

Quest Complete: The Holy Grail of Gradual Security

TIANYU CHEN and JEREMY G. SIEK, Indiana University, USA

Languages with gradual information-flow control combine static and dynamic techniques to prevent security leaks. Gradual languages should satisfy the gradual guarantee: programs that only differ in the precision of their type annotations should behave the same modulo cast errors. Unfortunately, Toro et al. [2018] identify a tension between the gradual guarantee and information security; they were unable to satisfy both properties in the language GSL_{Ref} and had to settle for only satisfying information-flow security. Azevedo de Amorim et al. [2020] show that by sacrificing type-guided classification, one obtains a language that satisfies both noninterference and the gradual guarantee. Bichhawat et al. [2021] show that both properties can be satisfied by sacrificing the no-sensitive-upgrade mechanism, replacing it with a static analysis.

In this paper we present a language design, λ_{IFC}^* , that satisfies both noninterference and the gradual guarantee without making any sacrifices. We keep the type-guided classification of GSL_{Ref} and use the standard no-sensitive-upgrade mechanism to prevent implicit flows through mutable references. The key to the design of λ_{IFC}^* is to walk back the decision in GSL_{Ref} to include the unknown label \star among the runtime security labels. We give a formal definition of λ_{IFC}^* , prove the gradual guarantee, and prove noninterference. Of technical note, the semantics of λ_{IFC}^* is the first gradual information-flow control language to be specified using coercion calculi (à la Henglein), thereby expanding the coercion-based theory of gradual typing.

CCS Concepts: • **Theory of computation**; • **Security and privacy** → **Formal security models**; • **Software and its engineering** → **Formal software verification**; **Semantics**;

Additional Key Words and Phrases: gradual typing, information flow security, machine-checked proofs, Agda

ACM Reference Format:

Tianyu Chen and Jeremy G. Siek. 2024. Quest Complete: The Holy Grail of Gradual Security. *Proc. ACM Program. Lang.* 8, PLDI, Article 212 (June 2024), 24 pages. <https://doi.org/10.1145/3656442>

1 INTRODUCTION

Information-flow control (IFC) ensures that information transfers within a program adhere to a security policy, e.g., by preventing high-security data from flowing to a low-security channel. This adherence can be enforced statically using a type system [Myers 1999; Myers and Liskov 1997; Volpano et al. 1996], or dynamically using runtime monitoring [Askarov and Sabelfeld 2009; Austin and Flanagan 2009; Austin et al. 2017; Devriese and Piessens 2010; Stefan et al. 2011; Xiang and Chong 2021], or using static analysis to pre-compute information that facilitates runtime monitoring [Chandra and Franz 2007; Le Guernic 2007; Le Guernic and Jensen 2005; Moore and Chong 2011; Russo and Sabelfeld 2010; Shroff et al. 2007]. The static and dynamic approaches have complementary strengths and weaknesses; the dynamic approach requires less effort from the programmer while the static approach provides stronger guarantees and less runtime overhead.

Taking inspiration from gradual typing [Siek and Taha 2006, 2007], researchers have explored how to give programmers seamless control over which parts of the program are secured statically versus dynamically. In general, gradually typed languages support the seamless transition between

Authors' address: Tianyu Chen, chen512@iu.edu; Jeremy G. Siek, jsiek@indiana.edu, Indiana University, Bloomington, Indiana, USA, 47408.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

ACM 2475-1421/2024/6-ART212

<https://doi.org/10.1145/3656442>

static and dynamic enforcement through the precision of type annotations. The programmer can choose when it is appropriate to increase the precision of the type annotations and put in the effort to pass the static checks and when it is appropriate to reduce the precision of type annotations, deferring the enforcement to runtime. The main property of gradually typed languages is the *gradual guarantee* [Siek et al. 2015], which states that removing type annotations should not change the runtime behavior. Adding type annotations should also result in the same behavior except that it may introduce more trapped errors because those new type annotations may contain mistakes.

The main challenge in the design of gradually typed languages is controlling the flow of values (and information) between the static and dynamic regions of code, which is accomplished using runtime casts. Typically source programs are compiled to an intermediate language, called a cast calculus, that includes explicit syntax for runtime casts. Disney and Flanagan [2011] design a cast calculus with IFC for a pure lambda calculus and prove noninterference. Fennell and Thiemann [2013] design a cast calculus named ML-GS with mutable references using the no-sensitive-upgrade (NSU) runtime checks of Austin and Flanagan [2009]. Fennell and Thiemann [2015] design a cast calculus for an imperative, object-oriented language.

Similar to gradual security-typed languages, Hybrid LIO (HLIO) [Buiras et al. 2015] also supports the choice of static or dynamic IFC in different parts of a single program. By default the checking is static, but a programmer can insert a defer clause to say that the security constraints should be checked at runtime. In gradual security-typed languages, the switch between static and dynamic is directed by types, so unlike in HLIO, the developer does not need to embed explicit casts or defer into the program. Moreover, the gradual guarantee relates the runtime behavior of programs that differ in the precision of type annotations, but there is no comparable theorem about adding and removing defer in HLIO.

Since the formulation of the gradual guarantee as a criterion for gradually typed languages [Siek et al. 2015], researchers have explored the feasibility of satisfying both the gradual guarantee and noninterference. Toro et al. [2018] identify a tension between the gradual guarantee and security enforcement. They analyze the semantics of runtime casts through the lens of Abstracting Gradual Typing [Garcia et al. 2016] and propose a type-driven semantics for gradual security. However, Toro et al. [2018] discover counterexamples to the gradual guarantee in the GSL_{Ref} language. They conjecture that it is not possible to enforce noninterference and satisfy the gradual guarantee.

Azevedo de Amorim et al. [2020] conjecture one possible source of the tension: the type-guided classification performed in GSL_{Ref} [Toro et al. 2018]. They propose a new gradually typed language, GLIO, which sacrifices type-guided classification. They prove that GLIO satisfies both noninterference and the gradual guarantee using a denotational semantics. Bichhawat et al. [2021] conjecture that *NSU checking* could be another possible source of the tension. As an alternative, they propose a hybrid approach that leverages static analysis ahead of program execution to determine the write effects in untaken branches. They study a simple imperative language with first-order stores and prove both noninterference and the gradual guarantee.

Contrary to the prior work, we show that one does not need to give up on type-guided classification or NSU checking to resolve the tension. Instead, the tension can be resolved by walking back a design choice in GSL_{Ref} , which was to allow \star as a runtime security label. For example, in GSL_{Ref} one can write a literal such as `true \star` in a program, and at runtime the literal becomes a value of unknown security level. That design was unusual because the unknown type \star is traditionally used in gradual languages to represent the lack of static information, not the lack of dynamic information. The design is also unusual when compared to dynamic systems for IFC, as those systems do not use an unknown security level [Askarov and Sabelfeld 2009; Austin and Flanagan 2009; Austin et al. 2017; Devriese and Piessens 2010; Stefan et al. 2011].

Table 1. Proposed sources of tension between security and the gradual guarantee

Language	Security (noninterference)	Gradual Guarantee	Type-guided classification	NSU checking	Runtime security labels
GSL _{Ref}	✓ Yes	✗ No	✓ Yes	✓ Yes	{low, high, ★}
GLIO	✓ Yes	✓ Yes	✗ No	✓ Yes	{low, high}
WHILE ^G	✓ Yes	✓ Yes	✓ Yes	✗ No	{low, high, ★}
λ_{IFC}^* (this paper)	✓ Yes	✓ Yes	✓ Yes	✓ Yes	{low, high}

One might think that allowing ★ as a label on literals and therefore on values is necessary so that programmers can run legacy code (without any security annotations) in a gradual language, by making ★ the default label for literals. However, prior information-flow languages use low security as the default security label for literals [Myers et al. 2006] and for good reasons. The security of a literal is something that only the programmer can know. That is, the identification of high-security data in a program must be considered as an input to an information flow system, and not something that can or should be inferred. When migrating legacy code into a system that supports secure information flow, a necessary part of the process for the programmer is to identify whether there is any high-security information in the legacy code. Our choice of low as the default label is because most literals (if not all) in real programs are low security. In fact, it is bad practice to embed high-security literals, such as passwords, in program text.

In our design, runtime security labels do not include ★, only low and high.¹ On the other hand, to support gradual typing, the security labels in a type annotation may include ★. Surprisingly, we find that removing ★ from the runtime labels is sufficient to reclaim the gradual guarantee, without sacrificing type-guided classification as in GLIO or NSU checking as in WHILE^G. This finding is the primary contribution of this paper. In our design, the security level of a literal defaults to low, similar to systems like Jif [Myers et al. 2006] and GLIO, but different from GSL_{Ref} and WHILE^G. We propose a new gradual, security-typed language λ_{IFC}^* , which (1) enforces information flow security, (2) satisfies the gradual guarantee, (3) enjoys type-guided classification, and (4) utilizes NSU checking to enforce implicit flows through the heap with no static analysis required.

The semantics of λ_{IFC}^* is given by translation to a new security cast calculus λ_{IFC}^c , for which we define a syntax, type system, and operational semantics. We compile λ_{IFC}^* into λ_{IFC}^c in a type-preserving way. In λ_{IFC}^c , security coercions serve as our runtime security monitor, in which we adapt ideas from the Coercion Calculus [Henglein 1994; Herman et al. 2010] to IFC.

Compared to prior work on gradual IFC languages, the λ_{IFC}^c cast calculus supports an additional feature called blame tracking [Findler and Felleisen 2002]. Blame tracking is important because it enables modular runtime error messages, e.g., they play an important role in production-quality languages such as Typed Racket [Tobin-Hochstadt and Felleisen 2008; Wilson et al. 2018].

Compared to prior work on gradual IFC based on abstracting gradual typing (AGT) [Toro et al. 2018], our use of a cast calculus makes it clear where in the program there is runtime overhead from dynamic checking (the casts). This is important for the programmer to know because one may wish to avoid runtime overhead in hot regions of a programs. In our translation from λ_{IFC}^* to λ_{IFC}^c , casts are only inserted where there is insufficient information during compilation to decide whether a security policy is enforced or not. In particular, casts are not inserted in statically typed regions. In contrast, the AGT mechanism for dynamic checking (called evidence) is attached to most nodes in the syntax tree.

¹Of course, any lattice of security labels could be used in place of low and high.

This paper makes the following technical contributions:

- Identify the real cause of the tension between information flow security and the gradual guarantee in GSL_{Ref} : the inclusion of \star in the runtime security levels (§ 2).
- A coercion calculus for security labels (§ 3) and a coercion calculus for secure values (§ 4). The two coercion calculi serve as our runtime IFC monitor. Our paper is the first work to apply the coercion calculus to an IFC setting.
- A cast calculus λ_{IFC}^c with IFC that defines the dynamic semantics of $\lambda_{\text{IFC}}^\star$ (§ 5). The proofs of (1) the gradual guarantee (§ 6.4), (2) compilation from $\lambda_{\text{IFC}}^\star$ to λ_{IFC}^c preserves types (§ 6.2), and (3) type safety of λ_{IFC}^c (§ 5.2.1) are mechanized in Agda.
- A proof of noninterference for $\lambda_{\text{IFC}}^\star$ (§ 6.3) through the simulation between its cast calculus λ_{IFC}^c and a dynamic IFC programming language (§ 5.2.3).
- The first design of a gradual security-typed language with type-guided classification that satisfies the gradual guarantee (§ 6).
- We mechanize the proof of the gradual guarantee in the Agda proof assistant.

The Appendix is in the supplementary material of this paper. The Agda code is available at:

<https://github.com/Gradual-Typing/LambdaIFCStar> at release v2.0 (PLDI 2024)

2 $\lambda_{\text{IFC}}^\star$ IN ACTION

In this section we present example programs that demonstrate how $\lambda_{\text{IFC}}^\star$ enables a gradual, smooth transition between static and dynamic information-flow control, while supporting type-based reasoning and satisfying the gradual guarantee. We briefly review the basics of gradual security typing in Section 2.1. In Section 2.2, we show that the tension between security and the gradual guarantee can be achieved by removing \star from the runtime security labels. In Section 2.3, we demonstrate that $\lambda_{\text{IFC}}^\star$ enables the same type-based reasoning capabilities as GSL_{Ref} .

For simplicity, we use the security lattice $\langle \{\text{high}, \text{low}\}, \preceq, \vee, \wedge \rangle$, where **high** is for private data and **low** is for publicly disclosable data. The ordering is standard: $\text{low} \preceq \text{high}$ and $\text{high} \not\preceq \text{low}$. So information is allowed to flow from public sources to private sinks but not the other way around. We refer to $\{\text{high}, \text{low}\}$ as *specific security labels*.

Types in $\lambda_{\text{IFC}}^\star$ have security labels associated with them, for example, $\text{Bool}_{\text{high}}$ is the type for booleans with high security, Unit_{low} is the type for the unit value with low security, and Bool_\star is the type of a boolean whose security level is unknown at compile time. We refer to $\{\text{high}, \text{low}, \star\}$ as *security labels*. We define a *precision* ordering \sqsubseteq on them, where $\star \sqsubseteq g$ for any label g and $\ell \sqsubseteq \ell$ for any specific security label ℓ . The precision ordering extends to types in a natural way, so for example, $\text{Bool}_\star \sqsubseteq \text{Bool}_{\text{low}}$. Figure 18 of the Appendix gives the definition of precision on types.

To enable information-flow control, $\lambda_{\text{IFC}}^\star$ allows the programmer to annotate constants, mutable references, and λ -abstractions with a specific security label and ensures that if a value is annotated with **high**, it will not flow into a sink that is **low** security. If the programmer does not annotate a value with a label, $\lambda_{\text{IFC}}^\star$ defaults the value's label to **low**. So `true` is shorthand for `truelow`.

We model I/O with two functions, `user-input` and `publish`: the former returns a high-security boolean that represents sensitive input information; the latter takes a low-security boolean and publishes it into a publicly visible channel.

2.1 Reviewing the Basics of Gradual Information Flow Security, in $\lambda_{\text{IFC}}^\star$

In this section we review the basic concepts of gradual information flow control using $\lambda_{\text{IFC}}^\star$. We start with fully static $\lambda_{\text{IFC}}^\star$ programs and show that $\lambda_{\text{IFC}}^\star$ can behave like a static security-typed language, guarding against both illegal explicit and implicit flows at compile time. We then replace some security label annotations in types with \star , so that the programs become partially typed and

the typing information alone is insufficient to enforce IFC. We show that coercions, our runtime security monitor, are able to capture both explicit flow and implicit flow violations at runtime, preventing information leakage and enforcing security.

Gradual IFC includes static IFC. For statically typed programs, λ_{IFC}^* behaves just like a statically typed IFC language. Consider the following well-behaved λ_{IFC}^* program that takes in a high-security user input, passes it to the function `fconst` that ignores the input and returns `false`, which is then published.

```

1 let fconst =  $\lambda$  b : Boolhigh. false in
2 let input = user-input () in
3 let result = fconst input in
4   publish result

```

The program type-checks and runs without error, with no need for runtime checks to enforce security. Indeed, a malicious party cannot infer anything about the high-security input because (1) the return value of `fconst` is always the same value `false` (2) the value `false` is of low security, so the explicit flow into `publish` is allowed.

If we replace `fconst` with the identity function `fid` with parameter type `Boollow`, the program becomes ill-typed because the type system of λ_{IFC}^* does not allow the explicit flow from the high-security input to `fid`, as is usual for a statically typed IFC language.

```

1 let fid =  $\lambda$  b : Boollow. b in
2 let input = user-input () in
3 let result = fid input in // static error, input is high security but fid expects low
4   publish result

```

Sometimes the observable behaviors of a program can depend on its branching structure. If some of the branch conditions have a data dependency on high-security input, a malicious party might be able to infer it from the observable behaviors, giving rise to illegal *implicit flows* [Denning 1976], which must be ruled out to guarantee security.

Consider the following program in which the function `flip` contains a conditional expression, whose condition is dependent on a high-security user input. Its two branches return different low-security booleans, creating a potential implicit flow from high to low:

```

1 let flip : Boolhigh  $\rightarrow$  Boollow =  $\lambda$  b : Boolhigh. if b then false else true in
2 let input = user-input () in
3 let result = flip input in
4   publish result

```

As is typical of statically typed IFC languages, the type system of λ_{IFC}^* rejects this program, thereby preventing an information leak through an implicit flow. To see why, note that the branch condition is of high security, so the type of the `if` expression as a whole is `Boolhigh`. In particular, the type checker computes the security level of an `if` to be the join of its branches (both `low`) and the condition (`high`), yielding `low \vee high = high`. The `flip` function is expected to return `Boollow` according to its type annotation, but returns `Boolhigh` because of the conditional, `high $\not\leq$ low`, so the program is ill-typed.

To summarize, λ_{IFC}^* behaves just like a static security-typed language in the above examples. When everything is statically typed, the type system of λ_{IFC}^* guards against illegal information flows, whether explicit or implicit.

Gradual IFC enables a mixture of static and dynamic IFC. We have seen that security labels (`low` and `high`) can appear in type annotations in a program, such as `Boollow` and `Boolhigh`. λ_{IFC}^* also

provides the *unknown security label*, written \star , for use in type annotations. We explain how the unknown security label works in the following discussion.

We return to the `fconst` example except this time the type of parameter `b` is `Bool \star` .

```

1 let fconst = λ b : Bool $\star$  . false in
2 let input  = user-input () in
3 let result = fconst input in
4   publish result

```

The type system of $\lambda_{\text{IFC}}^{\star}$ accepts this program because, in the call `fconst input`, it allows an implicit conversion from the type of `input`, which is `Boolhigh`, to the parameter type `Bool \star` . This program runs to completion and publishes `false`.

Now suppose we again replace `fconst` with `fid`, but keep the parameter type of `Bool \star` .

```

1 let fid    = λ b : Bool $\star$  . b in
2 let input  = user-input () in
3 let result = fid input in
4   publish result

```

The type system of $\lambda_{\text{IFC}}^{\star}$ still accepts this program. The type of result changes to `Bool \star` but in the call `publish result`, the type system allows an implicit conversion from `Bool \star` to `Boollow`. The security leak in this program is not caught statically; instead it is caught dynamically.

The dynamic semantics of $\lambda_{\text{IFC}}^{\star}$ is defined by compilation into λ_{IFC}^c by inserting casts. In λ_{IFC}^c , explicit casts are represented as *security coercions* that monitor the flow of information. We use the standard syntax for coercions [Henglein 1994] but with adaptations to handle IFC. A coercion whose target is \star (and source is not \star) is an *injection*, and is indicated by an exclamation mark. A coercion whose source is \star (and target is not \star) is a *projection*, and is indicated by a question mark. The projections perform runtime checks that may fail.

The translation from $\lambda_{\text{IFC}}^{\star}$ to λ_{IFC}^c inserts a security coercion wherever an implicit cast occurred in the type checking of the $\lambda_{\text{IFC}}^{\star}$ term. Here is the result of cast insertion on the above program:

```

1 let fid    = λ b . b in
2 let input  = user-input () in
3 let result = fid (input <high!>) in
4   publish (result <low?P>)

```

The coercion on Line 3 (`high!`) is an injection, casting from `high` to \star . The coercion on line 4 (`low?P`) is a projection, casting from \star to `low`. At runtime, a projection checks whether the incoming value has a security level that is less than or equal to the target security level. Now suppose we run the above example with `input true`. The injection on line 3 will create an injected value `truehigh <high!>`. This value is passed to and returned from `fid`, and then projected to `low`. Because `high` is greater than `low`, the projection fails.

Each projection is annotated with an identifier called a blame label (p). In case a projection fails, it raises a cast error, called *blame*, that contains its blame label. In this way, the programmer knows which cast is causing the problem. This feature is often referred to as *blame tracking* [Findler and Felleisen 2002; Wadler and Findler 2009]. Blame tracking is especially useful during the software development process, but in the context of IFC, one may not want blame to be observable in a production system as it could reveal information. This can be handled by causing the program to diverge whenever blame is detected, possibly sending a private error message to the software developer.

Next we return to the flip example to see how gradual IFC prevents implicit flows. Suppose that we change the parameter type of the λ from $\text{Bool}_{\text{high}}$ to Bool_{\star} . The return type remains Bool_{low} , to conform with the signature of `publish`. Line 1 thus becomes:

```
let flip :  $\text{Bool}_{\star}$   $\rightarrow$   $\text{Bool}_{\text{low}}$  =  $\lambda$  b :  $\text{Bool}_{\star}$ . if b then false else true in
```

This change makes the program well-typed in $\lambda_{\text{IFC}}^{\star}$. The IFC enforcement of the implicit flow is deferred until runtime because the branch condition now has type Bool_{\star} , with an unknown security level. Next we discuss the runtime behavior of this program.

The result of cast insertion on this program is the following λ_{IFC}^c term:

```
1 let flip =  $\lambda$  b. ((if $\star$  b then (false  $\langle$ low! $\rangle$ ) else (true  $\langle$ low! $\rangle$ ))  $\langle$ low? $\rangle^p$ ) in
2 let input = user-input () in
3 let result = flip (input  $\langle$ high! $\rangle$ ) in
4   publish result
```

where the `if` is changed to `if \star` because the condition expression has static security level \star . The type checking rule for `if \star` requires the two branches to have security level \star and the security level of the `if \star` as a whole is also \star . If we run the program with `true` or `false` as input, the λ_{IFC}^c term reduces to `blame p` with either input, thus capturing the illegal implicit flow. Consider running this program with input `true`. First, the λ is bound to `flip` and the input `true` is bound to the input variable. We then inject the input, producing the value (`true $_{\text{high}}$ \langle high! \rangle`). We call `flip` and the `if \star` branches on this boolean, evaluating the “then” branch to the result (`false $_{\text{low}}$ \langle low! \rangle`) and then, to protect against implicit flows, the `if \star` upgrades the result to `high` to match the runtime security level of the condition (`true $_{\text{high}}$ \langle high! \rangle`), producing (`false $_{\text{low}}$ \langle high! \rangle`). The subtype coercion \uparrow sends the security level of `false` from `low` to `high`. The last step in the body of `flip` is to apply the coercion `low? \rangle^p` to the value (`false $_{\text{low}}$ \langle high! \rangle`), which errors because `high` is greater than `low`.

2.2 Implicit Flow, NSU Checks, Unknown Security, and the Gradual Guarantee

The tension between information-flow security and the gradual guarantee arises from an interaction between implicit flows and the use of no-sensitive-upgrade checks to guard writes to mutable references. In brief, when \star is allowed as a runtime security label, some NSU checks have to conservatively trigger an error to preserve noninterference, even though no error would occur if the label was instead `high`. But the gradual guarantee says that if a program with a precise annotation runs without error, it should also run without error when that annotation is changed to \star .

In preparation to discuss this scenario in more detail, we review the no-sensitive-upgrade technique [Austin and Flanagan 2009] that protects against illegal implicit flows through writes to mutable references. We then show how allowing \star as a runtime security label leads to a situation where a language designer is forced to choose between noninterference and the gradual guarantee. Finally, we show how this problem is resolved in the $\lambda_{\text{IFC}}^{\star}$ language by walking back the choice of allowing \star as a runtime security label.

The main idea of no-sensitive-upgrades is to prevent data leaks through the mutable references by terminating execution whenever the program attempts to modify a low-security heap cell in a high-security execution context. In $\lambda_{\text{IFC}}^{\star}$, NSU checks happen at runtime when type information is insufficient to statically decide whether a heap-modifying operation is secure or not. Consider the following well-typed program in $\lambda_{\text{IFC}}^{\star}$:

```
1 let input :  $\text{Bool}_{\star}$  = user-input () in
2 let a      = ref low true in
3 let _     = if input then a := false else a := true in
4   publish (! a)
```

The assignments to a in the two branches try to write different low-security booleans into the cell at the address in a , depending on a branch condition whose security level is statically unknown. However, at runtime the user-input function returns a high-security boolean, so the branch condition is actually high security, and if the writes were successful, the program would leak information via an implicit flow. Fortunately, if we run this program, it reduces to blame regardless of the input. The way NSU checking works in λ_{IFC}^* is that a security level is associated with the current program counter. At the point of every write that requires an NSU, the system projects the program counter's security level to the level of the memory location, making sure that the later is at least as high as the former. In the above example, the NSU check fails because the program counter's security is **high** but the write is to a **low** security location.

In GSL_{Ref} [Toro et al. 2018], the dynamic enforcement of IFC through the heap is also based on NSU checks. Consider the following pair of programs adapted from Section 6.3 of their paper. The program on the left is derived from the program on the right by replacing some of the **high** annotations with the unknown label \star . Both variants of the program type check but the more precise variant runs to completion while the less precise variant triggers an error, thus violating the gradual guarantee. Let us examine their runtime behavior in further detail.

Left: less precise, more dynamic

```
let x = user-input () in
let y = ref Bool $\star$  true $\star$  in
  if x then (y := falsehigh) else ()
```

Right: more precise, more static

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

The program on the **right** runs without error in GSL_{Ref} because, at the assignment on line 3, variable y references a memory cell of **high** security and the PC's security level is also **high**, so the assignment is allowed.

In contrast, when the program on the **left** is run with input $\text{true}_{\text{high}}$, the assignment is conservatively rejected by the NSU check. This is because GSL_{Ref} considers \star as corresponding to the interval $[\text{low}, \text{high}]$, and the lower bound of this interval is not greater than or equal to the **high** PC label. So we see that this more precise program (**right**) runs successfully while the less precise one (**left**) errors in GSL_{Ref} .

In λ_{IFC}^* , the \star security label can be used in type annotations, as one would expect of a gradually-typed language, but \star is not allowed as a runtime security label and therefore also not allowed as a label on program literals and other introduction forms. The following adapts the above examples from GSL_{Ref} to λ_{IFC}^* . The fully static variant on the right is nearly identical to its GSL_{Ref} counterpart. To obtain the less-precise program on the left we change the type annotation on variable y to model the similar loss of precision in the GSL_{Ref} counterpart. We do not change the labels on the ref or true to \star because that is not allowed in λ_{IFC}^* as we just mentioned.

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

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

Branching on high-security input and assigning to a high-security memory location should be allowed. Indeed, both variants reduce to the unit value regardless of the input, thereby not violating the gradual guarantee. The fully annotated version (**right**) evaluates to unit without any overhead from runtime checking.

In the less-precise program (**left**), the type annotation $(\text{Ref Bool}\star)\star$ replaces $(\text{Ref Bool}_{\text{high}})_{\text{high}}$, meaning that both the security level of the memory and the security of the reference itself are statically unknown and should be checked at runtime. When executing the program, at the assignment on line 3, an NSU check happens and the assignment to high-security memory under high PC is allowed. As a result, the less-precise program also evaluates successfully to unit.

One might worry that the less precise program has a heavy annotation burden. However, as we mentioned, the default security label is `low` so the programmer does not have to label constants in λ_{IFC}^* . So we can remove the labels on constants to obtain the following program, which also reduces successfully to unit:

```

1 let x = user-input () in
2 let y : (Ref Bool★)★ = ref high true in
3   if x then (y := false) else ()

```

The low-security `true` is classified as high security during reference creation because the security level of the cell is `high` (line 2). Similarly, during assignment the `false` is also classified as high security because the security level of the cell (line 3). Assigning to a high-security memory cell is allowed under a high PC by the NSU check, so the program evaluates successfully to unit.

One might think that requiring the programmer to annotate the reference creation with `high` (line 2) is still a burden and that GSL_{Ref} is better in this regard. However, while GSL_{Ref} allows the unannotated version of this program to compile, it errors during program execution. We believe that it is better to require the programmer to annotate references during program development than to have the programs compile but then fail during program execution, perhaps only detected after the program is deployed.

2.3 Type-Based Reasoning in λ_{IFC}^*

Type-based reasoning in gradual IFC languages arises from two language design choices: vigilance and type-guided classification. Vigilance gives us type-based reasoning for explicit flows, while type-guided classification provides type-based reasoning about implicit flows. We show in this section that λ_{IFC}^* is both vigilant and performs typed-guided classification, so it enables type-based reasoning in the sense of Toro et al. [2018].

2.3.1 λ_{IFC}^* is Vigilant. A language with casts is *vigilant* if it checks whether all the casts that are applied to the same value are consistent with each other, and triggers an error if they are not.

Toro et al. [2018] present the following example to demonstrate how vigilance is needed for type-based reasoning and free theorems in the sense of Wadler [1989]. The example involves casts from `low` (line 3, the label annotation on 5_{low}) to `high` (line 1, the type annotation Int_{high} in the signature of `mix`) and then back to `low` via the unknown security level \star (line 2, the nested type annotations Int_{\star} and Int_{low}):

```

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

```

The type signature of `mix` should guarantee the free theorem that either (1) the result of `mix`, which is low security, never depends on the high-security `priv` argument or (2) `mix` produces a runtime error. In this case, the output of `mix` does depend on `priv` via an implicit flow, so the free theorem says that an error should be triggered at runtime. Let us focus on the three casts applied to 5_{low} , where the first cast sends the security level from `low` to `high` because of the type annotation on line 1, the second cast is an injection due to the first type annotation on line 2, and the third cast is a projection to `low` due to the second type annotation on line 2:

$$5_{\text{low}} \left\langle \text{low} \Rightarrow \text{high} \right\rangle \left\langle \text{high} \Rightarrow \star \right\rangle \left\langle \star \Rightarrow \text{low} \right\rangle$$

In λ_{IFC}^* , these casts produce the sequence of coercions: $5_{\text{low}} \left\langle \uparrow ; \text{high}! ; \text{low}?? \right\rangle$, which trigger an error when `high` collides with `low`, blaming label p .

Similarly, the interval refinement mechanism of GSL_{Ref} detects the conflict between the intermediate cast to **high** and the final cast to **low**. Surprisingly, in GLIO and in systems prior to GSL_{Ref} [Disney and Flanagan 2011; Fennell and Thiemann 2013], the program runs successfully and produces 1_{low} because they are forgetful [Greenberg 2014] regarding the intermediate cast of 5_{low} to **high** and only check that 5_{low} can be cast to low^2 .

2.3.2 λ_{IFC}^* *Performs Typed-Guided Classification.* A gradual IFC language employs *typed-guided classification* if the security-level of a value can be changed when the value flows through a cast.

The following example from Section 2 of Toro et al. [2018] demonstrates how type-guided classification interacts with implicit flow and type-based reasoning. Type-based reasoning tells us that the `smix` function below should either fail or return a value that does not depend on its high-security parameter `priv`. However, `smix` calls `mix` and there is an implicit flow from `priv` to the result value, so this program should fail.

```

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

```

In λ_{IFC}^* the program produces an error as type-based reasoning suggests. Security coercions in λ_{IFC}^c classify values, so when 5_{low} is passed into `smix` and then `mix`, it is wrapped in a coercion: $5_{\text{low}} \langle \uparrow; \text{high}! \rangle$ and is classified as high-security. Consequently, the `if` reduces to its then-branch protected with **high**. This implicit flow affects the result value of the then-branch, by (1) inserting a subtype coercion before the injection and (2) promoting the source of the injection to **high** to preserve types. So the result of the `if` is $1_{\text{low}} \langle \uparrow; \text{high}! \rangle$. The injection, whose source is **high**, collides with a projection to **low** (to match the Int_{low} return type of `mix`), causing the program to error as expected, blaming the projection.

3 A COERCION CALCULUS FOR SECURITY LABELS

In this section we present a coercion calculus on security labels. We show that we can use coercion composition and stamping to represent explicit flows and implicit flows respectively (Section 3.1). We define how these coercions act on security labels by defining a language of label expressions whose meaning is defined by a reduction relation (Section 3.2). Finally, we explore meta-theoretic properties of this coercion calculus (Section 3.3), establishing the intuition that the security coercion calculus can be used to enforce gradual IFC, while satisfying the gradual guarantee.

As we have seen in Section 2, gradual information flows can be modeled as casts. The cast sequence $\text{high} \Rightarrow \star \Rightarrow \text{low}$ should be statically accepted but dynamically rejected, while the sequence $\text{low} \Rightarrow \star \Rightarrow \text{high}$ should be statically and dynamically accepted, promoting the security of data to high. Such sequences of casts can be arbitrarily long (for example, $\text{low} \Rightarrow \text{high} \Rightarrow \star \Rightarrow \star \Rightarrow \text{low}$), which motivates us to model the casts on security labels as coercions. In λ_{IFC}^* , the source security label of a coercion sequence comes from literals, while the sink is whatever security level that the observer has: for example, the `publish` function of Section 2 is of **low**. Coercions can be easily sequenced and composed. Checking information flow at runtime is accomplished by reducing coercion sequences to their normal forms.

There are two noteworthy benefits of the coercion representation for IFC. First, coercions can be used to represent NSU checking while satisfying the gradual guarantee. In brief, whenever a memory location is written to, the current PC is coerced to the security level of that location. We are going to formally introduce label expressions as our representation for PC in Section 3.2 and

²Those systems still satisfy noninterference, because the labels *on values* track their security similar to fully dynamic IFC. It is just that the security labels on *type annotations* will not trigger errors at runtime, even if they are inconsistent.

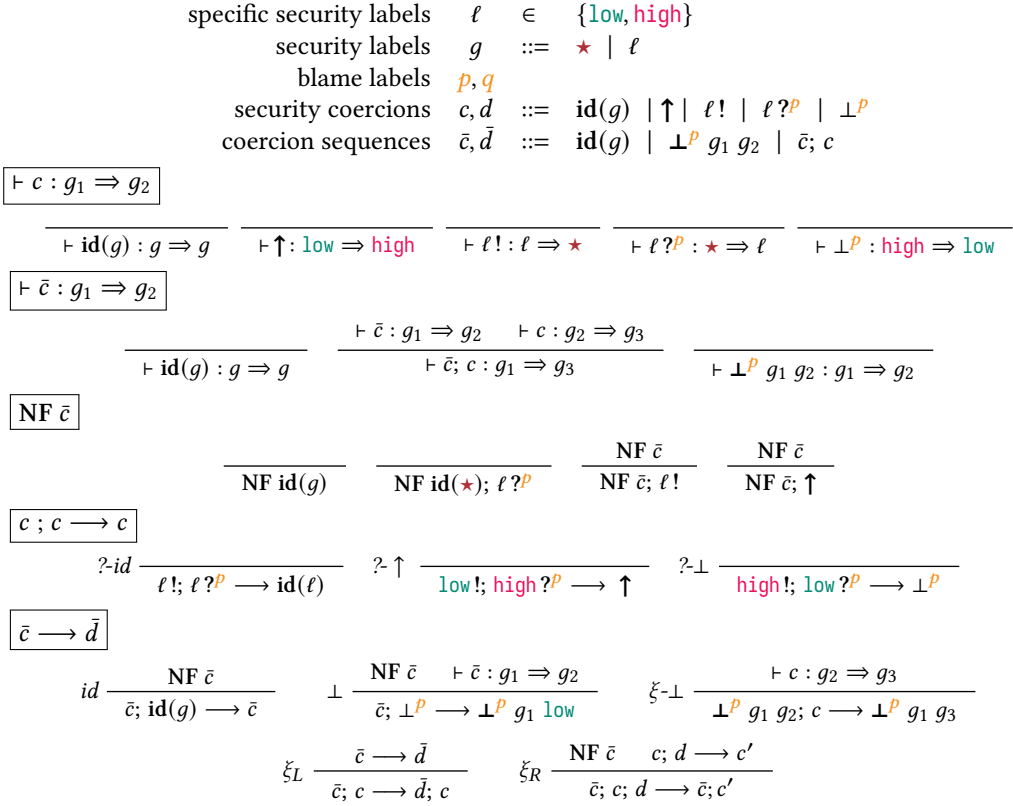


Fig. 1. Syntax, typing, normal forms, and semantics of security coercions and coercion sequences

discuss NSU in detail in Section 5.1.2. Second, the coercion representation benefits mechanization because it enables modular reasoning. The main simulation lemma (Lemma 13) depends on the simulation results of coercion sequences and label expressions, which are stated as separate lemmas and reasoned independently in our Agda code.

The syntax and typing for security coercions and coercion sequences is defined in Figure 1. A security label is either `low`, `high`, or statically unknown (\star). There are five security coercions: identity ($\text{id}(g)$), subtype (\uparrow), injection ($\ell!$), and projection ($\ell?^p$), and blame (\perp^p). Projection, which corresponds to the notion of a runtime check, is the only one responsible for blame, so it carries a blame label p . A coercion sequence \bar{c} starts with either success $\text{id}(g)$ or failure ($\perp^p g_1 g_2$). Each coercion has a source and target type $g_1 \Rightarrow g_2$. The $\text{id}(g)$ casts the label g to itself; \uparrow promotes security from `low` to `high`; injection casts to \star from a specific label ℓ and projection does the opposite. Appending a single coercion to a coercion sequence makes the target security label that of the single coercion.

Information flow is enforced in the reduction semantics of security coercions, shown in Figure 1. Injection followed by projection to the same label collapses to the identity ($?-\text{id}$). Flowing from `low` to `high` is allowed, so an injection from `low` followed by a projection to `high` collapses into the \uparrow coercion ($?-\uparrow$). An information flow from `high` to `low` is prohibited, so an injection to `high` followed by a projection to `low` triggers an error that blames the projection ($?-\perp$). The predicate NF that specifies the normal forms of coercion sequences. The reduction rules for coercion sequences are also defined in Figure 1. Appending $\text{id}(g)$ onto a coercion sequence reduces to the

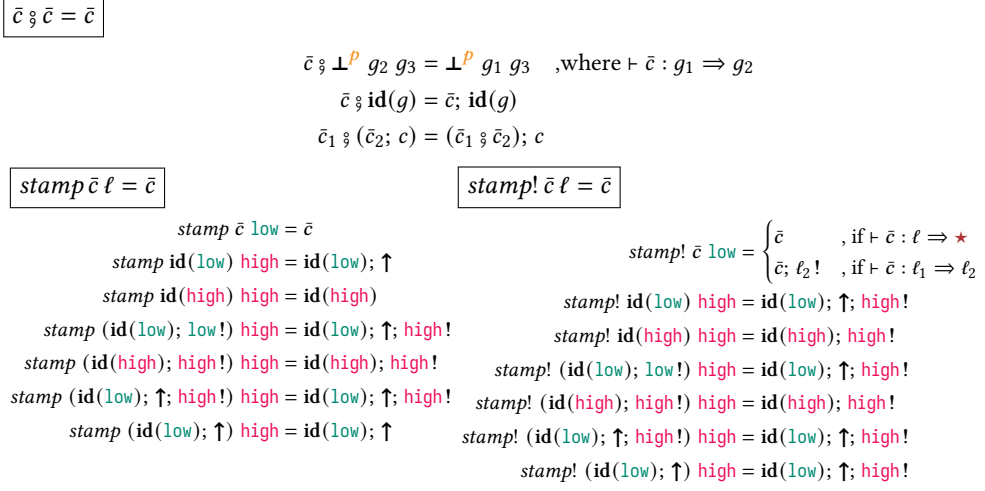


Fig. 2. Composing and stamping coercions

that sequence (*id*). The failure coercions annihilate the other coercions in the sequence (\perp and $\xi\text{-}\perp$). We choose the evaluation order in a coercion sequence to be from left to right (ξ_L and ξ_R), because that corresponds to the direction of information flow from source to sink: in the example above, $\text{low} \Rightarrow \text{high} \Rightarrow \star \Rightarrow \star \Rightarrow \text{low}$, we validate that low can flow to high before we check the flow from high through \star to low .

3.1 Monitoring Explicit and Implicit Flows

We model explicit flow using security coercions. We can compose two coercion sequences ($\bar{c} \circledast \bar{d}$), where $\bar{c} : g_1 \Rightarrow g_2$ and $\bar{d} : g_2 \Rightarrow g_3$, to form a flow from g_1 to g_3 , which is defined in Figure 2.

The stamping operation captures the intuition of an implicit flow from the security level ℓ' to a coercion sequence \bar{c} . We define the stamping operation in Figure 2 as two functions, $\text{stamp}(\bar{c}, \ell)$ and $\text{stamp!}(\bar{c}, \ell)$. Both function require \bar{c} to be in normal form and that its source label is not \star . The stamp! operator promotes the security of the coercion \bar{c} to be at least ℓ and then injects the coercion if necessary, while stamp only promotes the security but does not inject. These stamping operations satisfy the gradual guarantee, because when stamping on a more precise coercion sequence and a less precise coercion sequence, stamping preserves the precision relation (Section 3.3.2) between them (Lemma 9). The stamping operations of coercion sequences are used in the stamping operations of (1) label expressions, which are our representation of PC and (2) values in the cast calculus. Those three types of stamping together formalize the notion of implicit flow in λ_{IFC}^* .

3.2 Security Label Expressions

In this section we introduce security label expressions, which we use to model the security level of the PC. Security label expressions are crucial for implementing NSU checking in a way that satisfies the gradual guarantee.

A label expression is either (1) a specific security label, (2) blame (to signify an error), or (3) a coercion applied to a label expression (Figure 3). A label expression is in normal form (NF) if it is either (1) a specific security label or (2) an irreducible coercion applied to a specific security label. (A coercion is irreducible if it is a non-identity coercion in normal form). *PC* ranges over label expressions in normal form. The reduction relation for label expressions steps a label expression

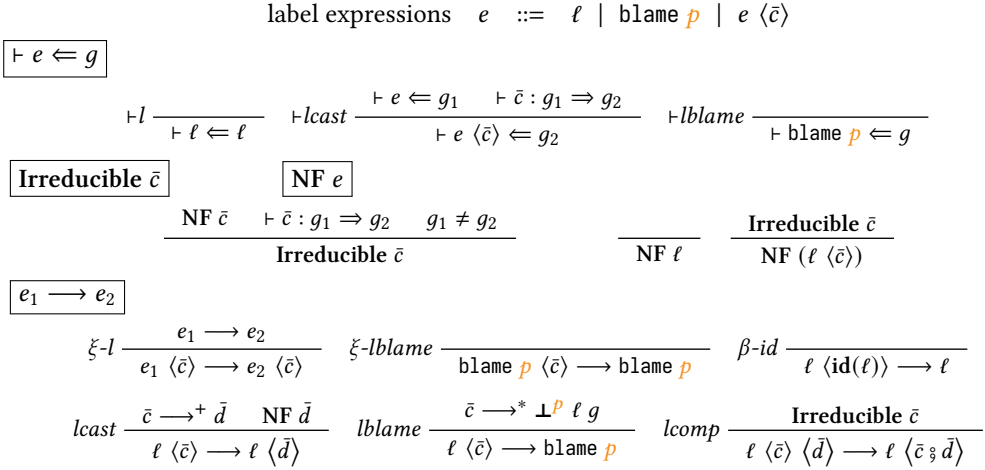


Fig. 3. Syntax, typing, normal forms, and semantics of label expressions

towards its normal form. The idea is that given a label expression of the form $e \langle \bar{d} \rangle$, we first reduce e to normal form and then apply the coercion \bar{d} . For example, if e reduces to a label wrapped in coercion $\ell \langle \bar{c} \rangle$, then the *lcomp* rule says to reduce by composing the two coercions, producing $\ell \langle \bar{c} \bar{d} \rangle$. Furthermore, in a label expression of the form $e \langle \bar{d} \rangle$, the coercion \bar{d} may also need to be reduced, which is accomplished by the *lcast* rule that refers to the reduction relation for coercion sequences (Figure 1). If the coercion reduces to an identity, then the coercion application goes away (*β -id*), whereas if the coercion reduces to a failure, then the label expression reduces to blame (*lblame*).

The stamping and security level operators for label expressions are defined in Figure 13 of the Appendix. They both require their input to be in normal form, which can be either (1) a specific security label ℓ , or (2) a label wrapped with an irreducible coercion sequence $\ell \langle \bar{c} \rangle$. For (1), stamping **low** with **high** results in **low** (\uparrow), otherwise the label expression remains unchanged; for (2), we directly stamp the coercion sequence using *stamp* for coercion sequences defined in Figure 2. The definition of *stamp!* is analogous, except that it turns to the *stamp!* operator of coercion sequences. The security level operator $|-|$ is defined such that (1) a specific security label indicates the security level for itself and (2) the security of the coercion sequence $|\bar{c}|$ records the security level for $\ell \langle \bar{c} \rangle$.

We describe in Section 5.1 how label expressions are used to implement NSU checks, which enforce the heap policy for write operations.

3.3 Properties of Coercion Calculus on Labels

3.3.1 Tracking Information Flow via Composition and Stamping. We show that composing coercions tracks explicit flows, while stamping tracks implicit flows. We first define of the security level of a coercion sequence with the $|-|$ operator³:

DEFINITION 1 (SECURITY LEVEL OF A COERCION). *Given a coercion sequence \bar{c} in normal form with type $\vdash \bar{c} : \ell \Rightarrow g$, then its security is given by the following $|-|$ operator:*

$$|\text{id}(\ell)| = \ell \quad |\text{id}(\ell); \ell!| = \ell \quad |\text{id}(\text{low}); \uparrow; \text{high}| = \text{high} \quad |\text{id}(\text{low}); \uparrow| = \text{high}$$

³In a larger lattice, the subtype coercion \uparrow would be parameterized by two labels: $\vdash (\uparrow_{\ell_1}^{\ell_2}) : \ell_1 \Rightarrow \ell_2$ where $\ell_1 < \ell_2$. The security would be the greater one: $|\uparrow_{\ell_1}^{\ell_2}| = \ell_2$. Here the lattice only consists of **low** and **high**, so $|\uparrow| = |\uparrow_{\text{low}}^{\text{high}}| = \text{high}$.

We reason about explicit flows first. If we compose one coercion sequence with another and then reduce the result to normal form, the security of the resulting coercion sequence should be greater than or equal to that of the first sequence.

LEMMA 2 (COMPOSITION MODELS EXPLICIT FLOW).

If $\text{NF } \bar{c}$ and $\bar{c} \bar{d} \longrightarrow^* \bar{c}'$ and $\text{NF } \bar{c}'$, then $|\bar{c}| \preceq |\bar{c}'|$.

Next we show that stamping models implicit flow correctly, promoting the security of the stamped coercion by joining it with the stamped label:

LEMMA 3 (STAMPING MODELS IMPLICIT FLOW).

If $\text{NF } \bar{c}$, then $|\text{stamp } \bar{c} \ell| = |\bar{c}| \vee \ell$ and $|\text{stamp! } \bar{c} \ell| = |\bar{c}| \vee \ell$.

3.3.2 Simulation Between More and Less Precise Coercion Sequences. Our goal is to prove the gradual guarantee for λ_{IFC}^* . The proof depends on a simulation lemma between more and less precise terms of the cast calculus λ_{IFC}^c . We use coercion sequences as the IFC monitor in λ_{IFC}^c . Reducing a coercion sequence can result in a blame which errors the program. So we would like to prove that the simulation lemma holds for the coercion calculus on security labels. The precision relation on security coercions is defined in Figure 20 of the Appendix. The precision relation between two coercion sequences \bar{c}, \bar{d} takes the form $\vdash \bar{c} \sqsubseteq \bar{d}$. Recall that the gradual guarantee states that replacing type annotations with \star (decreasing type precision) should result in the same value for a correctly running program while adding annotations (increasing type precision) may trigger more runtime errors. The precision relation is a syntactical characterization of the runtime behaviors of programs of different type precision. We explain the intuition with two examples.

EXAMPLE 4. Consider the following two programs that are related by precision because the first one has a \star annotation where the second one has **high**.

$$\text{true}_{\text{low}} : \text{Bool}_{\star} : \text{Bool}_{\star} \quad \text{and} \quad \text{true}_{\text{low}} : \text{Bool}_{\text{high}} : \text{Bool}_{\star}$$

At runtime, the less precise program on the left produces value ($\text{true } \langle \text{id}(\text{low}); \text{low}! \rangle$) and the more precise program on the right produces ($\text{true } \langle \text{id}(\text{low}); \uparrow; \text{high}! \rangle$). The trues are straightforwardly related; we need to show the two coercion sequences are also related:

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

Starting at the beginning of the two sequences, we have $\text{id}(\text{low}) \sqsubseteq \text{id}(\text{low})$ because \sqsubseteq is reflexive. Next we have $\text{low}! \sqsubseteq \uparrow$, which makes sense because the source and targets of the two coercions are related by precision: $\text{low} \sqsubseteq \text{low}$ and $\star \sqsubseteq \text{high}$. Finally, the coercion $\text{high}!$ can be added to the end of the more precise sequence because both its source and target type are more precise than the target of the left-hand sequence. That is, $\star \sqsubseteq \text{high}$ and $\star \sqsubseteq \star$. The injections at the ends of the two sequences, $\text{low}!$ and $\text{high}!$, cannot be directly related via precision because $\text{low} \not\sqsubseteq \text{high}$. Instead, $\text{low}!$ is related with \uparrow . This underlines the indispensability of explicit subtype coercion \uparrow for the purposes of proving the gradual guarantee.

EXAMPLE 5. Next we consider an example where the less precise program produces a value but the more precise program encounters an error. This situation is allowed by the gradual guarantee, but the opposite one is not. We extend the example with a cast to **low** on the more precise side.

$$\text{true}_{\text{low}} : \text{Bool}_{\star} : \text{Bool}_{\star} : \text{Bool}_{\star} \quad \text{and} \quad \text{true}_{\text{low}} : \text{Bool}_{\text{high}} : \text{Bool}_{\star} : \text{Bool}_{\text{low}}$$

The first program again produces value ($\text{true } \langle \text{id}(\text{low}); \text{low}! \rangle$). The second, on the other hand, reduces to ($\text{true } \langle \text{id}(\text{low}); \uparrow; \text{high}!; \text{low}^? \rangle$) \longrightarrow^* ($\text{true } \langle \perp^P \text{ low low} \rangle$), because of the contradicting annotations **high** and **low** (note that **high** is in the middle of the sequence and both the

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

$V \langle c \rangle \longrightarrow M$

$\text{cast} \frac{\bar{c} \longrightarrow^+ \bar{d} \quad \text{NF } \bar{d}}{V_r \langle c_r, \bar{c} \rangle \longrightarrow V_r \langle c_r, \bar{d} \rangle}$	$\text{cast-blame} \frac{\bar{c} \longrightarrow^* \perp^p g_1 g_2}{V_r \langle c_r, \bar{c} \rangle \longrightarrow \text{blame } p}$
$\text{cast-id} \frac{}{V_r \langle \text{id}(\iota), \text{id}(g) \rangle \longrightarrow V_r}$	$\text{cast-comp} \frac{\text{Irreducible } c}{V_r \langle c \rangle \langle d \rangle \longrightarrow V_r \langle c \ ; \ d \rangle}$

Fig. 4. Syntax and semantics of coercions on values.

source and target labels of the blame coercion are **low** so that types are preserved). The failure is then propagated out and the term further reduces to `blame p` . The precision of coercion sequences relates \perp on the right-hand side to any coercion sequence on the left so long as the respective source and target types are related via precision, in this case $\text{low} \sqsubseteq \text{low}$ and $\star \sqsubseteq \text{low}$.

Consider the security levels (Definition 1) of both sides of Example 4, which are **low** on the less precise side and **high** on the more precise side. We observe that λ_{IFC}^* programs related by precision may produce values of different security: a less precise value may have lower security than a more precise value. Indeed, we prove the following for coercions in normal form:

LEMMA 6 (SECURITY IS MONOTONIC WITH RESPECT TO PRECISION). *Suppose NF \bar{c} and NF \bar{d} . If $\vdash \bar{c} \sqsubseteq \bar{d}$, then $|\bar{c}| \preceq |\bar{d}|$.*

Next we prove a catch-up lemma for coercion sequences, where the less precise side catches up with a more precise sequence that is in normal form. The proof is by casing on NF \bar{d} first and then performing induction on the precision relation in each case.

LEMMA 7 (CATCHING UP TO A MORE PRECISE COERCION SEQUENCE). *If NF \bar{d} and $\vdash \bar{c} \sqsubseteq \bar{d}$, there exists \bar{c}' such that $\bar{c} \longrightarrow^* \bar{c}'$ and $\vdash \bar{c}' \sqsubseteq \bar{d}$.*

Using Lemma 7, we then prove the following simulation lemma for coercion sequences:

LEMMA 8 (SIMULATION BETWEEN RELATED COERCION SEQUENCES). *If $\vdash \bar{c} \sqsubseteq \bar{d}$ and $\bar{d} \longrightarrow \bar{d}'$, there exists \bar{c}' such that $\bar{c} \longrightarrow^* \bar{c}'$ and $\vdash \bar{c}' \sqsubseteq \bar{d}'$.*

We also prove that stamping on coercion sequences preserves precision:

LEMMA 9 (STAMPING PRESERVES PRECISION OF COERCION SEQUENCES). *If $\vdash \bar{c} \sqsubseteq \bar{d}$, then $\vdash \text{stamp } \bar{c} \ \ell \sqsubseteq \text{stamp } \bar{d} \ \ell$ and $\vdash \text{stamp! } \bar{c} \ \ell_1 \sqsubseteq \text{stamp! } \bar{d} \ \ell_2$ and $\vdash \text{stamp! } \bar{c} \ \ell_1 \sqsubseteq \text{stamp } \bar{d} \ \ell_2$ if $\ell_1 \preceq \ell_2$.*

4 A SECURITY COERCION CALCULUS ON VALUES

In this section we define a second coercion calculus whose purpose is to cast a program value from one type to another type. We use this coercion calculus as the representation of casts in the intermediate language λ_{IFC}^c . These coercions on values make use of the coercions on security labels that we defined in Section 3 because the types in λ_{IFC}^c are annotated with security labels, as is usual for a static and gradually-typed IFC languages.

We begin with the definition of types in Figure 4, which is standard for gradual security type systems: each type has a security label ascription on it, which is either a specific label ℓ or \star . Function types have one extra label g^c , which is a static approximation of the security level of the PC while executing the function body. In a reference type $(\text{Ref } T_{\hat{g}})_g$, the label \hat{g} of the referenced type also doubles as the security level of the memory cell.

The syntax and semantics for coercions on values is defined in Figure 4. Each coercion c consists of a raw coercion c_r that casts the type of the value and the label coercion \bar{c} that casts the security label of the value. There are three kinds of raw coercions: identity coercions $\text{id}(g)$, coercions between reference types $(\text{Ref } c \ d)$, and coercions between function types $(\bar{d}, c \rightarrow d)$. In the coercion on functions, the \bar{d} casts the PC of the function. (The syntax for values is not defined until the next section, so here we remark that V ranges over values, which can either be a raw value V_r (constant, address, or λ) or an irreducible coercion applied to a raw value: $V_r \langle c \rangle$, where there is no M such that $V_r \langle c \rangle \rightarrow M$. The definition of irreducible coercion is in Figure 15 of the Appendix.)

The reduction rules in Figure 4 apply a coercion to a value, yielding a value or triggering blame. The *cast* rule normalizes the coercion \bar{c} on the security label. If it normalizes to a failure coercion, the rule *cast-blame* triggers blame. We reduce identity coercions using rule *cast-id*. Finally, if the value is wrapped with an irreducible coercion, we compose the coercion with the coercion being applied (rule *cast-comp*). The composition operator $- \circ -$ is also defined in Figure 15 of the Appendix; the intuition of the composition operator is that $V_r \langle c \rangle \langle d \rangle$ must be contextual equivalent to $V_r \langle c \circ d \rangle$. There are no reduction rules specific to reference coercions $\text{Ref } c \ d$ or function coercions $(\bar{d}, c \rightarrow d)$ because they are irreducible coercions that wrap a value. Their action occurs when the value is used in an elimination form such as in a function call or a read or write to the reference, which we explain in the next section.

5 THE CAST CALCULUS λ_{IFC}^c

In this section we define the cast calculus λ_{IFC}^c (§ 5.1), prove that λ_{IFC}^c is type-safe (§ 5.2.1), and prove the main simulation lemma (§ 5.2.2) that is needed for the proof of the gradual guarantee. We conclude this section with a proof that λ_{IFC}^c satisfies noninterference (§ 5.2.3).

5.1 Syntax, Typing, and Operational Semantics of λ_{IFC}^c

5.1.1 Syntax and Typing of λ_{IFC}^c . As usual, the cast calculus λ_{IFC}^c is a statically-typed language that includes an explicit term for casts, written $M \langle c \rangle$, where M is a term and c is a coercion to be applied to the value of M . Furthermore, many of the operators in λ_{IFC}^c have two variants, a “static” one for when the pertinent security label is statically known and the “dynamic” one for when the security label is statically unknown. The operational semantics of the “dynamic” variants involve runtime checking. The syntax and typing rules for λ_{IFC}^c are shown in Figure 5 (excerpt of typing rules, full version in Figure 12 in the Appendix) and described in the following paragraphs.

A value is a raw value (constant, address or λ) or an irreducible coercion applied to a raw value.

The typing rules are syntax-directed. The typing judgment is of the form $\Gamma; \Sigma; g; \ell \vdash M \Leftarrow A$, which says that we are type-checking λ_{IFC}^c term M against the expected type A , where Γ is the typing context, Σ is the heap typing context, g is the security label that PC is typed at, and ℓ is the security level of PC. Both Σ and the security level ℓ play a role during runtime. The security level of the PC is constrained in rule $\vdash \text{prot}$, which is for the protection term that arises during reduction (we are going to discuss this rule momentarily). In the premises for sub-terms that do not immediately reduce, such as the body of a λ and the branches of an *if*, we universally quantify the security level (as in $\forall \ell$), which helps us prove that compilation preserves types (Theorem 15). The heap context Σ is mostly standard: looking up $\Sigma(\hat{\ell}, n)$, where n is the index in part of the heap

variables	x, y, z	
constants	k	$\in \{\text{unit}, \text{true}, \text{false}\}$
terms	L, M, N	$::= x \mid \$k \mid \text{addr } n \mid \lambda x. N \mid \text{app } L M A B \ell \mid \text{app}^\star L M A T$ $\mid \text{if } L A \ell M N \mid \text{if}^\star L T M N \mid \text{let } x=M:A \text{ in } N$ $\mid \text{ref } \ell M \mid \text{ref}^? P \ell M \mid !M A \ell \mid !^\star M T$ $\mid \text{assign } L M T \hat{\ell} \ell \mid \text{assign}^? P L M T \hat{g}$ $\mid \text{prot } P C \ell M A \mid M \langle c \rangle \mid \text{blame } P$
raw values	V_r, W_r	$::= \$k \mid \text{addr } n \mid \lambda x. N$
values	V, W	$::= V_r \mid V_r \langle c \rangle$, where Irreducible c

$\Gamma; \Sigma; g; \ell \vdash M \Leftarrow A$

$$\begin{array}{c}
 \frac{\Gamma; \Sigma; \ell'; \ell'' \vdash L \Leftarrow (A \xrightarrow{\ell' \vee \ell} B)_\ell \quad \Gamma; \Sigma; \ell'; \ell'' \vdash M \Leftarrow A \quad C = \text{stamp } B \ell}{\vdash \text{app} \frac{}{\Gamma; \Sigma; \ell'; \ell'' \vdash \text{app } L M A B \ell \Leftarrow C}} \quad \frac{\Gamma; \Sigma; g; \ell \vdash L \Leftarrow (A \xrightarrow{\star} (T_\star))_\star \quad \Gamma; \Sigma; g; \ell \vdash M \Leftarrow A}{\vdash \text{app}^\star \frac{}{\Gamma; \Sigma; g; \ell \vdash \text{app}^\star L M A T \Leftarrow T_\star}} \\
 \\
 \frac{\Gamma; \Sigma; \ell'; \ell'' \vdash L \Leftarrow (\text{Ref } T_{\hat{\ell}})_\ell \quad \Gamma; \Sigma; \ell'; \ell'' \vdash M \Leftarrow T_{\hat{\ell}} \quad \ell' \vee \ell \preceq \hat{\ell}}{\vdash \text{assign} \frac{}{\Gamma; \Sigma; \ell'; \ell'' \vdash \text{assign } L M T \hat{\ell} \ell \Leftarrow \text{Unit}_{\text{low}}}} \quad \frac{\Gamma; \Sigma; g; \ell \vdash L \Leftarrow (\text{Ref } T_{\hat{g}})_\star \quad \Gamma; \Sigma; g; \ell \vdash M \Leftarrow T_{\hat{g}}}{\vdash \text{assign}^? \frac{}{\Gamma; \Sigma; g; \ell \vdash \text{assign}^? P L M T \hat{g} \Leftarrow \text{Unit}_{\text{low}}}} \\
 \\
 \frac{\Gamma; \Sigma; g'; |PC| \vdash M \Leftarrow A \quad \vdash PC \Leftarrow g' \quad \ell' \vee \ell \preceq |PC| \quad B = \text{stamp } A \ell}{\vdash \text{prot} \frac{}{\Gamma; \Sigma; g; \ell' \vdash \text{prot } P C \ell M A \Leftarrow B}}
 \end{array}$$

Fig. 5. Syntax and selected typing rules of the cast calculus λ_{IFC}^c . The side condition that enforces the heap policy statically during assignment is **highlighted**

with security $\hat{\ell}$, gives us a raw type. Each memory cell is associated with a specific security label $\hat{\ell}$, which is specified by the programmer when that cell is created.

As the typing rules always stay in checking mode, constants, addresses, and λ s do not carry any label. The security of these raw values is in their types: for example, $\text{addr } n \Leftarrow (\text{Ref } T_{\hat{\ell}})_\ell$ says that the security of the address n itself is ℓ and it points to a memory cell labeled $\hat{\ell}$.

The typing rules for the static variants are similar to the typing rules in a static security type system. For example, in the static function application rule $\vdash \text{app}$, both top-level labels on the function type as well as the security label that the current PC is typed at are static and the rule mirrors one in a static system. On the other hand, in the dynamic version of application rule $\vdash \text{app}^\star$, both top-level labels as well as the label of the co-domain type are \star and PC is allowed to be typed at \star , indicating the presence of injections. As another example, in the static version of memory assignment rule $\vdash \text{assign}$, all labels to perform the heap policy check, including the security of the memory address itself (ℓ), the security of the memory cell that the address references ($\hat{\ell}$), and the security of PC (ℓ') are known statically and satisfy $\ell' \vee \ell \preceq \hat{\ell}$. At runtime, this static assignment can happen directly without any runtime overhead (as is shown in the example of Section 2.2). Its dynamic counterpart rule $\vdash \text{assign}^?$ does not maintain these static security invariants and thus requires runtime NSU checking, which is implemented as a projection on PC.

As we shall see in the reduction rules, the semantics of the protection term prot performs two things: (1) prot promotes the security of the value reduced from its body by level ℓ (2) it uses a new PC to reduce its body. However, the new PC cannot be any label expression. It has to be one with higher security than both the current PC and the security level ℓ . We capture this invariant

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

$$\begin{array}{c}
\text{prot-ctx} \frac{M \mid \mu \mid PC' \longrightarrow M' \mid \mu'}{\text{prot } PC' \ell MA \mid \mu \mid PC \longrightarrow \text{prot } PC' \ell M' A \mid \mu'} \quad \text{prot-val} \frac{}{\text{prot } PC' \ell VA \mid \mu \mid PC \longrightarrow \text{stamp-val } VA \ell \mid \mu} \\
\text{cast} \frac{V \langle c \rangle \longrightarrow M}{V \langle c \rangle \mid \mu \mid PC \longrightarrow M \mid \mu} \quad \beta \frac{}{\text{app } (\lambda x. N) V AB \ell \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp } PC \ell) \ell (N[x := V]) B \mid \mu} \\
\text{app-cast} \frac{\text{NF } \bar{c} \quad (\text{stamp } PC \ell) \langle \bar{d} \rangle \longrightarrow^* PC' \quad V \langle c \rangle \longrightarrow^* W}{\text{app } (\lambda x. N \langle \bar{d}, c \rightarrow d, \bar{c} \rangle) V CD \ell \mid \mu \mid PC \longrightarrow \text{prot } PC' \ell ((N[x := W]) \langle \bar{d} \rangle) D \mid \mu} \\
\text{app}\star\text{-cast} \frac{\text{NF } \bar{c} \quad (\text{stamp! } PC \mid \bar{c} \mid) \langle \bar{d} \rangle \longrightarrow^* PC' \quad V \langle c \rangle \longrightarrow^* W}{\text{app}\star (\lambda x. N \langle \bar{d}, c \rightarrow d, \bar{c} \rangle) V CT \mid \mu \mid PC \longrightarrow \text{prot } PC' \mid \bar{c} \mid ((N[x := W]) \langle \bar{d} \rangle) (T_\star) \mid \mu} \\
\text{if-true} \frac{}{\text{if } \$ \text{true } A \ell MN \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp } PC \ell) \ell MA \mid \mu} \\
\text{if}\star\text{-true-cast} \frac{\text{NF } \bar{c}}{\text{if}\star (\$ \text{true } \langle \text{id}(\text{Bool}), \bar{c} \rangle) T MN \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp! } PC \mid \bar{c} \mid) \mid \bar{c} \mid M (T_\star) \mid \mu} \\
\text{ref} \frac{n \text{ FreshIn } \mu(\ell)}{\text{ref } \ell V \mid \mu \mid PC \longrightarrow \text{addr } n \mid (\mu, \ell \mapsto n \mapsto V)} \quad \text{ref?} \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{ref?}\text{-blame} \frac{PC \langle \star \Rightarrow^P \ell \rangle \longrightarrow^* \text{blame } q}{\text{ref?}^P \ell V \mid \mu \mid PC \longrightarrow \text{blame } q \mid \mu} \quad \text{assign} \frac{}{\text{assign } (\text{addr } n) V T \hat{\ell} \ell \mid \mu \mid PC \longrightarrow \$ \text{unit} \mid [\hat{\ell} \mapsto n \mapsto V] \mu} \\
\text{assign?}\text{-cast} \frac{\text{NF } \bar{c} \quad \vdash c : T_g \Rightarrow S_{\hat{\ell}} \quad \vdash d : S_{\hat{\ell}} \Rightarrow T_g \quad (\text{stamp! } PC \mid \bar{c} \mid) \langle \star \Rightarrow^P \hat{\ell} \rangle \longrightarrow^* PC' \quad V \langle c \rangle \longrightarrow^* W}{\text{assign?}^P (\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) V T g \mid \mu \mid PC \longrightarrow \$ \text{unit} \mid [\hat{\ell} \mapsto n \mapsto W] \mu} \\
\text{deref} \frac{\mu(\hat{\ell}, n) = V}{!(\text{addr } n) T_{\hat{\ell}} \ell \mid \mu \mid PC \longrightarrow \text{prot_} \ell V T_{\hat{\ell}} \mid \mu} \\
\text{deref}\star\text{-cast} \frac{\text{NF } \bar{c} \quad \vdash c : S_\star \Rightarrow T_{\hat{\ell}} \quad \vdash d : T_{\hat{\ell}} \Rightarrow S_\star \quad \mu(\hat{\ell}, n) = V}{!\star (\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) S \mid \mu \mid PC \longrightarrow \text{prot_} \mid \bar{c} \mid (V \langle d \rangle) (S_\star) \mid \mu}
\end{array}$$

Fig. 6. Selected semantics rules for λ_{IFC}^c . NSU checks are represented using label expressions (highlighted)

in side condition $\ell' \vee \ell \approx |PC|$ in rule $\vdash \text{prot}$, where ℓ' is the security of the current PC and $|PC|$ is the security for the new PC. The invariant is used in the proof of noninterference.

5.1.2 Operational Semantics for λ_{IFC}^c . We show the interesting rules of the operational semantics of λ_{IFC}^c in Figure 6 (full version in Figure 16 and 17 of the Appendix). The reduction relation takes the form $M \mid \mu \mid PC \longrightarrow N \mid \mu'$, which reduces the configuration of term M and heap μ under the label expression PC to another configuration N and μ' . The heap is a map $(\ell, n) \mapsto V$, where a cell is indexed by its security level ℓ and by index n among the cells of ℓ . The predicate $(n \text{ FreshIn } \mu(\ell))$ means that the index n is fresh (not already in use) among all cells with security ℓ ; when performing a lookup, $\mu(\ell, n) = V$ retrieves the value V at index n whose security level is ℓ .

Protection terms. Following standard approaches to IFC, a protection term $\text{prot } PC' \ell MA$ has two functionalities: (1) it ensures that the reduction inside M does not leak information through heap write operations (2) it promotes the security level of the computation result of M to at least level ℓ . The first functionality is achieved by switching to PC' from the current PC when reducing the body M (rule prot-ctx). Recall Section 5.1.1, the typing of prot makes sure that PC' has the correct security that is at least as secure as both PC and ℓ . The second functionality is achieved by stamping the value produced by the body of prot (rule prot-val). The stamping of values in

λ_{IFC}^c (Figure 14 in the Appendix) is analogous to stamping of label expressions and turns to the stamping operation for coercions on labels (Figure 2) in a similar way. Again there are two cases, because the value is either (1) a raw value or (2) a coercion-wrapped value. Suppose the value is a raw value, if its type has **low** security and is stamped with **high**, the value becomes wrapped with a subtype coercion; otherwise the value stays unchanged. Otherwise, if the value is wrapped with an irreducible coercion, we stamp the top-level coercion sequence.

Function Application. The β rule is standard for IFC languages. It generates a prot term with the specific security label ℓ that comes from the label on the λ , preventing implicit flow from the function being applied through both the computation result and the heap. The *app-cast* rule applies a function wrapped in a function coercion to a value V . The application is “static”, so the security level of prot comes from the function type just like β . The function coercion is distributed into its domain coercion c , its co-domain coercion d , and the coercion on PC \bar{d} . The coercion \bar{c} is not used because the function type is fully static, so its security is already indicated by its type. The domain coercion c casts the input of the function V to W and W is substituted into the body of the λ . The substituted body goes through the co-domain cast d , and is then protected by ℓ using prot. The stamped PC casts to PC' by \bar{d} and PC' is used as the PC for prot. The rule *app \star -cast* is similar to *app-cast* except for two things: (1) the PC is stamped and then injected using *stamp!* to preserve types (2) the security level of the function proxy used in protection is indicated by $|\bar{c}|$ instead, because the top-level security label of the function is statically unknown (\star).

If-conditional. The static rule *if-true* is standard; the if term reduces a prot whose security ℓ comes from the type of the branch condition, guarding against implicit flow. The rationale of *if \star -true-cast* follows that of *app \star -cast*: (1) a *stamp!* is generated to stamp and then inject the PC and (2) the security of prot is retrieved from the coercion in the branch condition.

NSU and heap operations. Let us consider reference creation first. A “static” reference creation is secure because its typing (rule $\vdash\text{-ref}$, Figure 12 of the Appendix) already enforces the heap policy. Consequently, the allocation can happen directly (rule *ref*) without any runtime checking. Rule *ref?* does the same reference creation but with NSU checking, by casting the current PC to the security ℓ of the newly created memory cell. The coerce function ($- \Rightarrow^- -$) takes two security labels and a blame label to generate a coercion on labels. In this case, $\star \Rightarrow^p \ell$ generates $\text{id}(\star)$; $\ell ?^p$, which performs a projection whose target is ℓ . If the projected PC reduces to a blame, it means that NSU checking fails so we lift the blame to λ_{IFC}^c . Assignment follows the same pattern: a static assignment can happen directly, while *assign?* requires NSU checking, by stamping the current PC with the security indicated in the coercion and then projecting to $\hat{\ell}$, where $\hat{\ell}$ is the security of the heap cell. The input coercion c is applied before the value is stored into the cell.

The rule for static dereferencing *deref* looks up index n in all memory cells with security level $\hat{\ell}$. The value from the lookup is protected with ℓ , the top-level security label of the reference type. The PC of the prot does not matter, because V is already a value and will not reduce. The rule *deref \star -cast* dereferences a reference proxy. It looks up the value V in the heap, applies the output coercion d , and generates a prot with the security of the coercion $|\bar{c}|$. The PC of prot does not matter in this case either, because applying a coercion is pure and does not produce side effects.

5.2 Meta-theoretical Results of λ_{IFC}^c

We prove type safety for λ_{IFC}^c (§ 5.2.1) and the main simulation lemma for λ_{IFC}^c (§ 5.2.2) used in the proof of the gradual guarantee (§ 6.4). We also prove that λ_{IFC}^c satisfies noninterference (§ 5.2.3).

5.2.1 Type Safety. We show that λ_{IFC}^c is type safe by proving progress and preservation. Progress says that a well-typed λ_{IFC}^c term does not get stuck. The term is either a value or a blame, which does not reduce, or the term takes one reduction step. Heap well-typedness is defined point-wise.

THEOREM 10 (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 also well-typed: $\Sigma \vdash \mu$. Then either (1) M is a value or (2) M is a blame: $M = \text{blame } p$ or (3) M can take a reduction step: $M \mid \mu \mid PC \longrightarrow N \mid \mu'$ for some N and μ' .*

The operation semantics of λ_{IFC}^c preserves types and the well-typedness of heap:

THEOREM 11 (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 also 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 $\Sigma' \vdash \mu'$.*

5.2.2 Simulation Between λ_{IFC}^c Terms of Different Precision. The main simulation lemma says that if two terms are related by *precision* and the more precise side takes one step, then the less precise side is able to multi-step and get back in sync. The precision relation is in form $\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \sqsubseteq M' \Leftarrow A \sqsubseteq A'$, where $\Gamma; \Sigma; g; \ell$ corresponds to the typing context, heap context, type of PC , and security of PC of the less precise term M and $\Gamma'; \Sigma'; g'; \ell'$ is for those of the more precise term M' . The types of the two terms, A and A' , are related by precision between types. The intuition between this precision relation is that casts are allowed to appear in different places between the more precise and the less precise λ_{IFC}^c terms. Moreover, the casts must be in shapes that preserve the precision of λ_{IFC}^c (more or fewer static type annotations provided by the programmer). According to the gradual guarantee, the more precise side is allowed to signal more blames, so there is a rule (\sqsubseteq -blame) that relates $\text{blame } p$ to any term M on the less precise side as long as their types are in sync. We list selected precision rules in Figure 21 of the Appendix.

With the precision relation of λ_{IFC}^c defined, we first state the catch-up lemma, which catches up to a more-precise value by multi-stepping on the less-precise side:

LEMMA 12 (CATCHING UP TO MORE PRECISE). *If term M and value V' are related by precision:*

$$\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \sqsubseteq V' \Leftarrow A \sqsubseteq A'$$

then there exists value V s.t. $M \mid \mu \mid PC \longrightarrow^ V \mid \mu$ and $\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash V \sqsubseteq V' \Leftarrow A \sqsubseteq A'$.*

We prove the main simulation lemma using the catch-up lemma. Heap precision is defined point-wise similar to the definition of heap well-typedness.

LEMMA 13 (SIMULATION BETWEEN MORE PRECISE AND LESS PRECISE λ_{IFC}^c TERMS). *Suppose PC, PC' are related by precision: $\vdash PC \sqsubseteq PC' \Leftarrow g \sqsubseteq g'$. Moreover suppose M, M' are related by precision:*

$$\emptyset; \emptyset; \Sigma_1; \Sigma'_1; g; g'; |PC|; |PC'| \vdash M \sqsubseteq M' \Leftarrow A \sqsubseteq A'$$

heap contexts Σ_1, Σ'_1 are related by precision: $\Sigma_1 \sqsubseteq \Sigma'_1$, the initial heaps μ_1, μ'_1 are also related by precision: $\Sigma_1; \Sigma'_1 \vdash \mu_1 \sqsubseteq \mu'_1$.

If $M' \mid \mu'_1 \mid PC' \longrightarrow N' \mid \mu'_2$, there exists $\Sigma_2, \Sigma'_2, N, \mu_2$ s.t. $\Sigma_2 \supseteq \Sigma_1, \Sigma'_2 \supseteq \Sigma'_1, \Sigma_2 \sqsubseteq \Sigma'_2$,

$$M \mid \mu_1 \mid PC \longrightarrow^* N \mid \mu_2$$

the resulting terms are related by precision: $\emptyset; \emptyset; \Sigma_2; \Sigma'_2; g; g'; |PC|; |PC'| \vdash N \sqsubseteq N' \Leftarrow A \sqsubseteq A'$ and the resulting heaps are also related by precision: $\Sigma_2; \Sigma'_2 \vdash \mu_2 \sqsubseteq \mu'_2$.

5.2.3 Noninterference. We prove that λ_{IFC}^c satisfies termination-insensitive noninterference in Section 11 of the Appendix. The statement of termination-insensitive noninterference says that if we run a program with different high-security inputs in two executions, then their low-security output values should be the related (e.g. the same boolean):

LEMMA 14 (NONINTERFERENCE FOR λ_{IFC}^c). *If M is well-typed: $(x: \text{Bool}_{\text{high}}); \emptyset; \text{low}; \text{low} \vdash M : \text{Bool}_{\text{low}}$ and $M[x := \$b_1] \mid \emptyset \mid \text{low} \longrightarrow^* V_1 \mid \mu_1$ and $M[x := \$b_2] \mid \emptyset \mid \text{low} \longrightarrow^* V_2 \mid \mu_2$, then $V_1 = V_2$.*

$$\begin{array}{l}
 \text{terms } L, M, N ::= x \mid (\$ k)_\ell \mid (\lambda^g x:A. N)_\ell \mid (L M)^P \\
 \quad \quad \quad \mid (\text{if } L \text{ then } M \text{ else } N)^P \mid \text{let } x = M \text{ in } N \\
 \quad \quad \quad \mid (\text{ref } \ell M)^P \mid !^P M \mid (L := M)^P \mid (M : A)^P \\
 \\
 \boxed{\Gamma; g \vdash M : A} \\
 \\
 \vdash_{\text{lam}} \frac{(\Gamma, x:A); g \vdash N : B}{\Gamma; g' \vdash (\lambda^g x:A. N)_\ell : (A \xrightarrow{g} B)_\ell} \quad \vdash_{\text{assign}} \frac{\begin{array}{l} \Gamma; g' \vdash L : (\text{Ref } T_{\hat{g}})_g \quad \Gamma; g' \vdash M : A \\ A \lesssim T_{\hat{g}} \quad g \lesssim \hat{g} \quad g' \lesssim \hat{g} \end{array}}{\Gamma; g' \vdash (L := M)^P : \text{Unit}_{\text{low}}}
 \end{array}$$

Fig. 7. Syntax and selected typing rules of λ_{IFC}^* (highlighted security labels ℓ default to `low` if omitted)

6 THE λ_{IFC}^* LANGUAGE WITH GRADUAL INFORMATION-FLOW CONTROL

In this section, we first define the gradual language λ_{IFC}^* in Section 6.1. It is similar to GSL_{Ref} with respect to syntax and typing rules. The main syntactic difference is that in λ_{IFC}^* , the security labels of literals and newly created memory cells default to a specific label such as `low`, while in GSL_{Ref} they default to a runtime unknown security level \star . We show that λ_{IFC}^* can be compiled into our cast calculus λ_{IFC}^c and the compilation preserves types in Section 6.2. As a result, the semantics of λ_{IFC}^* can be defined by the operational semantics of λ_{IFC}^c . In Section 6.3, we prove the noninterference of λ_{IFC}^* as a corollary of the noninterference lemma for λ_{IFC}^c and compilation preserves types. Finally, we prove the gradual guarantee for λ_{IFC}^* as a corollary of the main simulation lemma of λ_{IFC}^c (Lemma 13), thus solving the tension discovered by Toro et al. [2018].

6.1 Syntax and Typing of the Surface Language λ_{IFC}^*

The syntax and selected typing rules of λ_{IFC}^* are shown in Figure 7 (full version in Figure 10 of the Appendix). They are directly adapted from those of GSL_{Ref} , by changing the security labels on values to disallow the \star label.

The rule \vdash_{assign} includes two side conditions $g \lesssim \hat{g}$ and $g' \lesssim \hat{g}$. If g , g' , and \hat{g} are all specific security labels, the heap policy is enforced statically, because the typing tells us that the security of the current PC as well as the memory address itself is lower than or equal to the security of the memory cell, thus no implicit flow through the heap. Indeed, we are going to see in the next section that an assignment where all three labels are statically known generates an static assign. The semantics of assign does not perform NSU because the static check on its typing rule suffices. Relating to the λ_{IFC}^* program at the end of Section 2.2, there is no runtime NSU check, because the program generates a static assign that enforces the heap policy statically.

6.2 Compiling λ_{IFC}^* to λ_{IFC}^c : Cast Insertion

The compile function takes the form $C M = M'$, where M is a λ_{IFC}^* program and M' is a λ_{IFC}^c term. Consider the case for assignment:

$$C (L := M)^P = \begin{cases} \text{assign } (C L) ((C M) \langle c_2 \rangle) T \hat{g} g & , \text{ if } g, g' \text{ and } \hat{g} \text{ are all specific} \\ \text{assign?}^P ((C L) \langle c_1 \rangle) ((C M) \langle c_2 \rangle) T \hat{g} & , \text{ otherwise} \end{cases}$$

$$\text{where } c_1 = (\text{Ref } T_{\hat{g}})_g \Rightarrow^P (\text{Ref } T_{\hat{g}})_{\star}, c_2 = A \Rightarrow^P T_{\hat{g}}, \Gamma; g' \vdash L : (\text{Ref } T_{\hat{g}})_g, \Gamma; g' \vdash M : A$$

If g , g' , and \hat{g} are specific, the check for heap policy can be statically justified. We recursively compile both L and M and generate a static assign. We cast M' using coercion c_2 , which casts from the type of M' to the type of the memory cell. The coercion is produced by the coerce function, $(- \Rightarrow^- -)$, which takes two types and a blame label, returning a coercion on values by calling the

coerce function of labels on each pair of security labels inside those two types. If at least one of the three security labels is \star , the typing information is insufficient to justify the assignment. The compilation produces an `assign?` instead, whose semantics performs NSU checking at runtime.

Compilation from $\lambda_{\text{IFC}}^{\star}$ to λ_{IFC}^c preserves types:

THEOREM 15 (COMPILATION PRESERVES TYPES). *If $\Gamma; g \vdash M : A$, then $\Gamma; \emptyset; g; \text{low} \vdash C M \Leftarrow A$.*

6.3 Noninterference for $\lambda_{\text{IFC}}^{\star}$

The noninterference theorem of $\lambda_{\text{IFC}}^{\star}$ is a straightforward corollary of the noninterference lemma for λ_{IFC}^c (Section 5.2.3) and compilation preserves types. The proof is in Section 11 of the Appendix.

THEOREM 16 (NONINTERFERENCE FOR $\lambda_{\text{IFC}}^{\star}$). *Suppose a $\lambda_{\text{IFC}}^{\star}$ program M is well-typed: $(x:\text{Bool}_{\text{high}}); \text{low} \vdash M : \text{Bool}_{\text{low}}$. If for any boolean inputs b_1, b_2*

$$(C M)[x := \$ b_1] \mid \emptyset \mid \text{low} \longrightarrow^* V_1 \mid \mu_1 \quad \text{and} \quad (C M)[x := \$ b_2] \mid \emptyset \mid \text{low} \longrightarrow^* V_2 \mid \mu_2$$

then the resulting values $V_1 = V_2$.

6.4 The Gradual Guarantee

Finally, we state the gradual guarantee of $\lambda_{\text{IFC}}^{\star}$ and prove it as a corollary of Lemma 13. The definition of term precision for $\lambda_{\text{IFC}}^{\star}$ is in Figure 19 of the Appendix.

THEOREM 17 (GRADUAL GUARANTEE). *Suppose M and M' are well-typed terms in $\lambda_{\text{IFC}}^{\star}$ that are related by precision, that is $\vdash M \sqsubseteq M'$. So $\emptyset; \text{low} \vdash M : A$, $\emptyset; \text{low} \vdash M' : A'$, and $A \sqsubseteq A'$. If the compilation of M' reduces to a value: $C M' \mid \emptyset \mid \text{low} \longrightarrow^* V' \mid \mu'$ there exists a value V and heap μ s.t. the compilation of M reduces to $V : C M \mid \emptyset \mid \text{low} \longrightarrow^* V \mid \mu$ and the resulting values are related by precision for some Σ, Σ' :*

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

PROOF. Compilation preserves precision, so $\emptyset; \emptyset; \emptyset; \text{low}; \text{low}; \text{low}; \text{low} \vdash C M \sqsubseteq C M' \Leftarrow A \sqsubseteq A'$. We then proceed by induction on the reduction of $C M'$ to a value V' , using Lemma 13 to show that $C M$ reduces to a corresponding term at each step. So we have $C M \longrightarrow^* N$ where $N \sqsubseteq V'$ for some N . We then apply Lemma 12 to show that N reduces to a value V where $V \sqsubseteq V'$. \square

7 CONCLUSION

We presented the design of a gradual information-flow language $\lambda_{\text{IFC}}^{\star}$ that satisfies both noninterference and the gradual guarantee while maintaining the principle of type-based reasoning. The key to the design of $\lambda_{\text{IFC}}^{\star}$ is to walk back the decision in GSL_{Ref} to include the unknown label \star among the runtime security labels. So $\lambda_{\text{IFC}}^{\star}$ takes a more standard approach to gradually-typed IFC: the \star label can be used in type annotations but not as the security level of a runtime value. The $\lambda_{\text{IFC}}^{\star}$ language is defined by translation to a cast calculus λ_{IFC}^c . This intermediate language employs a coercion calculus to express the implicit conversions between more-or-less precise parts of the program. We proved that $\lambda_{\text{IFC}}^{\star}$ satisfies termination-insensitive noninterference. We proved that $\lambda_{\text{IFC}}^{\star}$ satisfies the gradual guarantee and mechanized the result in Agda.

DATA-AVAILABILITY STATEMENT

The Agda code of this paper is in the supplementary material [Chen 2024].

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1763922.

REFERENCES

- Aslan Askarov and Andrei Sabelfeld. 2009. Tight Enforcement of Information-Release Policies for Dynamic Languages. In *2009 22nd IEEE Computer Security Foundations Symposium*. 43–59. <https://doi.org/10.1109/CSF.2009.22>
- Thomas H Austin and Cormac Flanagan. 2009. Efficient purely-dynamic information flow analysis. In *Proceedings of the ACM SIGPLAN Fourth Workshop on Programming Languages and Analysis for Security*. 113–124. <https://doi.org/10.1145/1554339.1554353>
- Thomas H. Austin, Tommy Schmitz, and Cormac Flanagan. 2017. Multiple Facets for Dynamic Information Flow with Exceptions. *ACM Trans. Program. Lang. Syst.* 39, 3, Article 10 (may 2017), 56 pages. <https://doi.org/10.1145/3024086>
- Arthur Azevedo de Amorim, Matt Fredrikson, and Limin Jia. 2020. Reconciling noninterference and gradual typing. In *Logic in Computer Science (LICS)*. <https://doi.org/10.1145/3373718.3394778>
- Abhishek Bichhawat, McKenna McCall, and Limin Jia. 2021. Gradual Security Types and Gradual Guarantees. In *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*. IEEE, 1–16. <https://doi.org/10.1109/CSF51468.2021.00015>
- Pablo Buiras, Dimitrios Vytiniotis, and Alejandro Russo. 2015. HLIO: Mixing Static and Dynamic Typing for Information-Flow Control in Haskell. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming (Vancouver, BC, Canada) (ICFP 2015)*. Association for Computing Machinery, New York, NY, USA, 289–301. <https://doi.org/10.1145/2784731.2784758>
- Deepak Chandra and Michael Franz. 2007. Fine-Grained Information Flow Analysis and Enforcement in a Java Virtual Machine. In *Twenty-Third Annual Computer Security Applications Conference (ACSAC 2007)*. 463–475. <https://doi.org/10.1109/ACSAC.2007.37>
- Tianyu Chen. 2024. *cty12/pldi2024-ae: PLDI Release 2*. <https://doi.org/10.5281/zenodo.10933110>
- Tianyu Chen and Jeremy G. Siek. 2022. Mechanized Noninterference for Gradual Security. arXiv:2211.15745 [cs.PL]
- Dorothy E Denning. 1976. A lattice model of secure information flow. *Commun. ACM* 19, 5 (1976), 236–243. <https://doi.org/10.1145/360051.360056>
- Dominique Devriese and Frank Piessens. 2010. Noninterference through Secure Multi-execution. In *2010 IEEE Symposium on Security and Privacy*. 109–124. <https://doi.org/10.1109/SP.2010.15>
- Tim Disney and Cormac Flanagan. 2011. Gradual Information Flow Typing. In *Workshop on Script to Program Evolution*.
- L. Fennell and P. Thiemann. 2013. Gradual Security Typing with References. In *2013 IEEE 26th Computer Security Foundations Symposium*. 224–239. <https://doi.org/10.1109/CSF.2013.22>
- Luminous Fennell and Peter Thiemann. 2015. LJGS: Gradual Security Types for Object-Oriented Languages. In *Workshop on Foundations of Computer Security (FCS)*. <https://doi.org/10.4230/LIPIcs.ECOOP.2016.9>
- Robert Bruce Findler and Matthias Felleisen. 2002. *Contracts for Higher-Order Functions*. Technical Report NU-CCS-02-05. Northeastern University. <https://doi.org/10.1145/2502508.2502521>
- Ronald Garcia, Alison M. Clark, and Éric Tanter. 2016. Abstracting Gradual Typing. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (St. Petersburg, FL, USA) (POPL 2016)*. ACM, New York, NY, USA, 429–442. <https://doi.org/10.1145/2837614.2837670>
- Michael Greenberg. 2014. Space-Efficient Manifest Contracts. CoRR abs/1410.2813 (2014). <http://arxiv.org/abs/1410.2813>
- Fritz Henglein. 1994. Dynamic typing: syntax and proof theory. *Science of Computer Programming* 22, 3 (June 1994), 197–230. [https://doi.org/10.1016/0167-6423\(94\)00004-2](https://doi.org/10.1016/0167-6423(94)00004-2)
- David Herman, Aaron Tomb, and Cormac Flanagan. 2010. Space-efficient gradual typing. *Higher-Order and Symbolic Computation* 23, 2 (2010), 167–189. <https://doi.org/10.1007/s10990-011-9066-z>
- Gurvan Le Guernic. 2007. Automaton-based confidentiality monitoring of concurrent programs. In *20th IEEE Computer Security Foundations Symposium (CSF'07)*. IEEE, 218–232. <https://doi.org/10.1109/CSF.2007.10>
- Gurvan Le Guernic and Thomas Jensen. 2005. Monitoring information flow. In *Proc. Workshop on Foundations of Computer Security*. 19–30.
- Peng Li and Steve Zdancewic. 2010. Arrows for secure information flow. *Theoretical Computer Science* 411, 19 (2010), 1974–1994. <https://doi.org/10.1016/j.tcs.2010.01.025>
- Mathematical Foundations of Programming Semantics (MFPS 2006).
- Scott Moore and Stephen Chong. 2011. Static analysis for efficient hybrid information-flow control. In *2011 IEEE 24th Computer Security Foundations Symposium*. IEEE, 146–160. <https://doi.org/10.1109/CSF.2011.17>
- Andrew C Myers. 1999. JFlow: Practical mostly-static information flow control. In *Proceedings of the 26th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. 228–241. <https://doi.org/10.1145/292540.292561>
- Andrew C. Myers and Barbara Liskov. 1997. A Decentralized Model for Information Flow Control. *SIGOPS Oper. Syst. Rev.* 31, 5 (oct 1997), 129–142. <https://doi.org/10.1145/269005.266669>
- Andrew C. Myers, Lantian Zheng, Steve Zdancewic, Stephen Chong, and Nathaniel Nystrom. 2006. *Jif 3.0: Java information flow*. <http://www.cs.cornell.edu/jif>
- Alejandro Russo and Andrei Sabelfeld. 2010. Dynamic vs. static flow-sensitive security analysis. In *2010 23rd IEEE Computer Security Foundations Symposium*. IEEE, 186–199. <https://doi.org/10.1109/CSF.2010.20>

- Paritosh Shroff, Scott Smith, and Mark Thober. 2007. Dynamic Dependency Monitoring to Secure Information Flow. In *20th IEEE Computer Security Foundations Symposium (CSF'07)*. 203–217. <https://doi.org/10.1109/CSF.2007.20>
- Jeremy G. Siek and Walid Taha. 2006. Gradual typing for functional languages. In *Scheme and Functional Programming Workshop*. 81–92.
- Jeremy G. Siek and Walid Taha. 2007. Gradual Typing for Objects. In *European Conference on Object-Oriented Programming (LCNS, Vol. 4609)*. 2–27. https://doi.org/10.1007/978-3-540-73589-2_2
- Jeremy G. Siek, Michael M. Vitousek, Matteo Cimini, and John Tang Boyland. 2015. Refined Criteria for Gradual Typing. In *SNAPL: Summit on Advances in Programming Languages (LIPIcs: Leibniz International Proceedings in Informatics)*. <https://doi.org/10.4230/LIPIcs.SNAPL.2015.274>
- Deian Stefan, David Mazières, John C. Mitchell, and Alejandro Russo. 2017. Flexible dynamic information flow control in the presence of exceptions. *Journal of Functional Programming* 27 (2017). <https://doi.org/10.1017/S0956796816000241>
- Deian Stefan, Alejandro Russo, John C Mitchell, and David Mazières. 2011. Flexible dynamic information flow control in Haskell. In *Proceedings of the 4th ACM symposium on Haskell*. 95–106. <https://doi.org/10.1145/2034675.2034688>
- Deian Stefan, Alejandro Russo, John C Mitchell, and David Mazières. 2012. Flexible dynamic information flow control in the presence of exceptions. *arXiv preprint arXiv:1207.1457* (2012).
- Sam Tobin-Hochstadt and Matthias Felleisen. 2008. The Design and Implementation of Typed Scheme. In *Symposium on Principles of Programming Languages*. <https://doi.org/10.1145/1328897.1328486>
- Matias Toro, Ronald Garcia, and Éric Tanter. 2018. Type-Driven Gradual Security with References. *ACM Trans. Program. Lang. Syst.* 40, 4, Article 16 (Dec. 2018), 55 pages. <https://doi.org/10.1145/3229061>
- Dennis Volpano, Cynthia Irvine, and Geoffrey Smith. 1996. A sound type system for secure flow analysis. *Journal of computer security* 4, 2-3 (1996), 167–187. <https://doi.org/10.3233/JCS-1996-42-304>
- Philip Wadler. 1989. Theorems for free!. In *FPCA '89: Proceedings of the fourth international conference on Functional programming languages and computer architecture* (Imperial College, London, United Kingdom). ACM, 347–359. <https://doi.org/10.1145/99370.99404>
- Philip Wadler and Robert Bruce Findler. 2009. Well-typed programs can't be blamed. In *European Symposium on Programming (ESOP)*. 1–16. https://doi.org/10.1007/978-3-642-00590-9_1
- Preston Tunnell Wilson, Ben Greenman, Justin Pombrio, and Shriram Krishnamurthi. 2018. The Behavior of Gradual Types: A User Study. In *Dynamic Languages Symposium*. <https://doi.org/10.1145/3393673.3276947>
- Jian Xiang and Stephen Chong. 2021. Co-Inflow: Coarse-grained Information Flow Control for Java-like Languages. In *Proceedings of the 2021 IEEE Symposium on Security and Privacy*. IEEE Press, Piscataway, NJ, USA. <https://doi.org/10.1109/SP40001.2021.00002>

APPENDIX

The Appendix is organized as follows. In Section 8 we present the complete definitions of λ_{IFC}^* and its cast calculus λ_{IFC}^c , and we define the precision relations needed to state the gradual guarantee for λ_{IFC}^* (Theorem 17). In Section 9, we show that vigilance *without* type-guided classification is insufficient to enable type-based reasoning by analyzing the semantics of λ_{SEC}^* [Chen and Siek 2022]. On the contrary, the coercion calculi of λ_{IFC}^* help us achieve both vigilance *and* type-guided classification, thus enabling type-based reasoning, while satisfying the gradual guarantee at the same time. In Section 10, we show that the normalization of coercion sequences is able to model information flow checks, because coercion sequences strongly normalize and the normalization is deterministic. Finally in Section 11, we present a proof for the noninterference result of λ_{IFC}^* , by establishing a simulation between its cast calculus λ_{IFC}^c and a dynamic IFC language λ_{SEC} derived from the calculus of Chen and Siek [2022].

8 SUPPLEMENTARY DEFINITIONS

We present the complete definition of the λ_{IFC}^* and λ_{IFC}^c languages in Sections 8.2 and 8.3, respectively. There are some definitions shared between the two languages, which we present in Section 8.1. The definitions needed to state and prove the gradual guarantee are presented in Section 8.4.

8.1 Shared Definitions for λ_{IFC}^* and λ_{IFC}^c

Figure 8 presents auxiliary operators on security labels and types that are used by the type systems of λ_{IFC}^* or its cast calculus λ_{IFC}^c . These operators include consistent join, join w.r.t precision, consistent meet, and the stamping operation on types.

$$\begin{array}{l}
\ell \sqcup \ell = \ell \quad \boxed{g \sqcup g} \\
\star \sqcup g = g \\
g \sqcup \star = g \\
\iota \sqcup \iota = \iota \quad \boxed{T \sqcup T} \\
(\text{Ref } A) \sqcup (\text{Ref } B) = \text{Ref } (A \sqcup B) \\
(A \xrightarrow{g_1} B) \sqcup (C \xrightarrow{g_2} D) = (A \sqcup C) \xrightarrow{g_1 \sqcup g_2} (B \sqcup D) \\
S_{g_1} \sqcup T_{g_2} = (S \sqcup T)_{g_1 \sqcup g_2} \quad \boxed{A \sqcup A} \\
\ell_1 \tilde{\vee} \ell_2 = \ell_1 \vee \ell_2 \quad \boxed{g \tilde{\vee} g} \\
-\tilde{\vee} \star = \star \\
\star \tilde{\vee} - = \star \\
\iota \tilde{\vee} \iota = \iota \quad \boxed{T \tilde{\vee} T} \\
(\text{Ref } A) \tilde{\vee} (\text{Ref } B) = \text{Ref } (A \sqcup B) \\
(A \xrightarrow{g_1} B) \tilde{\vee} (C \xrightarrow{g_2} D) = (A \tilde{\wedge} C) \xrightarrow{g_1 \tilde{\wedge} g_2} (B \tilde{\vee} D) \\
S_{g_1} \tilde{\vee} T_{g_2} = (S \tilde{\vee} T)_{g_1 \tilde{\vee} g_2} \quad \boxed{A \tilde{\vee} A} \\
\ell_1 \tilde{\wedge} \ell_2 = \ell_1 \wedge \ell_2 \quad \boxed{g \tilde{\wedge} g} \\
-\tilde{\wedge} \star = \star \\
\star \tilde{\wedge} - = \star \\
\iota \tilde{\wedge} \iota = \iota \quad \boxed{T \tilde{\wedge} T} \\
(\text{Ref } A) \tilde{\wedge} (\text{Ref } B) = \text{Ref } (A \sqcup B) \\
(A \xrightarrow{g_1} B) \tilde{\wedge} (C \xrightarrow{g_2} D) = (A \tilde{\vee} C) \xrightarrow{g_1 \tilde{\vee} g_2} (B \tilde{\wedge} D) \\
S_{g_1} \tilde{\wedge} T_{g_2} = (S \tilde{\wedge} T)_{g_1 \tilde{\wedge} g_2} \quad \boxed{A \tilde{\wedge} A} \\
\text{stamp } (T_{g_1}) g_2 = T_{g_1 \tilde{\vee} g_2}
\end{array}$$

Fig. 8. Auxiliary operators for security labels and types: join w.r.t precision (\sqcup), consistent join ($\tilde{\vee}$ for labels and $\tilde{\vee}$ for types), and consistent meet ($\tilde{\wedge}$ for labels and $\tilde{\wedge}$ for types). Stamping for types

8.2 Definition of λ_{IFC}^*

Figure 10 contains the type system of λ_{IFC}^* . To support the type system of λ_{IFC}^* , we first present the definitions of consistent subtyping for security labels and types in Figure 9.

$$\boxed{g_1 \lesssim g_2, S \lesssim T, \text{ and } A \lesssim B}$$

$$\begin{array}{c}
 \begin{array}{ccc}
 \lesssim^* \frac{}{g \lesssim^*} & \star \lesssim \frac{}{\star \lesssim g} & \lesssim^{-\ell} \frac{\ell_1 \preccurlyeq \ell_2}{\ell_1 \lesssim \ell_2} \\
 \\
 \lesssim^{-\iota} \frac{}{\iota \lesssim \iota} & \lesssim^{-\text{ref}} \frac{A \lesssim B \quad B \lesssim A}{\text{Ref } A \lesssim \text{Ref } B} & \\
 \\
 \lesssim^{-\text{fun}} \frac{g_2 \lesssim g_1 \quad C \lesssim A \quad B \lesssim D}{A \xrightarrow{g_1} B \lesssim C \xrightarrow{g_2} D} & \lesssim^{-\tau} \frac{g_1 \lesssim g_2 \quad S \lesssim T}{S_{g_1} \lesssim T_{g_2}} &
 \end{array}
 \end{array}$$

Fig. 9. Consistent subtyping for labels and types

$$\boxed{\Gamma; g \vdash M : A}$$

$$\begin{array}{c}
 \vdash \text{var} \frac{\Gamma \ni x : A}{\Gamma; g \vdash x : A} \quad \vdash \text{const} \frac{k : \iota}{\Gamma; g \vdash (\$ k)_\ell : \iota_\ell} \\
 \\
 \vdash \text{lam} \frac{(\Gamma, x:A); g_2 \vdash N : B}{\Gamma; g_1 \vdash (\lambda^{g_2} x:A. N)_\ell : (A \xrightarrow{g_2} B)_\ell} \quad \vdash \text{app} \frac{\Gamma; g \vdash L : (A \xrightarrow{g_2} B)_{g_1} \quad \Gamma; g \vdash M : A' \quad A' \lesssim A \quad g \lesssim g_2 \quad g_1 \lesssim g_2}{\Gamma; g \vdash (LM)^P : \text{stamp } B \ g_1} \\
 \\
 \vdash \text{let} \frac{\Gamma; g \vdash M : A \quad (\Gamma, x:A); g \vdash N : B}{\Gamma; g \vdash \text{let } x = M \text{ in } N : B} \quad \vdash \text{if} \frac{\Gamma; g_2 \vdash L : \text{Bool}_{g_1} \quad \Gamma; g_2 \tilde{\vee} g_1 \vdash M : A \quad \Gamma; g_2 \tilde{\vee} g_1 \vdash N : B \quad A \tilde{\vee} B = C}{\Gamma; g_2 \vdash (\text{if } L \text{ then } M \text{ else } N)^P : \text{stamp } C \ g_1} \\
 \\
 \vdash \text{ref} \frac{\Gamma; g_2 \vdash M : T_{g_1} \quad T_{g_1} \lesssim T_\ell \quad g_2 \lesssim \ell}{\Gamma; g_2 \vdash (\text{ref } \ell M)^P : (\text{Ref } T_\ell)_{\text{low}}} \quad \vdash \text{deref} \frac{\Gamma; g_2 \vdash M : (\text{Ref } A)_{g_1}}{\Gamma; g_2 \vdash !^P M : \text{stamp } A \ g_1} \\
 \\
 \vdash \text{assign} \frac{\Gamma; g_2 \vdash L : (\text{Ref } T_{\hat{g}})_{g_1} \quad \Gamma; g_2 \vdash M : A \quad A \lesssim T_{\hat{g}} \quad g_2 \lesssim \hat{g} \quad g_1 \lesssim \hat{g}}{\Gamma; g_2 \vdash (L := M)^P : \text{Unit}_{\text{low}}} \quad \vdash \text{ann} \frac{\Gamma; g \vdash M : A' \quad A' \lesssim A}{\Gamma; g \vdash (M : A)^P : A}
 \end{array}$$

Fig. 10. Typing rules of the surface language λ_{IFC}^* . Side conditions about the heap policy are highlighted

$$\boxed{\vdash \mathbf{c} : A \Rightarrow B}$$

$$\vdash_{cbase} \frac{\vdash \bar{c} : g_1 \Rightarrow g_2}{\vdash \mathbf{id}(i), \bar{c} : \iota_{g_1} \Rightarrow \iota_{g_2}} \quad \vdash_{cref} \frac{\vdash \mathbf{c} : B \Rightarrow A \quad \vdash \mathbf{d} : A \Rightarrow B \quad \vdash \bar{c} : g_1 \Rightarrow g_2}{\vdash \mathbf{Ref} \mathbf{c} \mathbf{d}, \bar{c} : (\mathbf{Ref} A)_{g_1} \Rightarrow (\mathbf{Ref} B)_{g_2}}$$

$$\vdash_{cfun} \frac{\vdash \bar{d} : g_4 \Rightarrow g_3 \quad \vdash \mathbf{c} : C \Rightarrow A \quad \vdash \mathbf{d} : B \Rightarrow D \quad \vdash \bar{c} : g_1 \Rightarrow g_2}{\vdash \bar{d}, \mathbf{c} \rightarrow \mathbf{d}, \bar{c} : (A \xrightarrow{g_3} B)_{g_1} \Rightarrow (C \xrightarrow{g_4} D)_{g_2}}$$

Fig. 11. Typing rules for coercions on values.

8.3 Definition of λ_{IFC}^c

8.3.1 *Type System.* To support the type system of λ_{IFC}^c , we define the typing rules for coercions on values in Figure 11. The full typing rules of the cast calculus λ_{IFC}^c are listed in Figure 12.

$$\boxed{\Gamma; \Sigma; g; \ell \vdash M \Leftarrow A}$$

$$\begin{array}{c}
 \vdash \text{var} \frac{\Gamma \ni x : A}{\Gamma; \Sigma; g; \ell \vdash x \Leftarrow A} \quad \vdash \text{const} \frac{k : \iota}{\Gamma; \Sigma; g; \ell \vdash \$k \Leftarrow \iota} \\
 \\
 \vdash \text{addr} \frac{\Sigma(\hat{\ell}, n) = T}{\Gamma; \Sigma; g; \ell' \vdash \text{addr } n \Leftarrow (\text{Ref } T_{\hat{\ell}})_{\ell}} \\
 \\
 \vdash \text{lam} \frac{\forall \ell''. (\Gamma, x:A); \Sigma; g; \ell'' \vdash N \Leftarrow B}{\Gamma; \Sigma; g'; \ell' \vdash \lambda x. N \Leftarrow (A \xrightarrow{g} B)_{\ell}} \quad \vdash \text{let} \frac{\Gamma; \Sigma; g; \ell \vdash M \Leftarrow A \quad \forall \ell'. (\Gamma, x:A); \Sigma; g; \ell' \vdash N \Leftarrow B}{\Gamma; \Sigma; g; \ell \vdash \text{let } x=M:A \text{ in } N \Leftarrow B} \\
 \\
 \vdash \text{app} \frac{\Gamma; \Sigma; \ell'; \ell'' \vdash L \Leftarrow (A \xrightarrow{\ell' \vee \ell} B)_{\ell} \quad \Gamma; \Sigma; \ell'; \ell'' \vdash M \Leftarrow A \quad C = \text{stamp } B \ell}{\Gamma; \Sigma; \ell'; \ell'' \vdash \text{app } L M A B \ell \Leftarrow C} \quad \vdash \text{app}\star \frac{\Gamma; \Sigma; g; \ell \vdash L \Leftarrow (A \xrightarrow{\star} (T_{\star}))_{\star} \quad \Gamma; \Sigma; g; \ell \vdash M \Leftarrow A}{\Gamma; \Sigma; g; \ell \vdash \text{app}\star L M A T \Leftarrow T_{\star}} \\
 \\
 \vdash \text{if} \frac{\Gamma; \Sigma; \ell'; \ell'' \vdash L \Leftarrow \text{Bool}_{\ell} \quad \forall \ell_1. \Gamma; \Sigma; \ell' \vee \ell; \ell_1 \vdash M \Leftarrow A \quad \forall \ell_2. \Gamma; \Sigma; \ell' \vee \ell; \ell_2 \vdash N \Leftarrow A \quad B = \text{stamp } A \ell}{\Gamma; \Sigma; \ell'; \ell'' \vdash \text{if } L A \ell M N \Leftarrow B} \quad \vdash \text{if}\star \frac{\Gamma; \Sigma; g; \ell \vdash L \Leftarrow \text{Bool}_{\star} \quad \forall \ell_1. \Gamma; \Sigma; \star; \ell_1 \vdash M \Leftarrow T_{\star} \quad \forall \ell_2. \Gamma; \Sigma; \star; \ell_2 \vdash N \Leftarrow T_{\star}}{\Gamma; \Sigma; g; \ell \vdash \text{if}\star L T M N \Leftarrow T_{\star}} \\
 \\
 \vdash \text{ref} \frac{\Gamma; \Sigma; \ell'; \ell'' \vdash M \Leftarrow T_{\ell} \quad \ell' \preceq \ell}{\Gamma; \Sigma; \ell'; \ell'' \vdash \text{ref } \ell M \Leftarrow (\text{Ref } T_{\ell})_{\text{low}}} \quad \vdash \text{ref}\star \frac{\Gamma; \Sigma; \star; \ell' \vdash M \Leftarrow T_{\ell}}{\Gamma; \Sigma; \star; \ell' \vdash \text{ref}\star^P \ell M \Leftarrow (\text{Ref } T_{\ell})_{\text{low}}} \\
 \\
 \vdash \text{deref} \frac{\Gamma; \Sigma; g; \ell' \vdash M \Leftarrow (\text{Ref } A)_{\ell} \quad B = \text{stamp } A \ell}{\Gamma; \Sigma; g; \ell' \vdash ! M A \ell \Leftarrow B} \quad \vdash \text{deref}\star \frac{\Gamma; \Sigma; g; \ell \vdash M \Leftarrow (\text{Ref } (T_{\star}))_{\star}}{\Gamma; \Sigma; g; \ell \vdash !\star M T \Leftarrow T_{\star}} \\
 \\
 \vdash \text{assign} \frac{\Gamma; \Sigma; \ell'; \ell'' \vdash L \Leftarrow (\text{Ref } T_{\hat{\ell}})_{\ell} \quad \Gamma; \Sigma; \ell'; \ell'' \vdash M \Leftarrow T_{\hat{\ell}} \quad \ell' \vee \ell \preceq \hat{\ell}}{\Gamma; \Sigma; \ell'; \ell'' \vdash \text{assign } L M T \hat{\ell} \ell \Leftarrow \text{Unit}_{\text{low}}} \\
 \\
 \vdash \text{assign}\star \frac{\Gamma; \Sigma; g; \ell \vdash L \Leftarrow (\text{Ref } T_{\hat{g}})_{\star} \quad \Gamma; \Sigma; g; \ell \vdash M \Leftarrow T_{\hat{g}}}{\Gamma; \Sigma; g; \ell \vdash \text{assign}\star^P L M T \hat{g} \Leftarrow \text{Unit}_{\text{low}}} \\
 \\
 \vdash \text{prot} \frac{\Gamma; \Sigma; g'; |PC| \vdash M \Leftarrow A \quad \vdash PC \Leftarrow g' \quad \ell' \vee \ell \preceq |PC| \quad B = \text{stamp } A \ell}{\Gamma; \Sigma; g; \ell' \vdash \text{prot } PC \ell M A \Leftarrow B} \\
 \\
 \vdash \text{cast} \frac{\Gamma; \Sigma; g; \ell \vdash M \Leftarrow A \quad \vdash c : A \Rightarrow B}{\Gamma; \Sigma; g; \ell \vdash M \langle c \rangle \Leftarrow B} \quad \vdash \text{blame} \frac{}{\Gamma; \Sigma; g; \ell \vdash \text{blame } p \Leftarrow A}
 \end{array}$$

Fig. 12. Typing rules of the cast calculus λ_{IFC}^c . The side conditions that enforce the heap policy statically during memory write operations are highlighted

$$\begin{array}{l}
\boxed{\text{stamp } e \ell = e} \\
\text{stamp } \ell \text{ low} = \ell \\
\text{stamp } \text{low } \text{high} = \text{low } \langle \text{id}(\text{low}); \uparrow \rangle \\
\text{stamp } \text{high } \text{high} = \text{high} \\
\text{stamp } (\ell \langle \bar{c} \rangle) \ell' = \ell \langle \text{stamp } \bar{c} \ell' \rangle
\end{array}
\qquad
\begin{array}{l}
\boxed{\text{stamp! } e \ell = e} \\
\text{stamp! } \ell \ell' = \ell \langle \text{stamp! } \text{id}(\ell) \ell' \rangle \\
\text{stamp! } (\ell \langle \bar{c} \rangle) \ell' = \ell \langle \text{stamp! } \bar{c} \ell' \rangle \\
\boxed{|e| = \ell} \\
|\ell| = \ell \\
|\ell \langle \bar{c} \rangle| = |\bar{c}|
\end{array}$$

Fig. 13. Stamping and security level operators for security label expressions

$$\begin{array}{l}
\boxed{\text{coerce-id } T = c_r \text{ and } \text{coerce-id } A = c} \\
\text{coerce-id } \iota = \text{id}(\iota) \\
\text{coerce-id } (\text{Ref } A) = \mathbf{Ref} (\text{coerce-id } A) (\text{coerce-id } A) \\
\text{coerce-id } (A \xrightarrow{g} B) = (\text{id}(g), \text{coerce-id } A \rightarrow \text{coerce-id } B) \\
\text{coerce-id } T_g = (\text{coerce-id } T, \text{id}(g)) \\
\boxed{\text{stamp } V A \ell = W} \\
\text{stamp } V - \text{low} = V \\
\text{stamp } V T_{\text{low}} \text{high} = V \langle (\text{coerce-id } T), \text{id}(\text{low}); \uparrow \rangle \\
\text{stamp } V T_{\text{high}} \text{high} = V \\
\text{stamp } (V \langle c_r, \bar{c} \rangle) - \ell = V \langle c_r, \text{stamp } \bar{c} \ell \rangle
\end{array}$$

Fig. 14. Stamping on values of λ_{IFC}^c

8.3.2 *Operational Semantics.* The operational semantics for λ_{IFC}^c is defined in Figure 16 and Figure 17. To support the operational semantics of λ_{IFC}^c , we define the following operators for label expressions, and values of λ_{IFC}^c :

- (1) The stamping operations, *stamp* and *stamp!*, for label expressions in normal form are defined in Figure 13.
- (2) The security level operator for label expressions in normal form is also defined in Figure 13.
- (3) The stamping operation for values in λ_{IFC}^c is defined in Figure 14.
- (4) The composition operator for coercions on values is defined in Figure 15.
- (5) The “irreducible” predicate for coercions on values is also defined in Figure 15.

$$\boxed{\mathbf{c} \circledast \mathbf{c} = \mathbf{c}}$$

$$\begin{aligned} &(\mathbf{id}(l), \bar{c}) \circledast (\mathbf{id}(l), \bar{d}) = (\mathbf{id}(l), \bar{c} \circledast \bar{d}) \\ &(\mathbf{Ref} \mathbf{c}_1 \mathbf{c}_2, \bar{c}) \circledast (\mathbf{Ref} \mathbf{d}_1 \mathbf{d}_2, \bar{d}) = (\mathbf{Ref} (\mathbf{d}_1 \circledast \mathbf{c}_1) (\mathbf{c}_2 \circledast \mathbf{d}_2), \bar{c} \circledast \bar{d}) \\ &(\bar{c}_1, \mathbf{c}_1 \rightarrow \mathbf{c}_2, \bar{c}_2) \circledast (\bar{d}_1, \mathbf{d}_1 \rightarrow \mathbf{d}_2, \bar{d}_2) = (\bar{d}_1 \circledast \bar{c}_1, (\mathbf{d}_1 \circledast \mathbf{c}_1) \rightarrow (\mathbf{c}_2 \circledast \mathbf{d}_2), \bar{c}_2 \circledast \bar{d}_2) \end{aligned}$$

$$\boxed{\text{Irreducible } \mathbf{c}}$$

$$\frac{\text{NF } \bar{c} \quad \ell \neq g}{\text{Irreducible } (\mathbf{id}(l), \bar{c})} \vdash \bar{c} : \ell \Rightarrow g \quad \frac{\text{NF } \bar{c}}{\text{Irreducible } (\mathbf{Ref} \mathbf{c} \mathbf{d}, \bar{c})} \quad \frac{\text{NF } \bar{c}}{\text{Irreducible } (\bar{d}, \mathbf{c} \rightarrow \mathbf{d}, \bar{c})}$$

Fig. 15. Composition of coercions on values.

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

$$\begin{array}{c}
\xi \frac{M \mid \mu \mid PC \longrightarrow M' \mid \mu'}{\text{plug } M F \mid \mu \mid PC \longrightarrow \text{plug } M' F \mid \mu'} \quad \xi\text{-blame} \frac{\text{plug } (\text{blame } p) F \mid \mu \mid PC \longrightarrow \text{blame } p \mid \mu}{} \\
\text{prot-ctx} \frac{M \mid \mu \mid PC' \longrightarrow M' \mid \mu'}{\text{prot } PC' \ell M A \mid \mu \mid PC \longrightarrow \text{prot } PC' \ell M' A \mid \mu'} \\
\text{prot-val} \frac{}{\text{prot } PC' \ell V A \mid \mu \mid PC \longrightarrow \text{stamp-val } V A \ell \mid \mu} \\
\text{prot-blame} \frac{}{\text{prot } PC' \ell (\text{blame } p) A \mid \mu \mid PC \longrightarrow \text{blame } p \mid \mu} \quad \text{cast} \frac{V \langle c \rangle \longrightarrow M}{V \langle c \rangle \mid \mu \mid PC \longrightarrow M \mid \mu} \\
\beta \frac{}{\text{app } (\lambda x. N) V A B \ell \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp } PC \ell) \ell (N[x := V]) B \mid \mu} \\
\text{app-cast} \frac{\text{NF } \bar{c} \quad (\text{stamp } PC \ell) \langle \bar{d} \rangle \longrightarrow^* PC' \quad V \langle c \rangle \longrightarrow^* W}{\text{app } (\lambda x. N \langle \bar{d}, c \rightarrow d, \bar{c} \rangle) V C D \ell \mid \mu \mid PC \longrightarrow \text{prot } PC' \ell ((N[x := W]) \langle \bar{d} \rangle) D \mid \mu} \\
\text{app-blame-pc} \frac{\text{NF } \bar{c} \quad (\text{stamp } PC \ell) \langle \bar{d} \rangle \longrightarrow^* \text{blame } p}{\text{app } (\lambda x. N \langle \bar{d}, c \rightarrow d, \bar{c} \rangle) V C D \ell \mid \mu \mid PC \longrightarrow \text{blame } p \mid \mu} \\
\text{app-blame} \frac{\text{NF } \bar{c} \quad (\text{stamp } PC \ell) \langle \bar{d} \rangle \longrightarrow^* PC' \quad V \langle c \rangle \longrightarrow^* \text{blame } p}{\text{app } (\lambda x. N \langle \bar{d}, c \rightarrow d, \bar{c} \rangle) V C D \ell \mid \mu \mid PC \longrightarrow \text{blame } p \mid \mu} \\
\text{app}\star\text{-cast} \frac{\text{NF } \bar{c} \quad (\text{stamp! } PC \mid \bar{c} \mid) \langle \bar{d} \rangle \longrightarrow^* PC' \quad V \langle c \rangle \longrightarrow^* W}{\text{app}\star (\lambda x. N \langle \bar{d}, c \rightarrow d, \bar{c} \rangle) V C T \mid \mu \mid PC \longrightarrow \text{prot } PC' \mid \bar{c} \mid ((N[x := W]) \langle \bar{d} \rangle) (T_\star) \mid \mu} \\
\text{app}\star\text{-blame-pc} \frac{\text{NF } \bar{c} \quad (\text{stamp! } PC \mid \bar{c} \mid) \langle \bar{d} \rangle \longrightarrow^* \text{blame } p}{\text{app}\star (\lambda x. N \langle \bar{d}, c \rightarrow d, \bar{c} \rangle) V C T \mid \mu \mid PC \longrightarrow \text{blame } p \mid \mu} \\
\text{app}\star\text{-blame} \frac{\text{NF } \bar{c} \quad (\text{stamp! } PC \mid \bar{c} \mid) \langle \bar{d} \rangle \longrightarrow^* PC' \quad V \langle c \rangle \longrightarrow^* \text{blame } p}{\text{app}\star (\lambda x. N \langle \bar{d}, c \rightarrow d, \bar{c} \rangle) V C T \mid \mu \mid PC \longrightarrow \text{blame } p \mid \mu} \\
\text{if-true} \frac{}{\text{if } (\$ \text{true}) A \ell M N \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp } PC \ell) \ell M A \mid \mu} \\
\text{if-true-cast} \frac{}{\text{if } (\$ \text{true}) (\text{id}(\text{Bool}), \text{id}(\text{low}); \uparrow) A \text{high} M N \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp } PC \text{high}) \text{high} M A \mid \mu} \\
\text{if}\star\text{-true-cast} \frac{\text{NF } \bar{c}}{\text{if}\star (\$ \text{true}) (\text{id}(\text{Bool}), \bar{c}) T M N \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp! } PC \mid \bar{c} \mid) \mid \bar{c} \mid M (T_\star) \mid \mu} \\
\text{if-false} \frac{}{\text{if } (\$ \text{false}) A \ell M N \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp } PC \ell) \ell N A \mid \mu} \\
\text{if-false-cast} \frac{}{\text{if } (\$ \text{false}) (\text{id}(\text{Bool}), \text{id}(\text{low}); \uparrow) A \text{high} M N \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp } PC \text{high}) \text{high} N A \mid \mu} \\
\text{if}\star\text{-false-cast} \frac{\text{NF } \bar{c}}{\text{if}\star (\$ \text{false}) (\text{id}(\text{Bool}), \bar{c}) T M N \mid \mu \mid PC \longrightarrow \text{prot } (\text{stamp! } PC \mid \bar{c} \mid) \mid \bar{c} \mid N (T_\star) \mid \mu} \\
\text{let} \frac{}{\text{let } x = V : A \text{ in } N \mid \mu \mid PC \longrightarrow N[x := V] \mid \mu}
\end{array}$$

Fig. 16. Operational semantics of λ_{IFC}^c (Part I).

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

$$\begin{array}{c}
\text{ref} \frac{n \text{ FreshIn } \mu(\ell)}{\text{ref } \ell V \mid \mu \mid PC \longrightarrow \text{addr } n \mid (\mu, \ell \mapsto n \mapsto V)} \\
\text{ref?} \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{ref?-blame} \frac{PC \langle \star \Rightarrow^P \ell \rangle \longrightarrow^* \text{blame } q}{\text{ref?}^P \ell V \mid \mu \mid PC \longrightarrow \text{blame } q \mid \mu} \\
\text{assign} \frac{}{\text{assign } (\text{addr } n) V T \hat{\ell} \ell \mid \mu \mid PC \longrightarrow \$\text{unit} \mid [\hat{\ell} \mapsto n \mapsto V] \mu} \\
\text{assign-cast} \frac{\text{NF } \bar{c} \quad \vdash c : T_{\hat{\ell}_2} \Rightarrow S_{\hat{\ell}_1} \quad \vdash d : S_{\hat{\ell}_1} \Rightarrow T_{\hat{\ell}_2} \quad V \langle c \rangle \longrightarrow^* W}{\text{assign } (\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) V T \hat{\ell}_2 \ell \mid \mu \mid PC \longrightarrow \$\text{unit} \mid [\hat{\ell}_1 \mapsto n \mapsto W] \mu} \\
\text{assign-blame} \frac{\text{NF } \bar{c} \quad \vdash c : T_{\hat{\ell}_2} \Rightarrow S_{\hat{\ell}_1} \quad \vdash d : S_{\hat{\ell}_1} \Rightarrow T_{\hat{\ell}_2} \quad V \langle c \rangle \longrightarrow^* \text{blame } p}{\text{assign } (\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) V T \hat{\ell}_2 \ell \mid \mu \mid PC \longrightarrow \text{blame } p \mid \mu} \\
\text{assign?-cast} \frac{\text{NF } \bar{c} \quad \vdash c : T_g \Rightarrow S_{\hat{\ell}} \quad \vdash d : S_{\hat{\ell}} \Rightarrow T_g \quad (\text{stamp! } PC \mid \bar{c} \mid) \langle \star \Rightarrow^P \hat{\ell} \rangle \longrightarrow^* PC' \quad V \langle c \rangle \longrightarrow^* W}{\text{assign?}^P (\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) V T g \mid \mu \mid PC \longrightarrow \$\text{unit} \mid [\hat{\ell} \mapsto n \mapsto W] \mu} \\
\text{assign?-cast-blame-pc} \frac{\text{NF } \bar{c} \quad \vdash c : T_g \Rightarrow S_{\hat{\ell}} \quad \vdash d : S_{\hat{\ell}} \Rightarrow T_g \quad (\text{stamp! } PC \mid \bar{c} \mid) \langle \star \Rightarrow^P \hat{\ell} \rangle \longrightarrow^* \text{blame } q}{\text{assign?}^P (\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) V T g \mid \mu \mid PC \longrightarrow \text{blame } q \mid \mu} \\
\text{assign?-cast-blame} \frac{\text{NF } \bar{c} \quad \vdash c : T_g \Rightarrow S_{\hat{\ell}} \quad \vdash d : S_{\hat{\ell}} \Rightarrow T_g \quad (\text{stamp! } PC \mid \bar{c} \mid) \langle \star \Rightarrow^P \hat{\ell} \rangle \longrightarrow^* PC' \quad V \langle c \rangle \longrightarrow^* \text{blame } q}{\text{assign?}^P (\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) V T g \mid \mu \mid PC \longrightarrow \text{blame } q \mid \mu} \\
\text{deref} \frac{\mu(\hat{\ell}, n) = V}{!(\text{addr } n) T_{\hat{\ell}} \ell \mid \mu \mid PC \longrightarrow \text{prot_} \ell V T_{\hat{\ell}} \mid \mu} \\
\text{deref-cast} \frac{\text{NF } \bar{c} \quad \vdash c : A \Rightarrow T_{\hat{\ell}} \quad \vdash d : T_{\hat{\ell}} \Rightarrow A \quad \mu(\hat{\ell}, n) = V}{!(\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) A \ell \mid \mu \mid PC \longrightarrow \text{prot_} \ell V \langle d \rangle A \mid \mu} \\
\text{deref}\star\text{-cast} \frac{\text{NF } \bar{c} \quad \vdash c : S_{\star} \Rightarrow T_{\hat{\ell}} \quad \vdash d : T_{\hat{\ell}} \Rightarrow S_{\star} \quad \mu(\hat{\ell}, n) = V}{!\star (\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) S \mid \mu \mid PC \longrightarrow \text{prot_} \mid \bar{c} \mid (V \langle d \rangle) (S_{\star}) \mid \mu}
\end{array}$$

Fig. 17. Operational semantics of λ_{IFC}^c (Part II). NSU checks are represented using label expressions (highlighted)

8.4 Definitions of Precision

To state and prove the gradual guarantee, we define the following precision relations:

- (1) The precision relations for security labels and types are defined in Figure 18.
- (2) The precision relation on terms of λ_{IFC}^* is defined in Figure 19.
- (3) The precision relation on the coercion calculus is defined in Figure 20.
- (4) We present noteworthy term precision rules of λ_{IFC}^c in Figure 21, with the full definition in `/src/CC2/Precision.agda` of the Agda code in the supplementary material.

$$\begin{array}{c}
 \boxed{g_1 \sqsubseteq g_2} \\
 \\
 \star \sqsubseteq \frac{}{\star \sqsubseteq g} \qquad \ell \sqsubseteq \ell \frac{}{\ell \sqsubseteq \ell} \\
 \\
 \boxed{S \sqsubseteq T} \\
 \\
 \sqsubseteq\text{-}\iota \frac{}{\iota \sqsubseteq \iota} \quad \sqsubseteq\text{-}\text{ref} \frac{A \sqsubseteq B}{\text{Ref } A \sqsubseteq \text{Ref } B} \quad \sqsubseteq\text{-}\text{fun} \frac{g_1 \sqsubseteq g_2 \quad A \sqsubseteq C \quad B \sqsubseteq D}{A \xrightarrow{g_1} B \sqsubseteq C \xrightarrow{g_2} D} \\
 \\
 \boxed{A \sqsubseteq B} \\
 \\
 \sqsubseteq\text{-}\tau \frac{g_1 \sqsubseteq g_2 \quad S \sqsubseteq T}{S_{g_1} \sqsubseteq T_{g_2}}
 \end{array}$$

Fig. 18. Precision of security labels and types

$$\begin{array}{c}
 \boxed{\vdash M \sqsubseteq M'} \\
 \\
 \frac{}{\vdash (\$ k)_\ell \sqsubseteq (\$ k)_\ell} \qquad \frac{g_1 \sqsubseteq g_2 \quad A_1 \sqsubseteq A_2 \quad \vdash N_1 \sqsubseteq N_2}{\vdash (\lambda^{g_1} x:A_1. N_1)_\ell \sqsubseteq (\lambda^{g_2} x:A_2. N_2)_\ell} \\
 \\
 \frac{\vdash M_1 \sqsubseteq M_2}{\vdash (\text{ref } \ell M_1)^P \sqsubseteq (\text{ref } \ell M_2)^P} \qquad \frac{\vdash M_1 \sqsubseteq M_2 \quad A_1 \sqsubseteq A_2}{\vdash (M_1 : A_1)^P \sqsubseteq (M_2 : A_2)^P}
 \end{array}$$

Fig. 19. Selected precision rules of λ_{IFC}^*

$$\boxed{\vdash c \sqsubseteq d, \vdash c \sqsubseteq g, \text{ and } \vdash g \sqsubseteq d}$$

$$\begin{array}{c} \sqsubseteq\text{-c} \frac{g_1 \sqsubseteq g'_1 \quad g_2 \sqsubseteq g'_2 \quad \vdash c : g_1 \Rightarrow g_2 \quad \vdash d : g'_1 \Rightarrow g'_2}{\vdash c \sqsubseteq d} \\ \sqsubseteq\text{-cl} \frac{g_1 \sqsubseteq g \quad g_2 \sqsubseteq g \quad \vdash c : g_1 \Rightarrow g_2}{\vdash c \sqsubseteq g} \quad \sqsubseteq\text{-cr} \frac{g \sqsubseteq g_1 \quad g \sqsubseteq g_2 \quad \vdash d : g_1 \Rightarrow g_2}{\vdash g \sqsubseteq d} \end{array}$$

$$\boxed{\vdash \bar{c} \sqsubseteq \bar{d}}$$

$$\begin{array}{c} \sqsubseteq\text{-id} \frac{g \sqsubseteq g'}{\vdash \text{id}(g) \sqsubseteq \text{id}(g')} \quad \sqsubseteq\text{-cast} \frac{\vdash \bar{c} \sqsubseteq \bar{d} \quad \vdash c \sqsubseteq d}{\vdash \bar{c}; c \sqsubseteq \bar{d}; d} \\ \sqsubseteq\text{-castl} \frac{\vdash \bar{c} \sqsubseteq \bar{d} \quad \vdash c \sqsubseteq g_2 \quad \vdash \bar{d} : g_1 \Rightarrow g_2}{\vdash \bar{c}; c \sqsubseteq \bar{d}} \quad \sqsubseteq\text{-castr} \frac{\vdash \bar{c} \sqsubseteq \bar{d} \quad \vdash g_2 \sqsubseteq d \quad \vdash \bar{c} : g_1 \Rightarrow g_2}{\vdash \bar{c} \sqsubseteq \bar{d}; d} \\ \sqsubseteq\text{-}\perp \frac{g_1 \sqsubseteq g_3 \quad g_2 \sqsubseteq g_4 \quad \vdash \bar{c} : g_1 \Rightarrow g_2}{\vdash \bar{c} \sqsubseteq \perp^P g_3 g_4} \end{array}$$

$$\boxed{\vdash \bar{c} \sqsubseteq_l g}$$

$$\sqsubseteq_l\text{-id} \frac{g \sqsubseteq g'}{\vdash \text{id}(g) \sqsubseteq_l g'} \quad \sqsubseteq_l\text{-cast} \frac{\vdash \bar{c} \sqsubseteq_l g \quad \vdash c \sqsubseteq_l g}{\vdash \bar{c}; c \sqsubseteq_l g}$$

$$\boxed{\vdash g \sqsubseteq_r \bar{d}}$$

$$\sqsubseteq_r\text{-id} \frac{g \sqsubseteq g'}{\vdash g \sqsubseteq_r \text{id}(g')} \quad \sqsubseteq_r\text{-cast} \frac{\vdash g \sqsubseteq_r \bar{d} \quad \vdash g \sqsubseteq_r d}{\vdash g \sqsubseteq_r \bar{d}; d} \quad \sqsubseteq_r\text{-}\perp \frac{g \sqsubseteq g_1 \quad g \sqsubseteq g_2}{\vdash g \sqsubseteq_r \perp^P g_1 g_2}$$

Fig. 20. Precision relation of the coercion calculus

$$\boxed{\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \sqsubseteq M' \Leftarrow A \sqsubseteq A'}$$

$$\begin{array}{c}
\sqsubseteq\text{-addr} \frac{\Sigma(\hat{\ell}, n) = T \quad \Sigma'(\hat{\ell}, n) = T'}{\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash \text{addr } n \sqsubseteq \text{addr } n \Leftarrow (\text{Ref } T_{\hat{\ell}})_{\ell} \sqsubseteq (\text{Ref } T'_{\hat{\ell}})_{\ell}} \\
\sqsubseteq\text{-ref?l} \frac{\Gamma; \Gamma'; \Sigma; \Sigma'; \star; \ell_1; \ell_2; \ell_3 \vdash M \sqsubseteq M' \Leftarrow T_{\ell} \sqsubseteq T'_{\ell} \quad \ell_1 \preceq \ell}{\Gamma; \Gamma'; \Sigma; \Sigma'; \star; \ell_1; \ell_2; \ell_3 \vdash \text{ref?}^P \ell M \sqsubseteq \text{ref } \ell M' \Leftarrow (\text{Ref } T_{\ell})_{\text{low}} \sqsubseteq (\text{Ref } T'_{\ell})_{\text{low}}} \\
\sqsubseteq\text{-prot!} \frac{\Gamma; \Gamma'; \Sigma; \Sigma'; g_1; g'_1; |PC|; |PC'| \vdash M \sqsubseteq M' \Leftarrow T_{\star} \sqsubseteq T'_{\star} \quad \vdash PC \sqsubseteq PC' \Leftarrow g_1 \sqsubseteq g'_1 \\ \ell_1 \vee \ell_2 \preceq |PC| \quad \ell'_1 \vee \ell'_2 \preceq |PC'| \quad \ell_2 \preceq \ell'_2}{\Gamma; \Gamma'; \Sigma; \Sigma'; g_2; g'_2; \ell_1; \ell'_1 \vdash \text{prot } PC \ell_2 M T_{\star} \sqsubseteq \text{prot } PC' \ell'_2 M' T'_{\star} \Leftarrow T_{\star} \sqsubseteq T'_{\star}} \\
\sqsubseteq\text{-cast} \frac{\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \sqsubseteq M' \Leftarrow A \sqsubseteq A' \quad c \sqsubseteq c' \quad \vdash c : A \Rightarrow B \quad \vdash c' : A' \Rightarrow B'}{\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \langle c \rangle \sqsubseteq M' \langle c' \rangle \Leftarrow B \sqsubseteq B'} \\
\sqsubseteq\text{-castl} \frac{\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \sqsubseteq M' \Leftarrow A \sqsubseteq A' \quad c \sqsubseteq A' \quad \vdash c : A \Rightarrow B}{\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \langle c \rangle \sqsubseteq M' \Leftarrow B \sqsubseteq A'} \\
\sqsubseteq\text{-castr} \frac{\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \sqsubseteq M' \Leftarrow A \sqsubseteq A' \quad A \sqsubseteq c' \quad \vdash c' : A' \Rightarrow B'}{\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \sqsubseteq M' \langle c' \rangle \Leftarrow A \sqsubseteq B'} \\
\sqsubseteq\text{-blame} \frac{\Gamma; \Sigma; g; \ell \vdash M \Leftarrow A \quad A \sqsubseteq A'}{\Gamma; \Gamma'; \Sigma; \Sigma'; g; g'; \ell; \ell' \vdash M \sqsubseteq \text{blame } p \Leftarrow A \sqsubseteq A'}
\end{array}$$

Fig. 21. Selected precision rules of λ_{FC}^c

9 VIGILANCE ALONE DOES NOT ENSURE TYPE-BASED REASONING

Chen and Siek [2022] study a language λ_{SEC}^* with gradual security that omits type-guided classification, similar to GLIO. However, λ_{SEC}^* differs from GLIO in that it is vigilant like traditional gradually-typed languages. Here we show that vigilance alone is not enough to ensure type-based reasoning, as the omission of type-guided classification in λ_{SEC}^* breaks type-based reasoning.

Returning to the `mix`/`smix` example by Toro et al. [2018] (§ 2.3), the cast insertion of λ_{SEC}^* produces:

```

1 let mix =
2   (λ pub priv . if (pub ⟨low ⇒ ★⟩) < priv then (1low ⟨low ⇒ ★⟩) else (2low ⟨low ⇒ ★⟩))
3   ⟨low → ★ → ★ ⇒ low → ★ → low⟩ in
4 let smix = λ pub priv . mix pub (priv ⟨high ⇒ ★⟩) in
5 smix 1low 5low

```

According to type-based reasoning, this program should fail at runtime. However, the program reduces to 1_{low} , violating the free theorem. Let us examine the reduction steps involving the if-expression that give rise to an implicit flow.

$$\longrightarrow^* \text{ (if (true}_{\text{low}} \langle \text{high} \Rightarrow \star \rangle) \text{ then } 1_{\text{low}} \langle \text{low} \Rightarrow \star \rangle \text{ else } \dots) \langle \star \Rightarrow \text{low} \rangle} \quad (1)$$

$$\longrightarrow^* \text{ (prot } \text{low} (1_{\text{low}} \langle \text{low} \Rightarrow \star \rangle)) \langle \star \Rightarrow \text{low} \rangle} \quad (2)$$

$$\longrightarrow^* 1_{\text{low}} \quad (3)$$

The if-expression branches on the value $\text{true}_{\text{low}} \langle \text{high} \Rightarrow \star \rangle$ (1) and this value is classified as `low` security because in λ_{SEC}^* , casts do not change the security classification of values. So the protection term gets security level `low` (2), and therefore $1_{\text{low}} \langle \text{low} \Rightarrow \star \rangle$ is allowed as the result of the protection term. Then the injection and projection collapse, producing 1_{low} (3).

10 THE NORMALIZATION OF COERCIONS CHECKS INFORMATION FLOW

We show that reducing coercion sequences to normal form models IFC checks because the normalization either succeeds or fails: if a coercion sequences successfully reduces to normal form, the IFC check succeeds and the flow is justified; if it reduces to a failure, then an illegal flow is detected and the program errors.

LEMMA 18 (STRONG NORMALIZATION OF COERCION SEQUENCES). *If $\vdash \bar{c} : g_1 \Rightarrow g_2$, then either (1) $\bar{c} \longrightarrow^* \bar{d}$ and $\text{NF } \bar{d}$ or (2) $\bar{c} \longrightarrow^* \perp^P g_1 g_2$.*

Normalization of coercion sequences is deterministic:

LEMMA 19 (REDUCTION OF COERCION SEQUENCES IS DETERMINISTIC). *If $\bar{c} \longrightarrow \bar{d}_1$ and $\bar{c} \longrightarrow \bar{d}_2$, then $\bar{d}_1 = \bar{d}_2$.*

LEMMA 20 (NORMALIZATION OF COERCION SEQUENCES IS DETERMINISTIC). *Suppose $\bar{c} \longrightarrow^* \bar{d}_1$ and $\bar{c} \longrightarrow^* \bar{d}_2$. If $\text{NF } \bar{d}_i$ or $\bar{d}_i = \perp^P g_1 g_2$, then $\bar{d}_1 = \bar{d}_2$.*

11 PROVING NONINTERFERENCE FOR λ_{IFC}^*

In this section we present a proof of noninterference for λ_{IFC}^* . The proof employs a three-step approach. First, we adapt an IFC language and noninterference proof from Chen and Siek [2022]. In particular, we define a dynamic IFC programming language named λ_{SEC} (Figure 22) and prove that it satisfies noninterference. This proof (Lemma 21) uses a standard erasure-based approach [Fennell and Thiemann 2013; Li and Zdancewic 2010; Stefan et al. 2017, 2011, 2012], where the high-security parts of a program are erased to an opaque value.

terms	L, M, N	$::=$	$x \mid (\$ k)_\ell \mid (\text{addr } n_\ell)_\ell \mid (\lambda x. N)_\ell$ $\mid L M \mid \text{if } L M N$ $\mid \text{ref}^? \ell M \mid ! M \mid L :=^? M \mid \text{prot } \ell M$
values	V, W	$::=$	$(\$ k)_\ell \mid (\text{addr } n_\ell)_\ell \mid (\lambda x. N)_\ell$
frames	F	$::=$	$\square M \mid V \square \mid \text{if } \square M N \mid \text{ref}^? \ell \square \mid ! \square \mid \square :=^? M \mid V :=^? \square$

$M \mid \mu \mid pc \longrightarrow N \mid \mu'$

$$\xi \frac{M \mid \mu \mid pc \longrightarrow M' \mid \mu'}{\text{plug } M F \mid \mu \mid pc \longrightarrow \text{plug } M' F \mid \mu'}$$

$$\text{prot-val} \frac{}{\text{prot } \ell V \mid \mu \mid pc \longrightarrow V \vee \ell \mid \mu} \quad \text{prot-ctx} \frac{M \mid \mu \mid pc \vee \ell \longrightarrow M' \mid \mu'}{\text{prot } \ell M \mid \mu \mid pc \longrightarrow \text{prot } \ell M' \mid \mu'}$$

$$\beta \frac{}{(\lambda x. N)_\ell V \mid \mu \mid pc \longrightarrow \text{prot } \ell (N[x := V]) \mid \mu}$$

$$\beta\text{-if-true} \frac{}{\text{if } (\$ \text{true})_\ell M N \mid \mu \mid pc \longrightarrow \text{prot } \ell M \mid \mu}$$

$$\beta\text{-if-false} \frac{}{\text{if } (\$ \text{false})_\ell M N \mid \mu \mid pc \longrightarrow \text{prot } \ell N \mid \mu}$$

$$\text{ref}^?\text{-ok} \frac{\text{pc} \preceq \ell \quad n \text{ FreshIn } \mu(\ell)}{\text{ref}^? \ell V \mid \mu \mid pc \longrightarrow (\text{addr } n_\ell)_{\text{low}} \mid (\mu, \ell \mapsto n \mapsto (V \vee \ell))}$$

$$\text{deref} \frac{\mu(\hat{\ell}, n) = V}{!(\text{addr } n_\ell)_\ell \mid \mu \mid pc \longrightarrow \text{prot } \ell V \mid \mu}$$

$$\text{assign}^?\text{-ok} \frac{\text{pc} \vee \ell \preceq \hat{\ell}}{(\text{addr } n_\ell)_\ell :=^? V \mid \mu \mid pc \longrightarrow (\$ \text{unit})_{\text{low}} \mid [\hat{\ell} \mapsto n \mapsto (V \vee \hat{\ell})] \mu}$$

Fig. 22. Syntax and reduction semantics (successful cases) of λ_{SEC}

Second, we prove a simulation lemma between λ_{IFC}^c and λ_{SEC} (Lemma 34). The main intuition of the simulation relation (Figure 23) is that a λ_{IFC}^c term always produces a value that is as secure as the one produced by its related λ_{SEC} term. We translate λ_{IFC}^c terms to λ_{SEC} by (1) getting rid of all the casts and (2) converting static heap enforcement (ref and assign) to dynamic enforcement (NSU) (Figure 24). The noninterference property of λ_{IFC}^c (Lemma 14) follows directly from the multi-step simulation lemma (Lemma 35) and the noninterference result of λ_{SEC} (Lemma 21).

Third, the noninterference theorem of λ_{IFC}^* (Theorem 16) is a corollary of the noninterference property of λ_{IFC}^c .

Similar to GSL_{Ref} and GLIO , the statements of noninterference for both λ_{SEC} and λ_{IFC}^c are termination-insensitive. Thus, we only consider successful executions that produce values. In other words, we do not consider reduction rules that trigger IFC monitor failures, such as NSU errors in λ_{SEC} and cast errors in λ_{IFC}^c . As we explained in Section 2.1, the programming language runtime can force the program to diverge whenever blame is detected, possibly sending a private error message to the software developer.

LEMMA 21 (NONINTERFERENCE OF λ_{SEC}). *If $M[x := (\$ b_1)_{\text{high}}] \mid \emptyset \mid \text{low} \longrightarrow^* (\$ b_3)_{\text{low}} \mid \mu_1$ and $M[x := (\$ b_2)_{\text{high}}] \mid \emptyset \mid \text{low} \longrightarrow^* (\$ b_4)_{\text{low}} \mid \mu_2$ then $b_3 = b_4$.*

PROOF. The proof is fully mechanized in Agda ⁴ and can be found in the supplementary material. The proof follows the noninterference proof of $\lambda_{\text{SEC}}^{\Rightarrow}$ [Chen and Siek 2022]. \square

DEFINITION 22 (SIMULATION BETWEEN HEAPS). $\Sigma \vdash \mu' \leq \mu \triangleq \forall \ell, n. \text{if } \mu(\ell, n) = V, \text{ then there exists } V' \text{ s.t } \mu'(\ell, n) = V' \text{ and } \text{low} \vdash V' \leq V \Leftarrow T_\ell \text{ where } T = \Sigma(\ell, n).$

LEMMA 23 (CASTING PRESERVES SIMULATION). *If $g^c \vdash W' \leq V \Leftarrow A, \vdash c : A \Rightarrow B,$ and $V \langle c \rangle \longrightarrow^* W,$ then $g^c \vdash W' \leq W \Leftarrow B.$*

PROOF. By induction on the multi-step cast reduction.

Zero step Directly proved by applying rule $\leq\text{-cast}$.

One or more steps Casing on the first reduction step yields three sub-cases:

cast

$$g^c \vdash W' \leq V_r \Leftarrow A \tag{4}$$

$$V_r \langle c_r, \bar{c} \rangle \longrightarrow V_r \langle c_r, \bar{d} \rangle \longrightarrow^* W \tag{5}$$

By induction hypothesis, $g^c \vdash W' \leq W \Leftarrow B.$

cast-id

$$g^c \vdash W' \leq V_r \Leftarrow A \tag{6}$$

$$V_r \langle \text{id}(\iota), \text{id}(g) \rangle \longrightarrow V_r \tag{7}$$

The goal is proved directly.

cast-comp

$$g^c \vdash W' \leq V_r \langle c \rangle \Leftarrow A \tag{8}$$

$$V_r \langle c \rangle \langle d \rangle \longrightarrow V_r \langle c \ ; \ d \rangle \longrightarrow^* W \tag{9}$$

We further reason about $V_r \langle c \ ; \ d \rangle \longrightarrow^* W.$ The coercion after composition $c \ ; \ d$ is reducible and not identity, so the reduction must take one step by *cast* and reduce the top-level coercion sequence to its normal form \bar{c}_n :

$$V_r \langle c_{r1}, \bar{c} \rangle \langle c_{r2}, \bar{d} \rangle \longrightarrow V_r \langle c_r, \bar{c} \ ; \ \bar{d} \rangle \longrightarrow V_r \langle c_r, \bar{c}_n \rangle \longrightarrow^* W$$

We know $|\bar{c}| \preceq |\bar{c}_n|$ because composition models explicit flow (Lemma 2).

- If $V_r = \$k$ and $\bar{c}_n = \text{id}(\ell).$ $W' = (\$k)_{\ell'}$ and $\ell' \preceq |\bar{c}|$ (Lemma 25). We know $g^c \vdash (\$k)_{\ell'} \leq \$k \Leftarrow \iota_\ell$ by rule $\leq\text{-const}$, because $\ell' \preceq |\bar{c}| \preceq |\bar{c}_n| = |\text{id}(\ell)| = \ell.$
- Otherwise, $V_r \langle c_r, \bar{c}_n \rangle$ is already a value. Apply rule $\leq\text{-wrapped-const}$, $\leq\text{-wrapped-lam}$, or $\leq\text{-wrapped-addr}$ depending on whether V_r is a constant, a λ , or an address.

\square

LEMMA 24 (SUBSTITUTION PRESERVES SIMULATION). *If $g \vdash N' \leq N \Leftarrow A, (\Gamma, x : B); \Sigma; g; \ell \vdash N \Leftarrow A,$ and $\forall g. g \vdash M' \leq M \Leftarrow B,$ then $g \vdash N'[x := M'] \leq N[x := M] \Leftarrow A.$*

PROOF. The proof is fully mechanized in Agda ⁵ and can be found in the supplementary material. \square

LEMMA 25 (SIMULATION WITH WRAPPED CONSTANT). *If $g^c \vdash M \leq (\$k) \langle \text{id}(\iota), \bar{c} \rangle \Leftarrow \iota_g$ and $(\text{id}(\iota), \bar{c})$ is irreducible, then there exists ℓ s.t $M = (\$k)_\ell,$ and $\ell \preceq |\bar{c}|.$*

PROOF. Inversion on the simulation relation:

⁴In /src/Dyn/Noninterference.agda.

⁵In substitution-pres- \leq of /src/Security/SubstPres.agda. The simulation relation is in /src/Security/SimRel.agda.

$$\boxed{g^c \vdash M \leq N \Leftarrow A}$$

$$\begin{array}{c}
\leq\text{-var} \frac{}{g^c \vdash x \leq x \Leftarrow A} \quad \leq\text{-const} \frac{\ell' \preceq \ell}{g^c \vdash (\$ k)_{\ell'} \leq \$ k \Leftarrow \iota_\ell} \\
\leq\text{-wrapped-const} \frac{\ell' \preceq |\bar{c}| \quad \vdash \bar{c} : \ell \Rightarrow g \quad \text{NF } \bar{c} \quad \ell \neq g}{g^c \vdash (\$ k)_{\ell'} \leq \$ k \langle \text{id}(\iota), \bar{c} \rangle \Leftarrow \iota_g} \\
\leq\text{-lam} \frac{g \vdash N' \leq N \Leftarrow B \quad \ell' \preceq \ell}{g^c \vdash (\lambda x. N')_{\ell'} \leq \lambda x. N \Leftarrow (A \xrightarrow{g} B)_\ell} \\
\leq\text{-wrapped-lam} \frac{g_3 \vdash N' \leq N \Leftarrow D \quad \ell \preceq |\bar{c}| \quad \vdash \bar{d} : g_1 \Rightarrow g_3 \quad \vdash c : A \Rightarrow C \quad \vdash d : D \Rightarrow B \quad \text{NF } \bar{c}}{g^c \vdash (\lambda x. N')_\ell \leq (\lambda x. N) \langle \bar{d}, c \rightarrow d, \bar{c} \rangle \Leftarrow (A \xrightarrow{g_1} B)_{g_2}} \\
\leq\text{-addr} \frac{\ell' \preceq \ell}{g^c \vdash (\text{addr } n_{\hat{\ell}})_{\ell'} \leq \text{addr } n \Leftarrow \text{Ref } (T_{\hat{\ell}})_\ell} \\
\leq\text{-wrapped-addr} \frac{\ell \preceq |\bar{c}| \quad \vdash c : T_{g_1} \Rightarrow S_{\hat{\ell}} \quad \vdash d : S_{\hat{\ell}} \Rightarrow T_{g_1} \quad \text{NF } \bar{c}}{g^c \vdash (\text{addr } n_{\hat{\ell}})_\ell \leq (\text{addr } n) \langle \text{Ref } c \ d, \bar{c} \rangle \Leftarrow (\text{Ref } (T_{g_1}))_{g_2}} \\
\leq\text{-app} \frac{\ell^c \vdash M' \leq M \Leftarrow (A \xrightarrow{\ell^c \vee \ell} B)_\ell \quad \ell^c \vdash N' \leq N \Leftarrow A}{\ell^c \vdash M' N' \leq \text{app } M N A B \ell \Leftarrow C} \\
\leq\text{-app } \star \frac{g^c \vdash M' \leq M \Leftarrow (A \xrightarrow{\star} T_\star)_\star \quad g^c \vdash N' \leq N \Leftarrow A}{g^c \vdash M' N' \leq \text{app} \star M N A T \Leftarrow T_\star} \\
\leq\text{-if} \frac{\ell^c \vdash L' \leq L \Leftarrow \text{Bool}_\ell \quad \ell^c \vee \ell \vdash M' \leq M \Leftarrow A \quad \ell^c \vee \ell \vdash N' \leq N \Leftarrow A}{\ell^c \vdash \text{if } L' M' N' \leq \text{if } L A \ell M N \Leftarrow B} \\
\leq\text{-if } \star \frac{g^c \vdash L' \leq L \Leftarrow \text{Bool}_\star \quad \star \vdash M' \leq M \Leftarrow T_\star \quad \star \vdash N' \leq N \Leftarrow T_\star}{g^c \vdash \text{if } L' M' N' \leq \text{if} \star L T M N \Leftarrow T_\star} \\
\leq\text{-ref} \frac{\ell^c \vdash M' \leq M \Leftarrow T_\ell}{\ell^c \vdash \text{ref}^? \ell M' \leq \text{ref } \ell M \Leftarrow (\text{Ref } T_\ell)_{\text{low}}} \quad \leq\text{-ref}^? \frac{\star \vdash M' \leq M \Leftarrow T_\ell}{\star \vdash \text{ref}^? \ell M' \leq \text{ref}^? \ell M \Leftarrow (\text{Ref } T_\ell)_{\text{low}}} \\
\leq\text{-deref} \frac{g^c \vdash M' \leq M \Leftarrow (\text{Ref } A)_\ell}{g^c \vdash ! M' \leq ! M A \ell \Leftarrow B} \quad \leq\text{-deref} \star \frac{g^c \vdash M' \leq M \Leftarrow (\text{Ref } T_\star)_\star}{g^c \vdash ! M' \leq ! \star M T \Leftarrow T_\star} \\
\leq\text{-assign} \frac{\ell^c \vdash L' \leq L \Leftarrow (\text{Ref } T_{\hat{\ell}})_\ell \quad \ell^c \vdash M' \leq M \Leftarrow T_{\hat{\ell}}}{\ell^c \vdash L' :=? M' \leq \text{assign } L M T \hat{\ell} \ell \Leftarrow \text{Unit}_{\text{low}}} \\
\leq\text{-assign}^? \frac{g^c \vdash L' \leq L \Leftarrow (\text{Ref } T_g)_\star \quad g^c \vdash M' \leq M \Leftarrow T_g}{g^c \vdash L' :=? M' \leq \text{assign}^? L M T g \Leftarrow \text{Unit}_{\text{low}}} \\
\leq\text{-cast} \frac{g^c \vdash M \leq N \Leftarrow A \quad \vdash c : A \Rightarrow B}{g^c \vdash M \leq N \langle c \rangle \Leftarrow B} \\
\leq\text{-prot} \frac{\ell' \preceq \ell \quad g_2 \vdash M' \leq M \Leftarrow A \quad \vdash PC \Leftarrow g_2}{g_1 \vdash \text{prot } \ell' M' \leq \text{prot } PC \ell M A \Leftarrow B}
\end{array}$$

Fig. 23. Simulation relation between λ_{IFC}^c and λ_{SEC}

Related by \leq -wrapped-const The goal is proved directly.

Related by \leq -cast We know $g^c \vdash M \leq \$k \Leftarrow \iota_\ell$ and $\bar{c} : \ell \Rightarrow g$. Note that $(\text{id}(\iota), \bar{c})$ is irreducible, so \bar{c} can be \uparrow , $\ell!$, or \uparrow ; **high!**, so $\ell \preceq |\bar{c}|$. By rule \leq -const, we know $M = (\$k)_{\ell'}$ for some ℓ' and $\ell' \preceq \ell$. Thus $\ell' \preceq \ell \preceq |\bar{c}|$. □

LEMMA 26 (SIMULATION WITH REFERENCE PROXY). *If $g^c \vdash M \leq (\text{addr } n) \langle \text{Ref } c \ d, \bar{c} \rangle \Leftarrow (\text{Ref } T_{g_1})_{g_2}$, $\vdash c : T_{g_1} \Rightarrow S_{\hat{c}}, \vdash d : S_{\hat{c}} \Rightarrow T_{g_1}$, and NF \bar{c} , then there exists ℓ s.t. $M = (\text{addr } n_\ell)_\ell$, and $\ell \preceq |\bar{c}|$.*

PROOF. Analogous to Lemma 25. The only extra case to consider is $\bar{c} = \text{id}(\ell)$, which also satisfies $\ell \preceq |\bar{c}|$. □

LEMMA 27 (SIMULATION WITH FUNCTION PROXY). *If $g^c \vdash M \leq (\lambda x. N) \langle \bar{d}, c \rightarrow d, \bar{c} \rangle \Leftarrow (A \xrightarrow{g_1} B)_{g_2}$, $\vdash \bar{d} : g_1 \Rightarrow g_3$, and NF \bar{c} , then there exists N', ℓ s.t. $M = (\lambda x. N')_\ell$, $g_3 \vdash N' \leq N \Leftarrow B$, and $\ell \preceq |\bar{c}|$.*

PROOF. Analogous to Lemma 25 and Lemma 26. □

LEMMA 28 (SIMULATION WITH λ_{IFC}^c VALUE). *If $g^c \vdash M \leq V \Leftarrow A$, then M is a value.*

PROOF. Inversion on the value V and the simulation relation $g^c \vdash M \leq V \Leftarrow A$ yields 6 cases. We consider the two cases for constants; the cases for λ s and addresses are analogous.

Related by \leq -const We know $M = (\$k)_\ell$ for some k, ℓ , so M is a value.

Related by \leq -wrapped-const By Lemma 25, $M = (\$k)_\ell$ for some k, ℓ , so M is a value. □

LEMMA 29 (STAMPING PRESERVES SIMULATION). *If $g^c \vdash V \leq W \Leftarrow A$ and $\ell' \preceq \ell$, then $g^c \vdash V \vee \ell' \leq \text{stamp } W \ \ell \Leftarrow \text{stamp } A \ \ell$.*

PROOF. Casing on $\ell' \preceq \ell$:

Case 1 $\ell' = \text{low}, \ell = \text{low}$: Straightforward because stamping low returns the same value.

Case 2 $\ell' = \text{low}, \ell = \text{high}$: Casing on value W and the simulation relation $g^c \vdash V \leq W \Leftarrow A$ produces 6 sub-cases. We consider the two cases for constants below (corresponding to \leq -const and \leq -wrapped-const), because the cases for λ s and addresses are analogous:

- $g^c \vdash (\$k)_{\ell'_1} \leq \$k \Leftarrow \iota_{\ell'_1}$
 - If $\ell'_1 = \text{low}$, $g^c \vdash (\$k)_{\ell'_1} \leq \$k \langle \text{id}(\iota), \uparrow \rangle \Leftarrow \iota_{\text{high}}$ because $\ell'_1 \preceq |\uparrow| = \text{high}$ (rule \leq -wrapped-const).
 - If $\ell'_1 = \text{high}$, $g^c \vdash (\$k)_{\ell'_1} \leq \$k \Leftarrow \iota_{\text{high}}$ because $\ell'_1 \preceq \text{high}$ (rule \leq -const).
- $g^c \vdash (\$k)_{\ell'_1} \leq \$k \langle \text{id}(\iota), \bar{c} \rangle \Leftarrow -$
 - $g^c \vdash (\$k)_{\ell'_1} \leq k \langle \text{id}(\iota), \text{stamp } \bar{c} \ \text{high} \rangle \Leftarrow -$ because $\ell'_1 \preceq |\text{stamp } \bar{c} \ \text{high}| = |\bar{c}| \vee \text{high} = \text{high}$ (Lemma 3 (Stamping models implicit flow) and rule \leq -wrapped-const).

Case 3 $\ell' = \text{high}, \ell = \text{high}$: Analogous to Case 2; the only difference is that instead of $(\$k)_{\ell'_1}$, the left side is always $(\$k)_{\text{high}}$ because it is stamped with $\ell' = \text{high}$. □

LEMMA 30 (CASTING A LABEL EXPRESSION MODELS EXPLICIT FLOW). *If NF e_1 and $e_1 \langle \bar{c} \rangle \longrightarrow^* e_2$ and NF e_2 , $|e_1| \preceq |e_2|$.*

PROOF. The proof is fully mechanized in Agda ⁶ and can be found in the supplementary material. The proof is by inversion on the multi-step reduction. □

⁶In cast-security of /src/LabelExpr/Security.agda.

LEMMA 31 (STAMPING A LABEL EXPRESSION MODELS IMPLICIT FLOW). *If* $\text{NF } e, |\text{stamp } e \ell| = |e| \vee \ell$ and $|\text{stamp! } e \ell| = |e| \vee \ell$.

PROOF. The proof is fully mechanized in Agda⁷ and can be found in the supplementary material. The proof is by casing on $\text{NF } e$. \square

LEMMA 32. *Suppose* $\Gamma; \Sigma; g; \ell \vdash V \Leftarrow S_{\ell_1}$ and $\vdash c : S_{\ell_1} \Rightarrow T_{\ell_2}$. *If* $V \langle c \rangle \longrightarrow^* W$, then $\ell_1 \preceq \ell_2$.

PROOF. By induction on the multi-step reduction.

Zero step $c = (c_r, \bar{c})$ is irreducible, so $\bar{c} : \ell_1 \Rightarrow \ell_2$ is in its normal form. Thus $\ell_1 \preceq \ell_2$.

One or more steps Case on the first reduction step:

- Rule *cast*: $V_r \langle c_r, \bar{c} \rangle \longrightarrow V_r \langle c_r, \bar{d} \rangle \longrightarrow^* W$. By preservation of coercion sequences, $\bar{d} : \ell_1 \Rightarrow \ell_2$. By induction hypothesis, $\ell_1 \preceq \ell_2$.
- Rule *cast-id*: $V_r \langle \text{id}(t), \text{id}(g) \rangle \longrightarrow V_r$. We directly know $\ell_1 = g = \ell_2$.
- Rule *cast-comp*: $V_r \langle c \rangle \langle d \rangle \longrightarrow V_r \langle c \circledast d \rangle \longrightarrow^* W$. Suppose $c = (-, \bar{c})$, $\bar{c} : \ell_0 \Rightarrow \ell_1$ and $d = (-, \bar{d})$, $\bar{d} : \ell_1 \Rightarrow \ell_2$. We know $V_r \langle -, \bar{c} \circledast \bar{d} \rangle \longrightarrow V_r \langle \bar{c}_n \rangle \longrightarrow^* W$. By Lemma 2 (Composition models explicit flow), $\ell_1 = |\bar{c}| \preceq |\bar{c}_n| = \ell_2$.

\square

LEMMA 33. *If* $g_1 \vdash M \leq V \Leftarrow A$ then $g_2 \vdash M \leq V \Leftarrow A$.

PROOF. The proof is fully mechanized in Agda and can be found in the supplementary material. The proof is by casing on value and the simulation relation. \square

LEMMA 34 (SIMULATION BETWEEN λ_{IFC}^c AND λ_{SEC}). *Suppose* M_1 is a well-typed λ_{IFC}^c term: $\Gamma; \Sigma_1; g^c; |PC| \vdash M_1 \Leftarrow A$, PC is a well-typed label expression: $\vdash PC \Leftarrow g^c$, and μ_1 is a well-typed heap: $\Sigma_1 \vdash \mu_1$. *Suppose* $g^c \vdash M'_1 \leq M_1 \Leftarrow A$, $\Sigma_1 \vdash \mu'_1 \leq \mu_1$, and $\ell \preceq |PC|$. *If* $M_1 | \mu_1 | PC \longrightarrow M_2 | \mu_2$, then there exists M_3, μ_3, M'_2, μ'_2 s.t $M_2 | \mu_2 | PC \longrightarrow^* M_3 | \mu_3, M'_1 | \mu'_1 | \ell \longrightarrow^* M'_2 | \mu'_2, g^c \vdash M'_2 \leq M_3 \Leftarrow A$, and $\Sigma_2 \vdash \mu'_2 \leq \mu_3$ for some Σ_2 .

PROOF. The proof is by induction on the reduction relation $M_1 | \mu_1 | PC \longrightarrow M_2 | \mu_2$ and inversion on the simulation relation $\vdash M'_1 \leq M_1 \Leftarrow A$. We only consider successful cases of the reduction (which do not produce errors), because the theorem statement of noninterference (Theorem 16) for λ_{IFC}^* is termination-insensitive.

Case ξ :

$$N_1 | \mu_1 | PC \longrightarrow N_2 | \mu_2 \quad (10)$$

$$\text{plug } N_1 F = M_1 \quad (11)$$

$$\text{plug } N_2 F = M_2 \quad (12)$$

We case on the frame F . We consider $F = \text{app } \square M A B \ell$ below; the cases for other frames all follow the same pattern. By inversion of the simulation relation, the left side must be a function application:

$$\ell^c \vdash N'_1 M' \leq \text{app } N_1 M A B \ell \Leftarrow C$$

We know $\ell^c \vdash N'_1 \leq N_1 \Leftarrow (A \xrightarrow{\ell^c \vee \ell} B)_\ell$ and $\ell^c \vdash M' \leq M \Leftarrow A$. By the induction hypothesis, there exists N_3, μ_3, N'_2, μ'_2 such that $N_2 | \mu_2 | PC \longrightarrow^* N_3 | \mu_3, N'_1 | \mu'_1 | \ell \longrightarrow^* N'_2 | \mu'_2, \ell^c \vdash N'_2 \leq N_3 \Leftarrow (A \xrightarrow{\ell^c \vee \ell} B)_\ell$, and $\Sigma_2 \vdash \mu'_2 \leq \mu_3$ for some Σ_2 . Choose $M_3 = \text{plug } N_3 F = \text{app } N_3 M A B \ell, \mu_3 = \mu_3, M'_2 = N'_2 M'$, and $\mu'_2 = \mu'_2$. We know $\text{app } N_2 M A B \ell | \mu_2 | PC \longrightarrow^* \text{app } N_3 M A B \ell | \mu_3, N'_1 M' | \mu'_1 | \ell \longrightarrow^* N'_2 M' | \mu'_2$, and $\Sigma_2 \vdash \mu'_2 \leq \mu_3$. We still need to show $\ell^c \vdash N'_2 M' \leq \text{app } N_3 M A B \ell \Leftarrow C$, which is proved by applying rule $\leq\text{-app}$.

⁷In `stampe-security` and `stamp!e-security` of `/src/LabelExpr/Security.agda`.

Case *prot-val*: By inversion on the simulation relation, the left side must be a protection term. By Lemma 28, the body V' must be a value:

$$g_1 \vdash \text{prot } \ell' V' \leq \text{prot } PC \ell V A \Leftarrow B \quad (13)$$

$$g_2 \vdash V' \leq V \Leftarrow A \quad (14)$$

$$\ell' \preceq \ell \quad (15)$$

$$\vdash PC \Leftarrow g_2 \quad (16)$$

We take a step by *prot-val* on the left side, we need to relate:

$$g_1 \vdash V' \vee \ell' \leq \text{stamp } V \ell \Leftarrow B$$

which is proved by applying Lemma 29 on (14) (with Lemma 33 applied) and (15).

Case *prot-ctx*:

$$M \mid \mu_1 \mid PC_2 \longrightarrow N \mid \mu_2$$

By inversion on the simulation relation, the left side must be protection:

$$g_1 \vdash \text{prot } \ell' M' \leq \text{prot } PC_2 \ell M A \Leftarrow B \quad (17)$$

$$g_2 \vdash M' \leq M \Leftarrow A \quad (18)$$

$$\ell' \preceq \ell \quad (19)$$

$$\vdash PC_2 \Leftarrow g_2 \quad (20)$$

By inversion on the typing of the protection term on the right:

$$|PC_1| \vee \ell \preceq |PC_2| \quad (21)$$

where PC_1 is the original PC to reduce the protection term on the right. To get the induction hypothesis, we need to show $\ell'' \vee \ell' \preceq |PC_2|$ where ℓ'' is the PC to reduce the protection term on the left. We know $\ell'' \preceq |PC_1|$, thus:

$$\ell'' \vee \ell' \preceq \ell' \vee \ell \preceq |PC_1| \vee \ell \preceq |PC_2|$$

The induction hypothesis shows that there exists L, μ_3, N', μ_2' s.t $N \mid \mu_2 \mid PC_2 \longrightarrow^* L \mid \mu_3, M' \mid \mu_1' \mid \ell'' \vee \ell' \longrightarrow^* N' \mid \mu_2', g_2 \vdash N' \leq L \Leftarrow A$, and $\Sigma_2 \vdash \mu_2' \leq \mu_3$ for some Σ_2 . We step the left side to $\text{prot } \ell' N'$ and step the right side to $\text{prot } PC_2 \ell L A$ using the congruence of *prot-ctx*. We then relate the protection terms on both sides using rule $\leq\text{-prot}$.

Case *cast*: There are three sub-cases: *cast*, *cast-id*, and *cast-comp*.

Sub-case *cast*: We know $g^c \vdash M' \leq V_r \langle c_r, \bar{c} \rangle \Leftarrow B$. $V_r \langle c_r, \bar{c} \rangle$ is not a value, so $g^c \vdash M' \leq V_r \Leftarrow A$. We need to relate $g^c \vdash M' \leq V_r \langle c_r, \bar{d} \rangle \Leftarrow B$, which is directly proved by applying $\leq\text{-cast}$.

Sub-case *cast-id*: We know $g^c \vdash M' \leq V_r \langle \text{id}(t), \text{id}(g) \rangle \Leftarrow A$. Again, note that $V_r \langle \text{id}(t), \text{id}(g) \rangle$ is not a value, so $g^c \vdash M' \leq V_r \Leftarrow A$.

Sub-case *cast-comp*: We know $g^c \vdash M' \leq V_r \langle c \rangle \langle d \rangle \Leftarrow B$. By inversion using rule $\leq\text{-cast}$, $g^c \vdash M' \leq V_r \langle c \rangle \Leftarrow A$; c is irreducible, by Lemma 28 M' is a value. We need to show there exists M such that $V_r \langle c \rangle \langle d \rangle \mid \mu_1 \mid PC \longrightarrow^* M \mid \mu_1$ and $g^c \vdash M' \leq M \Leftarrow B$. Casing on V_r :

- **If $V_r = \$k$.** Suppose $c \circledast d = d_r, \bar{c} \circledast \bar{d}$ for some d_r . We take a step using *cast* by reducing the composed coercion sequence to its normal form: $\bar{c} \circledast \bar{d} \longrightarrow^+ \bar{c}_n$. By Lemma 2 (Composition models explicit flow), $|\bar{c}| \preceq |\bar{c}_n|$. By Lemma 25, we know $M' = (\$k)_{\ell'}$ and $\ell' \preceq |\bar{c}|$. By transitivity, $\ell' \preceq |\bar{c}_n|$. (1) If \bar{c}_n is not an identity coercion, we relate $g^c \vdash (\$k)_{\ell'} \leq \$k \langle d_r, \bar{c}_n \rangle \Leftarrow B$ directly using rule $\leq\text{-wrapped-const}$. (2) If $\bar{c}_n = \text{id}(\ell)$, we take one additional step by *cast-id*. We know $\ell' \preceq |\text{id}(\ell)| = \ell$, thus $g^c \vdash (\$k)_{\ell'} \leq \$k \Leftarrow \iota_\ell$ by rule $\leq\text{-const}$.

- If $V_r = \lambda x. N$. Similar to the constant case above. The only difference is that we do not need to specially handle $\bar{c}_n = \mathbf{id}(\ell)$, because the function coercion $(\bar{d}, c \rightarrow \mathbf{d}, \mathbf{id}(\ell))$ is already irreducible.
- If $V_r = \text{addr } n$. Same as the lambda case above.

Case β : By inversion on the simulation relation, the left side must be a function application:

$$\ell^c \vdash L M \leq \text{app } (\lambda x. N) V A B \ell \Leftarrow C$$

where $C = \text{stamp } B \ell$. We know that L must be a λ by rule $\leq\text{-lam}$ and M must be a value by Lemma 28:

$$\ell^c \vdash ((\lambda x. N')_{\ell'}) V' \leq \text{app } (\lambda x. N) V A B \ell \Leftarrow C \quad (22)$$

$$\ell' \preceq \ell \quad (23)$$

$$\ell^c \vee \ell \vdash N' \leq N \Leftarrow B \quad (24)$$

$$\ell^c \vdash V' \leq V \Leftarrow A \quad (25)$$

We take one step by β on the left side of the simulation relation:

$$((\lambda x. N')_{\ell'}) V' \mid \mu' \mid \ell'' \longrightarrow \text{prot } \ell' (N'[x := V']) \mid \mu'$$

Now we need to show

$$\ell^c \vdash \text{prot } \ell' (N'[x := V']) \leq \text{prot } (\text{stamp } PC \ell) \ell (N[x := V]) B \Leftarrow C$$

Applying $\leq\text{-prot}$ yields three sub-goals: (1) $\ell' \preceq \ell$ is directly proved by (23) (2) $\ell^c \vee \ell \vdash N'[x := V'] \leq N[x := V] \Leftarrow B$ is proved by applying Lemma 24 on (24) and (25). (3) $\text{stamp } PC \ell \Leftarrow \ell^c \vee \ell$ because stamping is well-typed.

Case app-cast :

$$(\text{stamp } PC_1 \ell) \langle \bar{d} \rangle \longrightarrow^* PC_2 \quad (26)$$

$$V \langle c \rangle \longrightarrow^* W \quad (27)$$

By inversion on the simulation relation, the left side must be a function application:

$$\ell^c \vdash L M \leq \text{app } (\lambda x. N \langle \bar{d}, c \rightarrow \mathbf{d}, \bar{c} \rangle) V C D \ell \Leftarrow E$$

where $E = \text{stamp } D \ell, \vdash \bar{d} : \ell^c \vee \ell \Rightarrow g^c, \vdash c : C \Rightarrow A$, and $\vdash \mathbf{d} : B \Rightarrow D$.

We know that L must be a λ by Lemma 27 and M must be a value by Lemma 28:

$$\ell^c \vdash ((\lambda x. N')_{\ell'}) V' \leq \text{app } (\lambda x. N \langle \bar{d}, c \rightarrow \mathbf{d}, \bar{c} \rangle) V C D \ell \Leftarrow E \quad (28)$$

$$\ell' \preceq |\bar{c}| \quad (29)$$

$$g^c \vdash N' \leq N \Leftarrow B \quad (30)$$

$$\ell^c \vdash V' \leq V \Leftarrow C \quad (31)$$

We take a step by β on the left side of simulation relation:

$$((\lambda x. N')_{\ell'}) V' \mid \mu' \mid \ell'' \longrightarrow \text{prot } \ell' (N'[x := V']) \mid \mu'$$

Now we need to show

$$\ell^c \vdash \text{prot } \ell' (N'[x := V']) \leq \text{prot } PC_2 \ell ((N[x := W]) \langle \mathbf{d} \rangle) D \Leftarrow E$$

By rule $\leq\text{-prot}$, it is equivalent to showing:

$$\ell' \preceq \ell \quad (32)$$

$$g^c \vdash N'[x := V'] \leq (N[x := W]) \langle \mathbf{d} \rangle \Leftarrow D \quad (33)$$

$$\vdash PC_2 \Leftarrow g^c \quad (34)$$

We know $\vdash PC_2 \Leftarrow g^c$ because reducing label expressions preserves types. By (29) and $|\bar{c}| = \ell$ we prove (32). Apply Lemma 23 on (31) and (27), we get $\ell^c \vdash V' \leq W \Leftarrow A$. By Lemma 24, Lemma 33, and (30), $g^c \vdash N'[x := V'] \leq N[x := W] \Leftarrow B$. Apply rule $\leq\text{-cast}$ and we prove (33).

Case *app★-cast*: Similar to *app-cast*. The only noteworthy difference is that on the right side, we use $|\bar{c}|$ instead of the ℓ from the syntax of the function application, both in the protection term and when stamping the PC.

Cases *β-if-true* and *β-if-false*: By inversion on the simulation relation, the left side must also be an if-conditional:

$$\ell^c \vdash \text{if } (\$ \text{ true})_{\ell'} M' N' \leq \text{if } \$ \text{ true } A \ell M N \Leftarrow B$$

We know:

$$\ell' \preceq \ell \tag{35}$$

$$\ell^c \vee \ell \vdash M' \leq M \Leftarrow A \tag{36}$$

$$\ell^c \vee \ell \vdash N' \leq N \Leftarrow A \tag{37}$$

Take one step on the left side. We need to relate $\ell^c \vdash \text{prot } \ell' M' \leq \text{prot } (\text{stamp } PC \ell) \ell M A \Leftarrow B$. The goal is directly proved by applying rule $\leq\text{-prot}$. The case for *β-if-false* is analogous.

Cases *if-true-cast* and *if-false-cast*: Note that the label partial order trivially holds because $\ell' \preceq \text{high}$ for any ℓ' . The rest is analogous to *β-if-true* and *β-if-false*.

Cases *if★-true-cast* and *if★-false-cast*: By Lemma 25 we know $\ell' \preceq |\bar{c}|$. The rest is analogous to *β-if-true* and *β-if-false*.

Case *ref*: From the typing derivation (rule $\vdash\text{-ref}$), we also know:

$$\vdash PC \Leftarrow \ell^c \tag{38}$$

$$\ell^c \preceq \ell \tag{39}$$

which implies

$$|PC| \preceq \ell \tag{40}$$

We know $\ell' \preceq |PC| \preceq \ell$, where ℓ' is the PC that the left side is reduced with. We take a step by *ref?-ok* on the left side by choosing the same fresh address n . We need to show:

$$\ell^c \vdash (\text{addr } n_{\ell})_{\text{low}} \leq \text{addr } n \Leftarrow (\text{Ref } T_{\ell})_{\text{low}} \tag{41}$$

$$(\Sigma_1, \ell \mapsto n \mapsto T) \vdash (\mu', \ell \mapsto n \mapsto (V' \vee \ell)) \leq (\mu, \ell \mapsto n \mapsto V) \tag{42}$$

(41) is proved directly by $\leq\text{-addr}$. To prove (42) we need to show:

$$\text{low} \vdash V' \vee \ell \leq V \Leftarrow T_{\ell} \tag{43}$$

By inversion on the simulation relation and Lemma 33, we know $\text{low} \vdash V' \leq V \Leftarrow T_{\ell}$. From the definition of stamping we know $\text{stamp } V \ell = V$ if $V \Leftarrow T_{\ell}$; (43) is thus proved by Lemma 29.

Case *ref?*:

$$PC_1 \langle \star \Rightarrow^p \ell \rangle \longrightarrow^* PC_2$$

Reducing label expressions preserves types, so $\vdash PC_2 \Leftarrow \ell$. We know $|PC_2| = \ell$ and $\ell' \preceq |PC_1|$, where ℓ' is the PC that the left side is reduced with. Casting models explicit flow (Lemma 30), so $|PC_1| \preceq |PC_2|$. We thus have $\ell' \preceq |PC_1| \preceq |PC_2| = \ell$. Take one step by *ref?-ok* on the left side by choosing the same fresh address n . The rest of the proof follows that of *ref*.

Case *deref*:

$$\mu(\hat{\ell}, n) = V$$

By $\Sigma \vdash \mu' \leq \mu$ and Definition 22, there exists V' s.t $\mu'(\hat{\ell}, n) = V'$ and $\text{low} \vdash V' \leq V \Leftarrow T_{\hat{\ell}}$ where $T = \Sigma(\hat{\ell}, n)$. By inversion on the simulation relation, the left side must be a dereference:

$$g^c \vdash !(\text{addr } n_{\hat{\ell}})_{\ell'} \leq !(\text{addr } n) T_{\hat{\ell}} \ell \Leftarrow B$$

where $\ell' \preccurlyeq \ell$. Take one step on the left using rule *deref*, we need to relate:

$$g^c \vdash \text{prot } \ell' V' \leq \text{prot } \text{high } \ell V T_{\hat{\ell}} \Leftarrow B$$

which is directly proved by rule $\leq\text{-prot}$ and Lemma 33.

Case *deref-*cast*:**

$$\mu(\hat{\ell}, n) = V$$

By $\Sigma \vdash \mu' \leq \mu$ and Definition 22, there exists V' s.t $\mu'(\hat{\ell}, n) = V'$ and $\text{low} \vdash V' \leq V \Leftarrow S_{\hat{\ell}}$ where $S = \Sigma(\hat{\ell}, n)$. By inversion on the simulation relation, the left side must be a dereference:

$$g^c \vdash !(\text{addr } n_{\hat{\ell}})_{\ell} \leq !\star(\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) T \Leftarrow T_{\star}$$

where $\vdash c : T_{\star} \Rightarrow S_{\hat{\ell}}, \vdash d : S_{\hat{\ell}} \Rightarrow T_{\star}$. By Lemma 26, $\ell \preccurlyeq |\bar{c}|$. We take a step on the left side using *deref*. We need to relate:

$$g^c \vdash \text{prot } \ell V' \leq \text{prot } \text{high } |\bar{c}| (V \langle d \rangle) T_{\star} \Leftarrow B$$

which is proved directly by applying Lemma 33, $\leq\text{-prot}$, and then $\leq\text{-cast}$.

Case *deref*-*cast*: Analogous to *deref**-*cast*.**Case β -*assign*:** By inversion on the simulation relation, the left side is also an assignment:

$$\ell^c \vdash (\text{addr } n_{\hat{\ell}})_{\ell'} :=^? V' \leq \text{assign}(\text{addr } n) V T \hat{\ell} \ell \Leftarrow \text{Unit}_{\text{low}}$$

where $\ell' \preccurlyeq \ell$. From the typing derivation we know $\ell^c \vee \ell \preccurlyeq \hat{\ell}, \vdash PC \Leftarrow \ell^c$. We know $\ell'' \preccurlyeq |PC| = \ell^c$ where ℓ'' is the PC on the left, thus $\ell'' \vee \ell' \preccurlyeq \hat{\ell}$. We take one step on the left side using *assign?ok*. The units on both sides relate straightforwardly. We need to relate:

$$\Sigma \vdash [\hat{\ell} \mapsto n \mapsto (V' \vee \hat{\ell})] \mu' \leq [\hat{\ell} \mapsto n \mapsto V] \mu$$

Given $\Sigma(\hat{\ell}, n) = T$, we need to show:

$$\text{low} \vdash V' \vee \hat{\ell} \leq V \Leftarrow T_{\hat{\ell}}$$

which can be proved similar to (43) because $\vdash V \Leftarrow T_{\hat{\ell}}$.

Case *assign*-*cast*:

$$V \langle c \rangle \longrightarrow^* W$$

where $\vdash c : T_{\hat{\ell}_2} \Rightarrow S_{\hat{\ell}_1}$. By Lemma 32, $\hat{\ell}_2 \preccurlyeq \hat{\ell}_1$. From the typing derivation, $\vdash PC \Leftarrow \ell^c$ and $\ell^c \vee \ell \preccurlyeq \hat{\ell}_2$. By inversion on the simulation relation, the left side is also an assignment:

$$\ell^c \vdash (\text{addr } n_{\hat{\ell}_1})_{\ell'} :=^? V' \leq \text{assign}(\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) V T \hat{\ell}_2 \ell \Leftarrow \text{Unit}_{\text{low}}$$

We know $\ell' \preccurlyeq |\bar{c}| = \ell$. $\ell'' \preccurlyeq |PC| = \ell^c$ where ℓ'' is the PC that the left side is reduced with. Thus we have $\ell'' \vee \ell' \preccurlyeq \hat{\ell}_2 \preccurlyeq \hat{\ell}_1$. The check succeeds so we take one step on the left side by *assign?ok*. The units on both sides relate straightforwardly. We need to relate:

$$\Sigma \vdash [\hat{\ell}_1 \mapsto n \mapsto (V' \vee \hat{\ell}_1)] \mu' \leq [\hat{\ell}_1 \mapsto n \mapsto W] \mu$$

Need to show:

$$\text{low} \vdash V' \vee \hat{\ell}_1 \leq W \Leftarrow S_{\hat{\ell}_1}$$

which can be proved similar to (43).

Case *assign?-cast*:

$$(\text{stamp! } PC_1 \mid \bar{c}) \langle \star \Rightarrow^P \hat{\ell} \rangle \longrightarrow^* PC_2 \quad (44)$$

$$V \langle c \rangle \longrightarrow^* W \quad (45)$$

By inversion on the simulation relation, the left side is also an assignment:

$$g^c \vdash (\text{addr } n_{\hat{\ell}})_{\ell} :=^? V' \leq \text{assign?}^P (\text{addr } n \langle \text{Ref } c \ d, \bar{c} \rangle) VTg \Leftarrow \text{Unit}_{\text{low}}$$

where $\vdash c : T_g \Rightarrow S_{\hat{\ell}}, \vdash d : S_{\hat{\ell}} \Rightarrow T_g$. By Lemma 26, $\ell \preccurlyeq |\bar{c}|$. Stamping models implicit flow (Lemma 31), so $|\text{stamp } PC_1 \mid \bar{c}| = |PC_1| \vee |\bar{c}|$; casting models explicit flow (Lemma 30), so $|PC_1| \vee |\bar{c}| \preccurlyeq |PC_2|$. Reduction of label expressions preserves types, so $\vdash PC_2 \Leftarrow \hat{\ell}$. Thus $|PC_2| = \hat{\ell}, |PC_1| \vee |\bar{c}| \preccurlyeq \hat{\ell}$. We know $\ell' \preccurlyeq |PC_1|$ where ℓ' is the PC the left side reduces with. Thus $\ell' \vee \ell \preccurlyeq \hat{\ell}$; the check succeeds so we take one step on the left side by *assign?-ok*. The units on both sides relate straightforwardly. We need to relate:

$$\Sigma \vdash [\hat{\ell} \mapsto n \mapsto (V' \vee \hat{\ell})] \mu' \leq [\hat{\ell} \mapsto n \mapsto W] \mu$$

Need to show:

$$\text{low} \vdash V' \vee \hat{\ell} \leq W \Leftarrow S_{\hat{\ell}}$$

which again can be proved similar to (43). □

LEMMA 35 (MULTI-STEP SIMULATION). *Suppose M is a well-typed λ_{IFC}^c term $\Gamma; \Sigma; g^c; |PC| \vdash M \Leftarrow A$, PC is a well-typed label expression: $\vdash PC \Leftarrow g^c$, and μ_1 is a well-typed heap: $\Sigma \vdash \mu_1$. Suppose $g^c \vdash M' \leq M \Leftarrow A$, $\Sigma \vdash \mu'_1 \leq \mu_1$, and $\ell \preccurlyeq |PC|$. If $M \mid \mu_1 \mid PC \longrightarrow^* V \mid \mu_2$, then there exists V', μ'_2 such that $M' \mid \mu'_1 \mid \ell \longrightarrow^* V' \mid \mu'_2$ and $g^c \vdash V' \leq V \Leftarrow A$.*

PROOF. By induction on multi-step reduction $M \mid \mu_1 \mid PC \longrightarrow^* V \mid \mu_2$.

Zero step If M is already a value, choose V' to be M' . By Lemma 28, M' is also a value. The values are in sync because $g^c \vdash M' \leq M \Leftarrow A$.

One or more steps

$$M \mid \mu_1 \mid PC \longrightarrow N \mid \mu_3 \quad (46)$$

$$N \mid \mu_3 \mid PC \longrightarrow^* V \mid \mu_2 \quad (47)$$

Apply Lemma 34 and then use the induction hypothesis. □

LEMMA (NONINTERFERENCE FOR λ_{IFC}^c). *If M is well-typed: $(x:\text{Bool}_{\text{high}}); \emptyset; \text{low}; \text{low} \vdash M \Leftarrow \text{Bool}_{\text{low}}$ and*

$$M[x := \$ b_1] \mid \emptyset \mid \text{low} \longrightarrow^* V_1 \mid \mu_1 \quad \text{and} \quad M[x := \$ b_2] \mid \emptyset \mid \text{low} \longrightarrow^* V_2 \mid \mu_2$$

then $V_1 = V_2$.

PROOF. From the definition of ϵ (Figure 24), we know $\text{low} \vdash \epsilon(M)[x := (\$ b_i)_{\text{high}}] \leq M[x := \$ b_i] \Leftarrow \text{Bool}_{\text{low}}$. By Lemma 35, there exists V'_i, μ'_i s.t. $\epsilon(M)[x := (\$ b_i)_{\text{high}}] \mid \emptyset \mid \text{low} \longrightarrow^* V'_i \mid \mu'_i$ and $\text{low} \vdash V'_i \leq V_i \Leftarrow \text{Bool}_{\text{low}}$. The simulation relation must be of form $\text{low} \vdash (\$ a_i)_{\text{low}} \leq \$ a_i \Leftarrow \text{Bool}_{\text{low}}$ where a_i is the output boolean. By Lemma 21, $a_1 = a_2$, thus $V_1 = \$ a_1 = \$ a_2 = V_2$. □

THEOREM (NONINTERFERENCE FOR λ_{IFC}^*). *Suppose a λ_{IFC}^* program M is well-typed: $(x:\text{Bool}_{\text{high}}); \text{low} \vdash M : \text{Bool}_{\text{low}}$. If for any boolean inputs b_1, b_2*

$$(C M)[x := \$ b_1] \mid \emptyset \mid \text{low} \longrightarrow^* V_1 \mid \mu_1 \quad \text{and} \quad (C M)[x := \$ b_2] \mid \emptyset \mid \text{low} \longrightarrow^* V_2 \mid \mu_2$$

then the resulting values $V_1 = V_2$.

$$\boxed{\epsilon M A = M'}$$

$$\begin{aligned}
& \epsilon x - = x \\
& \epsilon (\$ k) (\iota_\ell) = (\$ k)_\ell \\
& \epsilon (\lambda x. N) ((A \xrightarrow{-} B)_\ell) = (\lambda x. \epsilon N B)_\ell \\
& \epsilon (\text{addr } n) (\text{Ref } (T_{\hat{\ell}})_\ell) = (\text{addr } n_{\hat{\ell}})_\ell \\
& \epsilon (\text{app } M N A B \ell) - = (\epsilon M (A \xrightarrow{\star} B)_\ell) (\epsilon N A) \\
& \epsilon (\text{app}\star M N A T) - = (\epsilon M (A \xrightarrow{\star} T_{\star})_{\star}) (\epsilon N A) \\
& \epsilon (\text{if } L A \ell M N) - = \text{if } (\epsilon L \text{Bool}_\ell) (\epsilon M A) (\epsilon N A) \\
& \epsilon (\text{if}\star L T M N) - = \text{if } (\epsilon L \text{Bool}_{\star}) (\epsilon M T_{\star}) (\epsilon N T_{\star}) \\
& \epsilon (\text{ref } \ell M) (\text{Ref } T_{\ell})_{\text{low}} = \text{ref}^? \ell (\epsilon M T_{\ell}) \\
& \epsilon (\text{ref?}^P \ell M) (\text{Ref } T_{\ell})_{\text{low}} = \text{ref}^? \ell (\epsilon M T_{\ell}) \\
& \epsilon (! M A \ell) - = ! (\epsilon M (\text{Ref } A)_\ell) \\
& \epsilon (!\star M T) - = ! (\epsilon M (\text{Ref } T_{\star})_{\star}) \\
& \epsilon (\text{assign } L M T \hat{\ell} \ell) - = (\epsilon L (\text{Ref } T_{\hat{\ell}})_\ell) :=^? (\epsilon M T_{\hat{\ell}}) \\
& \epsilon (\text{assign?}^P L M T \hat{g}) - = (\epsilon L (\text{Ref } T_{\hat{g}})_{\star}) :=^? (\epsilon M T_{\hat{g}}) \\
& \epsilon (N \langle c \rangle) - = \epsilon N A \quad , \text{ where } \vdash c : A \Rightarrow B \\
& \epsilon (\text{prot } P C \ell M A) - = \text{prot } \ell (\epsilon M A)
\end{aligned}$$

Fig. 24. Erasure from λ_{IFC}^c to λ_{SEC}

PROOF. By Theorem 15 (Compilation preserves types), $(x:\text{Bool}_{\text{high}}); \emptyset; \text{low}; \text{low} \vdash M' \Leftarrow \text{Bool}_{\text{low}}$.
By Lemma 14 (Noninterference for λ_{IFC}^c), $V_1 = V_2$. \square

Received 2023-11-16; accepted 2024-03-31