

Gradual Security Types and Gradual Guarantees

Abhishek Bichhawat
Carnegie Mellon University
Pittsburgh, USA
abichhaw@andrew.cmu.edu

McKenna McCall
Carnegie Mellon University
Pittsburgh, USA
mckennak@andrew.cmu.edu

Limin Jia
Carnegie Mellon University
Pittsburgh, USA
liminjia@andrew.cmu.edu

Abstract—Information flow type systems enforce the security property of noninterference by detecting unauthorized data flows at compile-time. However, they require precise type annotations, making them difficult to use in practice as much of the legacy infrastructure is written in untyped or dynamically-typed languages. Gradual typing seamlessly integrates static and dynamic typing, providing the best of both approaches, and has been applied to information flow control, where information flow monitors are derived from gradual security types. Prior work on gradual information flow typing uncovered tensions between noninterference and the dynamic gradual guarantee—the property that less precise security type annotations in a program should not cause more runtime errors.

This paper re-examines the connection between gradual information flow types and information flow monitors to identify the root cause of the tension between the gradual guarantees and noninterference. We develop runtime semantics for a simple imperative language with gradual information flow types that provides both noninterference and gradual guarantees. We leverage a proof technique developed for FlowML and reduce noninterference proofs to preservation proofs.

Index Terms—Information flow control, noninterference, gradual typing, gradual guarantees

I. INTRODUCTION

Information flow type systems combine types and security labels to ensure that well-typed programs do not leak secrets to attackers at compile-time [1]. However, purely statically-typed languages face significant adoption challenges. Most programmers are unfamiliar with and may be unwilling to use complex information flow type systems. Moreover, much of the legacy infrastructure is written in untyped and dynamically-typed languages without precise security type annotations.

Gradual typing is one promising technique to address these challenges [2]; it aims to seamlessly integrate statically-typed programs with dynamically-typed programs. At a high-level, gradual type systems introduce a dynamic type, often written as τ , to accommodate untyped portions of the program. The type system allows any program to be typed under τ . The type system enforces type safety on statically typed parts and the runtime semantics of gradual type systems monitor the interactions between parts typed as τ and statically typed parts to ensure type preservation.

Gradual typing has been applied to information flow types [3]–[7], where certain expressions have a dynamic security label τ (and are typed as, e.g., int^τ), which is determined at runtime. Information flow monitors are then derived from the runtime semantics of gradual information

flow types. Earlier adoptions of gradual typing to information flow types have developed ad hoc approaches to treat “gradual types”. For instance, Disney and Flanagan did not include a dynamic security label [3]. Instead, the programmer would insert type casts, which are checked at runtime, and are used to “gradually” make the programs more secure. Later, ML-GS [4] and LJGS [5] included the dynamic security label τ and a runtime monitor that performed checks on the dynamically typed parts of the program. Programmers still need to write annotations and casts in ML-GS and label constraints in LJGS. The dynamic label is instantiated as a single security label at run time in both ML-GS and LJGS.

At the same time, interests in the formal foundations of gradual types grew significantly. Formal properties related to gradual typing such as the *gradual guarantee* [8] were introduced, which says that loosening policies should not cause more type errors or runtime failures. Roughly, the gradual guarantee ensures that programs which type-check and run to completion with precise type annotations will also type-check and run to completion with less precise types, i.e., τ . This property ensures that programmers are not punished for not specifying type annotations if the program is safe. Such a guarantee is important for information flow type systems, as security annotations have been a road block for adoption. Without the gradual guarantee, the programmers’ burden of providing (unnecessary) security annotations is increased.

Garcia et al. developed the abstracting gradual typing (AGT) framework which provides a formal interpretation of gradual type systems [9]. In AGT, a principled interpretation of the dynamic type is that it represents the set of all possible types that are refined by the monitor to preserve type safety at runtime. By that interpretation, the semantics of the dynamic information flow label τ is the set of all possible labels. Early work on gradual information flow typing all instantiate the dynamic label as a single label at runtime [4], [5]. Recent work by Toro et al., GSL_{Ref} , aims to apply the AGT framework to information flow types [6]; however, it has to give up the dynamic gradual guarantee in favor of *noninterference*, the key information flow security property [10], when dealing with mutable references.

In this paper, we re-examine the connection between gradual information flow types and information flow monitors (c.f. [11]–[13]). We aim to identify the root cause of the tension between the dynamic gradual guarantee and security in systems that refine the set of possible labels for dynamically

labeled programs at runtime. To this end, we focus on a simple imperative language with first-order stores, which has been widely used to design information flow control systems [1], [13]–[17]. While simple, this language includes all the features to illustrate the problem of refining dynamic labels at runtime. We develop runtime semantics for this language with gradual information flow types that enjoy both noninterference and the dynamic gradual guarantee. We draw ideas from abstracting gradual typing, which advocates deriving runtime semantics for gradual types via the preservation proof [9].

We observe that as dynamic labels are updated, the semantics that only gradually refine the possible security labels during program execution *resemble* a naive flow-sensitive monitor and therefore inherit the problems with implicit leaks of flow-sensitive monitors [11]. To enforce noninterference and remove the implicit leaks caused by insecure writes in branches, the runtime semantics needs to take into consideration the variable and channel writes in the untaken branch [13]. Pure guessing which ignores information about the untaken branch like GSL_{Ref} 's runtime yields rigid semantics that break the gradual guarantee (more in Section III-D). The no-sensitive-upgrade (NSU) check [11] also doesn't solve the problem. Instead, a "hybrid" approach [13], [17] that leverages static analysis to obtain the write effects of the untaken branch and upgrade relevant references for both branches can be used to remove the implicit leaks and provide the gradual guarantees.

We leverage a proof technique developed for FlowML, which reduces noninterference proofs to preservation proofs [18]. The main idea is to extend the language with *pairs* of expressions and commands, representing two executions with different secrets in one program. Noninterference follows from preservation. This proof technique clearly illustrates the problem with purely dynamic flow-sensitive monitors and naturally suggests the hybrid approach [13], [17].

To summarize, we study the connection between gradual security types and information flow monitors and identify the conservative handling of implicit flows in GSL_{Ref} as the reason that it gives up dynamic gradual guarantee in favor of noninterference. Additionally, we show that the dynamic gradual guarantee can be recovered by using a hybrid approach that leverages the static phase to generate a list of variables that are written to in both the branches. Due to space constraints, we omit detailed definitions and proofs, which can be found in the full version of the paper [19].

II. OVERVIEW OF INFORMATION FLOW CONTROL

In information flow control systems, variables are annotated with a label from a security lattice, which have a partial-ordering (\preceq) and a well-defined join and meet operation. $\ell_1 \preceq \ell_2$ means information can flow from ℓ_1 to ℓ_2 . Consider a two-point security lattice with labels $\{L, H\}$ with $L \preceq H$ where L represents public and H represents secret. A variable x having type int^H contains a sensitive integer value.

Information flows can be broadly classified as *explicit* or *implicit* [10], [20]. Explicit flows arise from variable assign-

ments. For instance, the statement $x = y + z$ causes an explicit flow of values from y and z to x . Implicit flows arise from control structures in the program. For example, in the program $l = \text{false}; \text{if}(h)\{l = \text{true};\}$, there is an implicit flow of information from h to the final value of l (it is true iff h is true). Implicit flows are handled by maintaining a pc (program-context) label, which is an upper bound on the labels of all the predicates that have influenced the control flow thus far. In the example, the pc inside the branch is the label of h .

Information flow control systems aim to prevent leaks through these flows by either enforcing information flow typing rules and ruling out insecure programs at compile-time or dynamically monitoring programs and aborting the execution of insecure programs. In both systems, assignment to a variable is disallowed if either the pc label or the join of the label of the operands is not less than or equal to the label of the variable being assigned [1], [11]. Thus, in the above examples, if either the label of y or z is greater than the label of x or the label of h is greater than the label of l , the assignment does not type-check or the execution aborts at runtime. This guarantees a variant of noninterference, known as termination-insensitive noninterference [1], which we prove for our gradual type system. We assume that an adversary cannot observe or gain any information if a program's execution diverges or aborts and can only observe "public" outputs by the program.

We consider a flow-insensitive, fixed-label system in this paper and prove termination-insensitive noninterference for it.

III. GRADUAL SECURITY TYPING

Static information flow type systems do not scale up to scenarios where the security levels of some of the variables are not known at compile-time, while pure monitoring approaches cannot reject obviously insecure programs at compile-time. Gradual typing extends the reach of type-system based analysis by adding an imprecise (or dynamic) label, $?$, for variables whose labels are not known at compile-time. The runtime semantics then ensures that no information is leaked due to the relaxation of the type-system's handling of $?$ labels.

A. Imprecise Security Label: Interpretations and Operations

The label $?$ is not an actual element of the security lattice and its meaning is not universally agreed upon. Differences will manifest in the runtime monitoring semantics and proof of noninterference. For illustration, consider a variable x of type $\text{int}^?$. Semantically, in the literature $?$ has meant one of the following. (1) x 's label is dynamic (flow-sensitive) and can change at runtime (e.g., from $?$ representing L to representing H when x is assigned a secret). (2) the set of possible labels for x is refined at runtime and the set in a future state will be a subset of its the current state. Since we build on a flow-insensitive type system, we opt for the second meaning of $?$. Our runtime monitor will keep track of the set of possible labels for x . Note that a typical flow-sensitive IFC monitor (e.g. [11], [13]) takes an approach aligned with (1).

In the initial program state, $?$ could be interpreted as: (1) x could contain a secret or be observable to an adversary.

```

1  $y := \text{false}^H$ 
2 if ( $x$ ) then  $y := \text{true}^H$ 

```

Listing 1. Statically typed

```

1  $y := \text{false}^?$ 
2  $z := \text{false}^L$ 
3 if ( $x$ ) then  $y := \text{true}^?$ 
4 if ( $y$ ) then  $z := \text{true}^L$ 
5 output( $L, z$ )

```

Listing 3. NSU

```

1  $y := \text{false}^?$ 
2 if ( $x$ ) then  $y := \text{true}^?$ 

```

Listing 2. Dynamically typed

```

1  $y := \text{true}^?$ 
2  $z := \text{true}^L$ 
3 if ( $x$ ) then  $y := \text{false}^?$ 
4 if ( $y$ ) then  $z := \text{false}^L$ 
5 output( $L, z$ )

```

Listing 4. Implicit flows

Therefore, we should treat x conservatively as if it is both secret and public. (2) $?$ indicates indifference; the data x contains in the initial state is of no security value; otherwise, x should have been given the label H . We choose (2) again, as it is a cleaner interpretation. Note that this is only for the initial state. At runtime, the monitor maintains enough state to concretely know whether x contains a secret or not.

B. Gradually Refined Security Policy via Examples

Next, we describe how security labels can be gradually refined at runtime. Consider the program in Listing 1 and its variant with dynamic types in Listing 2. Suppose the lattice contains four elements $\{\perp, L, H, \top\}$ such that $\perp \preceq L \preceq H \preceq \top$. Assume that the initial type of x is bool^H in both examples while the type of y is bool^H in Listing 1 and $\text{bool}^?$ in Listing 2¹. The program in Listing 1 does not leak any information. With gradual typing, its variant in Listing 2 is also accepted by the type system. The runtime refines the set of possible labels for y as the program runs. With $x : \text{bool}^H$, y cannot be \perp or L as that would result in an implicit flow. Thus, the possible labels of y are refined to the set $\{H, \top\}$ when $x = \text{true}$. If, suppose, $x : \text{bool}^L$, then y is labeled $\{L, H, \top\}$ after executing line 2.

Suppose the program in Listing 2 is extended with another branch as shown in Listing 3 with $z : \text{bool}^L$. When $x : \text{bool}^L$, the possible labels for y on line 4 are $\{L, H, \top\}$. This allows the assignment on line 4 to succeed as the assignment is to a variable that has a label contained in the set of possible pc labels. When x has the type bool^H , then the possible labels for y on line 4 are $\{H, \top\}$. The assignment on line 4 is aborted as L is not equal or higher than any of the possible pc labels ($\{H, \top\}$). In this case, the monitor enforces NSU and prevents the implicit leak.

C. Gradual Guarantees

Desirable formal properties for gradual type systems are the *gradual guarantees*, proposed by [8]. The gradual guarantees relate programs that differ only in the precision of the type annotations. They state that changes that make the annotations of a gradually typed program less precise should not change the static or dynamic behavior of the program. In other words, if a program with more precise type annotations is well-typed in the static type system, and terminates in the runtime

semantics, then the same program with less precise terms is also well-typed, and terminates, respectively.

For illustration, consider the previous example from Listing 1. Assume that the program is well-typed under a gradual type system with $x : \text{bool}^H$ and $y : \text{bool}^H$ as secret variables. When x is `true`, the branch on line 2 is taken and y is assigned the value `true` and has the label H . When x is `false`, y remains `false`. This program is accepted by the security type system and dynamic information flow monitor, and runs to completion in all possible executions. As per the *static gradual guarantee*, the program should also be well-typed if, for instance, y had an imprecise $?$ security level as shown in Listing 2. By the *dynamic gradual guarantee*, the program in Listing 2 should run to completion at runtime, even with the imprecise label for y in all executions of the program.

The gradual guarantees are important in the context of information flow systems to show that the gradual security type system is strictly more permissive than the static security type system and the dynamic IFC monitor, while providing the same guarantees. They mean that programmers need not worry about insufficient annotations causing their safe programs to be rejected by the type system, or worse, at runtime, which may lead to undesirable behavior. Concerning programmers with unnecessary annotations defeats the purpose of gradual typing, which is meant to alleviate the burden of annotation.

D. Implicit Flows vs. Dynamic Gradual Guarantee

The example in Listing 4 illustrates how gradual typing semantics handle implicit flows. Assume that $x : \text{bool}^H$ is a secret variable while the security types of $y : \text{bool}^?$ and $z : \text{bool}^?$ are unknown at compile-time, and the security lattice is $\perp \preceq L \preceq H \preceq \top$. Consider two runs of the program with different initial values of x . When x is `true`, the branch on line 3 is taken and y is assigned the value `false`. With gradual typing, the labels of y are refined to $\{H, \top\}$. As y is `false`, the branch on line 4 is not taken and z remains `true` with the set of labels $\{\perp, L, H, \top\}$. As z 's value is visible at L , the output on line 5 succeeds. When x is `false`, the branch on line 3 is not taken and y remains `true` with the set of labels $\{\perp, L, H, \top\}$. As the pc on line 4 contains the set of labels $\{\perp, L, H, \top\}$, the assignment on line 4 succeeds, and z becomes `false` while the set of labels remains $\{\perp, L, H, \top\}$. Again, as z 's value is visible at L , the output on line 5 succeeds. Thus, in the two runs of the program, different values of z are output for different values of x , thereby leaking x to the adversary at level L . Here, the NSU mechanism does not apply, as the assignment to y on line 3 is merely refining, not “upgrading”, the label of y . If y 's label had been L , this program would have been rejected.

This program can be rejected by deploying a special monitoring rule that preemptively aborts an assignment statement if there is a possibility that the pc is not lower than or equal to the variable's label, as deployed by GSL_{Ref} [6]. In the above example, the assignment on line 3 will be aborted, because the pc is H , and y 's label could be \perp or L , which might leak information. This ensures noninterference, but unfortunately,

¹We will write x^ℓ to indicate that variable x has the type τ^ℓ .

```

1  $y := \text{true}^?$ 
2 if  $(x)$  then  $y := \text{false}^?$ 
3 output( $H, y$ )

```

Listing 5. Secure program violating gradual guarantee

the extra check does not retain the dynamic gradual guarantee. That is, enlarging the set of possible labels for the dynamic security type of y will cause the monitor to abort, which contradicts the dynamic gradual guarantee.

Consider the program in Listing 5 with the same security lattice as before such that the variable x is labeled H and y 's label is not specified at compile-time. As the pc on line 2 is H and the possible set of labels of y on line 2 is $\{\perp, L, H, \top\}$, the monitor aborts the execution of the program when x is true to satisfy noninterference. However, if y was labeled H or \top at compile-time instead of being $?$, the execution would have proceeded and output the value of y to H . In other words, the larger set of possible labels for y on line 2 (because of the unknown label) causes the monitor to abort while the precisely typed version of the program with $y : \text{bool}^H$ is accepted by the monitor, which violates the dynamic gradual guarantee.

To tackle this problem, we leverage the static phase of the gradual type system to determine the set of variables being written to in different branches and loops, and *refine* their possible security labels to implement a monitoring strategy that preserves the dynamic gradual guarantee. At the branch on line 2 in Listing 5, we know that y may be written to inside the branch, therefore, we narrow the possibility of the labels for y to $\{H, \top\}$ as the first step of executing the if statement, regardless of whether x is true or false. This is very similar to how hybrid monitors stop implicit leaks [13], [17], [21], [22]. We will discuss this further in Section IV-C.

IV. A LANGUAGE WITH GRADUAL SECURITY TYPES

The syntax and typing rules for the language with gradual security types (WHILE^G) are standard and provided in the full version of the paper [19]. In this section, we present the syntax and typing rules for our language with gradual information flow types and evidence ($\text{WHILE}_{\text{Evd}}^G$), define the translation from WHILE^G to $\text{WHILE}_{\text{Evd}}^G$, and explain the operational semantics for our monitor.

A. $\text{WHILE}_{\text{Evd}}^G$

The syntax of $\text{WHILE}_{\text{Evd}}^G$ is shown in Fig. 1. Gradual types, U , consist of a type (bool, or int) and a gradual security label, g . This label is either a static security label, denoted ℓ ; or an imprecise dynamic label, denoted $?$. As is standard, ℓ is drawn from $Labs$, a set of labels, which is a part of a security lattice $\mathcal{L} = (Labs, \preceq)$. Here \preceq is a partial order between labels in $Labs$. We commonly use the label H to indicate secret, L to indicate public data, and $L \preceq H$.

The partial-ordering (\preceq) and join operation (\vee) on security labels (ℓ) extends to consistent ordering (\preceq_c) and consistent-join (\vee_c) to account for $?$, as shown in Fig. 2. The consistent subtyping relation is written as $\tau^{g_1} \leq_c \tau^{g_2}$.

<i>Labels</i>	$\ell ::= L \mid \dots \mid H$
<i>Label-intervals</i>	$\iota ::= [\ell_{low}, \ell_{high}]$
<i>Raw values</i>	$u ::= n \mid \text{true} \mid \text{false}$
<i>Values</i>	$v ::= (\iota u)^g$
<i>Types</i>	$\tau ::= \text{bool} \mid \text{int}$
<i>Gradual labels</i>	$g ::= ? \mid \ell$
<i>Gradual types</i>	$U ::= \tau^g$
<i>Typing Context</i>	$\Gamma ::= \cdot \mid \Gamma, x : U$
<i>Cast evidence</i>	$E ::= (\iota_1, \iota_2)$
<i>Expressions</i>	$e ::= x \mid v \mid e_1 \text{ bop } e_2 \mid E^g e$
<i>Variable Set</i>	$X ::= \{x_1, \dots, x_n\}$
<i>Commands</i>	$c ::= \text{skip} \mid c_1; c_2 \mid x := e \mid \text{output}(\ell, e)$ $\quad \mid \text{if}^X e \text{ then } c_1 \text{ else } c_2 \mid \text{while}^X e \text{ do } c$

Fig. 1. Syntax for the language $\text{WHILE}_{\text{Evd}}^G$

$$\begin{array}{c}
\frac{\ell_1 \preceq \ell_2}{\ell_1 \preceq_c \ell_2} \qquad \frac{}{? \preceq_c g} \qquad \frac{}{g \preceq_c ?} \qquad \frac{g_1 \preceq_c g_2}{\tau^{g_1} \leq_c \tau^{g_2}} \\
\frac{}{\ell_1 \vee_c \ell_2 = \ell_1 \vee \ell_2} \qquad \frac{g \neq \top}{? \vee_c g = ?} \qquad \frac{g \neq \top}{g \vee_c ? = ?} \\
\frac{}{? \vee_c \top = \top} \qquad \frac{}{\top \vee_c ? = \top}
\end{array}$$

Fig. 2. Operations on gradual labels and types

Recall that examples in Section III-A use a set of possible security labels for preventing information leaks. This is *evidence* attesting to the validity of gradual labels. We use an interval of labels, representing the lowest and the highest possible label, as the refinement only narrows the interval, similar to GSL_{Ref} [6].

There are two types of evidence: a label-interval, ι , that justifies the dynamic label $?$; and a pair of intervals or the cast evidence, $E=(\iota_1, \iota_2)$, that justifies the consistent subtyping relation between two gradual types used in casts. Intuitively, ι represents the range of possible static labels that would allow the program to type-check. For static labels ℓ , the evidence is $[\ell, \ell]$. An interval $[\ell_l, \ell_r]$ is valid iff $\ell_l \preceq \ell_r$. The rest of the paper only considers valid intervals. Operations leading to an invalid interval are aborted.

A value in $\text{WHILE}_{\text{Evd}}^G$ is a raw value with an interval of possible security labels for the gradual label. Raw values are integer constants n , or boolean values. Expressions include values, variables, casts, and binary operations on expressions. An explicit type cast is written $E^g e$, where E is the evidence justifying the type cast and g is the label of the resulting type.

Commands include skip, sequencing, assignments, branches, loops, and outputs. This language does not have higher-order stores, so the left-hand side of the assignment is always a global variable. The output command outputs a value at a fixed security label ℓ . We include this command mainly to have a clear statement of the system's observable behavior, so we do not allow output at the imprecise label $?$. To prevent implicit leaks, we include a *write-set* of variables, X , which takes into account variable writes in both conditional branches and the loop body.

Label-interval operations: We first define functions and op-

$$\begin{array}{c}
\frac{}{\gamma(?) = [\perp, \top]} \quad \frac{}{\gamma(\ell) = [\ell, \ell]} \quad \frac{\ell_l \preceq \ell_r}{\text{valid}([\ell_l, \ell_r])} \quad \frac{\ell_2 \preceq \ell_1 \quad \ell'_1 \preceq \ell'_2}{[\ell_1, \ell'_1] \sqsubseteq [\ell_2, \ell'_2]} \quad \frac{\ell'_1 \preceq \ell_2}{[\ell_1, \ell'_1] \preceq [\ell_2, \ell'_2]} \\
\\
\frac{\ell_{1l} \preceq \ell_{1r} \wedge \ell_{2r} \quad \ell_{2l} \vee \ell_{1l} \preceq \ell_{2r}}{\text{refine}([\ell_{1l}, \ell_{1r}], [\ell_{2l}, \ell_{2r}]) = ([\ell_{1l}, \ell_{1r} \wedge \ell_{2r}], [\ell_{2l} \vee \ell_{1l}, \ell_{2r}])} \quad \frac{(\ell_{1l} \not\preceq \ell_{1r} \wedge \ell_{2r}) \vee (\ell_{2l} \vee \ell_{1l} \not\preceq \ell_{2r})}{\text{refine}([\ell_{1l}, \ell_{1r}], [\ell_{2l}, \ell_{2r}]) = \text{undef}} \\
\\
\frac{\iota_1 = [\ell_1, \ell'_1] \quad \iota_2 = [\ell_2, \ell'_2]}{\iota_1 \vee \iota_2 = [\ell_1 \vee \ell_2, \ell'_1 \vee \ell'_2]} \quad \frac{\iota_1 \sqsubseteq \gamma(g_1) \quad \iota_2 \sqsubseteq \gamma(g_2) \quad g_1 \preceq_c g_2}{(\iota_1, \iota_2) \vdash \tau^{g_1} \leq_c \tau^{g_2}} \\
\\
\frac{\iota_1 = [\ell_{1l}, \ell_{1r}] \quad \iota_2 = [\ell_{2l}, \ell_{2r}]}{\iota_1 \bowtie \iota_2 = [\ell_{1l} \vee \ell_{2l}, \ell_{1r} \wedge \ell_{2r}]} \quad \frac{\iota'' = (\iota' \bowtie \iota) \quad \iota'' \sqsubseteq \gamma(g)}{\iota' \bowtie (\iota u)^g = (\iota'' u)^g} \quad \frac{\text{refine}(\iota_1 \bowtie \iota, \iota_2) = (\iota'_1, \iota'_2)}{\iota \bowtie (\iota_1, \iota_2) = \iota'_2} \quad \frac{\text{refine}(\iota_1 \bowtie \iota, \iota_2) = \text{undef}}{\iota \bowtie (\iota_1, \iota_2) = \text{undef}}
\end{array}$$

Fig. 3. Label and label-interval operations

operations on the label-intervals that are used by the typing rules and operational semantics (shown in Fig. 3). The function $\gamma(g)$ returns the maximum possible label-interval for the gradual label g , assuming \perp and \top are the least and the greatest element in the lattice, respectively. Label-intervals form a lattice with the partial ordering defined as $\iota_1 \sqsubseteq \iota_2$. Here, ι_1 is said to be more precise than ι_2 . The label-intervals are refined throughout the execution of the program; i.e., they get more precise. Consider the example in Listing 3. Assume that the security lattice contains two elements L and H such that $L \preceq H$. Initially, y has a label $?$ with the evidence $[L, H]$ indicating that any of the two labels are possible. If $x : \text{bool}^H$, then the only possible label for y that allows assignment on line 3 is H . Thus, the evidence for the label on y is refined to $[H, H]$, which makes the program-context's evidence on line 4 $[H, H]$, disallowing assignments to L -labeled variables.

We define $\iota_1 \preceq \iota_2$ to mean for every security label in ι_1 , all labels in ι_2 are at higher or equal positions in the security lattice; and for every security label in ι_2 , all labels in ι_1 are at lower or equal positions in the security lattice. Even though this relation is not used in our typing rules or operational semantics, it is an invariant that must hold on the results of the binary label-interval operations used by the noninterference proofs. The function $\text{refine}(\iota_1, \iota_2)$ returns the largest sub-intervals of ι_1 and ι_2 (ι'_1 and ι'_2 , respectively) such that $\iota'_1 \preceq \iota'_2$. If the relation does not hold between ι'_1 and ι'_2 , the function returns undef .

The join of label-intervals is defined as $\iota_1 \vee \iota_2$. Note that this is not to be confused with the join operation in the lattice that the intervals form. The join of the label-intervals computes the interval corresponding to all possible joins of security labels in those intervals. The operation $\iota_1 \bowtie \iota_2$ computes the intersection of the intervals ι_1 and ι_2 . $\iota' \bowtie (\iota u)^g$ merges the labels for a value. $\iota \bowtie (\iota_1, \iota_2)$ refines ι_2 based on the intersection of ι and ι_1 .

Evidence-based consistent subtyping: Next, we define consistent subtyping relations for both gradual labels and types as supported by label-intervals. The consistent subtyping relation between two gradual types is written as $(\iota_1, \iota_2) \vdash \tau^{g_1} \leq_c \tau^{g_2}$ (defined in Fig. 3). In this relation, ι_1 , resp. ι_2 represents the set of possible labels for g_1 , resp. g_2 , and $g_1 \preceq_c g_2$.

The evidence (ι_1, ι_2) is to justify the consistent security label partial ordering relation between the labels of the gradual types. Note that we do not have $\iota_1 \preceq \iota_2$ in the premise. The reason is that $(\iota_1, \iota_2) \vdash \tau^{g_1} \leq_c \tau^{g_2}$ is used to type runtime terms; even though $\iota_1 \preceq \iota_2$ holds initially, as label-intervals are refined from ι_i to ι'_i , we cannot guarantee that $\iota'_1 \preceq \iota'_2$ hold. This will break preservation proofs. It is not the gradual type system's job to ensure all execution paths are secure. Instead, the runtime semantics will refine the intervals and abort the computation if necessary when a term is evaluated.

Typing rules: Expressions and commands with evidence are typed using rules shown in Fig. 4.

G-CAST casts an expression of type U_1 to U_2 , if the cast evidence E shows that U_1 is a consistent subtype of U_2 .

We augment the command typing with an interval for the gradual pc label; ι_{pc} is the range of possible static labels for the g_{pc} . The rules are similar to the ones in the original type-system except for the use of evidence for consistent ordering between the gradual labels. An exception is the use of $WtSet(c)$ that returns the set of variables being updated or assigned to in the command c . $WtSet$ is straightforwardly inductively defined over the structure of c and is shown below:

$$\begin{array}{l}
WtSet(\text{skip}) = \emptyset \quad WtSet(\text{output}(\ell, e)) = \emptyset \\
WtSet(x := e) = \{x\} \quad WtSet(\text{while } e \text{ do } c) = WtSet(c) \\
WtSet(c_1; c_2) = WtSet(c_1) \cup WtSet(c_2) \\
WtSet(\text{if } e \text{ then } c_1 \text{ else } c_2) = WtSet(c_1) \cup WtSet(c_2)
\end{array}$$

Further, G-ASSIGN and G-OUT do not consider expression subtyping and instead rely on the casts inserted by translation.

B. From WHILE^G to $\text{WHILE}_{\text{Evd}}^G$

The programs are written in WHILE^G , the language without evidence, which is then translated to $\text{WHILE}_{\text{Evd}}^G$. The explicit casts for ASSIGN and OUT are automatically inserted to account for the subtyping of expressions. We show the interesting rules for translating WHILE^G expressions and commands to $\text{WHILE}_{\text{Evd}}^G$ with evidence insertion in Fig. 5. T-ASSIGN inserts a cast for the expression e to have the same label as that of x . For example, $x^L := y^?$ is rewritten to $x := (\text{refine}([\perp, \top], [L, L])^L)y$. Similarly, T-OUT casts the expression e to the channel level ℓ .

$$\boxed{\Gamma \vdash e : U}$$

$$\frac{\iota \sqsubseteq \gamma(g)}{\Gamma \vdash (\iota u)^g : \Gamma(u)^g} \text{G-CONST} \quad \frac{}{\Gamma \vdash x : \Gamma(x)} \text{G-VAR} \quad \frac{\forall i \in \{1, 2\}, \Gamma \vdash e_i : \tau^{g_i}}{g = g_1 \curlywedge_c g_2} \text{G-BOP} \quad \frac{\Gamma \vdash e : \tau^{g_1}}{E \vdash \tau^{g_1} \leq_c \tau^g} \text{G-CAST}$$

$$\boxed{\Gamma; \iota_{pc} g_{pc} \vdash c}$$

$$\frac{}{\Gamma; \iota_{pc} g_{pc} \vdash \text{skip}} \text{G-SKIP} \quad \frac{\Gamma \vdash e : \text{bool}^g \quad X = \text{WtSet}(c) \quad \iota_c = \gamma(g) \quad \Gamma; \iota_{pc} \curlywedge \iota_c g_{pc} \curlywedge_c g \vdash c}{\Gamma; \iota_{pc} g_{pc} \vdash \text{while}^X e \text{ do } c} \text{G-WHILE}$$

$$\frac{\Gamma; \iota_{pc} g_{pc} \vdash c_1 \quad \Gamma; \iota_{pc} g_{pc} \vdash c_2}{\Gamma; \iota_{pc} g_{pc} \vdash c_1; c_2} \text{G-SEQ} \quad \frac{\Gamma \vdash x : \tau^g \quad \Gamma \vdash e : \tau^g \quad \iota_{pc} \sqsubseteq \gamma(g_{pc}) \quad g_{pc} \preceq_c g}{\Gamma; \iota_{pc} g_{pc} \vdash x := e} \text{G-ASSIGN} \quad \frac{\Gamma \vdash e : \tau^\ell \quad \iota_{pc} \sqsubseteq \gamma(g_{pc}) \quad g_{pc} \preceq_c \ell}{\Gamma; \iota_{pc} g_{pc} \vdash \text{output}(\ell, e)} \text{G-OUT}$$

$$\frac{\Gamma \vdash e : \text{bool}^g \quad \iota_c = \gamma(g) \quad X = \text{WtSet}(c_1) \cup \text{WtSet}(c_2) \quad \forall i \in \{1, 2\}, \Gamma; \iota_{pc} \curlywedge \iota_c g_{pc} \curlywedge_c g \vdash c_i}{\Gamma; \iota_{pc} g_{pc} \vdash \text{if}^X e \text{ then } c_1 \text{ else } c_2} \text{G-IF}$$

Fig. 4. Typing rules for $\text{WHILE}_{\text{Evd}}^G$. $\Gamma(u)$ maps a constant to its type (e.g. n to int , and true to bool)

$$\boxed{\Gamma \vdash e \rightsquigarrow e' : U}$$

$$\frac{\iota = \gamma(g)}{\Gamma \vdash b^g \rightsquigarrow (\iota b)^g : \text{bool}^g} \text{T-BOOL} \quad \frac{\iota = \gamma(g)}{\Gamma \vdash n^g \rightsquigarrow (\iota n)^g : \text{int}^g} \text{T-INT} \quad \frac{\Gamma(x) = \tau^g}{\Gamma \vdash x \rightsquigarrow x : \tau^g} \text{T-VAR}$$

$$\frac{\forall i \in \{1, 2\}, \Gamma \vdash e_i \rightsquigarrow e'_i : \tau^{g_i} \quad g = g_1 \curlywedge_c g_2}{\Gamma \vdash e_1 \text{ bop } e_2 \rightsquigarrow e'_1 \text{ bop } e'_2 : \tau^g} \text{T-BOP} \quad \frac{\Gamma \vdash e \rightsquigarrow e' : \tau^{g_1} \quad g_1 \preceq_c g \quad (\iota_1, \iota_2) = \text{refine}(\gamma(g_1), \gamma(g))}{\Gamma \vdash (e :: \tau^g) \rightsquigarrow (\iota_1, \iota_2)^g e' : \tau^g} \text{T-CAST}$$

$$\boxed{\Gamma; g_{pc} \vdash c \rightsquigarrow c'}$$

$$\frac{\Gamma(x) = \tau^g \quad \Gamma \vdash e \rightsquigarrow e' : \tau^{g'} \quad g' \preceq_c g \quad (\iota_1, \iota_2) = \text{refine}(\gamma(g'), \gamma(g))}{\Gamma; g_{pc} \vdash x := e \rightsquigarrow x := (\iota_1, \iota_2)^g e'} \text{T-ASSIGN} \quad \frac{\Gamma \vdash e \rightsquigarrow e' : \tau^g \quad g \preceq_c \ell \quad (\iota_1, \iota_2) = \text{refine}(\gamma(g), \gamma(\ell))}{\Gamma; g_{pc} \vdash \text{output}(\ell, e) \rightsquigarrow \text{output}(\ell, (\iota_1, \iota_2)^g e')} \text{T-OUT}$$

$$\frac{\Gamma \vdash e \rightsquigarrow e' : \text{bool}^g \quad \Gamma; g_{pc} \curlywedge_c g \vdash c \rightsquigarrow c' \quad X = \text{WtSet}(c')}{\Gamma; g_{pc} \vdash \text{while } e \text{ do } c \rightsquigarrow \text{while}^X e' \text{ do } c'} \text{T-WHILE} \quad \frac{\Gamma \vdash e \rightsquigarrow e' : \text{bool}^g \quad \Gamma; g_{pc} \curlywedge_c g \vdash c_1 \rightsquigarrow c'_1 \quad \Gamma; g_{pc} \curlywedge_c g \vdash c_2 \rightsquigarrow c'_2 \quad X = \text{WtSet}(c'_1) \cup \text{WtSet}(c'_2)}{\Gamma; g_{pc} \vdash \text{if } e \text{ then } c_1 \text{ else } c_2 \rightsquigarrow \text{if}^X e' \text{ then } c'_1 \text{ else } c'_2} \text{T-IF}$$

Fig. 5. Translation from WHILE^G to $\text{WHILE}_{\text{Evd}}^G$

The other interesting translation rules are T-IF and T-WHILE, which insert a write-set X that includes the set of all variables that might be written to in both the branches and the loop body. We prove that any well-typed term in WHILE^G is translated to another well-typed term in $\text{WHILE}_{\text{Evd}}^G$.

C. Operational Semantics

Runtime constructs: We define additional runtime constructs for our semantics, shown below. The store, δ , maps variables to values with their gradual labels and intervals. The gradual labels of the variables are suffixed on the values for the

purpose of evaluation. κ is a stack of pc labels, each of which is a gradual label, g_{pc} , with the corresponding interval, ι_{pc} .

$$\begin{aligned} \text{Store} \quad \delta &::= \cdot \mid \delta, x \mapsto v \\ \text{PC Stack} \quad \kappa &::= \emptyset \mid (\iota_{pc} g_{pc}) \mid \kappa_1 \triangleright \kappa_2 \\ \text{Actions} \quad \alpha &::= \cdot \mid (\ell, v) \\ \text{Commands} \quad c &::= \cdots \mid \{c\} \mid \text{if } e \text{ then } c_1 \text{ else } c_2 \end{aligned}$$

The stack is used for evaluating nested if statements. The operation $\kappa_1 \triangleright \kappa_2$ indicates that κ_1 is on top of κ_2 in the stack. α is an action, which may be silent or a labeled output. We add two runtime commands. $\{c\}$ is used in evaluating if statements. The curly braces help the monitor keep track of

$$\boxed{\delta / e \Downarrow v}$$

$$\frac{}{\delta / (\iota u)^g \Downarrow (\iota u)^g} \text{M-CONST} \quad \frac{}{\delta / x \Downarrow \delta(x)} \text{M-VAR}$$

$$\frac{\forall i \in \{1, 2\}, \delta / e_i \Downarrow (\iota_i u_i)^{g_i} \quad \iota = \iota_1 \curlywedge \iota_2 \quad g = g_1 \curlywedge g_2 \quad u = u_1 \text{ bop } u_2}{\delta / e_1 \text{ bop } e_2 \Downarrow (\iota u)^g} \text{M-BOP}$$

$$\frac{\delta / e \Downarrow (\iota u)^g \quad \iota' = \iota \boxtimes E}{\delta / E^{g'} e \Downarrow (\iota' u)^{g'}} \text{M-CAST}$$

$$\frac{\delta / e \Downarrow (\iota u)^g \quad \iota \boxtimes E = \text{undef}}{\delta / E^{g'} e \Downarrow \text{abort}} \text{M-CAST-ERR}$$

Fig. 6. Monitor semantics for expressions

the scope of a branch. The if statement without the write set is used in an intermediate evaluation state.

Expression monitoring semantics: Our monitoring semantics for expressions is of the form $\delta / e \Downarrow e'$ as shown in Fig. 6. Rules M-CONST and M-VAR are standard. To perform a binary operation on two values, the operation is performed on the raw values, and the join of their associated intervals and gradual labels is assigned to the computed value. M-CAST refines a value's interval according to the cast evidence. If the refinement is not valid, the execution aborts (M-CAST-ERR). Note that none of these operations modify the gradual label of the variable (the type of store locations remain the same); the operations only refine the intervals of the gradual label.

Commands monitoring semantics: Our monitoring semantics for commands is summarized in Fig. 7 and has the form $\kappa, \delta / c \xrightarrow{\alpha} \kappa', \delta' / c'$. Additional set of rules where the monitor aborts can be found in Fig. 16 in the Appendix. Rules M-PC and M-POP manage commands running in branches or loops. M-POP pops the top-most pc label from the stack, indicating the end of the branch or loop. We use braces around a command, $\{c\}$ to indicate that c is executing in a branch or loop. Such a command is run taking into account only the specific branch's pc stack. When the command execution finishes, the braces are removed and the current pc label is popped off the stack.

Rule M-ASSIGN updates the label-interval of the value being assigned based on the assignment's context ι_{pc} to prevent information leaks. The resulting label-interval is further restricted using the existing label-interval of the variable to ensure that we only refine the set of possible labels. The function $intvl(v)$ returns the label-interval of v . Formally:

$$intvl((\iota u)^g) = \iota$$

The function $refineLB$ refines the lower bound of a value's label-interval and raises it based on the interval ι_{pc} . The $updL$ function is defined in Fig. 8 and uses the interval $restrictLB$ operation. We raise the lower bound of the existing interval

$$\boxed{\kappa, \delta / c \xrightarrow{\alpha} \kappa', \delta' / c'}$$

$$\frac{\kappa, \delta / c_1 \xrightarrow{\alpha} \kappa', \delta' / c'_1}{\kappa, \delta / c_1; c_2 \xrightarrow{\alpha} \kappa', \delta' / c'_1; c_2} \text{M-SEQ}$$

$$\frac{\kappa, \delta / c \xrightarrow{\alpha} \kappa', \delta' / c'}{\kappa \triangleright \iota_{pc} g_{pc}, \delta / \{c\} \xrightarrow{\alpha} \kappa' \triangleright \iota_{pc} g_{pc}, \delta' / \{c'\}} \text{M-PC}$$

$$\frac{}{\iota_{pc} g_{pc} \triangleright \kappa, \delta / \{\text{skip}\} \longrightarrow \kappa, \delta / \text{skip}} \text{M-POP}$$

$$\frac{}{\kappa, \delta / \text{skip}; c \longrightarrow \kappa, \delta / c} \text{M-SKIP}$$

$$\frac{\delta / e \Downarrow v \quad v' = \text{refineLB}(\iota_{pc}, v) \quad v'' = \text{updL}(intvl(\delta(x)), v')}{\iota_{pc} g_{pc}, \delta / x := e \longrightarrow \iota_{pc} g_{pc}, \delta[x \mapsto v''] / \text{skip}} \text{M-ASSIGN}$$

$$\frac{\delta / e \Downarrow v \quad v' = \text{refineLB}(\iota_{pc}, v) \quad v'' = \text{updL}(intvl(\delta(x)), v')}{\iota_{pc} g_{pc}, \delta / \text{output}(\ell, e) \xrightarrow{(\ell, v'')} \iota_{pc} g_{pc}, \delta / \text{skip}} \text{M-OUT}$$

$$\frac{\iota'_{pc} = \iota_{pc} \curlywedge \iota \quad g'_{pc} = g_{pc} \curlywedge g \quad c_i = c_1 \text{ if } b = \text{true} \quad c_i = c_2 \text{ if } b = \text{false}}{\iota_{pc} g_{pc}, \delta / \text{if } (\iota b)^g \text{ then } c_1 \text{ else } c_2 \longrightarrow \iota'_{pc} g'_{pc} \triangleright \iota_{pc} g_{pc}, \delta / \{c_i\}} \text{M-IF}$$

$$\frac{\delta / e \Downarrow v \quad \delta' = \text{rfL}(\delta, X, \iota_{pc} \curlywedge intvl(v))}{\iota_{pc} g_{pc}, \delta / \text{if}^X e \text{ then } c_1 \text{ else } c_2 \longrightarrow \iota_{pc} g_{pc}, \delta' / \text{if } v \text{ then } c_1 \text{ else } c_2} \text{M-IF-REFINE}$$

$$\frac{}{\iota_{pc} g_{pc}, \delta / \text{while}^X e \text{ do } c \longrightarrow \iota_{pc} g_{pc}, \delta / \text{if}^X e \text{ then } (c; \text{while}^X e \text{ do } c) \text{ else skip}} \text{M-WHILE}$$

Fig. 7. Monitor semantics for commands

based on the interval of the newly computed value. Note that $restrictLB$ differs from $refine$ in the upper-bound of the computed interval. If either of these functions return an invalid interval, the execution aborts. The interval need not be checked against the reference's label-interval because the inserted cast would have ruled out unsafe programs earlier. Consider the following program, where $\delta = [x \mapsto ([L, L] 3)^L, y \mapsto ([H, H] 5)^H]$

$$x := ([\perp, \top], [L, L]) ([H, H], [\perp, \top]) y$$

This program tries to cast an H value to \perp , then back to L , which is accepted by the type system. The expression to be assigned to x is first evaluated to $([\perp, \top], [L, L]) ([H, H], [\perp, \top]) ([H, H] 5)^H$, then to $([\perp, \top], [L, L]) ([H, H] 5)^H$, then aborts, because $[H, H] \boxtimes ([\perp, \top], [L, L])$ evaluates to $refine([H, H], [L, L]) = \text{undef}$.

$$\begin{array}{c}
\frac{\text{refine}(\iota_c, \iota) = (_, \iota')}{\text{refineLB}(\iota_c, (\iota u)^g) = (\iota' u)^g} \quad \frac{}{\text{rfL}(\delta, \cdot, \iota) = \delta} \\
\\
\frac{\delta' = \text{rfL}(\delta, X, \iota) \quad v' = \text{refineLB}(\iota, v) \neq \text{undef}}{\text{rfL}((\delta, x \mapsto v), (X, x), \iota) = \delta', x \mapsto v'} \\
\\
\frac{\ell_{1l} \vee \ell_{2l} \preceq \ell_{1r}}{\text{restrictLB}([\ell_{1l}, \ell_{1r}], [\ell_{2l}, \ell_{2r}]) = [\ell_{1l} \vee \ell_{2l}, \ell_{1r}]} \\
\\
\frac{}{\text{updL}(\iota_o, (\iota_n u_n)^g) = (\text{restrictLB}(\iota_o, \iota_n) u_n)^g}
\end{array}$$

Fig. 8. Label-interval operations for the monitor

We explain the assignment rule via examples. Consider a three-point lattice $L \preceq M \preceq H$, the following command $x := [H, H] 5^?$, and two stores $\delta_1 = x \mapsto [M, H] 1^?$ and $\delta_2 = x \mapsto [L, M] 2^?$. Assume the following current pc -interval $\iota_{pc} = [L, H]$. Here, $v = [H, H] 5^?$. The second premise further refines the interval of v to make sure that the pc context is lower than or equal to the interval of the value to be written. This is to prevent low assignments in a high context. For this example, $\text{refine}([L, H], [H, H]) = ([L, H], [H, H])$, so the intervals remain the same. Next, we narrow down the possible label set using the existing value's interval. This is to adhere to our design choice that we do not change the type of the variables and only narrow down the label choices (Section III-A). For δ_1 , $v'' = [H, H] 5^?$, so now x stores a secret value with label H . For store δ_2 , $\text{restrictLB}([L, M], [H, H])$ is not defined and the monitor aborts.

When the old value in the store has a label-interval that is lower than the label-interval of the new value to be stored, the monitor aborts; for instance, under the store where $x \mapsto ([L, L]0)^?$, the monitor aborts the execution of $x := ([H, H]1)^?$, as $\text{restrictLB}([L, L], [H, H])$ is not defined. This is also consistent with our design choice to only refine the label set, not update it.

M-OUT makes similar comparisons as M-ASSIGN to ensure that the output is permitted. Rule M-IF is standard. The pc label is determined by joining the current pc with the gradual label and interval of the branch-predicate's value. Here, the pc stack grows and the branch is placed in the scoping braces. The rule for `while` reduces it to `if`. Rule M-IF-REFINE is the key for preventing implicit leaks. We refine the intervals for variables in both branches according to the write set, X , which contains the set of all variables being updated in either one of the two branches. Fig. 8 includes the auxiliary definitions for refining the intervals of variables in a write set. The function rfL refines the label-interval of values in the store and is defined inductively. Here, it is used to refine the intervals of the variables in the write set to be at least as high as the lower label in the interval of the current pc . When the functions rfL return `undef`, the execution aborts.

Example: Below, we define two initial memories, δ_t maps x

to true and δ_f maps x to false. Both y and z store true initially with y 's label being $?$ and z being L such that $L \sqsubseteq H$.

$$\begin{array}{l}
\delta_y = y \mapsto [L, H]\text{true}^? \\
\delta_z = z \mapsto [L, L]\text{true}^L \\
\delta_t = x \mapsto [H, H]\text{true}^H \\
\delta_f = x \mapsto [H, H]\text{false}^H \\
c_1 = \text{if}^{\{y\}} x \text{ then } y := [L, H]\text{false}^? \text{ else skip} \\
c_2 = \text{if}^{\{z\}} y \text{ then } z := [L, L]\text{false}^L \text{ else skip}
\end{array}$$

Below is the execution starting from the state where x is true.

$$\begin{array}{l}
[L, L] L, (\delta_t, \delta_z, \delta_y) / c_1; c_2 \\
\longrightarrow [L, L] L, (\delta_t, \delta_z, y \mapsto [H, H]\text{true}^?) / \\
\quad \text{if } x \text{ then } y := [L, H]\text{false}^? \text{ else skip}; c_2 \\
\longrightarrow [H, H] H \triangleright [L, L] L, (\delta_t, \delta_z, y \mapsto [H, H]\text{true}^?) / \\
\quad \{y := [L, H]\text{false}^?\}; c_2 \\
\longrightarrow [H, H] H \triangleright [L, L] L, (\delta_t, \delta_z, y \mapsto [H, H]\text{false}^?) / \\
\quad \{\text{skip}\}; c_2 \\
\longrightarrow [H, H] H \triangleright [L, L] L, (\delta_t, \delta_z, y \mapsto [H, H]\text{false}^?) / \\
\quad \text{skip}; c_2 \\
\longrightarrow [L, L] L, (\delta_t, \delta_z, y \mapsto [H, H]\text{false}^?) / c_2 \\
\longrightarrow \text{abort}
\end{array}$$

In the last step, rfL fails, because the operation $\text{refine}([H, H], [L, L])$ produces an invalid label-interval. Now let's see the execution starting from $x \mapsto \text{false}$.

$$\begin{array}{l}
[L, L] L, (\delta_f, \delta_z, \delta_y) / c_1; c_2 \\
\longrightarrow [L, L] L, (\delta_f, \delta_z, y \mapsto [H, H]\text{true}^?) / \\
\quad \text{if } x \text{ then } y := [L, H]\text{false}^? \text{ else skip}; c_2 \\
\longrightarrow [H, H] H \triangleright [L, L] L, (\delta_f, \delta_z, y \mapsto [H, H]\text{true}^?) / \\
\quad \{\text{skip}\}; c_2 \\
\longrightarrow [H, H] H \triangleright [L, L] L, (\delta_f, \delta_z, y \mapsto [H, H]\text{true}^?) / \\
\quad \text{skip}; c_2 \\
\longrightarrow [L, L] L, (\delta_f, \delta_z, y \mapsto [H, H]\text{true}^?) / c_2 \\
\longrightarrow \text{abort}
\end{array}$$

Notice that the label-intervals of y are changed the same way as when we start the execution from δ_t . Ultimately, the program aborts for the same reason.

V. NONINTERFERENCE

To prove noninterference, we extend $\text{WHILE}_{\text{Evd}}^{\text{C}}$ with pairs of values, expressions, and commands to simulate two executions which differ on secret values. This allows us to reduce our noninterference proof to a preservation proof [18].

A. Paired Execution

Syntax: The augmented syntax with pairs is shown below.

$$\begin{array}{l}
\text{Values } v ::= (\iota u)^g \mid \langle \iota_1 u_1 \mid \iota_2 u_2 \rangle^g \\
\text{Cmd. } c ::= \dots \mid \langle \kappa_1, \iota_1, c_1 \mid \kappa_2, \iota_2, c_2 \rangle_g
\end{array}$$

The store δ is extended to contain pairs of values. We also extend commands to be paired but do not allow pairs to be nested; an invariant maintained by our operational semantics. We only use pairs for values and commands whose values and effects are not observable by the adversary (are ‘‘high’’). Pairs of commands are part of the runtime statement, generated as

Expression Semantics: $\boxed{\delta /_i e \Downarrow v}$

$$\frac{}{\delta /_i (\iota u)^g \Downarrow (\iota u)^g} \text{P-CONST} \quad \frac{}{\delta /_i x \Downarrow \text{rd}_i \delta(x)} \text{P-VAR}$$

$$\frac{\delta /_i e \Downarrow v \quad v' = (E, g) \triangleright v}{\delta /_i E^g e \Downarrow v'} \text{P-CAST}$$

Command Semantics: $\boxed{\kappa, \delta /_i c \xrightarrow{\alpha} \kappa', \delta' /_i c'}$

$$\frac{\begin{array}{l} \kappa_i \triangleright \iota_{pc} \Upsilon \iota_i g_{pc} \Upsilon_c g, \delta /_i c_i \xrightarrow{\alpha} \\ \kappa'_i \triangleright \iota_{pc} \Upsilon \iota_i g_{pc} \Upsilon_c g, \delta' /_i c'_i \\ c_j = c'_j \quad \kappa_j = \kappa'_j \quad \{i, j\} = \{1, 2\} \end{array}}{\iota_{pc} g_{pc}, \delta / \langle \kappa_1, \iota_1, c_1 \mid \kappa_2, \iota_2, c_2 \rangle_g \xrightarrow{\alpha} \iota_{pc} g_{pc}, \delta' / \langle \kappa'_1, \iota_1, c'_1 \mid \kappa'_2, \iota_2, c'_2 \rangle_g} \text{P-C-PAIR}$$

$$\frac{\begin{array}{l} c_j = c_1 \text{ if } u_1 = \text{true} \quad c_j = c_2 \text{ if } u_1 = \text{false} \\ c_k = c_1 \text{ if } u_2 = \text{true} \quad c_k = c_2 \text{ if } u_2 = \text{false} \end{array}}{\iota_{pc} g_{pc}, \delta / \text{if } \langle \iota_1 u_1 \mid \iota_2 u_2 \rangle^g \text{ then } c_1 \text{ else } c_2 \xrightarrow{} \iota_{pc} g_{pc}, \delta / \langle \emptyset, \iota_1, c_j \mid \emptyset, \iota_2, c_k \rangle_g} \text{P-LIFT-IF}$$

$$\frac{}{\iota_{pc} g_{pc}, \delta / \langle \emptyset, \iota_1, \text{skip} \mid \emptyset, \iota_2, \text{skip} \rangle_g \xrightarrow{} \iota_{pc} g_{pc}, \delta / \text{skip}} \text{P-SKIP-PAIR}$$

$$\frac{\delta /_i e \Downarrow v \quad v' = \text{refineLB}(\iota_{pc}, v)}{\iota_{pc} g_{pc}, \delta /_i x := e \xrightarrow{} \iota_{pc} g_{pc}, \delta [x \mapsto \text{upd}_i \delta(x) v'] /_i \text{skip}} \text{P-ASSIGN}$$

$$\frac{\begin{array}{l} \delta /_i e \Downarrow v \\ v' = \text{refineLB}(\iota_{pc}, v) \quad v'' = \text{updL}[\ell, \ell] v' \end{array}}{\iota_{pc} g_{pc}, \delta /_i \text{output}(\ell, e) \xrightarrow{(i, \ell, v'')} \iota_{pc} g_{pc}, \delta /_i \text{skip}} \text{P-OUT}$$

Fig. 9. Selected rules of paired executions

a result of evaluating a branching statement. Each command represents an independent execution, capable of changing its own pc stack. As a result, we include local pc stacks in the pair with each command. The rationale behind additional pc stacks in command pairs is explained with the semantics.

Label-interval operations on pairs: The interval of a paired value is a pair of intervals, defined below.

$$\text{intvl}(\langle \iota_1 u_1 \mid \iota_2 u_2 \rangle^g) = \langle \iota_1 \mid \iota_2 \rangle$$

The intersection of an interval and a paired value is defined as follows. Other extensions to label-interval operations can be found in the full version of the paper [19].

$$\iota \bowtie \langle \iota_1 u_1 \mid \iota_2 u_2 \rangle^g = \langle \iota \bowtie \iota_1 u_1 \mid \iota \bowtie \iota_2 u_2 \rangle^g$$

Memory read and update operations: As we allow the intervals of values to be refined, the store read (rd) and update (upd) operations for paired values need to make sure that the correct paired value is read or updated. These functions are shown in Fig. 17 in the Appendix.

Operational semantics for pairs: The operational semantics are augmented with an index, i . The judgments now are of the form $\delta /_i e \Downarrow e'$ and $\kappa, \delta /_i c \xrightarrow{} \kappa', \delta' /_i c'$. The index i indicates which branch of a pair is executing (when $i \in \{1, 2\}$) or if it is a top-level execution (when i is omitted). Most of the rules can be directly obtained by adding the i to the monitor semantics shown in Fig. 6 and 7. Rules that deal with pairs, including read and write to the store need to be modified. We explain important rule changes (shown in Fig. 9).

Rule P-VAR uses the function $\text{rd}_i v$ to retrieve the value indexed by i within v . To evaluate a cast over a pair of values, we push the cast inside the pair (P-CAST).

Each command in the pair (P-C-PAIR) can make progress independently and the premise of the rule is indexed by the corresponding i . Here κ_i is the pc stack specific to c_i . Consider a command $c = \langle c_1 \mid c_2 \rangle$, where both c_1 and c_2 have nested if statements. The execution of c will create different κ_1 and κ_2 when executing c_1 and c_2 . Next, ι_i is the pc label-interval demonstrating that c is supposed to execute in a “high” context (unobservable by the adversary). The bottom pc in the stack is joined with ι_i . We will come back to this point when explaining the typing rules.

Rule P-LIFT-IF lifts the pair that appears as branch conditions to generate a paired command. The resulting commands on each side of the pair are determined by the value in the corresponding side of the branch condition. The branching context ι_i is the runtime interval of the branching condition. The initial local pc stack is empty.

Note that the individual branches do not contain pairs of commands. The only rule that generates paired command is P-LIFT-IF. To see how the semantics prevent nesting command pairs and how paired execution represents low runs with different secrets, consider the program in Listing 6. Assume that a and b are variables containing paired values such that $a = \langle u_{a1} \mid u_{a2} \rangle^H$ and $b = \langle u_{b1} \mid u_{b2} \rangle^H$, meaning both a and b contain secrets and u_{a1} and u_{b1} are values for the first execution and u_{a2} and u_{b2} are for the second. We ignore the intervals in this example for simplicity of exposition. On line 1, we use the P-LIFT-IF rule since we branch on a pair of values to create paired commands. In the first execution, if $u_{a1} = \text{true}$, we take the then branch. When evaluating b inside the branch we take the first part of the pair using the expression evaluation rules and rd_i operation (Fig. 17) for $i = 1$, i.e., $\text{rd}_1 b = [b]_1 = u_{b1}$. Thus, the branch on line 2 becomes: if u_{b1} then ... while the remaining parts remain the same. Similarly in the second execution, based on the value of u_{a2} , either the then branch or the else branch is chosen. If the then branch is chosen, the branch on line 2 becomes if u_{b2} then ... as we are in the second execution of the branch ($i = 2$) on line 1 and $\text{rd}_2 b = [b]_2 = u_{b2}$. Generating two different runs of the program is sufficient for reasoning about noninterference, which is what the projection semantics do.

Local pc refinements in pairs are forgotten when both sides of the pair finish executing in P-SKIP-PAIR. This is similar to the P-POP rule where pc for the branch or loop is forgotten.

Rule P-ASSIGN deals with the complexity of pairs updating

```

1 if a then
2   if b then y := [⊥, ⊤]1?
3   else skip
4 else y := [⊥, ⊤]2?

```

Listing 6. Example program to explain branching on pairs

the store in one branch with the helper function $\text{upd}_i v_o v_n$ (defined in Fig. 17). The refinement of labels during store updates is the same as the monitor semantics. When the update comes from a specific branch of execution ($i \in \{1, 2\}$), the value for the other branch should be preserved. If the value in the store is already a pair, only the i^{th} sub-expression is updated. Reconsider the example in Listing 6. The assignment on line 2 happens in either of the two branches, or both the branches depending on the values of u_{a1} to u_{b2} . If it happens in only the first projection, the first part of the value-pair in y is updated. Suppose that $y = \langle [H, \top]0 \mid [H, \top]42 \rangle^?$, initially, and $u_{a1} = u_{b1} = \text{true}$. Then, the value of y after the assignment on line 2 becomes $y = \langle [H, \top]1 \mid [H, \top]42 \rangle^?$. If $u_{a2} = \text{false}$, then the else branch is taken, and at the end of the assignment on line 4 the value of y is updated to $y = \langle [H, \top]1 \mid [H, \top]2 \rangle^?$. The first part of the pair is already updated through the then branch as we evaluate the two runs one after the other when branching on a pair of values.

If the store value is not a pair, the value becomes a pair where the i^{th} sub-expression is the updated value, and the other sub-expression is the old value. Considering the same example as above, if initially $y = ([H, \top]42)^?$, then at the end of then branch with $u_{a1} = \text{true}$, the updated value of y is $y = \langle [H, \top]1 \mid [H, \top]42 \rangle^?$. When updates happen at the top-level, the entire value in the store should be updated. The first rule applies when either the old or the new value is a pair and the second rule applies when none of them are pairs. Note that this is the reason why the intervals in a pair may differ.

The output rule is mostly the same. The event being output now includes the index to aid the statement and proof of noninterference. The P-IF-REFINE rule (for M-IF-REFINE) uses an augmented version of rfL , which only refines label-intervals for the i^{th} branch.

B. Semantic Soundness and Completeness

To connect the semantics of the extended language with pairs to the monitor semantics, we prove soundness and completeness theorems. These theorems depend on projections of the store, expression- and command-configurations. Similar to the value projection seen before, the goal of these projections is to obtain one execution from a paired execution.

The projection of a paired value, a paired interval, a normal value and interval are straightforward and defined in Fig. 17. The projection of stores (δ) and traces (\mathbb{T}) is inductively defined as shown in Fig. 10. The projection function only keeps the output events produced by the execution of concern and ignores output performed by the other execution. The projection function for expression configurations is $[\delta / e]_i = [\delta]_i / e$ and for command configurations is defined in Fig. 10.

Store projection:

$$[\cdot]_i = \cdot \quad [\delta, x \mapsto v]_i = [\delta]_i, x \mapsto [v]_i$$

Trace projection:

$$\begin{aligned} [\cdot]_i &= \cdot \\ [\mathbb{T}, (\ell, v)]_i &= [\mathbb{T}]_i, (\ell, [v]_i) \\ [\mathbb{T}, (j, \ell, v)]_i &= [\mathbb{T}]_i, (\ell, v) \quad \text{if } i = j \\ [\mathbb{T}, (j, \ell, v)]_i &= [\mathbb{T}]_i \quad \text{if } i \neq j \end{aligned}$$

Command-configuration projection:

$$\begin{aligned} &\frac{}{[\iota_{pc} g_{pc}, \delta / \text{skip}]_i = \iota_{pc} g_{pc}, [\delta]_i / \text{skip}} \\ &\frac{[\kappa, \delta / c_1]_i = \kappa', \delta' / c'_1}{[\kappa, \delta / c_1; c_2]_i = \kappa', \delta' / c'_1; c'_2} \\ &\frac{}{[\iota_{pc} g_{pc}, \delta / x := e]_i = \iota_{pc} g_{pc}, [\delta]_i / x := e} \\ &\frac{[\kappa, \delta / c]_i = \kappa', \delta' / c'}{[\kappa \triangleright \iota_{pc} g_{pc}, \delta / \{c\}]_i = \kappa' \triangleright \iota_{pc} g_{pc}, \delta' / \{c'\}} \\ &\frac{}{[\iota_{pc} g_{pc}, \delta / \text{output}(\ell, e)]_i = \iota_{pc} g_{pc}, [\delta]_i / \text{output}(\ell, e)} \\ &\frac{\forall \{i, j\} \in \{1, 2\}, c'_i = \begin{cases} \text{skip} & \text{if } c_i = \text{skip} \text{ and } c_j \neq \text{skip} \\ \{c_i\} & \text{else} \end{cases}}{[\iota_{pc} g_{pc}, \delta / \langle \kappa_1, \iota_1, c_1 \mid \kappa_2, \iota_2, c_2 \rangle_g]_i = \kappa_i \triangleright (\iota_{pc} \Upsilon \iota_i) (g_{pc} \Upsilon_c g) \triangleright \iota_{pc} g_{pc}, [\delta]_i / c'_i} \\ &\frac{}{[\iota_{pc} g_{pc}, \delta / \langle \emptyset, \iota_1, \text{skip} \mid \emptyset, \iota_2, \text{skip} \rangle_g]_i = \kappa_i \triangleright (\iota_{pc} \Upsilon \iota_i) (g_{pc} \Upsilon_c g) \triangleright \iota_{pc} g_{pc}, [\delta]_i / \{\text{skip}\}} \\ &\frac{}{[\iota_{pc} g_{pc}, \delta / \text{if } v \text{ then } c_1 \text{ else } c_2]_i = \iota_{pc} g_{pc}, [\delta]_i / \text{if } [v]_i \text{ then } c_1 \text{ else } c_2} \\ &\frac{}{[\iota_{pc} g_{pc}, \delta / \text{if}^X e \text{ then } c_1 \text{ else } c_2]_i = \iota_{pc} g_{pc}, [\delta]_i / \text{if}^X e \text{ then } c_1 \text{ else } c_2} \\ &\frac{}{[\iota_{pc} g_{pc}, \delta / \text{while}^X e \text{ do } c]_i = \iota_{pc} g_{pc}, [\delta]_i / \text{while}^X e \text{ do } c} \end{aligned}$$

Fig. 10. Projections

The interesting case is the projection of a command pair. We reassemble the pc stack and wrap c_i with curly braces to reflect the fact that these pairs only appear in an if branch.

The Soundness theorem ensures that if a configuration can transition to another configuration, then its projection can transition to the projection of the resulting configuration, generating the same trace modulo projection. The Completeness theorem ensures that if both projections of a configuration terminate, then the configuration terminates in an equivalent state. We write, $\vdash \kappa, \delta /_i c$ wf, to indicate that the configuration is

$$\boxed{\Gamma; \kappa \vdash_r c}$$

$$\frac{\Gamma; \kappa \vdash_r c}{\Gamma; \kappa \triangleright \iota_{g_{pc}} \vdash_r \{c\}} \text{R-POP} \quad \frac{\Gamma; \iota_{g_{pc}} \vdash c}{\Gamma; \iota_{g_{pc}} \vdash_r c} \text{R-END}$$

$$\frac{\Gamma; \kappa \triangleright \iota_{pc} g_{pc} \vdash_r c_1 \quad \Gamma; \iota_{pc} g_{pc} \vdash c_2 \quad \kappa \neq \emptyset}{\Gamma; \kappa \triangleright \iota_{pc} g_{pc} \vdash_r c_1; c_2} \text{R-C-SEQ}$$

$$\frac{\forall i \in 1, 2, \Gamma; \kappa_i \triangleright (\iota_{pc} \vee \iota_i) (g_{pc} \vee g) \vdash_r c_i \quad \iota_i \vdash g \in H(\ell_A)}{\Gamma; \iota_{pc} g_{pc} \vdash_r \langle \kappa_1, \iota_1, c_1 \mid \kappa_2, \iota_2, c_2 \rangle_g} \text{R-C-PAIR}$$

$$\frac{\Gamma \vdash e : \text{bool}^g \quad \iota_g = \gamma(g) \quad \forall i \in \{1, 2\}, \Gamma; \iota_{pc} \vee \iota_g g_{pc} \vee g \vdash_r c_i}{\Gamma; \iota_{pc} g_{pc} \vdash_r \text{if } e \text{ then } c_1 \text{ else } c_2} \text{R-C-IF}$$

Fig. 11. Typing rules for commands with pairs

well-formed (defined in Appendix A). Theorems 1 and 2 are the formal soundness and completeness theorem statements. The proofs can be found in the full version of the paper [19].

Theorem 1 (Soundness). *If $\kappa, \delta / c \xrightarrow{*} \kappa', \delta' / c'$ where $\vdash \kappa, \delta / c$ wf, then $\forall i \in \{1, 2\}, \llbracket \kappa, \delta / c \rrbracket_i \xrightarrow{[\mathbb{T}]_i} \llbracket \kappa', \delta' / c' \rrbracket_i$*

Theorem 2 (Completeness). *If $\forall i \in \{1, 2\}, \llbracket \kappa, \delta / c \rrbracket_i \xrightarrow{\mathbb{T}_i} \kappa_i, \delta_i / \text{skip}$ and $\vdash \kappa, \delta / c$ wf, then $\exists \kappa', \delta'$ s.t. $\kappa, \delta / c \xrightarrow{*} \kappa', \delta' / \text{skip}$, $\llbracket \kappa', \delta' / \text{skip} \rrbracket_i = \kappa_i, \delta_i / \text{skip}$ and $\mathbb{T}_i = \llbracket \mathbb{T} \rrbracket_i$.*

C. Preservation

Before we explain the typing rules for the extended configuration, we define another label relation. A gradual label is said to be “high” w.r.t an attacker, if the lower label in the interval is not lower than or equal to the level of the attacker.

$$\frac{\iota = [\ell_l, \ell_r] \quad \ell_l \not\leq \ell_A \quad \iota \sqsubseteq \gamma(g)}{\iota \vdash g \in H(\ell_A)}$$

All the pair typing rules are parameterized over attacker’s label ℓ_A , which we omit from the rules for simplicity. The typing rule for value-pairs is shown below. The second premise checks that the interval is representative of the gradual type U . The last premise checks if U ’s security label is high, meaning this pair of values is non-observable to the adversary.

$$\frac{\forall i \in \{1, 2\}, \Gamma \vdash \iota_i u_i : \tau^g \quad \iota \sqsubseteq \gamma(g) \quad \iota_i \vdash (g) \in H(\ell_A)}{\Gamma \vdash \langle \iota_1 u_1 \mid \iota_2 u_2 \rangle : \tau^g} \text{R-V-PAIR}$$

The judgement for typing commands with pairs is of the form $\Gamma; \kappa \vdash_r c$. Fig. 11 summarizes these typing rules.

Rule R-POP types the inner command with only the top part of the pc stack. When the pc stack contains only one element, R-END directly uses command typing. For pairs, R-C-PAIR first checks that each c_i is well-typing, using the pc context

Store typing:

$$\frac{}{\vdash \cdot \cdot} \text{T-S-EMP} \quad \frac{\vdash \delta : \Gamma \quad \Gamma \vdash v : U}{\vdash \delta, x \mapsto v : \Gamma, x : U} \text{T-S-IND}$$

Configuration typing:

$$\frac{\vdash \delta : \Gamma \quad \Gamma; \kappa \vdash_r c}{\vdash \kappa, \delta, c} \text{T-CONF}$$

Trace typing:

$$\frac{\vdash v : U \quad \vdash \text{intvl}(v) \preceq [\ell, \ell]}{\vdash (\ell, v)} \text{T-A-OUT}$$

$$\frac{\vdash v : U \quad \ell \not\leq \ell_A \quad \vdash \text{intvl}(v) \preceq [\ell, \ell]}{\vdash (i, \ell, v)} \text{T-A-OUTI}$$

$$\frac{}{\vdash \cdot} \text{T-T-EMP} \quad \frac{\vdash \alpha \quad \vdash \mathbb{T}}{\vdash \alpha, \mathbb{T}} \text{T-T-IND}$$

Fig. 12. Store, trace and configuration typing

assembled from the local pc context. The second premise makes sure that these commands are typed (executed) in a high context. Here ι_i is the witness for g , which demonstrates that c_i are high commands. The sequencing statement types the second command using only the last pc on the stack because the execution order is from left to right. We can only encounter branches in the first part of a sequencing statement and not the second part before beginning the execution of the second command in the sequence. The typing rule for if statements without a write set is straightforward.

We define store, trace and configuration typing in Fig. 12. The store δ types in the typing environment Γ if all variables in δ are mapped to their respective type and gradual label in Γ . We define top-level configuration typing as $\vdash \kappa, \delta, c$. To type traces and actions, the output value needs to be well-typed, and the label-interval of the value has to be lower than or equal to the channel label.

Using these definitions, we prove that our paired execution semantics preserve the configuration typing and generate a well-typed trace (Theorem 3). We write, $\vdash \kappa, \delta /_i c$ sf for $i \in \{1, 2\}$, to indicate that the configuration is safe. We say a configuration is safe if all of the following hold:

- 1) if $i \in \{1, 2\}$, then $\kappa \in H(\ell_A)$, $\forall x \in \text{WtSet}(c)$, $\text{intvl}(\delta(x)) \in H(\ell_A)$
- 2) if $c = \text{if } \langle _ \mid _ \rangle \text{ then } c_1 \text{ else } c_2$, then $\forall x \in \text{WtSet}(c)$, $\text{intvl}(\delta(x)) \in H(\ell_A)$
- 3) if $c = \langle \kappa_1, \iota_1, c_1 \mid \kappa_2, \iota_2, c_2 \rangle_g$, then $\forall i \in \{1, 2\}$, $\iota_i \vdash g \in H(\ell_A)$, and $\forall x \in \text{WtSet}(c)$, $\text{intvl}(\delta(x)) \in H(\ell_A)$

Theorem 3 (Preservation). *If $\kappa, \delta / c \xrightarrow{*} \kappa', \delta' / c'$ with $\vdash \kappa, \delta, c$ and $\vdash \kappa, \delta / c$ sf, then $\vdash \kappa', \delta', c'$ and $\vdash \mathbb{T}$*

Store equivalence:

$$\frac{g \preceq_c \ell_A \quad \iota \sqsubseteq \gamma(g) \quad \vdash (\iota u)^g : U}{\vdash (\iota u)^g \approx_{\ell_A} (\iota u)^g : U} \text{EQV-L}$$

$$\frac{\forall i \in \{1, 2\}, \iota_i \vdash g \in H(\ell_A) \quad \vdash (\iota_i u_i)^g : U}{\vdash (\iota_1 u_1)^g \approx_{\ell_A} (\iota_2 u_2)^g : U} \text{EQV-H}$$

$$\frac{}{\vdash \cdot \approx_{\ell_A} \cdot \dots} \text{EQS-EMP}$$

$$\frac{\vdash \delta_1 \approx_{\ell_A} \delta_2 : \Gamma \quad \vdash v_1 \approx_{\ell_A} v_2 : U}{\vdash \delta_1, x \mapsto v_1 \approx_{\ell_A} \delta_2, x \mapsto v_2 : \Gamma, x : U} \text{EQS-IND}$$

Trace equivalence:

$$\frac{}{\vdash \square \approx_{\ell_A} \square} \text{EQT-E}$$

$$\frac{\vdash \mathbb{T}_1 \approx_{\ell_A} \mathbb{T}_2 \quad \ell_1 = \ell_2 \preceq \ell_A \quad v_1 = v_2}{\vdash (\ell_1, v_1) :: \mathbb{T}_1 \approx_{\ell_A} (\ell_2, v_2) :: \mathbb{T}_2} \text{EQT-L}$$

$$\frac{\vdash \mathbb{T}_1 \approx_{\ell_A} \mathbb{T}_2 \quad \ell_1 \not\preceq \ell_A}{\vdash (\ell_1, v_1) :: \mathbb{T}_1 \approx_{\ell_A} \mathbb{T}_2} \text{EQT-HL}$$

$$\frac{\vdash \mathbb{T}_1 \approx_{\ell_A} \mathbb{T}_2 \quad \ell_2 \not\preceq \ell_A}{\vdash \mathbb{T}_1 \approx_{\ell_A} (\ell_2, v_2) :: \mathbb{T}_2} \text{EQT-HR}$$

Fig. 13. Equivalence definitions

D. Noninterference

We show that the gradual type system presented above satisfies termination-insensitive noninterference. We start by defining equivalence for values and stores (Fig. 13). Two values are said to be equivalent to an adversary at level ℓ_A if either they are both visible to the adversary and are the same, or neither are observable by the adversary. We also define equivalence of traces w.r.t an adversary at level ℓ_A in Fig. 13. The following noninterference theorem (Theorem 4) states that given a program and two stores equivalent for an adversary at level ℓ_A , if the program terminates in both the runs, then the ℓ_A -observable actions on both runs are the same.

Theorem 4 (Noninterference). *Given an adversary label ℓ_A , a program c , and two stores δ_1, δ_2 , s.t., $\vdash \delta_1 \approx_{\ell_A} \delta_2 : \Gamma$, and $\Gamma; [\perp, \perp] \perp \vdash c$, and $\forall i \in \{1, 2\}, [\perp, \perp] \perp, \delta_i / c \xrightarrow{\mathbb{T}_i}^* \kappa_i, \delta'_i, \text{skip}$, then $\vdash \mathbb{T}_1 \approx_{\ell_A} \mathbb{T}_2$.*

We know that only when $\ell \not\preceq \ell_A$, the individual runs can produce (i, ℓ, v) because the two runs diverge only when branching on pairs. Similarly, $(\ell, \langle v_1 | v_2 \rangle)$ can only be produced if $\ell \not\preceq \ell_A$, because pairs can only be typed if each individual interval in the pair is high and rule P-OUT makes sure ℓ is lower than or equal to the pair's interval. We prove a simple lemma that establishes that given a well-typed trace of actions $\mathbb{T}, \vdash [\mathbb{T}]_1 \approx_{\ell_A} [\mathbb{T}]_2$. By combining

Labels and intervals:

$$\frac{}{\ell \sqsubseteq ?} \quad \frac{}{\ell \sqsubseteq \ell} \quad \frac{\ell'_1 \preceq \ell_1 \quad \ell_2 \preceq \ell'_2}{[\ell_1, \ell_2] \sqsubseteq [\ell'_1, \ell'_2]}$$

Expressions:

$$\frac{\iota_1 \sqsubseteq \iota_2 \quad g_1 \sqsubseteq g_2}{(\iota_1 u)^{g_1} \sqsubseteq (\iota_2 u)^{g_2}} \quad \frac{}{x \sqsubseteq x} \quad \frac{e_1 \sqsubseteq e'_1 \quad e_2 \sqsubseteq e'_2}{e_1 \text{ bop } e_2 \sqsubseteq e'_1 \text{ bop } e'_2}$$

$$\frac{E \sqsubseteq E' \quad g_1 \sqsubseteq g_2 \quad e_1 \sqsubseteq e_2}{E_{g_1} e_1 \sqsubseteq E'_{g_2} e_2}$$

Commands:

$$\frac{}{\text{skip} \sqsubseteq \text{skip}} \quad \frac{c_1 \sqsubseteq c'_1 \quad c_2 \sqsubseteq c'_2}{c_1; c_2 \sqsubseteq c'_1; c'_2}$$

$$\frac{e_1 \sqsubseteq e_2}{x := e_1 \sqsubseteq x := e_2} \quad \frac{e_1 \sqsubseteq e_2}{\text{output}(\ell, e_1) \sqsubseteq \text{output}(\ell, e_2)}$$

$$\frac{e_1 \sqsubseteq e_2 \quad c_1 \sqsubseteq c'_1 \quad c_2 \sqsubseteq c'_2}{\text{if}^X e_1 \text{ then } c_1 \text{ else } c_2 \sqsubseteq \text{if}^X e_2 \text{ then } c'_1 \text{ else } c'_2}$$

$$\frac{e_1 \sqsubseteq e_2 \quad c_1 \sqsubseteq c_2}{\text{while}^X e_1 \text{ do } c_1 \sqsubseteq \text{while}^X e_2 \text{ do } c_2}$$

Store, Types and Typing-context:

$$\frac{g \sqsubseteq g'}{\tau^g \sqsubseteq \tau^{g'}} \quad \frac{\forall x \in \Gamma. \Gamma(x) \sqsubseteq \Gamma'(x)}{\Gamma \sqsubseteq \Gamma'}$$

PC-stack and Configurations:

$$\frac{}{\emptyset \sqsubseteq \emptyset} \quad \frac{\iota \sqsubseteq \iota' \quad g \sqsubseteq g'}{\iota g \sqsubseteq \iota' g'} \quad \frac{\kappa_1 \sqsubseteq \kappa'_1 \quad \kappa_2 \sqsubseteq \kappa'_2}{\kappa_1 \triangleright \kappa_2 \sqsubseteq \kappa'_1 \triangleright \kappa'_2}$$

$$\frac{\forall x \in \delta. \delta(x) \sqsubseteq \delta'(x)}{\delta \sqsubseteq \delta'} \quad \frac{\kappa \sqsubseteq \kappa' \quad \delta \sqsubseteq \delta' \quad c \sqsubseteq c'}{\kappa, \delta / c \sqsubseteq \kappa', \delta' / c'}$$

Fig. 14. Precision relations

the Preservation, Soundness, and Completeness Theorems, it follows that our gradual type system satisfies termination-insensitive noninterference.

VI. GRADUAL GUARANTEES

The gradual guarantees state that if a program with more precise labels type-checks and is accepted by the runtime semantics of the gradual type system, then the same program with less precise labels is also accepted by the gradual type system. To establish these guarantees, we define a precision relation between labels, expressions, and commands. The precision relations are shown in Fig. 14.

Our type system with gradual labels satisfies the static gradual guarantee. The dynamic gradual guarantee is also ensured by our calculus, i.e., if a command takes a step under

a store and *pc* stack, then a less precise command can also take a step under a less precise store and *pc* stack.

Theorem 5 (Static Guarantee). *If $\Gamma_1; g_1 \vdash c_1$, $\Gamma_1 \sqsubseteq \Gamma_2$, $g_1 \sqsubseteq g_2$, and $c_1 \sqsubseteq c_2$, then $\Gamma_2; g_2 \vdash c_2$.*

Theorem 6 (Dynamic Guarantee). *If $\kappa_1, \delta_1 / c_1 \xrightarrow{\alpha_1} \kappa'_1, \delta'_1 / c'_1$ and $\kappa_1, \delta_1 / c_1 \sqsubseteq \kappa_2, \delta_2 / c_2$, then $\kappa_2, \delta_2 / c_2 \xrightarrow{\alpha_2} \kappa'_2, \delta'_2 / c'_2$ such that $\kappa'_1, \delta'_1 / c'_1 \sqsubseteq \kappa'_2, \delta'_2 / c'_2$ and $\alpha_1 = \alpha_2$.*

VII. DISCUSSION

Monitor comparison: Our monitor is a hybrid monitor. The differences between our monitor and traditional hybrid monitors (c.f., [13], [17], [21]–[23]) are that (1), we update the memory before executing the branch, while other hybrid monitors update the labels at the merge points; (2), our monitor will not upgrade variables with fixed static labels (the monitor will abort) and only refine label intervals for dynamically labeled variables. Fig. 15 highlights the differences between our monitor and other information flow monitors using the example program in Listing 4 and a two point lattice ($L \preceq H$). We show two cases for our monitor, differing in y 's security label ($?$, or L). Note that hybrid monitors including ours have the same behavior regardless of x 's value. When y has a dynamic label, our monitor is more precise than NSU, as precise as permissive-upgrade [12], and less precise than a traditional hybrid monitor [13]. Ours can be as precise as a traditional hybrid monitor if z 's initial value is $([L, H]\text{true})^?$. This would allow us to refine z 's label interval, and only abort at the output. When y has a static label, or an effectively static label (i.e., $([L, L]\text{true})^L$), our monitor is the least precise and will abort at the first branch. This rigidity is due to our decision to not update variables' label-intervals, except by refinement.

Implicit leaks manifested in noninterference proofs: Let's revisit the example at the end of Section IV-C to see how the implicit leak manifests in the paired execution and why it leads to our current design; where we do not use the write sets in the if statements and simply refine the label-intervals. Because x is H , it's initialized with a paired value.

$$\begin{aligned} \delta &= x \mapsto \langle [H, H]\text{true}^H \mid [H, H]\text{false}^H \rangle, \\ &\quad y \mapsto [L, H]\text{true}^?, z \mapsto [L, L]\text{true}^L \\ c_1 &= \text{if } x \text{ then } y := [L, H]\text{false}^? \text{ else skip} \\ c_2 &= \text{if } y \text{ then } z := [L, L]\text{false}^L \text{ else skip} \end{aligned}$$

$$\begin{aligned} [L, L] L, \delta / c_1; c_2 &\longrightarrow \\ \dots / \text{if } \langle [H, H]\text{true}^H \mid [H, H]\text{false}^H \rangle \text{ then} & \\ \quad y := [L, H]\text{false}^? \text{ else skip}; c_2 &\longrightarrow \\ \dots / \langle \dots, y := [L, H]\text{false}^?; c_2 \mid \dots, \text{skip}; c_2 \rangle &\longrightarrow \end{aligned}$$

The variable y is updated only in the left branch. To prove soundness and completeness of the paired semantics, the two executions should be independent. Therefore, we try to update y in the store as $\langle [H, H]\text{false} \mid [L, H]\text{true} \rangle^?$. However, this pair is not well-formed because pairs are only well-typed if both intervals are in H . Clearly, the right branch of y does not satisfy this requirement. Therefore, we cannot

prove preservation for the assignment case. For preservation to succeed, we would need to refine the right branch to be $[H, H]\text{true}$ when assigning to y in the left execution. But then the two executions are no longer independent, which breaks soundness (i.e., the projected execution is not guaranteed to make progress or stay in the same state).

With these constraints in place, we need to refine y before the branch, which ultimately leads to our final design.

VIII. RELATED WORK

Static Information Flow Type Systems: Quite a few type-systems have been proposed to statically enforce noninterference by annotating variables with labels. Volpano et al. [1] present the first type-system with information flow labels that satisfies a variant of noninterference, also known as termination-insensitive noninterference. If all variables are annotated with concrete security labels, our type system behaves the same as a flow-insensitive information flow type system. Being a gradual type system, we can additionally accept programs with no security labels and enforce termination-insensitive noninterference at runtime. Our formalization borrows the proof-technique from FlowML, presented by Pottier and Simonet [18], for enforcing noninterference using pairs.

Static type systems that resemble gradual typing: JFlow [24] (and later Jif) includes polymorphic labels, for which programmers can specify the upper bound of a polymorphic label. Polymorphic labels are essentially (bounded) universally quantified labels and these labels are instantiated by concrete labels at runtime. There is no label refinement associated with polymorphic labels. Jif also allows run-time labels (also called dynamic labels). These are runtime representation of label objects that users can generate and perform tests on. This is not to be confused with the dynamic label ($?$) in gradual typing. Runtime labels do not mean unknown security labels, nor are they refined at runtime. Finally, label inference is a widely used compile-time algorithm to reduce programmers' annotation burdens. A flow-insensitive type system with the most powerful inference algorithm is less permission than our system. Ill-typed programs, rejected by a type system, can be accepted by our type system. Their safety is ensured by our runtime monitors.

Purely dynamic monitors: Dynamic approaches use a runtime monitor to track the flow of information through the program. The labels are mostly flow-sensitive in nature. Austin and Flanagan [11] present a purely dynamic information flow monitoring approach that disallows assignments to public values in secret contexts. Our monitor semantics follows a similar approach to prevent information leaks at runtime. Subsequent work presents approaches to make the analysis more permissive and amenable to dynamic languages [12], [25]–[27]. A recent paper shows the equivalence between coarse-grained and fined-grained dynamic monitors [28]. Detailed comparisons between our monitor and dynamic monitors can be found in Fig. 15.

$y = \text{true}^L, z = \text{true}^L$ Program	$x = \text{false}^H$ NSU/Permissive	$x = \text{true}^H$		Hybrid	WHILE ^G _{Evd}	
		NSU	Permissive		$y = [L, H]\text{true}^L$	$y = [L, L]\text{true}^L$
if x then $y := \text{false}^L$ if y then $z := \text{false}^L$ output(L, z)	branch not taken branch taken output(L, false)	branch taken $pc = H, \text{abort}$	branch taken $y \uparrow P$ abort	$y \uparrow H$ $z \uparrow H$ abort	$y \uparrow [H, H]$ try $z \uparrow [H, H], \text{abort}$	try $y \uparrow [H, H], \text{abort}$

Fig. 15. Comparison of monitor behavior. $y \uparrow \ell$ denotes monitor’s attempt to update y ’s label (interval).

Hybrid monitors: To leverage the benefits of static and dynamic approaches for precision and permissiveness, researchers have also proposed hybrid approaches to enforce noninterference [13], [17], [21]–[23], [29]–[31]. We demonstrate that the hybrid monitoring approach is suitable for generating runtime behavior of gradual types that rely on refining label intervals. Gradual typing has the added benefit of allowing programmers to reject ill-typed programs. As shown in Fig. 15, our monitor aborts earlier than a typical hybrid monitor because of the lack of support for label updating. We support termination-insensitive noninterference, while others support progress-sensitive noninterference [22], [31].

Gradual information flow type systems: More closely related to our work are works on gradual security types. Disney and Flanagan [3] study gradual security types for a pure lambda calculus, and Fennell and Thiemann [4] present a gradual type system for a calculus with ML-style references. However, these works are based on adding explicit programmer-provided checks and casts to the code. Fennell and Thiemann [5] extend their prior work to object-oriented programs in a flow-sensitive setting for a Java-like language. They use a hybrid approach to perform effect analysis that upgrades the labels of variables similar to the write set used in our analysis. At runtime, these systems cast the dynamic label to a fixed label, rather than a set of possible labels, and the monitors updates labels of memory locations, which we do not do. On the other hand, our approach has fixed gradual labels and refines only the label-intervals associated with the value to satisfy dynamic gradual guarantee. More recently, Toro et al. [6] presented a type-driven gradual type system for a higher-order language with references based on abstract gradual typing [9]. Their formalization satisfies the static gradual guarantee, but sacrifices the dynamic gradual guarantee for noninterference. They briefly discuss the idea of using hybrid approaches and faceted evaluation for regaining the dynamic gradual guarantee. The language presented in this paper is simpler than their language but has mutable global variables and hence, a similar issue with proving noninterference while satisfying the dynamic gradual guarantee.

GLIO [7] presents another interpretation of gradual information flow types that enjoys both noninterference and gradual guarantees. *GLIO* is the most expressive among the above-mentioned projects; it includes higher-order functions, general references, coarse-grained information flow control, and first-class labels. *GLIO*’s monitor decides the concrete label for a dynamically labeled reference at allocation time. While

avoiding problems stemmed from refining label intervals, the concrete label results in a less permissive approach.

Our work explores yet another design space of gradual information flow types and highlights the necessity of a hybrid approach for a system that refines label intervals to ensure both noninterference and the gradual guarantees. Extending our static type system to include higher-order functions and references would require a precise static analysis to determine the write set accurately, as pointed out by prior work [17]. This is common for hybrid approaches. For instance, LJGS [5] uses a sophisticated points-to analysis. Moore and Chong have identified sufficient conditions for safely incorporating memory abstractions and static analyses into a hybrid information-flow monitor [17]. An interesting future direction is to investigate such conditions and abstractions for a higher-order language. Another possible direction for handling languages with first-class functions and references can be using the ideas proposed by Nielson et al. [32] and Foster et al. [33], who use regions and side-effect analysis to determine aliases.

IX. CONCLUSION

We presented a gradual information flow type system for a simple imperative language that enforces termination-insensitive noninterference and ensures the gradual guarantee at the same time. We demonstrated that our hybrid monitor can stop implicit flows by refining the labels for references in the write-sets of both branches, regardless of which branch is taken. The non-conventional proof technique of noninterference that we used helps us identify the conditions for ensuring the gradual guarantees.

ACKNOWLEDGMENT

We would like to thank the anonymous reviewers for their insightful comments and feedback. This work was supported in part by the National Science Foundation via grant CNS1704542 and the CyLab Presidential Fellowship at Carnegie Mellon University.

REFERENCES

- [1] D. Volpano, C. Irvine, and G. Smith, “A sound type system for secure flow analysis,” *Journal of Computer Security*, vol. 4, no. 2-3, pp. 167–187, Jan. 1996.
- [2] J. G. Siek and W. Taha, “Gradual typing for functional languages,” in *In Scheme And Functional Programming Workshop*, 2006, pp. 81–92.
- [3] T. Disney and C. Flanagan, “Gradual information flow typing,” in *Proceedings of the 2nd International Workshop on Scripts to Programs Evolution*, 2011.
- [4] L. Fennell and P. Thiemann, “Gradual security typing with references,” in *Proceedings of the 2013 IEEE 26th Computer Security Foundations Symposium*, 2013, pp. 224–239.

- [5] —, “LJGS: gradual security types for object-oriented languages,” in *30th European Conference on Object-Oriented Programming*, 2016, pp. 9:1–9:26.
- [6] M. Toro, R. Garcia, and E. Tanter, “Type-driven gradual security with references,” *ACM Trans. Program. Lang. Syst.*, vol. 40, no. 4, pp. 16:1–16:55, Dec. 2018.
- [7] A. A. de Amorim, M. Fredrikson, and L. Jia, “Reconciling noninterference and gradual typing,” in *Proceedings of the 35th ACM/IEEE Symposium on Logic in Computer Science*, 2020, pp. 116–129.
- [8] J. G. Siek, M. M. Vitousek, M. Cimini, and J. T. Boyland, “Refined Criteria for Gradual Typing,” in *1st Summit on Advances in Programming Languages*, ser. Leibniz International Proceedings in Informatics, vol. 32, 2015, pp. 274–293.
- [9] R. Garcia, A. M. Clark, and E. Tanter, “Abstracting gradual typing,” in *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2016, pp. 429–442.
- [10] J. A. Goguen and J. Meseguer, “Security policies and security models,” in *Proc. IEEE Symposium on Security and Privacy*, 1982, pp. 11–20.
- [11] T. H. Austin and C. Flanagan, “Efficient purely-dynamic information flow analysis,” in *Proc. ACM SIGPLAN Fourth Workshop on Programming Languages and Analysis for Security*, 2009, pp. 113–124.
- [12] —, “Permissive dynamic information flow analysis,” in *Proc. 5th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security*, 2010, pp. 3:1–3:12.
- [13] A. Russo and A. Sabelfeld, “Dynamic vs. static flow-sensitive security analysis,” in *Proceedings of the 2010 23rd IEEE Computer Security Foundations Symposium*, 2010, pp. 186–199.
- [14] S. Hunt and D. Sands, “On flow-sensitive security types,” in *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2006, pp. 79–90.
- [15] D. Volpano and G. Smith, “Eliminating covert flows with minimum typings,” in *Proceedings 10th Computer Security Foundations Workshop*, June 1997, pp. 156–168.
- [16] T. Terauchi and A. Aiken, “Secure information flow as a safety problem,” in *Proceedings of the 12th International Conference on Static Analysis*, 2005, pp. 352–367.
- [17] S. Moore and S. Chong, “Static analysis for efficient hybrid information-flow control,” in *Proceedings of the 2011 IEEE 24th Computer Security Foundations Symposium*, 2011, pp. 146–160.
- [18] F. Pottier and V. Simonet, “Information flow inference for ML,” in *Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2002, pp. 319–330.
- [19] A. Bichhawat, M. McCall, and L. Jia, “First-order gradual information flow types and gradual guarantees,” <https://arxiv.org/abs/2003.12819>, 2021.
- [20] D. E. Denning and P. J. Denning, “Certification of programs for secure information flow,” *Commun. ACM*, vol. 20, no. 7, pp. 504–513, Jul. 1977.
- [21] D. Hedin, L. Bello, and A. Sabelfeld, “Value-sensitive hybrid information flow control for a JavaScript-like language,” in *Proceedings of the 2015 IEEE 28th Computer Security Foundations Symposium*, 2015, pp. 351–365.
- [22] A. Bedford, S. Chong, J. Desharnais, E. Kozyri, and N. Tawbi, “A progress-sensitive flow-sensitive inlined information-flow control monitor,” *Computers & Security*, vol. 71, pp. 114 – 131, 2017.
- [23] P. Buiras, D. Vytiniotis, and A. Russo, “HLIO: Mixing static and dynamic typing for information-flow control in Haskell,” in *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming*, 2015, pp. 289–301.
- [24] A. C. Myers, “JFlow: Practical mostly-static information flow control,” in *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 1999, pp. 228–241.
- [25] D. Hedin and A. Sabelfeld, “Information-flow security for a core of JavaScript,” in *Proceedings of the 2012 IEEE 25th Computer Security Foundations Symposium*, 2012, pp. 3–18.
- [26] A. Bichhawat, V. Rajani, D. Garg, and C. Hammer, “Generalizing permissive-upgrade in dynamic information flow analysis,” in *Proceedings of the Ninth Workshop on Programming Languages and Analysis for Security*, 2014, pp. 15:15–15:24.
- [27] J. F. Santos and T. Rezk, “An information flow monitor-inlining compiler for securing a core of JavaScript,” in *ICT Systems Security and Privacy Protection*, 2014, pp. 278–292.
- [28] M. Vassena, A. Russo, D. Garg, V. Rajani, and D. Stefan, “From fine- to coarse-grained dynamic information flow control and back,” *Proc. ACM Program. Lang.*, vol. 3, no. POPL, Jan. 2019.
- [29] D. Chandra and M. Franz, “Fine-grained information flow analysis and enforcement in a Java virtual machine,” in *23rd Annual Computer Security Applications Conference*, 2007, pp. 463–475.
- [30] S. Just, A. Cleary, B. Shirley, and C. Hammer, “Information flow analysis for JavaScript,” in *Proceedings of the 1st ACM SIGPLAN International Workshop on Programming Language and Systems Technologies for Internet Clients*, 2011, pp. 9–18.
- [31] A. Askarov, S. Chong, and H. Mantel, “Hybrid monitors for concurrent noninterference,” in *2015 IEEE 28th Computer Security Foundations Symposium*, 2015, pp. 137–151.
- [32] F. Nielson, H. R. Nielson, and C. Hankin, *Principles of Program Analysis*. Berlin, Heidelberg: Springer-Verlag, 1999.
- [33] J. S. Foster, T. Terauchi, and A. Aiken, “Flow-sensitive type qualifiers,” in *Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation*, 2002, pp. 1–12.

APPENDIX

A. Well-formedness

- 1) $\vdash v$ wf, if
 - a) $v = \langle v_1 \mid v_2 \rangle$, then $\forall i \in \{1, 2\}. v_i = (\iota_i u_i)$
 - b) $v = (\iota u)^g$
- 2) $\vdash \delta$ wf, if $\forall x \in \delta, \vdash \delta(x)$ wf
- 3) $\vdash \delta /_i e$ wf for $i \in \{1, 2\}$ if $\vdash \delta$ wf
- 4) $\vdash c$ wf, when the following hold:
 - a) if $c = \langle \kappa_1, \iota_1, c_1 \mid \kappa_2, \iota_2, c_2 \rangle_g$, then $\vdash c_1$ wf, $\vdash c_2$ wf, and c_1 and c_2 do not contain pairs
 - b) if $c = \text{if } e \text{ then } c_1 \text{ else } c_2$, then $\vdash c_1$ wf, $\vdash c_2$ wf, and c_1 and c_2 do not contain pairs or braces

$$\begin{array}{c}
 \boxed{\kappa, \delta / c \xrightarrow{\alpha} \kappa', \delta' / c'} \\
 \\
 \frac{\kappa, \delta / c_1 \longrightarrow \text{abort}}{\kappa, \delta / c_1; c_2 \longrightarrow \text{abort}} \text{ M-SEQ-ERR} \\
 \\
 \frac{\delta / e \Downarrow v \quad \text{rfL}(\delta, X, \iota_{pc} \Upsilon \text{intvl}(v)) = \text{undef}}{\iota_{pc} g_{pc}, \delta / \text{if}^X e \text{ then } c_1 \text{ else } c_2 \longrightarrow \text{abort}} \text{ M-IF-REFINE-ERR} \\
 \\
 \frac{\delta / e \Downarrow v \quad \text{refineLB}(\iota_{pc}, v) = \text{undef}}{\iota_{pc} g_{pc}, \delta / x := e \longrightarrow \text{abort}} \text{ M-ASSIGN-ERR} \\
 \\
 \frac{\delta / e \Downarrow v \quad v' = \text{refineLB}(\iota_{pc}, v) \quad \text{updL}(\text{intvl}(\delta(x)), v') = \text{undef}}{\iota_{pc} g_{pc}, \delta / x := e \longrightarrow \text{abort}} \text{ M-ASSIGN-ERR2} \\
 \\
 \frac{\delta / e \Downarrow v \quad \text{refineLB}(\iota_{pc}, v) = \text{undef}}{\iota_{pc} g_{pc}, \delta / \text{output}(\ell, e) \longrightarrow \text{abort}} \text{ M-OUT-ERR} \\
 \\
 \frac{\delta / e \Downarrow v \quad \text{refineLB}(\iota_{pc}, v) = v' \quad \text{updL}([\ell, \ell], v') = \text{undef}}{\iota_{pc} g_{pc}, \delta / \text{output}(\ell, e) \longrightarrow \text{abort}} \text{ M-OUT-ERR2}
 \end{array}$$

Fig. 16. Monitor abort cases for commands

Read operations:

$$\text{rd } v = v \quad \text{rd}_1 v = [v]_1 \quad \text{rd}_2 v = [v]_2$$

Write operations:

$$\frac{v_n = \langle \iota_1 u_1 \mid \iota_2 u_2 \rangle^g \quad \forall i \in 1, 2, \text{restrictLB}(\iota_o, \iota_i) = \iota'_i}{\text{updL } \iota_o v_n = \langle \iota'_1 u_1 \mid \iota'_2 u_2 \rangle^g}$$

$$\frac{\text{restrictLB}(\iota_o, \iota_n) = \iota}{\text{updL } \iota_o (\iota_n u_n)^g = (\iota u_n)^g}$$

$$\frac{\exists j \in [1, 2], \text{restrictLB}(\iota_o, \iota_j) = \text{undef}}{\text{updL } \iota_o v_n = \text{undef}}$$

$$\frac{\text{restrictLB}(\iota_o, \iota_n) = \text{undef}}{\text{updL } \iota_o (\iota_n u_n)^g = \text{undef}}$$

- c) if $c = \text{while}^X e \text{ do } c$, then $\vdash c$ wf and c does not contain pairs or braces
- d) if $c = c_1; c_2$ then $\vdash c_1$ wf, $\vdash c_2$ wf and c_2 does not contain pairs or braces
- e) if $c = \{c_1\}$, then $\vdash c_1$ wf
- 5) $\vdash \kappa, \delta /_i c$ wf for $i \in \{1, 2\}$ if all of the following hold
- $\vdash c$ wf and $\vdash \delta$ wf
 - if $i \in \{1, 2\}$, then c does not contain pairs

B. Additional definitions

We define the following constraints on configurations to facilitate proofs related to paired values and commands. We start by defining $\iota \in H(\ell_A)$, $\Pi \in H(\ell_A)$ and $\kappa \in H(\ell_A)$ for any observer at level ℓ_A .

$$\frac{\iota = [\ell_l, \ell_r] \quad \ell_l \not\leq \ell_A}{\iota \in H(\ell_A)} \iota\text{-H} \quad \frac{\Pi = \langle \iota_1 \mid \iota_2 \rangle \quad \iota_i \in H(\ell_A), i \in \{1, 2\}}{\Pi \in H(\ell_A)} \Pi\text{-H}$$

$$\frac{\kappa = \iota g \triangleright \kappa' \quad \iota \in H(\ell_A) \quad (\kappa' \in H(\ell_A) \vee \kappa' = \emptyset)}{\kappa \in H(\ell_A)} \kappa\text{-H}$$

We say a configuration is safe (written $\vdash \kappa, \delta /_i c$ sf for $i \in \{1, 2\}$) if all of the following hold

- if $i \in \{1, 2\}$, then $\kappa \in H(\ell_A)$, $\forall x \in \text{WtSet}(c)$, $\text{intvl}(\delta(x)) \in H(\ell_A)$
- if $c = \text{if } \langle _ \mid _ \rangle$ then c_1 else c_2 , then $\forall x \in \text{WtSet}(c)$, $\text{intvl}(\delta(x)) \in H(\ell_A)$
- if $c = \langle \kappa_1, \iota_1, c_1 \mid \kappa_2, \iota_2, c_2 \rangle_g$, then $\forall i \in \{1, 2\}$, $\iota_i \vdash g \in H(\ell_A)$, and $\forall x \in \text{WtSet}(c)$, $\text{intvl}(\delta(x)) \in H(\ell_A)$

We first define when a gradual label of an initial store location is not observable by the attacker. Formally: $g \in H(\ell_A)$ iff $g = \ell$ and $\ell \not\leq \ell_A$.

Simple write/output operations:

$$\frac{\text{intvl}(v_o) = \langle \iota_1 \mid \iota_2 \rangle \text{ or } v_n = \langle \iota'_1 u_1 \mid \iota'_2 u_2 \rangle^g \quad \forall i \in 1, 2, \text{restrictLB}(\lfloor \text{intvl}(v_o) \rfloor_i, \lfloor \text{intvl}(v_n) \rfloor_i) = \iota''_i}{\text{upd } v_o v_n = \langle \iota''_1 u_1 \mid \iota''_2 u_2 \rangle^g}$$

$$\frac{\text{intvl}(v_o) = \langle \iota_1 \mid \iota_2 \rangle \text{ or } v_n = \langle \iota'_1 u_1 \mid \iota'_2 u_2 \rangle^g \quad \exists i \in 1, 2, \text{restrictLB}(\lfloor \text{intvl}(v_o) \rfloor_i, \lfloor \text{intvl}(v_n) \rfloor_i) = \text{undef}}{\text{upd } v_o v_n = \text{undef}}$$

$$\frac{\text{restrictLB}(\iota_o, \iota_n) = \iota}{\text{upd } (\iota_o u_o)^g (\iota_n u_n)^g = (\iota u_n)^g}$$

$$\frac{\text{restrictLB}(\iota_o, \iota_n) = \text{undef}}{\text{upd } (\iota_o u_o)^g (\iota_n u_n)^g = \text{undef}}$$

$$\frac{[v_n]_1 = (\iota_1 u_n)^{g'} \quad \text{restrictLB}(\lfloor \text{intvl}(v_o) \rfloor_1, \iota_1) = \iota'_1 \quad [v_o]_2 = (\iota_2 u_2)^g}{\text{upd}_1 v_o (\iota_n u_n)^g = \langle \iota'_1 u_n \mid \iota_2 u_2 \rangle^g}$$

$$\frac{[v_n]_2 = (\iota_2 u_n)^{g'} \quad \text{restrictLB}(\lfloor \text{intvl}(v_o) \rfloor_2, \iota_2) = \iota'_2 \quad [v_o]_1 = (\iota_1 u_1)^g}{\text{upd}_2 v_o (\iota_n u_n)^g = \langle \iota_1 u_1 \mid \iota'_2 u_n \rangle^g}$$

$$\frac{[v_n]_i = (\iota_i u_n)^{g'} \quad \text{restrictLB}(\lfloor \text{intvl}(v_o) \rfloor_i, \iota_i) = \text{undef}}{\text{upd}_i v_o (\iota_n u_n)^g = \text{undef}}$$

Fig. 17. Operations with pairs

$$\boxed{\kappa, \delta /_i c \xrightarrow{\alpha} \kappa', \delta' /_i c'}$$

$$\frac{\kappa_i \triangleright \iota_{pc} \vee \iota_i g_{pc} \vee_c g, \delta /_i c_i \longrightarrow \text{abort} \quad \{i, j\} = \{1, 2\}}{\iota_{pc} g_{pc}, \delta / \langle \kappa_1, \iota_1, c_1 \mid \kappa_2, \iota_2, c_2 \rangle_g \longrightarrow \text{abort}} \text{P-C-PAIR-ERR}$$

Fig. 18. Paired executions abort

We define merging of two stores $(\delta_1 \bowtie \delta_2)$ as below:

$$\frac{}{\Gamma \vdash \cdot \bowtie \cdot = \cdot} \text{MGS-EMP}$$

$$\frac{\Gamma \vdash \delta_1 \bowtie \delta_2 = \delta \quad \text{lab}(\Gamma(x)) \in H(\ell_A) \quad v_i = (\iota_i u_i)^g (i \in \{1, 2\})}{\Gamma \vdash \delta_1, x \mapsto v_1 \bowtie \delta_2, x \mapsto v_2 = \delta, x \mapsto \langle \iota_1 u_1 \mid \iota_2 u_2 \rangle^g} \text{MGS-H}$$

$$\frac{\Gamma \vdash \delta_1 \bowtie \delta_2 = \delta \quad \text{lab}(\Gamma(x)) \notin H(\ell_A) \quad v_1 = v_2 = v}{\Gamma \vdash \delta_1, x \mapsto v_1 \bowtie \delta_2, x \mapsto v_2 = \delta, x \mapsto v} \text{MGS-L}$$