



The Holy Grail of Gradual Security

PL Wonks Seminar - Logic Seminar, Spring 2024

Tianyu Chen

Computer Science, Indiana University



Road Map

Background:

- Information flow properties and noninterference
 - Information flow control: static, dynamic, gradual
 - The gradual guarantee
 - The tension between noninterference and the gradual guarantee
- ▶ λ_{IFC}^* in Action
- ★ Solving the Tension Between Noninterference and the Gradual Guarantee
 - Type-Based Reasoning in λ_{IFC}^*
- ▶ Coercion-based Semantics for Gradual Security
- ▶ Meta-theoretical Results of λ_{IFC}^*

Information-Flow and Noninterference

Consider boolean negation. Can we infer output from input?

```
let output = ¬input
```

Yes! ✓

- ▶ Requires witnessing at least two executions:
input = true, output = false and input = false, output = true
- ▶ An information-flow property is a hyperproperty: a predicate on sets of executions
- ▶ Noninterference:¹ two successful executions of a program produce the same value for different input.

¹Specifically, termination-insensitive noninterference (TINI) for a functional programming language

Information-Flow Control (IFC)

- ▶ Security labels. Label input as **high**; output as **low**.
Track and check the security labels
- ▶ IFC in a programming language, traditionally
 - Static: using a type system
 - Dynamic: using runtime monitoring
- ▶ A *gradual* security programming language
 - embeds both **static** and **dynamic** IFC
 - enables seamless transition between static and dynamic
 - security label annotations $\left\{ \begin{array}{l} \text{specific: } \mathbf{low, high} \\ \text{statically unknown: } \star \end{array} \right.$

Review: Static IFC Using a Type System

Modeling input and output:

`user-input` : $\text{Unit}_{\text{low}} \rightarrow \text{Bool}_{\text{high}}$ `publish` : $\text{Bool}_{\text{low}} \rightarrow \text{Unit}_{\text{low}}$

Consider a statically-typed program:

```
1 let fconst = λ b : Boolhigh. false in
2 let input  = user-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 $\left\{ \begin{array}{l} \text{always false} \\ \text{of low-security} \end{array} \right.$
- ▶ No runtime check!
- ▶ IFC enforced by the type system alone

Guarding Against Illegal Explicit Flows, **Statically**

Consider another **fully static** program:

```
1 let fid      = λ b : Boollow . b in
2 let input    = user-input () in
3 let result = fid input in    // error
4   publish result
```

✗ **Ill-typed**. Why?

- ▶ *Illegal explicit flow* from the **high**-security input to fid
 - fid expects **low** argument
- * Program rejected by type-checker. Illegal explicit flow ruled out **at compile time**

Guarding Against Illegal Implicit Flows, **Statically**

Different observable behaviors in different branches:

```
1 let flip : Boolhigh → Boollow =  
2   λ b : Boolhigh. if b then false else true in  
3 let input = user-input () in  
4 let result = flip input in  
5   publish result
```

✗ Ill-typed

- ▶ Security label on the type of `if` is the join of its branches (both `low`) and the branch condition (`high`).
 - Expected: `low` (from type annotation `Boollow`)
 - Actual: `high` (because of conditional)
- * `high` $\not\leq$ `low`, thus rejected by type-checker. Illegal implicit flow ruled out **at compile time**

Guarding Against Illegal Explicit Flows, **Dynamically**

Consider the following dynamically-typed fid example that could potentially leak information through explicit flow:

```
1 let fid      = λ b : Bool*. b in
2 let input   = user-input () in
3 let result  = fid input in    // error
4   publish result
```

✓ Well-typed *but* ✗ Fails at runtime

- ▶ The program errors regardless of input
- ▶ A runtime happens before the call to `publish` and checks whether **high** can flow to **low** (of course, no)
- * Illegal explicit flow ruled out **at runtime**

¹We annotate `Bool*` explicitly, which conforms with the syntax of λ_{IFC}^*

Against Illegal Implicit Flows, **Dynamically**

Consider the following dynamically-typed flip example that could potentially leak information through implicit flow:

```
1 let flip : Bool* → Bool* =  
2   λ b : Bool*. if b then false else true in  
3 let input = user-input () in  
4 let result = flip input in  
5   publish result
```

- ✓ **Well-typed** *but* ✗ **Fails at runtime**
 - ▶ The program (again) errors regardless of input
 - ▶ flip produces a **high** value because of **high** branch condition
 - ▶ A runtime happens before the call to **publish** and checks whether **high** can flow to **low** (of course, no)
- * **Illegal implicit flow ruled out at runtime**

Gradual Typing Bridges Static and Dynamic IFC

Consider the following **partially annotated** version of `flip`. The return value must be `low`, because we intend to output the result:

```
1 let flip : Bool* → Boollow =  
2   λ b : Bool*. if b then false else true in  
3 let input = user-input () in  
4 let result = flip input in  
5   publish result
```

- ✓ **Well-typed** *but* ✗ **Fails at runtime** (for both true and false)
 - thus preventing the leak through implicit flow
- ▶ The information flow violation is detected *earlier* than the dynamic version, as `flip` returns
- * Checking happens on the boundaries between statically- and dynamically-typed code fragments

The Gradual Guarantee

- * **Removing annotations** from a correctly running program:

Example: $(42 : \text{low} : \text{high}) \sqsupseteq (42 : * : \text{high}) \sqsupseteq (42 : * : *)$
→ ... results in the same runtime behavior (42)

- * **Adding annotations** may introduce more errors:

Example: $(42 : * : * : *) \sqsubseteq (42 : \text{high} : * : \text{low})$
○ $(42 : * : * : *) \Downarrow 42$ but $(42 : \text{high} : * : \text{low}) \Downarrow \text{error}$

Satisfying Noninterference and the Gradual Guarantee in One Programming Language

... is **hard** according to the literature:

“We believe that there might be an **inherent incompatibility** between the strictness required to enforce a hyper-property like noninterference, and the optimistic flexibility dictated by the dynamic gradual guarantee.”

Matías Toro, Ronald Garcia, and Éric Tanter. 2018. Type-Driven Gradual Security with References

“There is some recent evidence that the dynamic gradual guarantee – which some see as essential to gradual typing – is **incompatible** with various hyperproperties, like noninterference and parametricity.”

Michael Greenberg. 2019. The Dynamic Practice and Static Theory of Gradual Typing

Review: No-Sensitive-Upgrade Checking

- ▶ No-sensitive-upgrade (NSU) (Austin and Flanagan 2009) prevents implicit flow leaks through writes to mutable references
- ▶ For gradual typing, NSU happens at **runtime**, when type information is insufficient in deciding if a heap write is secure
- ▶ Program that potentially leaks information through the heap:

```
1 let input : Bool* = user-input () in
2 let a      = ref low true in
3   if input then a := false else a := true ;
4   publish (! a)
```

- ✓ **Well-typed** *but* ✗ **Fails at runtime** (for both true and false)
- ▶ NSU checking **terminates** this program, because it attempts to write to a **low** memory location under a **high** execution context (PC), thus preventing the leak through heap

The Tension (in a Nutshell)

Toro et al. [2018] discover a tension between noninterference and the gradual guarantee in their language design, GSL_{Ref} .

Counterexample of the gradual guarantee in GSL_{Ref} :

Left: less precise, more dynamic

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

Right: more precise, more static

```
let x = user-input () in
let y = ref Boolhigh truehigh in
  if x then (y := falsehigh)
            else ()
```

✓ Both are well-typed

✓ The more precise (Right) program runs **successfully** to unit

✗ The less precise (Left) program **errors!**

- In GSL_{Ref} , \star corresponds to the interval $[\text{low}, \text{high}]$


✗ Violates **the gradual guarantee!**

Possible Sources of the Tension

Lang.	Noninterference	Gradual Guarantee	Type-guided classification	NSU	Runtime security labels
GSL _{Ref}	✓	✗	✓	✓	{low, high, ★}
GLIO	✓	✓	✗	✓	{low, high}
WHILE ^G	✓	✓	✓	✗	{low, high, ★}
λ_{IFC}^* (ours)	✓	✓	✓	✓	{low, high}

Road Map

- ▶ Background

- ▶  λ_{IFC}^* in Action:

- ★ Solving the Tension Between Noninterference and the Gradual Guarantee

- Type-Based Reasoning in λ_{IFC}^*

- ▶ Coercion-based Semantics for Gradual Security

- ▶ Meta-theoretical Results of λ_{IFC}^*

Solution to the Tension, in λ_{IFC}^*

Left: less precise, more dynamic

```
1 let x = user-input () in
2 let y : (Ref Bool*)* =
3   ref high truehigh in
4   if x then (y := falsehigh)
5     else ()
```

Right: more precise, more static

```
let x = user-input () in
let y : (Ref Boolhigh)high =
  ref high truehigh in
if x then (y := falsehigh)
  else ()
```

- ✓ Both are well-typed
- ✓ The more precise (Right) program runs **successfully** to unit
- ✓ The less precise (Left) one also runs **successfully** to unit.
- * Does *not* violate the gradual guarantee!

Problem solved!

But why?



Less precise in GSL_{Ref} :

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

Less precision in λ_{IFC}^* :

```
1 let x = user-input () in
2 let y : (Ref Bool*)* =
3   ref high truehigh in
4   if x then (y := falsehigh)
5             else ()
```

More precise in GSL_{Ref} :

```
let x = user-input () in
let y = ref Boolhigh truehigh in
  if x then (y := falsehigh)
            else ()
```

More precise in λ_{IFC}^* :

```
let x = user-input () in
let y : (Ref Boolhigh)high =
  ref high truehigh in
  if x then (y := falsehigh)
            else ()
```

In λ_{IFC}^* , Security labels on type annotations can be **specific** or **★**,
but those on literals and memory locations **stay specific**.

Omitted security label annotations on literals default to **low**:

Less precise in GSL_{Ref} :

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

Less precision in λ_{IFC}^* :

```
let x = user-input () in
let y : (Ref Bool*)* =
  ref high true in
  if x then (y := false)
  else ()
```

Solving the Tension in λ_{IFC}^* (Summary)

Design choices of GSL_{Ref} :

- ▶ Security labels on both types and literals can be \star
- ▶ Runtime security labels can also be \star (due to \star on literals)
- ▶ Runtime has to “guess” conservatively
 - more runtime errors when moving toward less precise
 - violates the gradual guarantee!

Design choices of λ_{IFC}^* :

- ▶ Security labels on **type annotations** may decrease in precision

$$(\text{Ref Bool}_\star)_\star \sqsubseteq (\text{Ref Bool}_{\text{high}})_{\text{high}}$$

- NSU checking happens. Heap IFC policy enforced at runtime
- ▶ Labels on **literals and memory locations** remain specific
 - security of data: only the programmer knows; must not be inferred
 - runtime security levels remain specific during program execution

Security Coercions as Runtime IFC Monitor

Revisit the dynamically-typed λ_{IFC}^* program:

```
1 let flip : Bool* → Bool* =  
2   λ b : Bool*. if b then false else true in  
3 let input = user-input () in  
4 let result = flip input in  
5   publish result
```

Compile the λ_{IFC}^* program to the following cast calculus λ_{IFC}^c term,
by making all casts explicit:

```
1 let flip = λ b . if b then (false <low!> )  
2                               else (true <low!> ) in  
3 let input = user-input () in  
4 let result = flip (input <high!> ) in  
5   publish (result <low?P> )
```

Reducing the λ_{IFC}^c term blames the projection (before calling publish):

→* `let result = ((λ b. if b then (false <low!>) else ...)
 (true <high!>)) in
 publish (result <low?P>)` (1)

→* `let result = prot low (if (true <high!>)
 then (false <low!>) else ...) in
 publish (result <low?P>)` (2)

→* `let result = prot low (prot high (false <low!>)) in
 publish (result <low?P>)` (3)

→* `let result = prot low (false <↑; high!>) in
 publish (result <low?P>)` (4)

→* `publish (false <↑; high!; low?P>)` (5)

→* `blame p` (6)

Sequencing models explicit flow. Stamping models implicit flow.
Checking by reducing coercion sequences

Type-Based Reasoning in λ_{IFC}^*

- ▶ Type-based reasoning: Toro et al. [2018] observe that security typing induces “free theorems” about noninterference
- ▶ Type-based reasoning is the synergy of two design choices:
 1. Vigilance
 2. Type-Guided Classification
- ▶ GLIO (Azevedo de Amorim et al. 2020) satisfies the gradual guarantee by sacrificing type-guided classification, which they claim to be the reason GSL_{Ref} (Toro et al. 2018) violates the gradual guarantee
- ▶ λ_{IFC}^* supports type-based reasoning just like GSL_{Ref}

Vigilance: Type-Based Reasoning for Explicit Flows

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

```
1 let mix : Intlow → Inthigh → Intlow =  
2   λ pub priv .  
3     if pub < (priv : Int* : Intlow) then 1 else 2 in  
4   mix 1low 5low
```

Free theorem: Either ① the **low** result of `mix` never depends on the **high** `priv` argument or ② **mix produces a runtime error**.

(GLIO: not vigilant → does not produce an error → violates the free theorem)

In λ_{IFC}^* , $5 \langle \uparrow; \text{high}!; \text{low}?^p \rangle \Downarrow \text{blame } p$

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

Another example from Toro et al. [2018]:

```
1 let mix : Intlow → Int* → Intlow =  
2   λ pub priv. if pub < priv then 1 else 2 in  
3 let smix : Intlow → Inthigh → Intlow =  
4   λ pub priv. mix pub priv in  
5   smix 1low 5low
```

Free theorem: The `smix` function either ① returns a value that does not depend on `priv` or ② produces a runtime error

(GLIO: ① not vigilant ② does not classify values using types → does not produce an error → violates the free theorem)

```

1 let mix = λ pub priv.
2   (if (pub <low!>) < priv
3     then (1 <low!>)
4     else (2 <low!>)) <low?P> in
5 let smix = λ pub priv. mix pub (priv <high!>) in
6   smix 1 (5 <↑>)

```

→* (if (1 <low!> < 5 <↑; high!>) then 1 <low!> else ...) <low?^P> (7)

→* (if (true <↑; high!>) then 1 <low!> else ...) <low?^P> (8)


→* (prot high (1 <low!>)) <low?^P> (9)

→* 1 <↑; high!> <low?^P> (10)

→* blame *p* (11)

In λ_{IFC}^* , the program errors, thus satisfying the free theorem

Road Map

- ▶ Background
- ▶ λ_{IFC}^* in Action
- ▶  Coercion-based Semantics for Gradual Security
- ▶ Meta-theoretical Results of λ_{IFC}^*

Coercion Calculus for Security Labels

Syntax and typing for security coercions and coercion sequences:

specific security labels	ℓ	\in	$\{\text{low}, \text{high}\}$
gradual security labels	g	$::=$	$\star \mid \ell$
blame labels	p, q		
security coercions	c, d	$::=$	$\text{id}(g) \mid \uparrow \mid \ell! \mid \ell?^p \mid \perp^p$
coercion sequences	\bar{c}, \bar{d}	$::=$	$\text{id}(g) \mid \perp^p g_1 g_2 \mid \bar{c}; c$

$\vdash c : g_1 \Rightarrow g_2$

 $\vdash \text{id}(g) : g \Rightarrow g$

 $\vdash \uparrow : \text{low} \Rightarrow \text{high}$

 $\vdash \ell! : \ell \Rightarrow \star$

 $\vdash \ell?^p : \star \Rightarrow \ell$

 $\vdash \perp^p : \text{high} \Rightarrow \text{low}$

Reduction semantics and normal forms of the coercion calculus on security labels:

$\boxed{\text{NF } \bar{c}}$

$$\frac{}{\text{NF id}(g)} \quad \frac{}{\text{NF id}(\star); \ell ?^P} \quad \frac{\text{NF } \bar{c}}{\text{NF } \bar{c}; \ell !} \quad \frac{\text{NF } \bar{c}}{\text{NF } \bar{c}; \uparrow}$$

$\boxed{c; c \longrightarrow c}$

$$\begin{array}{c} ?\text{-id} \frac{}{\ell !; \ell ?^P \longrightarrow \text{id}(\ell)} \quad ?\text{-}\uparrow \frac{}{\text{low} !; \text{high} ?^P \longrightarrow \uparrow} \\ \\ ?\text{-}\perp \frac{}{\text{high} !; \text{low} ?^P \longrightarrow \perp^P} \end{array}$$

$\boxed{\bar{c} \longrightarrow \bar{d}}$

$$\begin{array}{c} \text{id} \frac{\text{NF } \bar{c}}{\bar{c}; \text{id}(g) \longrightarrow \bar{c}} \quad \perp \frac{\text{NF } \bar{c} \quad \vdash \bar{c} : g_1 \Rightarrow g_2}{\bar{c}; \perp^P \longrightarrow \perp^P g_1 \text{ low}} \\ \\ \xi\text{-}\perp \frac{\vdash c : g_2 \Rightarrow g_3}{\perp^P g_1 g_2; c \longrightarrow \perp^P g_1 g_3} \\ \\ \xi_L \frac{\bar{c} \longrightarrow \bar{d}}{\bar{c}; c \longrightarrow \bar{d}; c} \quad \xi_R \frac{\text{NF } \bar{c} \quad c; d \longrightarrow c'}{\bar{c}; c; d \longrightarrow \bar{c}; c'} \end{array}$$

(A Glimpse of) the Cast Calculus λ_{IFC}^c

- ▶ Representation of PC: label expressions

$e, PC ::= \ell \mid \text{blame } p \mid e \langle \bar{c} \rangle$

- ▶ Coercions on values of λ_{IFC}^c :

base types	ι	$::=$	Unit Bool
raw types	T, S	$::=$	$\iota \mid A \xrightarrow{g^c} B \mid \text{Ref } (T_g)$
types	A, B	$::=$	T_g
raw coercions	c_r, d_r	$::=$	$\text{id}(\iota) \mid \text{Ref } \mathbf{c} \mathbf{d} \mid (\bar{d}, \mathbf{c} \rightarrow \mathbf{d})$
coercions	\mathbf{c}, \mathbf{d}	$::=$	c_r, \bar{c}

- ▶ NSU checking: reducing label expressions

$$\frac{
 \frac{
 n \text{ FreshIn } \mu(\ell) \quad PC \langle \star \Rightarrow^P \ell \rangle \longrightarrow^* PC'
 }{
 \text{ref?}^P \ell V \mid \mu \mid PC \longrightarrow \text{addr } n \mid (\mu, \ell \mapsto n \mapsto V)
 }
 }{
 \text{NF } \bar{c} \quad (\text{stamp! } PC \mid \bar{c} \mid) \langle \star \Rightarrow^P \hat{\ell} \rangle \longrightarrow^* PC' \quad V \langle \mathbf{c} \rangle \longrightarrow^* W
 }
 \text{assign?}^P (\text{addr } n \langle \text{Ref } \mathbf{c} \mathbf{d}, \bar{c} \rangle) V T g \mid \mu \mid PC \longrightarrow \$ \text{unit} \mid [\hat{\ell} \mapsto n \mapsto W] \mu$$

$$\vdash \mathbf{c} : T_g \Rightarrow S_{\hat{\ell}}, \vdash \mathbf{d} : S_{\hat{\ell}} \Rightarrow T_g$$

Road Map

- ▶ Background
- ▶ λ_{IFC}^* in Action
- ▶ Coercion-based Semantics for Gradual Security
- ▶ Meta-theoretical Results of λ_{IFC}^*

Theorem (Compilation preserves types)

If $\Gamma; g \vdash M : A$, then $\Gamma; \emptyset; g; \text{low} \vdash \mathcal{C} M : A$.

Theorem (Progress)

Suppose PC is well-typed: $\vdash PC \Leftarrow g$, M is well-typed:

$\emptyset; \Sigma; g; |PC| \vdash M \Leftarrow A$, and the heap μ is well-typed: $\Sigma \vdash \mu$.

Then either (1) M is a value or (2) M is a blame or (3) M can take a reduction step: $M \mid \mu \mid PC \longrightarrow N \mid \mu'$ for some N and μ' .

Theorem (Preservation)

Suppose PC is well-typed: $\vdash PC \Leftarrow g$, M is well-typed:

$\emptyset; \Sigma; g; |PC| \vdash M \Leftarrow A$ and the heap μ is well-typed: $\Sigma \vdash \mu$.

If $M \mid \mu \mid PC \longrightarrow N \mid \mu'$, there exists Σ' s.t $\Sigma' \supseteq \Sigma$,

$\emptyset; \Sigma'; g; |PC| \vdash N \Leftarrow A$, and $\Sigma' \vdash \mu'$.

Theorem (The gradual guarantee)

Suppose M, M' are related by precision:

$$\emptyset; \emptyset; \emptyset; \emptyset; \text{low}; \text{low}; \text{low}; \text{low} \vdash M \sqsubseteq M' \Leftarrow A \sqsubseteq A'$$

If M' evaluates to a value:

$$M' \mid \emptyset \mid \text{low} \longrightarrow^* V' \mid \mu'$$

there exists V and μ s.t. M evaluates to V :

$$M \mid \emptyset \mid \text{low} \longrightarrow^* V \mid \mu$$

and the resulting values are related by precision for some Σ, Σ' :

$$\emptyset; \emptyset; \Sigma; \Sigma'; \text{low}; \text{low}; \text{low}; \text{low} \vdash V \sqsubseteq V' \Leftarrow A \sqsubseteq A'$$

The noninterference of λ_{IFC}^* is conjectured by that of λ_{SEC}^* :

λ_{IFC}^* performs type-guided classification but λ_{SEC}^* does not, so the value that a λ_{IFC}^* program produces is at least as secure as the value produced by the same program in λ_{SEC}^* .

Theorem (Noninterference of λ_{SEC}^*)

If M is well-typed $(x:\text{Bool}_{\text{high}}); \emptyset; \text{low}; \text{low} \vdash M : \text{Bool}_{\text{low}}$ and

$$\emptyset \mid \text{low} \vdash M[x := (b_1)_{\text{high}}] \Downarrow V_1 \mid \mu_1$$

$$\emptyset \mid \text{low} \vdash M[x := (b_2)_{\text{high}}] \Downarrow V_2 \mid \mu_2$$

then $V_1 = V_2$.

Code and Data Availability

<https://github.com/Gradual-Typing/LambdaSecStar>

```
tianyu@beluga:~/workspace/agda/LambdaSecStar
> tianyu @ beluga in ~/workspace/agda/LambdaSecStar on git:master v [17:15:00]
$ polyglot --exclude plot/ --exclude notes/ --exclude bin/ --exclude _build/ --exclude src/CCExpSub .
-----
Language      Files    Lines   Code   Comments  Blanks
-----
Agda           201     22618  19774   288       2556
Makefile       1         20     15      0         5
Markdown       4        182    148     0        34
-----
Total          206     22820  19937   288       2595
-----

> tianyu @ beluga in ~/workspace/agda/LambdaSecStar on git:master v [17:15:06]
$ |
```

Main Takeaways

1. It is possible to satisfy both noninterference and the gradual guarantee in a gradual security-typed language, provided that the security level of data remains specific at runtime
2. Gradual information flow can be represented as coercions. In particular, NSU checking is a special projection that casts PC to the security of the memory location to modify
3. The key to the semantics design of of a gradual security-typed language is identifying injections ($\ell !$) and projections ($\ell ?^p$)

Thank you for your attention!