add the primitive

$$\frac{\Gamma \vdash t : R_{\text{public}}(T_{\text{public}} A)}{\Gamma \vdash \text{distr}(t) : T_{\text{public}}(R_{\text{public}} A)}$$

to the language.

More generally, we obtain a prog. language for IFC with printing to multiple channels by *simply* combining DCC (or SC or ...) with EF_E (Katsumata [2]) via a primitive

$$\frac{\Gamma \vdash t : R_l(W_C A)}{\Gamma \vdash \text{distr}(t) : W_C(R_l A)} \quad \forall c \in C. l \sqsubseteq \text{label}(c)$$

Noninterference proofs à la Kavvos [3]

Noninterference proofs à la Kavvos [3]

“Explain” previous approaches to IFC with effects

Noninterference proofs à la Kavvos [3]

“Explain” previous approaches to IFC with effects

Study other effects: e.g. exceptions or global store

Thank you for your attention!

- [1] Martín Abadi et al. “A Core Calculus of Dependency”. In: *POPL '99, Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, TX, USA, January 20-22, 1999*. Ed. by Andrew W. Appel and Alex Aiken. ACM, 1999, pp. 147–160. DOI: 10.1145/292540.292555. URL: <https://doi.org/10.1145/292540.292555>.
- [2] Shin-ya Katsumata. “Parametric effect monads and semantics of effect systems”. In: *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*. Ed. by Suresh Jagannathan and Peter Sewell. ACM, 2014, pp. 633–646. DOI: 10.1145/2535838.2535846. URL: <https://doi.org/10.1145/2535838.2535846>.
- [3] G. A. Kavvos. “Modalities, cohesion, and information flow”. In: *Proc. ACM Program. Lang.* 3.POPL (2019), 20:1–20:29. DOI: 10.1145/3290333. URL: <https://doi.org/10.1145/3290333>.
- [4] Eugenio Moggi. “Notions of Computation and Monads”. In: *Inf. Comput.* 93.1 (1991), pp. 55–92. DOI: 10.1016/0890-5401(91)90052-4. URL: [https://doi.org/10.1016/0890-5401\(91\)90052-4](https://doi.org/10.1016/0890-5401(91)90052-4).
- [5] Naokata Shikuma and Atsushi Igarashi. “Proving Noninterference by a Fully