

Static enforcement of security in runtime systems

Mathias V. Pedersen
Department of Computer Science
Aarhus University
mvp@cs.au.dk

Aslan Askarov
Department of Computer Science
Aarhus University
aslan@cs.au.dk

Abstract—Underneath every modern programming language is a runtime environment (RTE) that handles features such as automatic memory management and thread scheduling. In the information-flow control (IFC) literature, the RTE is often part of the trusted computing base (TCB), and there has been little focus on applying IFC to the implementation of the RTE itself. In this paper we address this problem by designing an IFC language, Zee, for implementing secure RTEs, thereby removing the RTE from the TCB. We implement Zee and design and implement secure versions of garbage collectors and thread schedulers using Zee. We also prove that a faithful calculus of Zee satisfies a strong variant of timing-sensitive noninterference.

I. INTRODUCTION

Modern programming languages offer many abstraction mechanisms to simplify the development of programs, increase their throughput, or reduce their resource consumption at runtime. Such abstractions are often implemented by compiler writers in a specialized program called the runtime environment (RTE). A runtime environment is a program (possibly written in a different language) that is running alongside programs written by the user, and may use knowledge about the implementation details of the language to perform its tasks. Examples of features commonly associated with the RTE include thread scheduling and automatic memory management. For a user of the language, implementing such functionality is difficult as it often requires breaking abstractions enforced by the language designer, and careful reasoning must be done by the developers of the runtime environment to ensure that the guarantees offered by the language are not violated by the runtime environment. For instance, a garbage collector must determine which heap allocations are reachable by following pointers through the heap, starting from a set of *root* pointers. Root pointers typically include the local variables stored on the call stack, and thus the RTE must traverse every stack frame currently stored on the call stack, potentially breaking local state encapsulation [30]. Even worse is the situation from a security perspective, where operations carried out by an RTE might reveal confidential information about the data handled by the user-written programs through storage- or timing channels [24], [29], [32], [37].

Language-based information-flow control (IFC) is a popular approach to solve the problem of ensuring integrity and confidentiality of data [27]. This approach uses programming language techniques to analyze information flows in potentially untrusted programs before and/or during their execution, and verifies that the execution does not leak sensitive information.

In this paper we design and implement Zee: an IFC programming language for implementing secure runtime environments. Zee supports secure and type-safe programming on heterogeneous data (e.g., data at multiple security levels). We also define a faithful calculus of Zee and prove that well-typed programs satisfy timing-sensitive noninterference.

The calculus and its implementation are defined over an abstract *instantiation language*, which allows for Zee to be extended with additional features without redoing many of the formal proofs. Instead, each instantiation must only prove that a certain semantic interface is satisfied.

In summary, this paper makes the following contributions:

- It identifies access to the call stack and runtime type analysis as core features necessary for a practical language for programming runtime environments.
- It presents the design and implementation of Zee: an extensible language for implementing RTEs, and proves that the combination of Zee’s type system and runtime semantics enforces timing-sensitive noninterference.
- It uses Zee’s extensible semantics to design and implement a secure garbage collector, and a secure thread scheduler.

The rest of the paper is structured as follows: Section II gives an introduction to Zee through examples, and Section III formally defines the syntax and semantics of Zee. Section IV formalizes the attacker model and the security guarantees enforced by Zee’s type system. Section V presents two case studies on how Zee can be used to implement well-typed (and hence secure) garbage collectors and thread schedulers. Section VI describes the implementation of Zee, and Section VII discusses related work. The Appendix contains full definitions and the accompanying technical report contains full proofs.

II. PROGRAMMING IN ZEE

This section introduces Zee using three example programs. The first example demonstrates how Zee uses existentially quantified type variables to compute securely on values with different security levels located in the same data structure.¹ The second example demonstrates how a similar technique can be used to securely inspect the call-stack at runtime. Finally, the third example demonstrates how Zee uses implicit revocation of expired pointers to prevent reuse of invalidated memory.

¹This example uses a specific instantiation language defined in Section V-A.

A. Computing on heterogeneous values

The first example uses a simple two-point lattice consisting of two elements \mathbf{L} and \mathbf{H} with the partial order $\mathbf{L} \sqsubseteq \mathbf{H}$, and $\mathbf{L} \sqcup \mathbf{L} = \mathbf{L}$ and $\ell_1 \sqcup \ell_2 = \mathbf{H}$ otherwise. We call \mathbf{L} the “public” level and \mathbf{H} the “secret” level. We use a level-partitioned heap [24] and so, in the two-point lattice, the heap consists of two partitions that we denote as the \mathbf{L} partition and the \mathbf{H} partition. As is standard for type systems for security, types are augmented with security levels, i.e., the type $\text{int}_{\mathbf{H}}$ is the type of integers at security level \mathbf{H} .

The type system also tracks which partitions pointers point into. The type $(\ell_1 \mapsto s)_{\ell_2}$ describes pointers that point into the ℓ_1 section, containing data of type s , and where the size of the allocation depends on information up to level ℓ_2 .

Zee uses existential types [36] and runtime type analysis [12] for secure handling of heterogeneous data. Traditional existential types are written as $\exists : \text{type}. s$, and a value of type $\exists : \text{type}. s$ is a pair (τ, v) where τ is a runtime representation of a type,² and v is a value of type $s[\tau/ \]$. That is, v is of type s , but where the free type variable has been replaced by τ . A value of an existential type can be introduced using the expression $\text{pack}(s, e)$ as $\exists : \text{type}. s$, which evaluates s to τ and e to v before returning a pair (τ, v) . Dually, given an expression e of type $\exists : \text{type}. s$, it can be eliminated using the command $\text{let}(\ : \text{type}, x : s) := e$ in c , that brings the type variable $\$, and the variable x into scope for the evaluation of command c . For information-flow control, the existential type is augmented with security levels, and a value (τ, v) is of type $(\exists : \text{type}_{\ell_1}. s)_{\ell_2}$ if τ depends only on information up to level ℓ_1 . Introduction and elimination rules are augmented accordingly as $\text{pack}(s, e)$ as $\exists : \text{type}_{\ell_1}. s$, and $\text{let}(\ : \text{type}_{\ell_1}, x : s) := e$ in c respectively. We call the latter command an *unpacking* command.

Figure 1 shows a Zee function `compute` that, given an array `xs` with elements of different security types, computes the sum of all the public integers in the array, revealing no information about the secret values in `xs`. The array `xs` is annotated with the type $(\mathbf{L} \mapsto (\exists : \text{type}_{\mathbf{L}}.)_{\mathbf{L}})_{\mathbf{L}}$, representing an array of heterogeneous data in the \mathbf{L} partition of public length.

Function `compute` also specifies two additional labels (both of which are \mathbf{L} in this example): the bottom label is a lower bound on the program counter label, which is classic in IFC literature [21]. The top label is novel: it represents an upper bound on the sensitivity of the information that can be learned by knowing the type of a local variable in the current activation frame. We defer the discussion of this label until Section III-C.

On lines 5 to 9, function `compute` loops over the elements of `xs`, and on lines 6 and 7 it extracts the witness $\$ and the value of type $\$. On line 8 the code performs runtime type analysis on the value of $\$, and it matches the pattern $\text{int}_{\mathbf{L}}$ (i.e., the type of public integers). In this branch on line 8 the value y is known to be of type $\text{int}_{\mathbf{L}}$, and can be added to the public variable `sum`, without leaking sensitive information. Finally,

```

1 compute(xs : (L ↦ (∃ : typeL. )L)L) =L
2   let n : intL := length xs in
3   let i : intL := 0 in
4   let sum : intL := 0 in
5   while i < n do
6     let x : (∃ : typeL. )L := *(xs + i) in
7     let ( : typeL, y : ) := unpack x in
8     match with intL ) sum := sum + y
9     | - ) skip;
10    i := i + 1

```

Fig. 1: Zee program demonstrating computations on heterogeneous values.

if $\$ is of any other type, this value is omitted from the final sum.

This example demonstrates how Zee securely computes on heterogeneous values using existential types and runtime type analysis. In the next example, we extend such use of existential types to access the call stack while guaranteeing type-safety and security. We do this by treating the frame pointer as a pointer to an array of existentially quantified types, similar to the type of `xs` in Figure 1. For the remaining examples in the paper we elide some type annotations, but all of the missing type annotations are inferred by our prototype implementation of Zee.

B. Computing on the call-stack

Zee allows fine-grained reasoning about the call stack, which is important for operations such as stack traversal for many garbage collection algorithms, or unwinding the stack to handle exceptions. Figure 2 shows the structure of the call stack during an execution with two activation frames belonging to functions g and f respectively. Function g pushes two arguments on the stack: a value of type $(\mathbf{H} \mapsto \text{int}_{\mathbf{H}})_{\mathbf{H}}$, and a value of type $\text{int}_{\mathbf{H}}$, and then invokes f . Function f then establishes a new frame by pushing the value of the old frame pointer onto the stack, so that the previous frame can be restored upon returning from f . Furthermore, f modifies the frame pointer to point to its local variables, and modifies the stack pointer to point to the next free address on the stack. Finally, f allocates its local variables and proceeds with computation in the newly established activation frame.

As Zee features runtime type analysis, which is crucial for the implementation of many useful programming language features, the types themselves must be protected using security labels. To accomplish this, we introduce the notion of a frame label fr , which represents an upper bound on the sensitivity of the information that can be learned by knowing the type of a local variable in the current frame.

In Zee the expression `FP` returns a pointer to the beginning of the current activation frame. The type assigned to `FP` is a recursive type,³ that reflects the layout of the stack (cf. Figure 2) and contains the types of the function parameters and local variables.

³We will explain the precise typing of the frame pointer in detail in Section III-C.

²The value τ is often called the *witness* of the type s .

```

1  f(p : (H ↗ intH)H, h : intH) =H
2  let a : intL := 0 in
3  let b : (L ↗ intH)L := null in
4  let c : intH := h in skip
5  g() =H ... ; f(null, 42); ...
6

```




Fig. 2: Bottom: the structure of the call stack just before executing skip in f . Top: two functions g and f .


To understand the need for the frame label, consider the program in Figure 3, showing a function inspect that inspects its own frame. Unlike the compute function, the inspect function is parameterized by a secret label ℓ , which the caller provides when invoking inspect. In this example, the first element of the frame layout consists of an integer x with the security level ℓ , as declared on line 2. Using the FP construct on line 3, the type variable args is assigned a tuple type representing the types of the arguments passed to inspect,⁴ and on line 4, locals is assigned the types of the local variables of inspect.

As no arguments have been passed to inspect, the value of args during execution is the unit type ε . The value of locals is $\text{int}_{\ell'} \cdot \tau_{\text{tail}}$, where ℓ' is the value of ℓ that has been provided when inspect was called, and τ_{tail} is the rest of the frame layout. After obtaining the type of the local variables, line 5 performs a pattern match on locals , and if the pattern $\text{int}_{\ell'} \cdot _$ matches the value of locals , it follows that ℓ' is equal to ℓ , which is information classified at H . So to track the indirect flow from the label ℓ to the match command, we introduce the frame label $f\ell$, which assigns locals the security level $f\ell$ (which is equal to H as declared by the inspect function), properly tracking the dependency of ℓ in locals .

C. Fail-stop revocation of expired pointers

This section demonstrates how Zee uses a dynamic enforcement technique to prevent the reuse of pointers that point to “expired” data. Expired data include local variables of functions that have returned control to its caller, or heap data that have been reclaimed by a garbage collector. To do this we introduce a technique similar to identifier-based temporal checking [22] which we simply refer to as *versioning*. At runtime, every pointer is assigned a natural number ν called the version number. Similarly, every stack frame is assigned a version number. When a pointer to a local variable on a stack frame is created, the pointer is assigned the version number of that stack frame. When reading data pointed to by a pointer with

⁴The unroll expression is needed as Zee uses isorecursive types to represent the type isomorphism between $\mu \text{ : type. } s$ and $s[\mu \text{ : type. } s/_]$ (i.e., between a recursive type and its unrolling). We will omit unroll expressions throughout the paper.



```

1 inspect h : levelH i() : =H
2 let x : int := 42 in
3 let ( args, e ) := unpack (unroll FP) in
4 let ( locals, _ ) := unpack e in
5 match locals with intℓ * _ ) ... // = ℓ
6 | _ ) skip

```

Fig. 3: Bottom: Indirect information flow from ℓ to locals when inspect inspects its own stack frame. Top: The activation frame for function inspect.

version ν , the system checks that the stack frame, which is read from, has a version number γ such that $\gamma \leq \nu$, ensuring that this stack frame was “live”⁵ when the pointer was created, and a similar check is done for writes through pointers.

Figure 4 demonstrates how an illegal flow could be constructed without versioning: on line 6, x is allocated on the stack and initialized to null.⁶ Then, function f is called and writes the address of its local variable y into x through the passed pointer ρx on line 2. When f returns, the value of x is a pointer, pointing to the local variable y from the “dead” stack frame of f . Then, when g is called, the value of secret is stored in h , which might be located in the same offset from the base pointer as y was in function f (cf. lower part of Figure 4). So when g reads x , the value of secret is stored in low , violating the type ascribed to the variable.

Versioning prevents this leak by assigning a version number $\nu \in \mathbb{N}$ to x and the stack frame allocated on line 6. The stack frames of f and g are assigned versions $\nu + 1$ and $\nu + 2$ respectively. When g reads from x , this variable has version $\nu + 1$, while the data it points to is located on g ’s stack frame, which has version $\nu + 2$. This violates the version check and execution stops, successfully preventing the leak.

III. LANGUAGE

This section presents a formalization of Zee. We assume a lattice \mathcal{L} of security levels with a least element \perp and let ℓ range over the elements of \mathcal{L} .

A. Syntax

Figure 5 defines the grammar for Zee. We give an informal description of each syntactic category, and delay the formal semantics until Section III-B.

1) *Commands*: Commands c include standard constructs for imperative languages. Non-standard commands include commands $*e := e$ and $x := *e$ that respectively write to, and read from, a location on the stack. Command at $k e c$ raises the program counter label for command c during type checking,

⁵A live stack frame refers to a frame of a computation that is still ongoing.

⁶For now the type $@ s$ can be read as “pointer to a value of type s ”, but we will define a more general version of this type in Section III.

```

1 f(px : (@(@ intL)L)L) =  $\dagger$ 
2   let y : intL := 0 in *px := &y
3 g(z : (@ intL)L) =  $\dagger$ 
4   let h : intH := secret in
5   let low : intL := *z in skip
6 let x : (@ intL)L := null in f(&x); g(x)

```




Fig. 4: Attempting to read x in function g will fail: The pointer has been implicitly revoked when f no longer was live.

and also provides lightweight⁷ predictive mitigation [4], [41], [42]. The match command allows for runtime type analysis [1], [12], which is crucial for implementing many useful tasks associated with the RTE. We call the \bar{p} being matched on the *scrutinee*, and the patterns are defined in the syntactic category p . These include integer type patterns (int), stack pointer type patterns ($(p @ p)$), heap pointer type patterns ($(\mapsto p)$), product type patterns (\bar{p}) and a “catch all” pattern that names the type value $(\)$. To facilitate traversing the stack at runtime, the language exposes the frame pointer, which can be obtained using the $x := \text{fp}$ command. A function can be invoked by supplying the function name with label⁸ parameters \bar{k} , type parameters \bar{s} , and expression parameters \bar{e} . Finally, both existentially quantified labels and types can be eliminated using the unpacking commands $\text{let } (\ : \mathcal{X}_k, x : s) := e \text{ in } c$ for $\mathcal{X} \in \{\text{type}, \text{level}\}$.

2) *Expressions*: The meta-variable e ranges over expressions which include numbers, variables, binary operations, a special pointer value null , and pack expressions for introducing existentially quantified labels and types. We include recursive types to give the language type-safe access to the call stack, and expressions $\text{unroll } e$ and $\text{roll } e$ allow for simple (i.e., isorecursive) typing rules for recursive types [26]. The size of a type can be calculated using the expression $\text{sizeof } s$, which is also used to facilitate stack traversal. Finally, given a variable x the address of x on the stack can be obtained as $\&x$.

3) *Security labels and types*: The meta-variable k ranges over security labels. As labels can be existentially and universally quantified, the category includes variables $_$. Finally, one can form the join \sqcup or meet \sqcap of two security labels, representing the least upper bound, or greatest lower bound of two labels respectively.

The meta-variable s ranges over security types and include base types t with a security label k , written t_k , type variables $_$, or product types \bar{s} . Base types are integers, heap pointers ($k \mapsto s$) representing a pointer into the heap partition associated

⁷In particular, the predicted time for mitigation commands is given by the programmer, nor is it not automatically updated by the semantics.

⁸We distinguish between levels (i.e., elements of \mathcal{L}) and labels (i.e., expressions that evaluate to elements of \mathcal{L})

```

c ::= skip | let x : s := e in c | if e c c
    | while e c | c; c | x := e | *e := e
    | x := *e | at k e c | if (k  $\sqsubseteq$  k) c c
    | match  $\bar{p} \Rightarrow \bar{c} \mid x := \text{fp} \mid f(\bar{k})(\bar{s})(\bar{e})$ 
    | let ( : typek, x : s) := e in c
    | let ( : levelk, x : s) := e in c
e ::= n | x | e  $\oplus$  e | null | unroll e | roll e
    | pack (s, e) as  $\exists \ : \text{type}_k. s \mid \text{sizeof } s$ 
    | pack (k, e) as  $\exists \ : \text{level}_k. s \mid \&x$ 
k ::=  $\ell \mid \_ \mid k \sqcup k \mid k \sqcap k$ 
s ::= tk |  $\_ \mid \bar{s}$ 
t ::= int | k  $\mapsto$  s | s @ s |  $\exists \ : \text{type}_k. s$ 
    |  $\exists \ : \text{level}_k. s \mid \mu \ : \text{type}_k. s \mid \text{size}[s]$ 
p ::= int | (p @ p) | (  $\mapsto$  p) |  $\bar{p} \mid$ 

```

Fig. 5: The syntax of Zee.

with label k [24], stack pointers ($s_1 @ s_2$) representing pointers to a value of type s_2 that, on the stack, are located “above” a value of type s_1 [2], [25]. Base types also include the type of existentially quantified security types and labels, recursive types and singleton types [34] $\text{size}[s]$ describing the size of the type s . A function definition is given as $f \langle _ : \text{level}_{k_1} \rangle \langle _ : \text{type}_{k_2} \rangle \langle \bar{x} : \bar{s} \rangle =_{\rho c}^{fr} c$ that defines a function f with label parameters $\bar{_}$, type parameters $\bar{_}$, and value parameters \bar{x} . Label parameter $_i$ can depend on the previous label parameters $_1, \dots, _i$, and type parameter $_i$ can depend on all label parameters and type parameters $_1, \dots, _i$. Finally, the types \bar{s} for the expression parameters can depend on all label and type parameters. Furthermore, ρc is a lower bound on the side effects produced by f , and fr is an upper bound on the type of any local variable declaration. We revisit this label in Section III-C.

B. Semantics

The semantics of Zee is given by a small-step relation \rightarrow on configurations C of the form $\langle c, M, P, q \rangle_\nu$. We first define each component of the configuration before describing the small-step relation.

1) *Values*: Figure 7 describes the syntax of values. A value v is either a number n , an address a with a *version number* $\nu \in \mathbb{N}$, a pair consisting of either a level and a value (ℓ, v) , or a security type value and a value (τ, v) . The meta-variable τ ranges over security type values and is either a base type value π with a security level ℓ , or a product of security type values.

Security type values also include a *nonsense* [19] type value $_$. To motivate the need for nonsense type values, consider the program in Figure 6. When evaluating FP on line 3, the variable x with value 0 is in scope, but none of the variables e , p or y have entered the scope, and their value on the stack are “garbage” values. So locals is a product type $\text{int}_L \cdot \dots$ containing three nonsense type values $_$ as the variables e , p and y have yet to enter the scope and be assigned a value. The type variables args and locals live on a different stack (cf. top right of Figure 6) where extracting type-information is not possible, and therefore technically do not need a nonsense




Fig. 6: Left: the value of locals contains three nonsense type values, corresponding to the three variables e , p and y that has yet to enter the scope upon evaluating FP on line 3. Right: the stack layout when FP is evaluated.

$$\begin{aligned}
v &::= n \mid a_\nu \mid (\ell, v) \mid (\tau, v) \\
\tau &::= \pi_\ell \mid \bar{\tau} \\
\pi &::= \text{int} \mid \ell \mapsto \tau \mid \tau @ \tau \mid \exists : \text{type. } s \\
&\quad \mid \exists : \text{level. } s \mid \mu : \text{type. } s \mid \text{size}[\tau]
\end{aligned}$$

Fig. 7: Values in Zee.

type value, as we do not track “the type of type variables” at runtime.

The meta variable π ranges over base type values, and contain the same constructs as the base types, but where security labels not under binders \exists and μ are fully evaluated, and labels on types and labels are erased (as they are only needed for type checking).

2) *Exposed and private stack frames*: An exposed stack frame m is a triple $(l, |m|, \nu)$ where $|m| : l \rightarrow \nu$ is a partial function from a nonempty interval $l \subset \mathbb{N}$ to values, and a frame version $\nu \in \mathbb{N}$. We call $\min(l)$ the frame pointer, written $\text{fp}(m)$, and $\max(l) + 1$ the stack pointer, written $\text{sp}(m)$. Intuitively, $\text{fp}(m)$ is the minimum address in the stack frame, corresponding to the usual notion of a frame pointer, and similarly $\text{sp}(m)$ is the address of the next available stack location. Given an exposed stack frame $m = (l, |m|, \nu)$ we write $m[a \mapsto v]$ to mean $(l, |m|[a \mapsto v], \nu)$, and $m(a)$ to mean $|m|(a)$. We call a list of exposed stack frames an exposed stack M .

A private stack frame p is a triple of partial functions $(p_{\text{var}}, p_{\text{arg}}, p_{\text{local}})$ where $p_{\text{var}} : \text{Var} \rightarrow \mathcal{L} \cup \tau$ and $p_{\text{arg}}, p_{\text{local}} : \text{Var} \rightarrow \tau$. Function p_{var} maps label- and type variables to levels and type values, p_{arg} maps function arguments to their security type, and p_{local} maps local variable names to their security type. The name “private” refers to the fact that this stack cannot be traversed and inspected at runtime (unlike the exposed stack). Finally, a private stack P is a list of private stack frames. We call a pair of an exposed and a private stack (M, P) a stack.

We distinguish between exposed and private stacks because Zee allows type-safe traversal of the exposed stack, but does not directly expose the private stack to programs.

As exposed stack frames include mappings from (subsets of) natural numbers, a translation from local variable names to addresses is needed. This translation is usually performed by a compiler, and many techniques exist for such translations

```

1 fib(n : intL, r : (@ intL)L) =  $\perp$ 
2   if n <= 1 then *r := n
3   else let r1 : intL := 0 in
4         let r2 : intL := 0 in
5         fib(n-1, &r1); fib(n-2, &r2);
6         *r := r1 + r2

```




Fig. 8: Top: an implementation of a function `fib` for computing a public Fibonacci number given a public input. Bottom right: the stack frame layout for `fib`. Bottom left: the mapping δ between local variable names and stack frame offsets.

[3]. We abstract away the specifics of such translations by assuming a global mapping $\delta : \text{Var} \rightarrow \mathbb{N}$ from variables names to offsets. Figure 8 demonstrates how a compiler might generate a mapping δ that maps variable names to activation frame offsets. At the moment we assume that the compiler does not coalesce stack locations when the lifetimes of two variables do not overlap. So given two local variables $x \neq y$ it holds that $\delta(x) \neq \delta(y)$. That is, different variables are stored at different locations on the stack. The expression $\delta(x) + \text{fp}(m)$ computes the address of the local variable x in the stack frame $\text{fp}(m)$.

3) *Model of time*: We represent time as a number $q \in \mathbb{N}$ that counts the number of operational steps in the computation. This simple model is sufficient to demonstrate that runtime environment tasks can be computed in a timing-sensitive security setting. Naturally, a realistic implementation would need to soundly relate the operational steps with the wall-clock, but that is outside of the scope of the current work.

4) *Version counter*: Finally, the configuration contains a version counter that keeps track of the next free version number. This is needed when constructing new stack frames, as each new frame is given a fresh version number.

5) *Big-step evaluation for expressions*: The big-step evaluation for expressions is defined on configurations of the form $\langle e, m, p \rangle$, where m is an exposed stack frame and p is a private stack frame. The evaluation of expressions need both the exposed, and the private stack frame, as both expression variables and type- and level variables might appear in expressions. Figure 9 shows excerpts of the big-step evaluation semantics for expressions. Rule E-NUM evaluates a literal, and E-VAR evaluates a variable by looking up its value on the stack using the global mapping δ . Finally, rule E-PACK-TY evaluates a pack expression containing a security type s and an expression e to a pair (τ, v) . The remaining rules are found in the technical appendix [23].

6) *Small-step relation for commands*: Figure 10 shows an excerpt of the small-step relation, and the full semantics is

$$\begin{array}{c}
\text{E-NUM} \\
\langle n, m, p \rangle \Downarrow n \\
\\
\text{E-VAR} \\
\frac{m(\delta(x) + \text{fp}(m)) = v}{\langle x, m, p \rangle \Downarrow v} \\
\\
\text{E-PACK-TY} \\
\frac{\langle s, p \rangle \Downarrow_{\text{type}} \tau \quad \langle e, m, p \rangle \Downarrow v}{\langle \text{pack}(s, e) \text{ as } \exists : \text{type}_k. s', m, p \rangle \Downarrow (\tau, v)}
\end{array}$$

Fig. 9: Excerpts of the big-step evaluation of expressions.

in Figure 22 in the Appendix. Rule S-ASGN evaluates an expression e and inserts the resulting value v in the memory at offset $\delta(x)$ of the current frame pointer. Rule S-FP stores the frame pointer in the variable x . In addition to the value of the frame pointer $\text{fp}(m)$, the value v contains the the list of types $\text{cod}(p.\text{arg})$ of the arguments passed to the current function, and the list of types $\text{cod}(p.\text{local})$ of the local variables. Rule S-LET declares a new local variable x and (1) updates the stack location to contain the initial value of x , and (2) updates the private stack frame to contain type information about the type of x . This causes the type value of the variable x to be updated from a nonsense type value to a meaningful type value τ , which is the result of evaluating the type s . After executing c , a command $\text{unscope}(x)$ removes the type information of x from the private stack p .⁹ Rule S-UNPACK-TY unpacks an existential value containing a security type value and a regular value. Several maps are updated: the private stack frame is updated to contain the security type value and the type information about the newly allocated local variable. Finally, the exposed stack frame is updated to contain the regular value. Rule S-MATCH evaluates the scrutinee to a security type value and performs type analysis according to a list of patterns \bar{p} using the relation $\tau - p$ containing rules such as $\text{int}_\ell - \text{int}$ and $\tau - \text{int}$ ¹⁰ (i.e., integer patterns matches integer types, and name patterns matches any security type value). Upon finding the first (due to argmin) matching pattern the private stack frame is updated by binding relevant parts of the security type value to type variables, and execution proceeds with the command associated with the pattern.

Rule S-READ reads from a stack location a_γ . The version number γ is used to prevent attacks based on pointer reuse as was described in Section II-C. Dually, S-WRITE writes a value v to an address a , requiring the same relation between the versions of the target stack frame and the address being read from.

Rules S-AT and S-DELAY implement simple predictive mitigation of direct timing channels, i.e., channels represented directly in the control-flow of the program: S-AT reduces to the underlying command c , but ensures that the command terminates in exactly n steps, where n is the result of evaluating expression e , by delaying further commands until the command delay n has terminated.

⁹The semantics of unscope is defined in the Appendix.

¹⁰The complete definition of matching is found in the Appendix.

$$\begin{array}{c}
\text{S-ASGN} \\
\frac{\langle e, m, P \rangle \Downarrow v \quad m' = m[\delta(x) + \text{fp}(m) \mapsto v]}{\langle x := e, m \cdot M, P, q \rangle_\nu \rightarrow \langle \text{stop}, m' \cdot M, P, q + 1 \rangle_\nu} \\
\\
\text{S-FP} \\
\frac{v = (\text{cod}(p.\text{arg}), (\text{cod}(p.\text{local}), \text{fp}(m)_\nu)) \quad m = (l, |m|, \nu) \quad m' = m[\delta(x) + \text{fp}(m) \mapsto v]}{\langle x := \text{fp}, m \cdot M, p \cdot P, q \rangle_\nu \rightarrow \langle \text{stop}, m' \cdot M, p \cdot P, q + 1 \rangle_\nu} \\
\\
\text{S-LET} \\
\frac{\langle s, p \rangle \Downarrow_{\text{type}} \tau \quad \langle e, m, p \rangle \Downarrow v \quad c' = c; \text{unscope}(x) \quad m' = m[\delta(x) + \text{fp}(m) \mapsto v] \quad p' = p[p_{\text{local}} \mapsto p_{\text{local}}[x \mapsto \tau]]}{\langle \text{let } x : s := e \text{ in } c, m \cdot M, p \cdot P, q \rangle_\nu \rightarrow \langle c', m' \cdot M, p' \cdot P, q + 1 \rangle_\nu} \\
\\
\text{S-UNPACK-TY} \\
\frac{\langle e, m, p \rangle \Downarrow (\tau_1, v_2) \quad \langle s, p' \rangle \Downarrow_{\text{type}} \tau \quad p' = p[\text{var} \mapsto p.\text{var}[\mapsto \tau_1]] \quad p'' = p'[\text{local} \mapsto p'.\text{local}[x \mapsto \tau]] \quad m' = m[\delta(x) + \text{fp}(m) \mapsto v_2]}{\langle \text{let } (\exists : \text{type}_k, x : s) := e \text{ in } c, m \cdot M, p \cdot P, q \rangle_\nu \rightarrow \langle c; \text{unscope}(x), m' \cdot M, p' \cdot P, q + 1 \rangle_\nu} \\
\\
\text{S-READ} \\
\frac{m_i = (l, |m_i|, \nu_i) \in m \cdot M \quad a \in l \quad \nu_i \leq \gamma \quad \langle e, m, P \rangle \Downarrow a_\gamma \quad m' = m[\delta(x) + \text{fp}(m) \mapsto m_i(a)]}{\langle x := *e, m \cdot M, P, q \rangle_\nu \rightarrow \langle \text{stop}, m' \cdot M, P, q + 1 \rangle_\nu} \\
\\
\text{S-WRITE} \\
\frac{m_i = (l, |m_i|, \nu) \quad M = M_1 \cdot m_i \cdot M_2 \quad a \in l \quad \nu_i \leq \gamma \quad \langle e_1, M, P \rangle \Downarrow a_\gamma \quad \langle e_2, M, P \rangle \Downarrow v \quad m'_i = m_i[a \mapsto v]}{\langle *e_1 := e_2, M, P, q \rangle_\nu \rightarrow \langle \text{stop}, M_1 \cdot m'_i \cdot M_2, P, q + 1 \rangle_\nu} \\
\\
\text{S-AT} \quad \text{S-DELAY} \\
\frac{\langle e, M, P \rangle \Downarrow n}{\langle \text{at } k \ e \ c, M, P, q \rangle_\nu \rightarrow \langle c; \text{delay } n, M, P, q + 1 \rangle_\nu} \quad \frac{n \leq q}{\langle \text{delay } n, M, P, q \rangle_\nu \rightarrow \langle \text{delay } n, M, P, n + 1 \rangle_\nu} \\
\\
\text{S-MATCH} \\
\frac{\text{argmin}_{i=1, \dots, n}(\tau - p_i) = j \quad \langle s, p \rangle \Downarrow_{\text{type}} \tau \quad \text{Jp}_j \text{K}(p, \tau) = p'}{\langle \text{match } (p_i \Rightarrow c_i)_{i=1, \dots, n}, M, p \cdot P, q \rangle_\nu \rightarrow \langle c_j, M, p' \cdot P, q + 1 \rangle_\nu}
\end{array}$$

Fig. 10: Semantics of Zee: commands.

C. Type system

We now describe a type system for Zee, which we will show ensures secure information-flow in Section IV. The type system integrates previous work on type safety in stack-based languages [2], [25] with dynamic security labels [44] and existential types for information-flow control [36] into a single language that is able to express complex indirect data dependencies. The typing judgment for commands has the form $\Gamma, \Pi, \phi, \rho c, fr \vdash c$, and the typing judgment for expressions has the form $\Gamma, \Pi, \phi \vdash e : s$. We explain each component of the judgment before presenting the judgment rules.

Function $\Gamma : \text{Var} \rightarrow s$ maps regular variable names to security types, and similarly $\Pi : \text{Var} \rightarrow \{\text{type}, \text{level}\} \times k$ maps type variables to type_k and label variables to level_k . Formula ϕ is a finite conjunction of (possibly negated) flow relations such as $1 \sqcup 2 \sqsubseteq 3 \wedge 3 \not\sqsubseteq 4$. These formulas are gathered during type-checking in such a way that the constructed formulae always represents flows that will be true at runtime. Adding such formulae to the typing relation improves the expressiveness of static information-flow control in the presence of dynamic security labels [16], [43], [44]. Finally, the type system tracks two labels: the program counter label ρc and the frame label fr . We now explain the typing judgments involved in typing Zee programs.

1) *Typing judgment for expressions*: Figure 11 shows excerpts of the typing rules for expressions. Rules T-NUM and T-VAR are standard rules for literals and variables, and T-PACK-TY is the standard rule for introducing an existentially quantified type [36]. Finally, rule T-SIZEOF assigns a singleton-type $\text{size}[s]$ to an expression $\text{sizeof } s$, representing that the expression will evaluate to the size of the type s at runtime. Such expressions are crucial for secure and type-safe operations in a language with heterogeneous data like in Zee: they allow the type system to track how pointer arithmetic changes the type of the pointer. This becomes clear in rule T-BINOP, which assigns types to the result of binary expressions using the relation $s_1 \text{J}\oplus\text{K } s_2 \rightarrow s$. This states that applying operator \oplus to expressions of type s_1 and s_2 results in an expression of type s . The full judgment is shown in the Appendix, and excerpts of this relation include

$$\text{int}_{k_1} \text{J}\oplus\text{K } \text{int}_{k_2} \rightarrow \text{int}_{(k_1 \sqcup k_2)} \quad (1)$$

$$(s_1 @ s \cdot s_2)_{k_1} \text{J}+\text{K } \text{size}[s]_{k_2} \rightarrow (s_1 \cdot s @ s_2)_{(k_1 \sqcup k_2)} \quad (2)$$

$$(s_1 \cdot s @ s_2)_{k_1} \text{J}-\text{K } \text{size}[s]_{k_2} \rightarrow (s_1 @ s \cdot s_2)_{(k_1 \sqcup k_2)} \quad (3)$$

In words, performing a binary operation on two integers results in an integer labeled with the join of the two labels (1), adding the size of a type s to a pointer pointing to a value of type $s_1 @ s \cdot s_2$ results in a pointer to a value of type $s_1 \cdot s @ s_2$, and the labels are raised accordingly. Dually, one can subtract the size of a type s from a pointer of type $s_1 \cdot s @ s_2$, and obtain a value of type $s_1 @ s \cdot s_2$.

2) *Typing judgment for commands*: Figure 12 shows excerpts of the typing rules for commands. Rule T-LET states that a variable declaration $\text{let } x : s := e \text{ in } c$ is well-typed if the type of e is a subtype of s . The subtype relation is standard for imperative languages for information-flow [21]. To prevent implicit flows, the program counter label should also flow to s . Finally, x is added to the typing context Γ , and the frame label fr is raised to reflect the fact that the frame layout has been influenced by the variable declaration.

Rule T-IF states that a command $\text{if } e \text{ } c_1 \text{ } c_2$ is well-typed when e is an expression of type $\text{int}_{\rho c}$, and both branches can be shown to be well-typed. Readers familiar with IFC type systems may wonder why it is necessary to restrict the label on the type of e to be ρc . This is done to facilitate predictive mitigation of direct timing channels: the ρc must be explicitly

<p>T-NUM</p> $\frac{}{\Gamma, \Pi, \phi \vdash n : \text{int}_{\perp}}$	<p>T-VAR</p> $\frac{}{\Gamma, \Pi, \phi \vdash x : \Gamma(x)}$
<p>T-PACK-TY</p> $\frac{\Pi, \phi \vdash_{\text{type}} s_2 \quad \Gamma, \Pi, \phi \vdash e : s_2[s_1/x] \quad \Pi, \phi \vdash_{\text{type}} s_1 : k_1 \quad t = \exists x : \text{type}_{k_1} \cdot s_2}{\Gamma, \Pi, \phi \vdash \text{pack}(s_1, e) \text{ as } t : t_{\perp}}$	
<p>T-SIZEOF</p> $\frac{\Pi, \phi \vdash_{\text{type}} s : k}{\Gamma, \Pi, \phi \vdash \text{sizeof } s : \text{size}[s]_k}$	<p>T-BINOP</p> $\frac{\Gamma, \Pi, \phi \vdash e_i : s_i \quad s_1 \text{J}\oplus\text{K } s_2 \rightarrow s}{\Gamma, \Pi, \phi \vdash e_1 \oplus e_2 : s}$

Fig. 11: Excerpts of the typing rules for expressions.

raised using an `at` command. Rule T-AT states that a command `at k e c` is well-typed when the label k and the computation time e only depends on information up to ρc . Furthermore, the command must not lower the program counter label, and the command c must be well-typed under the raised program counter label k .

The command $x := \text{fp}$ is well-typed when x is a subtype of the type $\text{T}_{\text{st}}(\rho c, fr, k)$, which abbreviates the type

$$(\mu : \text{type}_k. (\exists : \text{type}_{fr}. (\exists \gamma : \text{type}_{fr}. (\cdot @ \gamma)_{\rho c})_{\perp})_{\perp})_{\perp}$$

This type reflects the layout of the stack at runtime (cf. Figure 2). Each frame consists of some type representing the type of the arguments given to the function, followed by a pointer to the previous stack frame (which is represented as the recursive type variable μ), and finally the type γ representing the types of the local variables. By assigning each existentially quantified type the label fr we ensure that no type leaks information, as fr represents the upper bound of the information that can be learned from knowing the value of the types.

Rule T-MATCH states when a command $\text{match } \overline{p} \Rightarrow \overline{c}$ is well-typed. First, \overline{p} must be a type variable with a label that flows to ρc , as the direct timing channels must be controlled using `at` commands. Judgment $\Pi \vdash p_i \quad k \quad \Pi_i : s_i$ generates a type variable environment Π_i for type-checking the command for the i 'th pattern, and the type s_i to assign \overline{c} in the command. Furthermore, any type and label variable bound by the new typing type variable environment Π_i is bound to the label k . Most of the rules of this judgment are of the form

$$\Pi \vdash \text{int}_{\cdot} \quad k \quad \Pi[\mapsto \text{level}_k] : \text{int}$$

which expresses that, in environment Π , when the pattern is int_{\cdot} and the scrutinee depends on information up to label k , the environment is updated to $\Pi[\mapsto \text{level}_k]$ and the type of the scrutinee can be assumed to have type int_{\cdot} in the command guarded by the pattern int_{\cdot} . The full judgment can be found in the technical report. Finally, each command c_i is type-checked in the generated type- and variable environment.

Rule T-UNPACK-TY states when an elimination of an

$$\begin{array}{c}
\text{T-LET} \\
\frac{\Gamma, \Pi, \phi \vdash e : r \quad \Pi, \phi \vdash_{\text{type}} s : k \quad \phi \vdash r^{\rho_C} <: s \quad fr' = fr \sqcup k \quad \Gamma[x \mapsto s], \Pi, \phi, \rho_C, fr' \vdash c}{\Gamma, \Pi, \phi, \rho_C, fr \vdash \text{let } x : s := e \text{ in } c} \\
\\
\text{T-AT} \\
\frac{\Gamma, \Pi, \phi \vdash e : \text{int}_{\rho_C} \quad \Pi; \phi \vdash_{\text{lab}} k : \rho_C \quad \phi \vdash \rho_C \sqsubseteq k \quad \Gamma, \Pi, \phi, k, fr \vdash c}{\Gamma, \Pi, \phi, \rho_C, fr \vdash \text{at } k \text{ e } c} \\
\\
\text{T-IF} \quad \text{T-FP} \\
\frac{\Gamma, \Pi, \phi \vdash e : \text{int}_{\rho_C} \quad \Gamma, \Pi, \phi, \rho_C, fr \vdash c_i \quad i = 1, 2}{\Gamma, \Pi, \phi, \rho_C, fr \vdash \text{if } e \text{ c}_1 \text{ c}_2} \quad \frac{\Pi; \phi \vdash_{\text{lab}} fr : k \quad \phi \vdash \text{T}_{\text{st}}(\rho_C, fr, k)^{\rho_C} <: \Gamma(x)}{\Gamma, \Pi, \phi, \rho_C, fr \vdash x := \text{fp}} \\
\\
\text{T-MATCH} \\
\frac{\Pi() = \text{type}_{k_i} \quad \phi \vdash k \sqsubseteq \rho_C \quad \Pi \vdash p_i \quad k \quad \Pi_i : s_i \quad \Gamma[s_i / _], \Pi_i[s_i / _], \phi, \rho_C, fr \vdash c_i[s_i / _]}{\Gamma, \Pi, \phi, \rho_C, fr \vdash \text{match } \overline{p} \Rightarrow \overline{c}} \\
\\
\text{T-UNPACK-TY} \\
\frac{\Gamma, \Pi, \phi \vdash e : (\exists _ : \text{type}_{k_1}. r)_{\rho_C} \quad \phi \vdash r^{\rho_C} <: s \quad \Gamma' = \Gamma[x \mapsto s] \quad \Pi' = \Pi[_ \mapsto \text{type}_{k_1}] \quad \Pi', \phi \vdash_{\text{type}} r : k_2 \quad fr' = fr \sqcup k_1 \sqcup k_2 \quad \Gamma', \Pi', \phi, \rho_C, fr' \vdash c}{\Gamma, \Pi, \phi, \rho_C, fr \vdash \text{let } (_ : \text{type}_{k_1}, x : s) := e \text{ in } c} \\
\\
\text{T-FLOWS TO} \\
\frac{\Pi; \phi \vdash_{\text{lab}} k_i : \rho_C \quad \Gamma, \Pi, \phi \wedge k_1 \sqsubseteq k_2, \rho_C, fr \vdash c_1 \quad \Gamma, \Pi, \phi \wedge k_1 \not\sqsubseteq k_2, \rho_C, fr \vdash c_2}{\Gamma, \Pi, \phi, \rho_C, fr \vdash \text{if } (k_1 \sqsubseteq k_2) \text{ c}_1 \text{ c}_2}
\end{array}$$

Fig. 12: Excerpts of the typing rules for commands.

existentially quantified type is well-typed. The rule follows previous work on existential types for security-typed languages [36]: the type r , in which $_$ may appear free, must be a subtype of the declared type s , and to prevent implicit flows the program counter label must also flow to the label on s . Furthermore, the frame label is raised to reflect that two new variables, each of which has a type that may depend on sensitive information, is now part of the frame layout. Finally, the command is type-checked in the updated environments with the raised frame label. Rule T-FLOWS TO branches on the runtime relation between the two labels k_1 and k_2 . Each command is checked in the extended formula capturing whether $k_1 \sqsubseteq k_2$ holds at runtime.

3) *Typing judgment for types and labels*: The typing judgments for security types and labels are straightforward. Figure 13 shows an excerpt of the typing judgment for security types, and the judgment for labels is similar. Rule T-INT says that, if the label k depends on information up to label k' then int_k depends on information up to label k' as well. Rule T-MU states that a recursive type $(\mu _ : \text{type}_{k_1}. s)_{k_2}$ depends on information up to label k if, assuming $_$ depends on information up to label k , the type s can be shown to depend on information up to k and finally, both k_1 and k_2 must also not depend on information above k .

$$\begin{array}{c}
\text{T-MU} \\
\frac{\phi \vdash k_1 \sqsubseteq k \quad \Pi[_ \mapsto \text{type}_{k_1}]; \phi \vdash_{\text{lab}} s : k}{\Pi; \phi \vdash_{\text{lab}} k_1 : k \quad \Pi; \phi \vdash_{\text{lab}} k_2 : k} \\
\\
\text{T-INT} \\
\frac{\Pi; \phi \vdash_{\text{lab}} k : k' \quad \Pi, \phi \vdash_{\text{type}} \text{int}_k : k'}{\Pi, \phi \vdash_{\text{type}} (\mu _ : \text{type}_{k_1}. s)_{k_2} : k}
\end{array}$$

Fig. 13: Excerpts of the typing relation for security types.

$$\begin{array}{c}
\text{S-LIFT} \\
\frac{\langle c, M, P, q \rangle_{\nu} \rightarrow \langle c', m', P', q' \rangle_{\nu'}}{\langle c, M, P, h, q \rangle_{\nu} \rightarrow \langle c', m', P', h, q' \rangle_{\nu'}} \\
\\
\text{S-INST} \quad \text{T-INST} \\
\frac{\langle c, M, P, h, q \rangle_{\nu} \rightarrow \langle c', M', P', h', q' \rangle_{\nu'}}{\langle c, M, P, h, q \rangle_{\nu} \rightarrow \langle c', M', P', h', q' \rangle_{\nu'}} \quad \frac{\Gamma, \Pi, \phi, \rho_C, fr \vdash c}{\Gamma, \Pi, \phi, \rho_C, fr \vdash c}
\end{array}$$

Fig. 14: Extending Zee with rules for modular extensions of the reduction semantics and the typing judgment.

D. An extensible language

To allow the specification of additional operations in Zee, we include a *hole* $[_]$ command:

$$c ::= \dots \mid [_]$$

We call the language without the hole construct the *base* language, and the additional commands the *instantiation* language. We let c range over commands in the instantiation language, and write \mathcal{C} for the set of commands in the base language. Given an instantiation language \mathcal{D} we write $\mathcal{C}[\mathcal{D}]$ for the set containing commands from both the base language and the instantiation language.¹¹

We add a heap to the configuration, which can be modified by the instantiation language. A heap is a partial mapping $h : A \rightarrow v$ from addresses to values. We write $\text{dom}(h)$ for the set of addresses currently allocated in h . We add an additional rule, S-LIFT, that lifts the semantics of base commands to configurations that include a heap.

Finally, we add a rule for specifying semantics of commands in $\mathcal{C}[\mathcal{D}]$: rule S-INST delegates reduction steps to the small-step semantics of the instantiation language. The new rules are shown in Figure 14. We extend the typing judgment with an additional rule T-INST that delegates typing to the typing relation for the instantiation language using the typing judgment \vdash provided by the instantiation language.

IV. SECURITY GUARANTEES

In this section we formalize the security guarantees obtained by adhering to the type system described in Section III-C. Section IV-A defines the attacker model, and Section IV-B specifies the semantic interface that each instantiation language must satisfy. Finally, Section IV-C defines *termination-insensitive timing-sensitive noninterference* (TINI) [8], [24] and shows that well-typed programs satisfy TINI.

¹¹Section IV-B formally defines the notion of an instantiation language.

$$\begin{aligned}
ev ::= & \varepsilon \mid \text{asgn}(x \leftarrow v, q) \mid \text{rd}(x \leftarrow v, q) \mid [\cdot] \\
& \mid \text{unp}(\ell, x : \tau \leftarrow v, q) \mid \text{let}(x : \tau \leftarrow v, q)
\end{aligned}$$

Fig. 15: Grammar for events.

A. Attacker model

To precisely define the security condition we introduce an augmented semantics that adds observable events to the reduction rules. We associate the attacker with a fixed level $\mathcal{A} \in \mathcal{L}$, and now define what an attacker at level \mathcal{A} can observe and which values \mathcal{A} can distinguish.

1) *Events and event semantics:* The grammar of events is shown in Figure 15. We assume that only commands that modify the stack generates an event, but nothing fundamental prevents adding a more fine-grained syntax of events. Event $\text{asgn}(x \leftarrow v, q)$ contains the variable x assigned to, along with the value v assigned to x , and the time q of when the assignment happened. Similarly, $\text{rd}(x \leftarrow v, q)$ describes obtaining a value v from the stack by reading a pointer, and assigning the value to variable x at time q . Event $\text{unp}(\ell, x : \tau \leftarrow v, q)$ describes an unpack command that declares a type (or level) variable at security level ℓ , and a variable of type τ with the initial value v at time q . Event $\text{let}(x : q \leftarrow v, q)$ describes the declaration of a regular variable x of type τ with initial value v at time q .

As our events capture the time at which the events are emitted, our definition of noninterference is timing-sensitive. Finally, like commands, the language of events can be extended with *instantiation events* using a hole construct $[\cdot]$, and we write the events of the instantiation language as ev .

We denote by \widetilde{ev} an *event tuple* of the form (ev, Γ, P) where Γ is the typing environment and P is the private stack, and we define an event semantics $\xrightarrow{\widetilde{ev}}$ over configurations that emits event tuples. Finally, we write $\xrightarrow{t^*}$ for the reflexive, transitive closure of the event semantics relation that concatenates all event tuples into a trace t .

2) *Attacker observability:* Given some level $\mathcal{A} \in \mathcal{L}$ we say that τ_ℓ is *observable* to \mathcal{A} if $\ell \sqsubseteq \mathcal{A}$, and *invisible* to \mathcal{A} otherwise. We lift observability to events as follows: given an event ev we write $\Gamma, P \vdash ev \sqsubseteq \mathcal{A}$ if \mathcal{A} can observe event ev given typing environment Γ and private stack P . We write $\widetilde{ev} \sqsubseteq \mathcal{A}$ if $\widetilde{ev} = (ev, \Gamma, P)$ and $\Gamma, P \vdash ev \sqsubseteq \mathcal{A}$. Section IV-B places restrictions on the instantiation events which are necessary for the noninterference theorem to hold.

3) *Attacker equivalence:* We say two values $v_i : \tau_i$ for $i = 1, 2$ are \mathcal{A} -equivalent given private stacks frames p_1 and p_2 , written $p_1, p_2 \vdash v_1 =_{\mathcal{A}} v_2 : \tau_1 \times \tau_2$, if an attacker \mathcal{A} is unable to distinguish them.

We lift \mathcal{A} -equivalence of values to \mathcal{A} -equivalence of events and write $\Gamma \mid p_1, p_2 \vdash ev_1 =_{\mathcal{A}} ev_2$ when events ev_1 and ev_2 are \mathcal{A} -equivalent. The typing environment Γ is needed in the judgment to associate a type with the variable x in the case of an assignment event $\text{asgn}(x \leftarrow v, q)$. All judgments are spelled out in the technical appendix.

Given a trace t , Figure 16 defines the \mathcal{A} -projected trace $[t]_{\mathcal{A}}$ containing only \mathcal{A} -observable events. We say two traces t_1 and

$$\begin{aligned}
[\varepsilon]_{\mathcal{A}} = \varepsilon \quad & [\widetilde{ev} \cdot t]_{\mathcal{A}} = \begin{cases} \widetilde{ev} \cdot [t]_{\mathcal{A}} & \widetilde{ev} \sqsubseteq \mathcal{A} \\ [t]_{\mathcal{A}} & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 16: \mathcal{A} -projected trace.

t_2 are \mathcal{A} -equivalent, written $t_1 =_{\mathcal{A}} t_2$, if $[t_1]_{\mathcal{A}}$ and $[t_2]_{\mathcal{A}}$ are pairwise \mathcal{A} -equivalent. Finally, given two pairs of exposed and private stack frames (p_i, m_i) we write $\Gamma \vdash (p_1, m_1) =_{\mathcal{A}} (p_2, m_2)$ when an attacker \mathcal{A} is unable to distinguish their content.

We extend this judgment pointwise and obtain an \mathcal{A} -equivalence on exposed and private stacks, which we write as $\Gamma \vdash (P_1, M_1) =_{\mathcal{A}} (P_2, M_2)$.

B. Specification of an instantiation language

In this section we describe the specification that an instantiation language must satisfy. To define the requirements of the instantiation language, we define an augmented small-step semantics $\Gamma, \Pi, \phi, \rho c, fr \vdash C \rightarrow C' : \Gamma', \Pi', \phi', \rho c', fr'$. This relation specifies that C steps to C' and updates the typing environments Γ, Π , constraints ϕ and labels ρc and fr to $\Gamma', \Pi', \phi', \rho c'$ and fr' respectively. We lift this relation to the event semantics relation and write $\Gamma, \Pi, \phi, \rho c, fr \vdash C \xrightarrow{\widetilde{ev}} C' : \Gamma', \Pi', \phi', \rho c', fr'$ when event tuple \widetilde{ev} is emitted when evaluating $\Gamma, \Pi, \phi, \rho c, fr \vdash C \rightarrow C' : \Gamma', \Pi', \phi', \rho c', fr'$. A well-formedness relation $\Gamma, \Pi, \phi \vdash C$ in the technical report formalizes well-formed configurations, and given a private stack P and a constraint formula ϕ relation $P \vdash \phi$ specifying that ϕ is true when evaluating all labels in ϕ in the private stack P .

Formally, an instantiation language \mathcal{D} is a tuple 6-tuple $(c, \rightarrow, \vdash, ev, =_{\mathcal{A}}, \sqsubseteq)$ where c is a set of syntactically valid commands, relation \rightarrow is a small-step relation on configurations, and \vdash is a typing judgment. The set ev contains syntactically valid events, and relations $\Gamma \mid p_1, p_2 \vdash ev_1 =_{\mathcal{A}} ev_2$ and $\phi, P \vdash ev \sqsubseteq \mathcal{A}$ defines when two events ev_1 and ev_2 are considered equivalent by an attacker at level \mathcal{A} , and when an event is observable to a \mathcal{A} respectively.

For the following properties, let c be a command c such that $\Gamma, \Pi, \phi, \rho c, fr \vdash c$, and let (M, P) and h be a stack and a heap such that $\Gamma, \Pi, \phi \vdash \langle c, M, h, P, q \rangle_{\nu}$ and $P \vdash \phi$. The following three properties must then be satisfied:

Property 1 (Single-run reduction properties). *If*

$$\begin{aligned}
& \Gamma, \Pi, \phi, \rho c, fr \vdash \langle c, M, h, P, q \rangle_{\nu} \rightarrow \\
& \quad \langle c', M', h', P', q' \rangle_{\nu'} : \Gamma', \Pi', \phi', \rho c', fr'
\end{aligned}$$

it holds that $\Gamma \sqsubseteq \Gamma', \Pi \sqsubseteq \Pi', \nu \leq \nu'$, and $\phi' \implies \phi$. Finally it holds that $\Gamma', \Pi', \phi' \vdash \langle c', M', h', P', q' \rangle_{\nu'}$, $P' \vdash \phi'$ and $\Gamma', \Pi', \phi', \rho c', fr' \vdash c'$.

Property 1 formalizes the type-safety requirements of the instantiation language. Intuitively, the semantics of the instantiation language should preserve well-formedness of configurations (i.e., $\Gamma', \Pi', \phi' \vdash \langle c', M', h', P', q' \rangle_{\nu'}$). Furthermore, the semantics should not prevent future use of variables that

are already in scope by removing them from the typing environments Γ' or Π' (i.e., $\Gamma \subseteq \Gamma'$ and $\Pi \subseteq \Pi'$). To prevent the possibility of reusing version numbers the version counter ν' should not decrease (i.e., $\nu \leq \nu'$), and finally the semantics must not weaken the constraint formula ϕ' ($\phi' \implies \phi$), but should also not strengthen the formula to the point where it is not guaranteed to hold at runtime (i.e., $P' \vdash \phi'$).

Property 2 (Single-step noninterference). *If $\Gamma \vdash (P_1, M_1) =_{\mathcal{A}} (P_2, M_2)$ and $\phi \vdash \rho c \sqsubseteq \mathcal{A}$ and*

$$\Gamma, \Pi, \phi, \rho c, fr \vdash \langle c, M_i, h_i, P_i, q \rangle_{\nu} \xrightarrow{\overline{ev}} \langle c'_i, M'_i, h'_i, P'_i, q' \rangle_{\nu'} : \Gamma', \Pi', \phi', \rho c', fr'$$

for $i = 1, 2$ then $\Gamma'_1 = \Gamma'_2$, $\Pi'_1 = \Pi'_2$, $\Gamma'_1 \vdash (P'_1, M'_1) =_{\mathcal{A}} (P'_2, M'_2)$, $q'_1 = q'_2$ and $\Gamma'_1 \mid P'_1, P'_2 \vdash \overline{ev}_1 =_{\mathcal{A}} \overline{ev}_2$.

Property 2 ensures that a command c in \mathcal{A} -equivalent environments results in \mathcal{A} -equivalent observations for a single-step.

Property 3 (Confinement). *If $\phi \vdash \rho c \not\sqsubseteq \mathcal{A}$ and*

$$\Gamma, \Pi, \phi, \rho c, fr \vdash \langle c, M, h, P, q \rangle_{\nu} \xrightarrow{\overline{ev}} \langle c', M', h', P', q' \rangle_{\nu'} : \Gamma', \Pi', \phi', \rho c', fr'$$

then $\Gamma \vdash (P, M) =_{\mathcal{A}} (P', M')$ and $\phi, P \vdash \overline{ev} \not\sqsubseteq \mathcal{A}$.

Property 3 ensures that the semantics does not leak sensitive information through indirect flows. That is, when the reachability of a program point depends on sensitive information (i.e., $\phi \vdash \rho c \not\sqsubseteq \mathcal{A}$), no \mathcal{A} -observable event is emitted, and \mathcal{A} is unable to distinguish the environments before and after the execution of c .

When properties 1, 2 and 3 are satisfied we say that the instantiation language $(c, \rightarrow, \vdash, \overline{ev}, =_{\mathcal{A}}, \sqsubseteq)$ is a well-formed instantiation language.

C. Security guarantees

Finally, we show that well-typed Zee programs satisfy termination-insensitive timing-sensitive noninterference. This definition permits attackers to learn information by observing the termination-behavior of the program, but it does not permit an attacker to learn information due to the timing-behavior of terminating programs. This distinction between termination and timing is unusual compared to previous literature where timing-sensitivity implies termination-sensitivity [17] but in a setting like ours, where a program can fail to terminate in many different ways, i.e., by attempting to read invalid memory or by non-exhaustive pattern matching, this definition is suitable [24].

Theorem 1 (Soundness). *Let \mathcal{D} be a well-formed instantiation language and let $c \in \mathcal{C}[\mathcal{D}]$, and let Γ, Π be typing environments. Assume $\Gamma, \Pi \vdash c$ and for all function definitions*

$$f(\overline{\cdot} : k_1) \langle \overline{\cdot} : k_2 \rangle (\overline{x} : \overline{s}) =_{\rho c}^{fr} c_f$$

it holds that $\Gamma_f, \Pi_f, \top, \rho c, fr \vdash c_f$, where

$$\Gamma_f = \{ \overline{x} \mapsto \overline{s} \} \quad \Pi_f = \{ \overline{\cdot} \mapsto \text{level}_{k_1} \} \cup \{ \overline{\cdot} \mapsto \text{type}_{k_2} \}.$$

If $\Gamma \vdash (P_1, M_1) =_{\mathcal{A}} (P_2, M_2)$ and $\langle c, M_i, P_i \rangle \xrightarrow{t_i}^*$ for $i = 1, 2$ then $t_1 =_{\mathcal{A}} t_2$.

Theorem 1 states that, if each function definition is well-typed, and the command c is well-typed, then executing c with two \mathcal{A} -equivalent stacks will result in \mathcal{A} -equivalent traces t_1 and t_2 .

V. CASE STUDIES

This section presents two case studies demonstrating realistic activities of a runtime system for a modern programming language. Both case studies, and all programs presented in the paper, are executable using our prototype implementation.

The first case study is an implementation of a segregated garbage collector (GC) that splits the heap into partitions indexed by security levels from a fixed lattice [24]. The algorithm is a modified version of a mark-and-sweep collector [13], and to the best of our knowledge, is the first GC algorithm that implements the abstract semantics formally proven secure in [24]. The security property obtained from the well-typedness (Theorem 1) of the GC implementation implies that the timing behavior of the garbage collector does not depend on the memory operations caused by handling sensitive information.

The second case study is an implementation of a simple cooperative thread scheduling algorithm. The security property obtained from the well-typedness of the thread scheduler implies that the scheduling of “public threads” is independent from the presence of threads spawned due to handling of sensitive information.

The programs in the case studies use a syntax more suitable for programming compared to the formal language, but it can be desugared into the core calculus presented in Section III.

A. Secure garbage collection

The job of a GC is to reclaim memory that will not be used in the future by the program. This property is in general undecidable, and GC algorithms instead only reclaim memory that is not *reachable* by the program. The GC implementation is split into two phases: The *marking* phase and the *sweeping* phase. The marking phase starts when the RTE decides that a GC is needed. Our GC implementation is a stop-the-world collector: it stops the execution of the program and marks every heap allocation that is currently reachable by the program. Guaranteeing security and type-safety for such an operation is nontrivial as data of different types, and with different security policies, must be traversed and handled differently depending on the base type of the value, and its security label.

We proceed by defining the syntax of an instantiation language MS : a language for implementing secure mark-and-sweep garbage collectors. We then define the small-step semantics and the typing relation of MS . Finally, we show that MS is a well-formed instantiation language (cf. Section IV-B).

1) *The instantiation language MS* : We instantiate Zee with the instantiation language MS , whose commands are shown in Figure 17. We assume an operation on ℓ -indexed partitioning of the heap and write \mathbb{A}^{ℓ} for the set of addresses belonging to security level ℓ .

```


c ::= mark e | x := is_marked e | free e | x := length e
    | x := alloc(e, e, s, k, k) | x := start k | x := next e
    | x := read(e) | write(e1, e2)

```

Fig. 17: Instantiation language for mark-and-sweep garbage collection.

Command $\text{mark } e$ marks a heap address, representing the information that the address is reachable in the heap. Command $x := \text{is_marked } e$ checks if an address, given by the evaluation of e , has previously been marked, and stores this information in variable x . Command $\text{free } e$ reclaims the memory pointed to by e , and $x := \text{alloc}(e_1, e_2, s, k_1, k_2)$ allocates e_1 number of entries, all initialized to the value e_2 , in the heap partition associated with the security level k_2 , and where the label k_1 denotes the sensitivity of the size of the allocation. Command $x := \text{length } e$ stores the length of an allocation pointed to by e in variable x . The final two commands are used in the sweep phase, and implement an abstract notion of heap parsability [13]: command $x := \text{start } k$ stores a pointer to the first allocation in the heap partition associated with the security level k in variable x , and $x := \text{next } e$ stores the next pointer in the same heap partition as e (i.e., the allocation with the smallest address that is larger e) in variable x .

Figure 19 shows the small-step semantics for allocation in MS . A heap allocation of size n is structured as follows:



That is, the first address stores the *mark* of the allocation, which is used during the marking phase of garbage collection to denote that the allocation is reachable by the program. Next to the mark is the *version* entry, which ensures that it is not possible to read stale values from the heap when addresses are reused. This is similar to the technique described in Section II-C for values stored on the exposed stack. The *meta* field stores type and label information using existentially quantified labels and types: Specifically, a value of type $\exists \text{ : level}_\ell. \exists \text{ : type}_{e_\ell}. \text{int}$ is stored. Here, the level ℓ is the index of the partition in which the allocation is stored on the heap. The label int stores the security label on the size of the array, which is used during the sweep phase, and the type int stores the type information about the type of elements in the array, which is used to traverse the heap during the marking phase, and the integer with security label int stores the length of the allocation. Finally, the data of the array is stored at the end.

Rule S-ALLOC allocates such a data structure on the heap. First, an address a is found such that the range $a, \dots, a+n+2$ is free (i.e., not part of the domain of the heap). A pointer to the first element of the array is then stored in stack variable x , and the address is given a fresh version count. Finally, the allocation structure is stored on the heap, and the version counter is incremented.

The version number in each allocation prevents leaks caused by dangling pointers and aliasing. To see an example of this,

```

1 let p : (L ∇ intL)L := null in
2 p := alloc(10, 0, intL, L, L);
3 free(p);
4 let q : (L ∇ intH)L := null in
5 q := alloc(10, h, intH, L, L)

```




Fig. 18: Top: a program that attempts to read the secret h through the pointer p when the RTE decides to reuse the recently freed memory for the allocation on line 4. Bottom: stack- and heap layout after executing line 5 of the program.

S-ALLOC

$$a = \operatorname{argmin}([a, \dots, a+n+2] \cap \operatorname{dom}(h) = ?)$$

$$a \in A \text{ s.t. } a, \dots, a+n+2 \in A \setminus \ell_2$$

$$c = x := \text{alloc}(e_1, e_2, s, k_1, k_2) \quad \langle e_1, m, p \rangle \Downarrow n$$

$$P = p \cdot P' \quad \langle e_2, m, p \rangle \Downarrow v \quad \langle s, p \rangle \Downarrow_{\text{type}} \tau$$

$$\langle k_i, p \rangle \Downarrow_{\text{lab}} \ell_i \quad m' = m[\delta(x) + \text{fp}(m) \mapsto (a+3)]_\nu$$

$$h' = h[a \mapsto 0, a+1 \mapsto \nu, a+2 \mapsto (\ell_1, (\tau, n))] \cup \{a+i+3 \mapsto v \mid i \in 0, \dots, n-1\}$$

$$\frac{\langle c, m \cdot M, P, h, q \rangle_\nu \rightarrow \langle \text{stop}, m' \cdot M, P, h', q+1 \rangle_{\nu+1}}$$

T-ALLOC

$$\frac{\Gamma, \Pi, \phi \vdash e_1 : \text{int}_{k_1} \quad \Gamma, \Pi, \phi \vdash e_2 : s \quad \Pi; \phi \vdash_{\text{lab}} k_1 : k_2 \quad \Pi, \phi \vdash_{\text{type}} s : k_2 \quad \Pi; \phi \vdash_{\text{lab}} k_2 : \perp \quad \phi \vdash (k_2 \mapsto s)_{k_1 \sqcup \rho C} <: \Gamma(x)}{\Gamma, \Pi, \phi, \rho C, \text{fr} \vdash x := \text{alloc}(e_1, e_2, s, k_1, k_2)}$$

Fig. 19: Static and dynamic semantics for allocation. The remaining judgments are defined in the technical report [23].

consider the program in Figure 18. On line 2 a pointer to an allocation containing public data is stored in p with a version number ν . On line 3 the memory is freed, making it possible for S-ALLOC to reuse the memory from the allocation on line 2. The memory is reused on line 5, causing p to point to secret data, even though the type of p specifies that it points to public data. However, an updated version number $\nu+1$ is stored when the memory is being reused, so any attempt to access the secret data through p will fail the version check.

Figure 19 also defines the typing judgment: first, the expression e_1 , which denotes the size of the allocation, must be of integer type with label at most k_1 . This ensures that the label argument k_1 correctly captures the security of the size of the allocation at runtime. Similarly, the expression e_2 , which denote the initial value of the array entries, must be typeable at type s , ensuring that the type argument s correctly captures the type information of the array at runtime. Finally, the partition label k_2 must be typeable at the label \perp . This guarantees that no information can be learned by knowing which partition the allocation happens in.

The events for MS is defined in the technical report, as well as the rest of the instantiation language. We conclude this section with the following lemma showing that MS satisfies

all the requirements stated in Section IV-B.

Lemma 1. *The instantiation language MS is well-formed.*

We now describe how MS can be used to implement a secure mark-and-sweep garbage collector. This case study is developed in a setting of the two-point lattice $L \sqsubseteq H$ from Section II-A.

2) *Using the instantiation language:* Figure 20 presents two functions representing the beginning of the marking phase. Function `mark_frames` is invoked by the `gc` function, which is invoked by the runtime.

```

1 gc hpc : L, fr : Li()  $\stackrel{pc, fr}{=}_{pc}$ 
2 let (  $\overset{2}{args}, e$  ) := unpack FP in
3 let (  $\overset{2}{local\ s}, p$  ) := unpack e in
4 let (  $\overset{2}{args}, e_2$  ) := unpack
5   * (  $p - \text{sizeof } T_{st}(pc, fr, L)$  ) in
6 let (  $\overset{2}{local\ s}, p_2$  ) := unpack e2 in
7 mark_frames hpc, fr ih  $\overset{2}{args}, \overset{2}{local\ s} i(p_2); \dots$ 

```

Function `gc` starts by reading the frame pointer `FP`, on line 2, to obtain an existentially quantified pointer e of type $(\exists \text{ locals} : \text{type}_{fr}. (\text{arg} \cdot T_{st}(pc, fr, L) @ \text{locals})_{pc})L$. On line 3 e is unpacked, revealing the pointer p of type $(\text{arg} \cdot T_{st}(pc, fr, L) @ \text{locals})_{pc}$, pointing to the beginning of `gc`'s stack frame. Lines 4 to 6 then follow the same procedure to obtain a pointer p_2 to the beginning of the previous stack frame (i.e., the function that was executing before the GC occurred). This pointer is then passed to a recursive function `mark_frames` on line 7, which traverses each stack frame. This function is shown in Figure 20. On line 3 the function checks if the pointer p is non-zero (i.e., we are not at the last stack frame). It marks allocations reachable from the stack frame starting at p using the function `mark_frame` on line 4, and computes the pointer to the previous stack frame on lines 5 to 7. On line 8 the function then invokes itself recursively to mark the previous frame starting at p_2 .

Function `mark_frame` is also recursively defined, as it traverses each entry in a single stack frame. On line 3 the function performs runtime type analysis on the type locals . If the runtime representation of the type is a product type with a pointer type $(p \mapsto \cdot)$ at its head, lines 5 to 10 are executed. Line 5 performs a dynamic “flows to” check¹² to ensure that it is secure to reclaim this allocation [24]. If so, the pointer is read off the stack on line 6, and line 7 marks the pointer if it is non-null. For simplicity we elide the code that marks objects recursively reachable from this object, but the full implementation is available in our technical appendix, and the code is executable using our prototype implementation of Zee.

After having marked all allocations reachable from q on line 7, `mark_frame` calculates the offset n to the address of the next entry in the stack frame on line 9, and then invokes itself recursively on line 10 with the new address $p + n$ as its argument.

```

1 mark_frames hpc : L, fr : Lih args : L, locals : Li
2 (p : ( args * Tst(pc, fr, L) @ locals)pc)  $\stackrel{fr}{=}_{pc}$ 
3 if p then
4   mark_frame hpc, fr ih args, locals i(p);
5   let (  $\overset{2}{args}, e$  ) := unpack
6     * (  $p - \text{sizeof } T_{st}(pc, fr, L)$  ) in
7   let (  $\overset{2}{local\ s}, p_2$  ) := unpack e in
8   mark_frames hpc, fr ih  $\overset{2}{args}, \overset{2}{local\ s} i(p_2)$ 
9 else skip

```

```

1 mark_frame hpc : L, fr : Lih args : L, locals : Li
2 (p : ( args * Tst(pc, fr, L) @ locals)pc)  $\stackrel{fr}{=}_{pc}$ 
3 match locals with
4   (  $\overset{2}{p} \mapsto \cdot$  ) * )
5   if  $\overset{2}{p}$  pc then
6     let q := *p in
7     if q then mark(q); ... else skip
8   else skip;
9   let n := sizeof (  $\overset{2}{p} \mapsto \cdot$  ) in
10  mark_frame hpc, fr ih args, i(p + n)
11  | * )
12  | let n := sizeof in
13  mark_frame hpc, fr ih args, i(p + n)
14  | _ ) skip

```

Fig. 20: Snippets from the GC case study: The stack frames are traversed (top) and each frame is traversed looking for pointers into the heap (bottom).

If locals is not of the form $(p \mapsto \cdot)$, but is still a product type \cdot , lines 12 and 13 are executed. Line 12 computes the number of addresses that must be skipped in order to skip past the current entry in the stack frame, and line 13 then calls the function recursively with the next address $p + n$ to inspect. Finally, if local is not a product type, the frame has been completely traversed and line 14 is executed, and the function returns.

B. Secure thread scheduling

Once we have the possibility of allocating memory on the heap, we can use the same instantiation language MS to implement a thread scheduler. Concurrency has received a lot of attention in the literature on language-based security [15], [28], [29], [31], [32], [40], especially in the context of timing-channels. Several authors [15], [28], [32] propose special-purpose thread schedulers designed to close such timing-channels, and in this section we present an implementation of a secure cooperative thread scheduling algorithm. For the purpose of this case study, each function written by the user is assumed to have been rewritten into continuation passing style (CPS), as is standard for many compilers for functional programming languages [3], [14], and defunctionalized into a form that contains no higher-order functions, i.e., closure is an identifier followed by a heterogeneous array of local variables. Each security level from some fixed lattice is associated with a queue of closures, and a thread schedules a function f to be invoked by enqueueing its closure in the queue associated with the program counter label of f . The scheduler is a function `schedul e` that receives a queue for each security level, and

¹²The expression $\text{flows } 1 \equiv 2$ is shorthand for $1 \sqsubseteq 2 \wedge 2 \sqsubseteq 1$.

a bound for how long to run sensitive computation. We have implemented a small security-typed queue datastructure in Zee that supports operations such as checking if the queue is empty, as well as queueing and dequeuing elements. We use a pseudocode-style description of the scheduling algorithm, and refer interested readers to the implementation for a precise description.

```

1 schedule(n : intL,
2 schedL : (L ↯ (9 : typeL . )L)L,
3 schedH : (L ↯ (9 : typeL . )H)H) =H
4 while nonempty(schedL) do
5   let ( , proc) := unpack dequeue(&schedL)
6   in runL/H i(proc);
7   at H with bound n do
8     if nonempty(schedH) then
9       let ( , proc) :=
10        unpack dequeue(&schedH)
11        in runH/H i(proc);
12     else skip

```

For simplicity, the initial program counter label is L , meaning that threads with program counter L do not need a bound on their computation time, but the scheduler can only be called when the program counter label is L . The frame label, on the other hand, is set to H allowing any sensitive information to flow to the types of the data, but any attempts to compute information based on types will be assigned the label H .

The thread scheduler executes a *public quanta* (i.e., the execution of one closure in the sched_L queue), followed by one *secret quanta* (i.e., the execution of one closure in the sched_H). There are positives and negative things to point out about the design of this scheduler: on the positive side, the bound on secret computations n only needs to bound one quanta, and the function that calls `schedule` does not need to consider the number of closures in sched_H . On the negative side, to guarantee timing-sensitive noninterference, each public quanta must be followed by n steps of computation, no matter if there is any secret threads to execute or not. Furthermore, in order to run secret threads, there must also be public threads available, as the while loop terminates when the sched_L queue is empty. An alternative strategy would be a scheduler that provides a bound on the total computation time on high threads. With this approach secret threads can run without the presence of public threads. Having developed this prototype, we leave the design of a more practical thread scheduler as future work.

VI. IMPLEMENTATION

We have implemented a type checker and interpreter for Zee in Haskell in about 3600 lines of code, and the case studies consists of about 500 lines each. Providing an instantiation language corresponds to an implementation of a particular type class, and the type checking and evaluation of instantiation language constructs is delegated to the relations provided by the instantiation language, similar to how the judgments in the paper are defined. The implementation can be found at: <https://www.dropbox.com/s/bl2juns8nqkqhu/zee.zip>.

VII. RELATED WORK

Our work on securing runtime environments combines previous efforts of memory safety for unsafe programming languages, extensible reasoning about type systems and information-flow control. We review the relevant literature in each of these classes separately.

A. Stack typing and memory safety

The work on typed assembly languages initiated by Morrisett *et al.* [20] paved the way for type systems for low-level programming languages. As the target language was expressed in continuation-passing style, there was no need for a stack. Morrisett *et al.* [19] introduced local stack variables, but as the goal of that work is type preserving compilation it does not support reasoning about stack traversal, and so the “previous frame pointer” is not available on the stack for accessing the previous stack frames. For this reason Morrisett *et al.* do not consider runtime type analysis.

Our stack typing discipline is inspired by the bunched adjacency logic of Ahmed and Walker [2]. They use logic formulae more^{\leftarrow} and $\text{more}^{\rightarrow}$ to describe the type of an infinite sequence of locations that increases “to the left” and “to the right” respectively, similar to our use of stack pointer types.

Our version-based enforcement mechanism is inspired by CETS’s [22] identifier-based temporal checking. CETS is a program transformation that adds temporal memory safety checking capable of detecting dangling pointer dereferences and double frees errors at runtime. Our version-based enforcement mechanism could be replaced with static reasoning about regions. Regions were introduced by Tofte and Talpin [35] and later used to provide memory safety for a safe dialect of C [10]. We believe the use of regions is orthogonal to our choice of a dynamic enforcement mechanism.

B. Attacks on runtimes

The work on observational determinism by Zdancewic and Myers [40] contains a detailed collection of common scheduler-related attacks. Other parts of the RTE that has been attacked include garbage collectors [24]. Pedersen and Askarov [24] present a series of attacks on the garbage collectors of the Java virtual machine and the V8 JavaScript engine, and design a type system and a small-step semantics for a high-level language with automatic memory management for which they prove a noninterference result similar to ours. Finally, Vassena *et al.* [37] present attacks that combine concurrent execution and lazy evaluation for leaking sensitive information. They propose a new construct for Haskell called `lazydup`, which lazily duplicates thunks on the heap when entering secret contexts (i.e., when the program counter label is H , as they only consider a two-point lattice).

C. Securing runtimes

Vassena *et al.* [38] present a new foundation for a dynamic information flow control parallel runtime system. The goal of their work is securing the execution platform of LIO [33], a dynamic information flow control library for Haskell. Similar

to our work, Vassena *et al.* [38] consider a setting in which an attacker can obtain the current global time as a natural number counting execution steps, and (unlike our model) the current size of the heap. They design a system for hierarchically managing space and time resources with some amount of burden on the programmer: a parent thread has to manually kill their child thread to reclaim resources. Their end goal is an implementation of a modified GHC runtime system, but such a modified runtime has yet to be implemented.

The work by Sabelfeld and Sands [28] contains an interesting observation:

Abstractly we will take a scheduler to be a mechanism for selecting threads which itself satisfies some noninterference property, i.e., its behaviour is independent of high data.

This is exactly the approach we have taken: a thread scheduler is a program written in our language, and Theorem 1 proves that, since this program is well-typed, its public observable behavior is independent of high data.

D. Static information flow control

There is a large body of literature focusing on static information flow control, starting with the seminal work by Denning and Denning [6] and later formulated as a type system by Volpano *et al.* [39]. Sabelfeld and Myers [27] survey the different enforcement techniques and security definitions. Zheng and Myers [44] introduce the technique of including a formulae expressing which flows are guaranteed to hold at specific program points, allowing for static reasoning about information flow policies that vary at runtime. The use of existentially quantified labels is introduced by Tse and Zdancewic [36], and we follow the same typing discipline for such values. Dependent type systems for IFC has also been explored by Lourenço and Caires [16], Zhang *et al.* [43] and Gregersen *et al.* [9].

E. Verified runtimes

There has been much work on verifying runtime system components such as garbage collectors [5], [7], [18] and thread schedulers [11] using program logics. We view our work complementary to these efforts. The constructs needed for implementing secure runtime environments, that we identify in this work, may serve as guidelines when applying enforcement techniques different from our type system. Additionally, a program logic may be used to verify the requirements of the instantiation languages used in Zee.

REFERENCES

- [1] M. Abadi, L. Cardelli, B. Pierce, and G. Plotkin, Dynamic Typing in a Statically-typed Language, in *Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM, 1989,
- [2] A. Ahmed and D. Walker, The Logical Approach to Stack Typing, in *Proceedings of the 2003 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation*, ACM, 2003,
- [3] A. W. Appel, *Compiling with continuations*. Cambridge University Press, 2006.
- [4] A. Askarov, D. Zhang, and A. C. Myers, Predictive Black-box Mitigation of Timing Channels, in *Proceedings of the 17th ACM Conference on Computer and Communications Security*, ACM, 2010,
- [5] L. Birkedal, N. Torp-Smith, and J. C. Reynolds, Local Reasoning About a Copying Garbage Collector, in *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM, 2004,
- [6] D. E. Denning and P. J. Denning, Certification of Programs for Secure Information Flow, *Commun. ACM*, Jul. 1977.
- [7] P. Gammie, A. L. Hosking, and K. Engelhardt, Relaxing Safely: Verified On-the-fly Garbage Collection for x86-TSO, in *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ACM, 2015,
- [8] J. A. Goguen and J. Meseguer, Security Policies and Security Models, in *1982 IEEE Symposium on Security and Privacy*, IEEE, Apr. 1982.
- [9] S. Gregersen, S. E. Thomsen, and A. Askarov, A Dependently Typed Library for Static Information-Flow Control in Idris, *CoRR*, 2019. arXiv: 1902.06590.
- [10] D. Grossman, G. Morrisett, T. Jim, M. Hicks, Y. Wang, and J. Cheney, Region-Based Memory Management in Cyclone, *ACM SIGPLAN Notices*, May 2002.
- [11] Y. Guo, X. Feng, Z. Shao, and P. Shi, Modular Verification of Concurrent Thread Management, in *Programming Languages and Systems*, R. Jhala and A. Igarashi, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2012,
- [12] R. Harper and G. Morrisett, Compiling Polymorphism Using Intensional Type Analysis, in *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM, 1995,
- [13] R. Jones, A. Hosking, and E. Moss, *The Garbage Collection Handbook: The Art of Automatic Memory Management*, 1st. Chapman & Hall/CRC, 2011.
- [14] S. L. P. Jones, Implementing lazy functional languages on stock hardware: the Spineless Tagless G-machine, *Journal of Functional Programming*, Apr. 1992.
- [15] A. Karbyshev, K. Svendsen, A. Askarov, and L. Birkedal, Compositional Non-interference for Concurrent Programs via Separation and Framing, in *Principles of Security and Trust*, Cham: Springer International Publishing, 2018,
- [16] L. Lourenço and L. Caires, Dependent Information Flow Types, in *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM, 2015,
- [17] H. Mantel and A. Sabelfeld, A Generic Approach to the Security of Multi-Threaded Programs, in *Proceedings of the 14th IEEE Workshop on Computer Security Foundations*, IEEE Computer Society, 2001,
- [18] A. McCreight, Z. Shao, C. Lin, and L. Li, A General Framework for Certifying Garbage Collectors and Their Mutators, in *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ACM, 2007,
- [19] G. Morrisett, K. Crary, N. Glew, and D. Walker, Stack-based Typed Assembly Language, in *Types in Compilation*, X. Leroy and A. Ohori, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 1998,
- [20] J. G. Morrisett, D. Walker, K. Crary, and N. Glew, From System F to Typed Assembly Language, in *POPL*, 1998.
- [21] A. C. Myers and A. C. Myers, JFlow: Practical Mostly-static Information Flow Control, in *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM, 1999,

- [22] S. Nagarakatte, J. Zhao, M. M. Martin, and S. Zdancewic, CETS: Compiler enforced temporal safety for C, in *Proceedings of the 2010 International Symposium on Memory Management*, ACM, 2010.
- [23] Pedersen, Mathias V., Askarov, Aslan, “Static enforcement of secure runtime systems: technical report,” Tech. Rep., 2018.
- [24] M. V. Pedersen and A. Askarov, From Trash to Treasure: Timing-Sensitive Garbage Collection, in *2017 IEEE Symposium on Security and Privacy (SP)*, IEEE, May 2017.
- [25] F. Perry, C. Hawblitzel, and J. Chen, Simple and Flexible Stack Types, Jul. 2007,
- [26] B. C. Pierce, Types and Programming Languages, 1st. The MIT Press, 2002.
- [27] A. Sabelfeld and A. C. Myers, Language-based Information-flow Security, *IEEE J.Sel. A. Commun.*, Sep. 2006.
- [28] A. Sabelfeld and D. Sands, Probabilistic noninterference for multi-threaded programs, in *Proceedings 13th IEEE Computer Security Foundations Workshop. CSFW-13*, IEEE Comput. Soc.
- [29] A. Sabelfeld, The Impact of Synchronisation on Secure Information Flow in Concurrent Programs, in *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2001,
- [30] L. Skorstengaard, D. Devriese, and L. Birkedal, Reasoning About a Machine with Local Capabilities, in *Programming Languages and Systems*, A. Ahmed, Ed., Cham: Springer International Publishing, 2018,
- [31] G. Smith and D. Volpano, Secure Information Flow in a Multi-threaded Imperative Language, in *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM, 1998,
- [32] D. Stefan, A. Russo, P. Buiras, A. Levy, J. C. Mitchell, and D. Mazières, Addressing Covert Termination and Timing Channels in Concurrent Information Flow Systems, in *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming*, ACM, 2012,
- [33] D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières, Flexible Dynamic Information Flow Control in Haskell, in *Proceedings of the 4th ACM Symposium on Haskell*, ACM, 2011,
- [34] C. A. Stone and R. Harper, Extensional Equivalence and Singleton Types, *ACM Trans. Comput. Logic*, Oct. 2006.
- [35] M. Tofte and J.-P. Talpin, Region-Based Memory Management, *Inf. Comput.*, Feb. 1997.
- [36] S. Tse and S. Zdancewic, Run-time Principals in Information-flow Type Systems, *ACM Trans. Program. Lang. Syst.*, Nov. 2007.
- [37] M. Vassena, J. Breitner, and A. Russo, Securing Concurrent Lazy Programs Against Information Leakage, in *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, IEEE, Aug. 2017.
- [38] M. Vassena, G. Soeller, P. Amidon, M. Chan, and D. Stefan, Towards parallel information flow control foundations, in *Principles of Security and Trust*, 2019.
- [39] D. Volpano, C. Irvine, and G. Smith, A Sound Type System for Secure Flow Analysis, *J. Comput. Secur.*, Jan. 1996.
- [40] S. Zdancewic and A. Myers, Observational determinism for concurrent program security, in *16th IEEE Computer Security Foundations Workshop, 2003. Proceedings.*, IEEE Comput. Soc.
- [41] D. Zhang, A. Askarov, and A. C. Myers, Language-based Control and Mitigation of Timing Channels, in *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation*, ACM, 2012,

$$\frac{}{\text{int}_\ell - \text{int}} \quad \frac{\tau - p}{(\ell_1 \mapsto \tau)_{\ell_2} - (\ell_1 \mapsto p)_2}$$

$$\frac{\tau_i - p_i \quad i = 1, 2}{(\tau_1 @ \tau_2)_\ell - (p_1 @ p_2)} \quad \frac{}{\tau - p}$$

$$\frac{|\bar{\tau}| = n \quad |\bar{p}| = m \quad m \leq n \quad \forall i \in \{1, \dots, m-1\} . \tau_i - p_i \quad \tau_{m\dots n} - p_m}{\bar{\tau} - \bar{p}}$$

(a) Pattern Matching relation.

$$\text{Jint } \llbracket (p, \text{int}_\ell) \rrbracket = p[\ell \mapsto \ell]$$

$$\frac{\text{J}p_1 \llbracket (p, \tau_1) \rrbracket = p' \quad \text{J}p_2 \llbracket (p', \tau_2) \rrbracket = p''}{\text{J}(p_1 @ p_2) \llbracket (p, (\tau_1 @ \tau_2)_\ell) \rrbracket = P''[\ell \mapsto \ell]}$$

$$\frac{\text{J}p \llbracket (p, \tau) \rrbracket = p' \quad p'' = p'[\ell_1 \mapsto \ell_1, \ell_2 \mapsto \ell_2]}{\text{J}(\ell_1 \mapsto p)_2 \llbracket (p, \ell_1 \mapsto \tau_{\ell_2}) \rrbracket = p''} \quad \text{J} \llbracket (p, \tau) \rrbracket = p[\ell \mapsto \tau]$$

$$\frac{|\bar{\tau}| = n \quad |\bar{p}| = m \quad m \leq n \quad p_0 = p \quad \forall i \in \{1, \dots, m-1\} . \text{J}p_i \llbracket (p_{i-1}, \tau_i) \rrbracket = p_i \quad \text{J}p_m \llbracket (p_{m-1}, \tau_{m\dots n}) \rrbracket = p'}{\text{J}\bar{p} \llbracket (p, \bar{\tau}) \rrbracket = p'}$$

(b) Binding free variables in the pattern p by deconstructing the type value τ .

$$\Pi \vdash \text{int} \quad k \quad \Pi[\ell \mapsto \text{level}_k] : \text{int}$$

$$\Pi \vdash k \quad \Pi[\ell \mapsto \text{type}_k] :$$

$$\frac{\Pi \vdash p_1 \quad k \quad \Pi_1 : s_1 \quad \Pi_1 \vdash p_2 \quad k \quad \Pi_2 : s_2}{\Pi \vdash (p_1 @ p_2) \quad k \quad \Pi_2[\ell \mapsto \text{level}_k] : (s_1 @ s_2)}$$

$$\frac{\Pi \vdash p \quad k \quad \Pi' : s}{\Pi \vdash (\ell_1 \mapsto p)_2 \quad k \quad \Pi'[\ell_1 \mapsto \text{level}_k, \ell_2 \mapsto \text{level}_k] : (\ell_1 \mapsto s)_2}$$

$$\frac{\Pi_0 = \Pi \quad \Pi_{i-1} \vdash p_i \quad k \quad \Pi_i : s_i \quad i = 1, \dots, n}{\Pi \vdash p_1, \dots, p_n \quad k \quad \Pi_n : s_1, \dots, s_n}$$

(c) Relation for closing the free type variables in the pattern p , and computing the type of the scrutinee when it matches pattern p .

- [42] —, Predictive Mitigation of Timing Channels in Interactive Systems, in *Proceedings of the 18th ACM Conference on Computer and Communications Security*, ACM, 2011,

- [43] D. Zhang, Y. Wang, G. E. Suh, and A. C. Myers, A Hardware Design Language for Timing-Sensitive Information-Flow Security, in *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems*, ACM, 2015,

- [44] L. Zheng and A. C. Myers, Dynamic Security Labels and Static Information Flow Control, *Int. J. Inf. Secur.*, Mar. 2007.

APPENDIX

$$\begin{array}{c}
\text{S-IF-T} \\
\frac{\langle e, M, P \rangle \Downarrow n \quad n \neq 0}{\langle \text{if } e \ c_1 \ c_2, M, P, q \rangle_\nu \rightarrow \langle c_1, M, P, q + 1 \rangle_\nu} \\
\\
\text{S-IF-F} \\
\frac{\langle e, M, P \rangle \Downarrow 0}{\langle \text{if } e \ c_1 \ c_2, M, P, q \rangle_\nu \rightarrow \langle c_2, M, P, q + 1 \rangle_\nu} \\
\\
\text{S-WHILE-F} \\
\frac{\langle e, M, P \rangle \Downarrow 0}{\langle \text{while } e \ c, M, P, q \rangle_\nu \rightarrow \langle \text{stop}, M, P, q + 1 \rangle_\nu} \\
\\
\text{S-WHILE-T} \\
\frac{\langle e, M, P \rangle \Downarrow n \quad n \neq 0}{\langle \text{while } e \ c, M, P, q \rangle_\nu \rightarrow \langle c; \text{while } e \ c, M, P, q + 1 \rangle_\nu} \\
\\
\text{S-SKIP} \\
\langle \text{skip}, M, P, q \rangle_\nu \rightarrow \langle \text{stop}, M, P, q + 1 \rangle_\nu \\
\\
\text{S-WRITE} \\
\frac{m_i = (m_1, \nu_i) \in M \quad a \in \text{I} \quad \nu_i \leq \gamma \quad \langle e_1, M, P \rangle \Downarrow a_\gamma \quad \langle e_2, M, P \rangle \Downarrow v}{\langle *e_1 := e_2, M, P, q \rangle_\nu \rightarrow \langle \text{stop}, M[a \mapsto v], P, q + 1 \rangle_\nu} \\
\\
\text{S-READ} \\
\frac{M = m \cdot M' \quad m_i = (m_1, \nu_i) \in M \quad a \in \text{I} \quad \nu_i \leq \gamma \quad \langle e, M, P \rangle \Downarrow a_\gamma \quad m' = m[\delta(x) + \text{fp}(m) \mapsto M(n)]}{\langle x := *e, m \cdot M, P, q \rangle_\nu \rightarrow \langle \text{stop}, m' \cdot M', P, q + 1 \rangle_\nu} \\
\\
\text{S-AT} \\
\frac{\langle e, m, P \rangle \Downarrow n}{\langle \text{at } k \ e \ c, m, P, q \rangle_\nu \rightarrow \langle c; \text{delay } n, m, P, q + 1 \rangle_\nu} \\
\\
\text{S-DELAY} \\
\frac{n \leq q}{\langle \text{delay } n, m, P, q \rangle_\nu \rightarrow \langle \text{delay } n, m, P, n + 1 \rangle_\nu} \\
\\
\text{S-FP} \\
\frac{v = (\text{cod}(p.\text{arg}), (\text{cod}(p.\text{local}), \text{fp}(m)_\nu)) \quad m = (\text{I}, |m|, \nu) \quad m' = m[\delta(x) + \text{fp}(m) \mapsto v]}{\langle x := \text{fp}, m \cdot M, p \cdot P, q \rangle_\nu \rightarrow \langle \text{stop}, m' \cdot M, p \cdot P, q + 1 \rangle_\nu} \\
\\
\text{S-LET} \\
\frac{M = m \cdot M' \quad \langle s, p \rangle \Downarrow_{\text{type}} \tau \quad \langle e, m, p \rangle \Downarrow v \quad m' = m[\delta(x) + \text{fp}(m) \mapsto v] \quad p' = p[p_{\text{local}} \mapsto p_{\text{local}}[x \mapsto \tau]]}{\langle \text{let } x : s := e \text{ in } c, M, p \cdot P, q \rangle_\nu \rightarrow \langle c; \text{unscope}(x), m' \cdot M', p' \cdot P, q + 1 \rangle_\nu} \\
\\
\text{S-UNSCOPE} \\
\frac{p' = p[\text{local} \mapsto p.\text{local}[x \mapsto \]]}{\langle \text{unscope}(x), M, p \cdot P, q \rangle_\nu \rightarrow \langle \text{stop}, M, p' \cdot P, q + 1 \rangle_\nu} \\
\\
\text{S-UNPACK-LEV} \\
\frac{P = p \cdot P' \quad \langle s, p' \rangle \Downarrow_{\text{type}} \tau \quad \langle e, m, p \rangle \Downarrow (\ell_1, \nu_2) \quad p' = p[p_{\text{var}} \mapsto p_{\text{var}}[\] \mapsto \ell_1], p_{\text{local}} \mapsto p_{\text{local}}[x \mapsto \tau]] \quad M = m \cdot M' \quad m' = m[\delta(x) + \text{fp}(m) \mapsto \nu_2]}{\langle \text{let } (\ : \text{level}_k, x : s) := e \text{ in } c, M, P, q \rangle_\nu \rightarrow \langle c; \text{unscope}(x), m' \cdot M', p' \cdot P', q + 1 \rangle_\nu} \\
\\
\text{S-UNPACK-TY} \\
\frac{P = p \cdot P' \quad \langle s, p' \rangle \Downarrow_{\text{type}} \tau \quad \langle e, m, p \rangle \Downarrow (\tau_1, \nu_2) \quad p' = p[p_{\text{var}} \mapsto p_{\text{var}}[\] \mapsto \tau_1], p_{\text{local}} \mapsto p_{\text{local}}[x \mapsto \tau]] \quad M = m \cdot M' \quad m' = m[\delta(x) + \text{fp}(m) \mapsto \nu_2]}{\langle \text{let } (\ : \text{type}_k, x : s) := e \text{ in } c, M, P, q \rangle_\nu \rightarrow \langle c; \text{unscope}(x), m' \cdot M', p' \cdot P', q + 1 \rangle_\nu} \\
\\
\text{S-EPILOGUE} \\
\frac{}{\langle \text{epilogue}, (\ell_1, |m_1|, \nu_1) \cdot (\ell_2, |m_2|, \nu_2) \cdot M, p \cdot P, q \rangle_\nu \rightarrow \langle \text{stop}, (\ell_2, |m_2|, \max(\nu_1, \nu_2) + 1) \cdot M, P, q + 1 \rangle_\nu} \\
\\
\text{S-MATCH} \\
\frac{\text{argmin}_{i=1, \dots, n}(\tau - p_i) = j \quad \langle \ , p \rangle \Downarrow_{\text{type}} \tau \quad \text{J}p_j \text{K}(p, \tau) = p'}{\langle \text{match } (p_i \Rightarrow c_i)_{i=1, \dots, n}, M, p \cdot P, q \rangle_\nu \rightarrow \langle c_j, M, p' \cdot P, q + 1 \rangle_\nu} \\
\\
\text{S-CALL} \\
\frac{\text{F}(f) = \langle \ell_1, \dots, \ell_n \rangle \langle \tau_1, \dots, \tau_m \rangle (x_1 : s'_1, \dots, x_r : s'_r) = c \quad \langle k_i, P \rangle \Downarrow_{\text{lab}} \ell_i \quad \langle s_i, P \rangle \Downarrow_{\text{type}} \tau_i \quad \langle e_i, M, P \rangle \Downarrow v_i \quad \langle s'_i, P' \rangle \Downarrow_{\text{type}} \tau'_i \quad M = m \cdot M' \quad m' = (\ell', |m'|, \nu) \quad P = (p_{\text{var}}, p_{\text{args}}, p_{\text{local}}) \cdot P' \quad p' = (p'_{\text{var}}, p'_{\text{arg}}, p'_{\text{local}}) \quad p'_{\text{var}} = \{ i \mapsto \ell_i \mid i = 1, \dots, n \} \cup \{ i \mapsto \tau_i \mid i = 1, \dots, m \} \quad p'_{\text{arg}} = \{ x_i \mapsto \tau'_i \mid i = 1, \dots, r \} \quad p'_{\text{local}} = \{ x \mapsto \perp \mid x \in c \} \quad \ell' = \{ \delta(x_i) + \text{sp}(m) \mid i = 1, \dots, r \} \cup \{ \text{sp}(m) \} \cup \{ \delta(z) + \text{sp}(m) \mid z \in c \} \quad |m'| = \{ \text{sp}(m) \mapsto (\text{cod}(p_{\text{arg}}), (\text{cod}(p_{\text{local}}), \text{fp}(m)_\nu)) \} \cup \{ \delta(x_i) + \text{sp}(m) \mapsto v_i \mid i = 1, \dots, r \}}{\langle f \langle k_1, \dots, k_n \rangle \langle s_1, \dots, s_m \rangle (e_1, \dots, e_r), M, P, q \rangle_\nu \rightarrow \langle c; \text{epilogue}, m' \cdot M, p' \cdot P, q + 1 \rangle_{\nu+1}} \\
\\
\text{S-SEQ-CONT} \\
\frac{c'_1 \neq \text{stop}}{\langle c_1, m, P, q \rangle_\nu \rightarrow \langle c'_1, m', P', q' \rangle_\nu} \\
\frac{}{\langle c_1; c_2, m, P, q \rangle_\nu \rightarrow \langle c'_1; c_2, m', P', q' \rangle_\nu} \\
\\
\text{S-SEQ-STOP} \\
\frac{}{\langle c_1, m, P, q \rangle_\nu \rightarrow \langle \text{stop}, m', P', q' \rangle_\nu} \\
\frac{}{\langle c_1; c_2, m, P, q \rangle_\nu \rightarrow \langle c_2, m', P', q' \rangle_\nu} \\
\\
\text{S-INST} \\
\frac{}{\langle c, m, P, h, q \rangle_\nu \rightarrow \langle c', m', P', h', q' \rangle_\nu} \\
\frac{}{\langle c, m, P, h, q \rangle_\nu \rightarrow \langle c', m', P', h', q' \rangle_\nu}
\end{array}$$