

### The Holy Grail of Gradual Security

Final Examination for Doctor of Philosophy in Computer Science

### Tianyu Chen

#### Indiana University



<sup>°</sup> The Summons. The Holy Grail tapestries by Morris & Co. Birmingham Museum and Art Gallery

### Tianyu's Thesis Statement

It *is possible* to design a gradual IFC programming language that satisfies both noninterference and the gradual guarantee while supporting type-based reasoning, by excluding the unknown label **\*** from runtime security labels and using security coercions to represent casts.

# Road Map

### 🖙 Background

- Explicit flow and implicit flow
- $\circ~$  Information flow control (IFC): static, dynamic, and gradual
- $\circ~$  The gradual guarantee and its tension with IFC
- Source of the tension: including **\*** in runtime labels
- $\lambda_{\rm IFC}^{\star}$ : a gradual IFC calculus
  - $\circ~\lambda^{\star}_{ ext{IFC}}$  enforces IFC while satisfying gradual guarantee
  - $\circ \lambda_{\rm IFC}^{\star}$  supports type-based reasoning (free theorems)
- Technical development
  - Formal definition of  $\lambda_{\text{IFC}}^{\star}$
  - Coercion calculi for IFC
  - $\circ$  The IFC cast calculus  $\lambda^c_{ ext{IFC}}$
- Meta-theoretic results
  - $\circ~$  Type safety for  $\lambda^{\star}_{\rm IFC}$
  - $\circ$  Gradual guarantee for  $\lambda^{\star}_{\mathrm{IFC}}$
  - Noninterference for  $\lambda_{\text{IFC}}^{\star}$

### Explicit Information Flow

```
Can we infer input from output in the following program?
```

```
let input = private-input () in
    publish (¬ input)
```

### Explicit Information Flow

Can we infer input from output in the following program?

```
let input = private-input () in
    publish (¬ input)
```

✓ Yes!

- Witness at least two executions
- Output is the negation of input
- Explicit flow

### Implicit Information Flow

Can we infer input from output in the following program?

let input = private-input () in
 publish (if input then false else true)

### Implicit Information Flow

Can we infer input from output in the following program?

let input = private-input () in
 publish (if input then false else true)

- 🗸 Also yes
  - ► Again, output is the negation of input
  - Implicit flow: input influences output through branching

### Information-Flow Control (IFC)

- Ensures that information transfers adhere to a security policy ►
- For example, high input must not flow to low output
- Propagate and check the security labels

• IFC in PL  $\begin{cases} static using a type system \\ dynamic using runtime monitoring \end{cases}$ 

### Static IFC Accepts Legal Explicit Flow

(Static IFC using a type system)

- 1 let fconst = λ b : Bool<sub>high</sub>. false in
  2 let input = private-input () in
  3 let result = fconst input in
  4 publish result
- ✓ Well-typed and runs successfully to unit
- Why? The return value of fconst is { always false of low-security
- ► Accepted by type-checker. No runtime check

 $<sup>^{\</sup>circ}$ private-input : Unit<sub>low</sub>  $\rightarrow$  Bool<sub>high</sub> and publish : Bool<sub>low</sub>  $\rightarrow$  Unit<sub>low</sub>

### Static IFC Rejects Illegal Explicit Flow

(Replace fconst with flip)

1 let flip = λ b : Bool<sub>low</sub>. ¬ b in
2 let input = private-input () in
3 let result = flip input in // compilation error
4 publish result

✗ Ill-typed. Illegal explicit flow:

- input is **high**
- flip expects low argument
- ► Rejected by type-checker. Again no runtime check

### Dynamic Enforcement of Explicit Flow

(Revisit flip with dynamic IFC)

I	let	flip	=	λb.	- b	in			
2	let	input	=	priva	ate-	inp	ut	()	in
3	let	result	=	flip	inp	ut	in		
4	рι	ublish r	esι	ılt	//	rur	ntime	e er	ror

**×** Errors at runtime (regardless of input)

• A runtime check happens before calling publish

In dynamic IFC, runtime values are tagged with their security level. The labels can originate from

- primitive operations
- annotations on literals
- ► the security level of the execution context

### Static Enforcement of Implicit Flow

(Different behavior in different branches)

### **✗** Ill-typed

- Security label on the type of if is the join (least upper bound) of its branches (low) and the branch condition (high).
- ► Rejected by type-checker. No runtime check

### Dynamic Enforcement of Implicit Flow

(Enforcing implicit flow with dynamic IFC)

- let flip =  $\lambda$  b. if b then false else true in
- 2 let input = private-input () in
- 3 let result = flip input in
- publish result
- **×** Errors at runtime (regardless of input)
- flip produces a high value because of high branch condition
- ► A runtime check happens before calling publish
- Illegal implicit flow ruled out at runtime

### Static IFC Is Hard to Use

"Another myth spread by security researchers is that the planet Earth contains more than six programmers who can correctly use security labels and information flow control (IFC)."<sup>1</sup>

<sup>&</sup>lt;sup>1</sup>James Mickens. This World of Ours. Usenix ;login: 2014

### Gradual IFC Comes to the Rescue!



In gradual IFC, static type information is optional:

- Go static in parts of the application where performance matters
- ► Go dynamic when performance matters less, and ease-of-use matters more

Principle of Locality: 90% of execution time in 10% of the code <sup>2</sup>

<sup>&</sup>lt;sup>2</sup>Dr. Ranjani Parthasarathi. Computer Architecture: Engineering And Technology

Gradual Typing Bridges Static and Dynamic IFC

Partially-annotated flip:

 $let flip : Bool_* → Bool_{low} =$  $λ b : Bool_*. if b then false else true in$ let input = private-input () inlet result = flip input in

publish result

- ► Well-typed but errors at runtime
- Checking happens on the boundaries between static and dynamic fragments
- ► The information flow violation is detected earlier than the dynamic version, as flip returns.

<sup>&</sup>lt;sup>2</sup>Blue: types with statically-known security. Red: types with unknown security.

### The Gradual Guarantee

#### less precise

#### more precise

- ► In the absense of errors, adding or removing security annotations does not change the result of the program.
- Adding security annotations may trigger errors.
- Removing security annotations may not trigger errors.

### The Gradual Guarantee

#### less precise

#### more precise

- ► In the absense of errors, adding or removing security annotations does not change the result of the program.
- Adding security annotations may trigger errors.
- Removing security annotations may not trigger errors.

There is a tension between the gradual guarantee and IFC!

# Static Enforcement of Flows Through Mutable References

```
1 let a = ref low true in
2 let input = private_input () in
3 if input then
4 a := false
6 a := true
7 publish (! a)
```

- ► The reference has type Ref (Bool<sub>low</sub>). It points to a low memory location.
- The type of the branch condition is Bool<sub>high</sub>
- X Writing to low memory under a high branch condition

### Dynamic Enforcement of Flows Through Mutable References

```
1 let a = ref low true in
2 let input = private_input () in
3 if input then
4 a := false
6 a := true
7 publish (! a)
```

The assignments fail at runtime because the no-sensitive-upgrade (NSU) mechanism <sup>3</sup> prevents writing to a low security pointer in a high security branch.

<sup>&</sup>lt;sup>3</sup>Austin and Flanagan. Efficient purely-dynamic information flow analysis. PLAS 2009.

### Counterexample of Gradual Guarantee in GSL<sub>Ref</sub>

#### less precise

```
1 let x = private-input () in
2 let a = ref * true* in
3 if x then (a := falsehigh)
4 else ()
```

#### more precise

- ✓ The more precise program (right) runs successfully
- ✗ But the less precise version (left) errors in GSL<sub>Ref</sub>⁴
- The assignment fails because it is in a high-security branch and GSL<sub>Ref</sub> conservatively treats the reference's label (\*) as if it were low

<sup>&</sup>lt;sup>4</sup>Toro, Garcia, Tanter. Type-Driven Gradual Security with References. TOPLAS 2018.

### But wait... GSL<sub>Ref</sub> allows \* labels on values?

The counterexample depends on labeling a reference with unknown security  $(\star)$ :

```
1 let x = private-input () in
2 let a = ref * true* in
3 if x then (a := falsehigh)
4 else ()
```

- ► Dynamic IFC languages don't use ★ as a runtime security label.
- ► Gradual languages traditionally use ★ for type checking, but not for categorizing runtime values.
- The inputs to an information flow system are the user's choices regarding what data is high or low security.

### Sources of the Tension with the Gradual Guarantee

Lang.	Noninter- ference	Gradual Guarantee	Type-guided classification	NSU	Runtime security labels
GSL <sub>Ref</sub>		×	<b>~</b>	1	<pre>{low, high, *}</pre>
GLIO	✓	<pre></pre>	×	1	<pre>{low, high}</pre>
WHILE <sup>G</sup>	✓			× – – – – – – – – – – – – – – – – – – –	<pre>{low, high, *}</pre>
$\lambda^{\star}_{ t{IFC}}$ (ours)	✓	<ul> <li></li> </ul>	<ul> <li>Image: A start of the start of</li></ul>	<ul> <li>Image: A second s</li></ul>	<pre>{low, high}</pre>

Removing  $\star$  from the runtime labels enables the gradual guarantee.

# Road Map

- Background
- ${}^{\tiny \hbox{\tiny ISC}} \ \lambda^{\star}_{\rm IFC}$ : a gradual IFC calculus
  - $\circ~\lambda^{\star}_{ ext{IFC}}$  enforces IFC while satisfying gradual guarantee
  - $\circ \lambda_{IFC}^{\star}$  supports type-based reasoning (free theorems)
  - Technical development
  - Meta-theoretic results

### $\lambda^{\star}_{\mathtt{IFC}}$ Excludes $\star$ From Runtime Labels

#### less precise

```
1 let x = private-input () in
2 let a : (Ref Bool<sub>*</sub>)<sub>*</sub> =
3 ref high true<sub>high</sub> in
4 if x then (a := false<sub>high</sub>)
5 else ()
```

#### more precise

```
let x = private-input () in
let a : (Ref Bool<sub>high</sub>)<sub>high</sub> =
    ref high true<sub>high</sub> in
    if x then (a := false<sub>high</sub>)
        else ()
```

The more precise program runs successfully to unit
 The less precise program also runs successfully to unit

### $\lambda^{\star}_{\mathtt{IFC}}$ Excludes $\star$ From Runtime Labels

#### less precise

```
1 let x = private-input () in
2 let a : (Ref Bool<sub>*</sub>)<sub>*</sub> =
3 ref high true<sub>high</sub> in
4 if x then (a := false<sub>high</sub>)
5 else ()
```

#### more precise

```
let x = private-input () in
let a : (Ref Bool<sub>high</sub>)<sub>high</sub> =
    ref high true<sub>high</sub> in
    if x then (a := false<sub>high</sub>)
        else ()
```

The more precise program runs successfully to unit
 The less precise program also runs successfully to unit
 Problem solved!

# Comparing $\lambda_{\rm IFC}^{\star}$ With ${\rm GSL}_{\rm Ref}$

- ► The default security label in \u03c6<sup>\*</sup><sub>IFC</sub> is low, so the programmer does not have to label constants
- Remove the labels on constants to obtain the following program, which also reduces successfully to unit:

```
let x = private-input () in
let a : (RefBool<sub>*</sub>)<sub>*</sub> = ref high true in
if x then (a := false)
        else ()
```

Comparing with the program in  $\ensuremath{\mathsf{GSL}_{\mathsf{Ref}}}$  , which errors:

### Vigilance: Type-Based Reasoning for Explicit Flows

Consider the example from Toro et al. [2018]:

<u>Free theorem</u>: The mix function either ① returns a value that does not depend on priv or ② produces a runtime error.

```
In \lambda_{\text{IFC}}^{\star}, 5 \langle \uparrow; \text{high}!; \text{low}?^p \rangle \longrightarrow \text{blame } p
```

# Type-Guided Classification: Type-Based Reasoning for Implicit Flows

Consider another example from Toro et al. [2018]:

<u>Free theorem</u>: The smix function either (1) returns a value that does not depend on priv or (2) produces a runtime error.

### Type-Based Reasoning for Implicit Flows in $\lambda^{\star}_{\text{IFC}}$

- let mix =  $\lambda$  pub priv. (if ((pub  $\langle 1ow! \rangle) < priv)$ then (1  $\langle 1ow! \rangle$ )  $\Rightarrow$  else (2  $\langle 1ow! \rangle$ ))  $\langle 1ow?^{p} \rangle$  in let smix =  $\lambda$  pub priv. mix pub (priv  $\langle high! \rangle$ ) in smix 1 (5  $\langle \uparrow \rangle$ )

- $\longrightarrow^*$  (if  $(1\langle low! \rangle < 5\langle \uparrow; high! \rangle)$  then  $1\langle low! \rangle$  else ...) $\langle low?^p \rangle$
- $\longrightarrow^*$  (if (true  $\langle \uparrow; high! \rangle$ ) then  $1 \langle low! \rangle$  else ...)  $\langle low?^p \rangle$
- $\longrightarrow^*$  (prot high  $(1\langle low! \rangle))\langle low?^p \rangle$
- $\longrightarrow^*$  1( $\uparrow$ ; high! )( low?<sup>p</sup> )
- $\longrightarrow^*$  blame p

# Road Map

- Background
- $\lambda_{\text{IFC}}^{\star}$ : a gradual IFC calculus
- Technical development
  - $\circ$  Formal definition of  $\lambda^{\star}_{ ext{IFC}}$
  - Coercion calculi for IFC
  - $\circ~$  The IFC cast calculus  $\lambda_{\rm IFC}^c$
  - Meta-theoretic results

### Syntax of $\lambda^{\star}_{\mathrm{IFC}}$

We define a gradual IFC calculus  $\lambda_{\text{IFC}}^{\star}$  with mutable references, first-class functions, and conditionals. Highlighted security labels default to low if omitted:

$$\ell \in \{\text{low}, \text{high}\}\$$

$$g \in \{\text{low}, \text{high}\}, \star\}$$

$$\iota ::= \text{Unit} \mid \text{Bool}$$

$$T ::= \iota \mid A \xrightarrow{g} A \mid \text{Ref}(T_g)$$

$$A ::= T_g$$

$$M ::= x \mid k \underset{\ell}{\ell} \mid (\lambda^g x : A \cdot M) \underset{\ell}{\ell} \mid (M M)^p$$

$$\mid (\text{if } M \text{ then } M \text{ else } M)^p$$

$$\mid (\text{ref } \ell M)^p \mid !^p M \mid (M := M)^p$$

### Semantics of $\lambda^{\star}_{\text{IFC}}$ The semantics of $\lambda^{\star}_{\text{IFC}}$ is by translation C to a cast calculus.

evaluation result r := k | fun | addr | diverge | stuck

obs(V) = r

$$\begin{array}{l} \textit{obs}(k) = k\\ \textit{obs}(k \ \langle \ \bm{c} \ \rangle) = k\\ \textit{obs}(\lambda x. \ N) = \texttt{fun}\\ \textit{obs}((\lambda x. \ N) \ \langle \ \bm{c} \ \rangle) = \texttt{fun}\\ \textit{obs}((\texttt{addr} \ n) = \texttt{addr}\\ \textit{obs}((\texttt{addr} \ n) \ \langle \ \bm{c} \ \rangle) = \texttt{addr} \end{array}$$

Let M be a well-typed  $\lambda_{\text{IFC}}^{\star}$  term:  $(x:\text{Bool}_{\text{high}}); \text{low} \vdash M : \text{Bool}_{\text{low}}$ 

 $\begin{array}{l} \textit{eval}(M,b) = r \\ \hline \textit{eval}(M,b) = \textit{obs}(V) & \text{if } (\mathcal{C} \ M)[x := b] \mid \varnothing \Downarrow V \mid \mu \\ \textit{eval}(M,b) = \textit{diverge} & \text{if } (\mathcal{C} \ M)[x := b] \mid \varnothing \Downarrow \textit{blame } p \mid \mu \\ & \text{or } (\mathcal{C} \ M)[x := b] \mid \varnothing \Uparrow \\ \textit{eval}(M,b) = \textit{stuck} & \text{otherwise} \end{array}$ 

### The Cast Calculus $\lambda_{ extsf{IFC}}^c$

The casts in  $\lambda_{\text{IFC}}^c$  are represented by coercions on types (a la Henglein) and coercions on security labels.

$$c ::= id(g) | \uparrow | \ell! | \ell?^{p} | \perp^{p}$$

$$\bar{c} ::= id(g) | \perp^{p} | \bar{c}; c$$

$$e ::= \ell | blame p | e \langle \bar{c} \rangle$$

$$c_{r} ::= id(\iota) | \operatorname{Ref} c c | (\bar{c}, c \rightarrow c)$$

$$c ::= (c_{r}, \bar{c})$$

$$M ::= x | k | \lambda x. M | let x=M:A in M$$

$$| M \langle c \rangle$$

$$| M \land M | M \land^{\star} M$$

$$| if M then M else M | if^{\star} M then M else M$$

$$| ref \ell M | ref?^{p} \ell M$$

$$| M := M | M :=?^{p} M$$

$$| addr n | prot e \ell M A | blame p$$

# Road Map

- Background
- $\lambda_{\rm IFC}^{\star}$ : a gradual IFC calculus
- Technical development
- Meta-theoretic results
  - $\circ~$  Type safety for  $\lambda^{\star}_{\rm IFC}$
  - $\circ$  Gradual guarantee for  $\lambda^{\star}_{ ext{IFC}}$
  - $\circ~$  Noninterference for  $\lambda^{\star}_{\rm IFC}$

$$\vdash r:A$$

$$\begin{array}{c|c} k:\iota \\ \hline \vdash k:\iota_{\ell} \end{array} & \begin{array}{c} \vdash \mathsf{addr}:(\mathsf{Ref}\;A)_g \end{array} \\ \hline \vdash \mathsf{fun}:(A \xrightarrow{g_2} B)_{g_1} & \begin{array}{c} \vdash \mathsf{diverge}:A \end{array} \end{array}$$

Theorem (Type safety of  $\lambda_{IFC}^{\star}$ ) If M is a well-typed  $\lambda_{IFC}^{\star}$  term:  $(x:Bool_{high}); low \vdash M : Bool_{low}$ and eval(M, b) = rthen the evaluation result is well-typed  $\vdash r : Bool_{low}$ . Theorem (Gradual guarantee for  $\lambda_{IFC}^{\star}$ ) Suppose M and M' are well-typed  $\lambda_{IFC}^{\star}$  terms:

> $(x:\operatorname{Bool}_{\operatorname{high}}); \operatorname{low} \vdash M : \operatorname{Bool}_{\operatorname{low}}$  $(x:\operatorname{Bool}_{\operatorname{high}}); \operatorname{low} \vdash M' : \operatorname{Bool}_{\operatorname{low}}$

and they are related by precision:  $\vdash M \sqsubseteq M'$ . If  $eval(M', b_1) = b_2$ 

then

$$eval(M, b_1) = b_2$$

### Theorem (Noninterference for $\lambda_{IFC}^{\star}$ ) If M is a well-typed $\lambda_{IFC}^{\star}$ term: $(x:Bool_{high}); low \vdash M : Bool_{low}$ and

$$eval(M, b_1) = b'_1$$

and

$$eval(M, b_2) = b'_2$$

*then*  $b'_1 = b'_2$ .

### Conclusion

- ► It is possible to satisfy noninterference and the gradual guarantee while supporting type-based reasoning.
- The security labels on constants and memory locations should default to low or high so that \* is not included in runtime security labels.
- Security checks can be modeled using coercions.

Show me the code! The Agda mechanization of  $\lambda^{\star}_{\rm IFC}$  is at

https://github.com/Gradual-Typing/LambdaIFCStar

# Thank you! ©

### NSU Checking

$$M \mid \mu \mid \textit{PC} \longrightarrow N \mid \mu'$$

$$\begin{array}{c|c} n \operatorname{FreshIn} \mu(\ell) & PC\langle \star \Rightarrow^{p} \ell \rangle \longrightarrow^{*} PC' \\ \hline & \operatorname{ref}?^{p} \ell V \mid \mu \mid PC \longrightarrow \operatorname{addr} n \mid (\mu, \ell \mapsto n \mapsto V) \\ \\ \operatorname{NF} \bar{c} & \vdash c : T_{g} \Rightarrow S_{\hat{\ell}} & \vdash d : S_{\hat{\ell}} \Rightarrow T_{g} \\ & (stamp! \ PC \mid \bar{c} \mid) \langle \star \Rightarrow^{p} \hat{\ell} \rangle \longrightarrow^{*} PC' \\ & V \langle c \rangle \longrightarrow^{*} W \\ \\ (\operatorname{addr} n \langle \operatorname{Ref} c d, \bar{c} \rangle) :=^{p} V \mid \mu \mid PC \longrightarrow \operatorname{unit} \mid [\hat{\ell} \mapsto n \mapsto W] \mu \end{array}$$

### Simulation Between More and Less Precise Coercions

Consider the following  $\lambda^{\star}_{\rm IFC}$  terms related by precision:

 $true_{low} : Bool_{\star} : Bool_{\star}$  and  $true_{low} : Bool_{high} : Bool_{\star}$ 

We need to show the two coercion sequences are related:

 $\vdash$  id(low); low!  $\sqsubseteq$  id(low);  $\uparrow$ ; high!