

Language Support for Secure Software Development with Enclaves

Aditya Oak
TU Darmstadt

Amir M. Ahmadian
KTH Royal Institute of Technology

Musard Balliu
KTH Royal Institute of Technology

Guido Salvaneschi
University of St.Gallen

Abstract—Confidential computing is a promising technology for securing code and data-in-use on untrusted host machines, e.g., the cloud. Many hardware vendors offer different implementations of Trusted Execution Environments (TEEs). A TEE is a hardware protected execution environment that allows performing confidential computations over sensitive data on untrusted hosts. Despite the appeal of achieving strong security guarantees against low-level attackers, two challenges hinder the adoption of TEEs. First, developing software in high-level managed languages, e.g., Java or Scala, taking advantage of existing TEEs is complex and error-prone. Second, partitioning an application into components that run inside and outside a TEE may break application-level security policies, resulting in an insecure application when facing a realistic attacker.

In this work, we study both these challenges. We present J_E , a programming model that seamlessly integrates a TEE, abstracting away low-level programming details such as initialization and loading of data into the TEE. J_E only requires developers to add annotations to their programs to enable the execution within the TEE. Drawing on information flow control, we develop a security type system that checks confidentiality and integrity policies against realistic attackers with full control over the code running outside the TEE. We formalize the security type system for the J_E core and prove it sound for a semantic characterization of security. We implement J_E and the security type system, enable Java programs to run on Intel SGX with strong security guarantees. We evaluate our approach on use cases from the literature, including a battleship game, a secure event processing system, and a popular processing framework for big data, showing that we correctly handle complex cases of partitioning, information flow, declassification, and trust.

Index Terms—Information Flow Control, Trusted Execution Environment, Robust Declassification, Security Type System

I. INTRODUCTION

Confidential computing includes recent technologies to protect data-in-use through isolating computations to a hardware-based Trusted Execution Environment (TEE). TEEs provide hardware-supported enclaves to protect data and code from the system software. Over the past few years, an array of TEE designs has been developed, including Intel’s Software Guard Extensions (SGX) [1], [2], ARM TrustZone [3], MultiZone [4] and others [5], [6], [7], [8], [9]. Using TEEs, data can be loaded securely in plain text and processed at native speed within an enclave even on a third-party machine. SGX is a TEE implementation from Intel which has been successfully used in a number of industry products [10], [11].

The issue with confidential computing Supporting confidential computing in a way that is both accessible for developers and technically secure is still an open problem.

First, **seamless integration** of enclave programming into software applications remains challenging. For example, Intel provides a C/C++ interface to the SGX enclave but no direct support is available for managed languages. As managed languages like Java and Scala are extensively used for developing distributed applications, developers need to either interface their programs with the C++ code executing in the enclave (e.g., using the Java Native Interface [12]) or compile their programs to native code (e.g., using Java Native [13]) relinquishing many advantages of managed environments.

A second aspect concerns security with **realistic attackers**. Standard security analysis of code protects against passive attackers, as common for untrusted/buggy code executing in a single trusted host [14]. Yet, with enclaves, programs run in a trusted environment within an untrusted host: the attacker can control the untrusted environment to cause additional leaks of sensitive information through the interface between the two environments. An active attacker may force the enclave program to violate the security policy by compromising the integrity of inputs at the interface or by controlling the execution order of interface components, e.g., to trigger execution paths and side effects that were not possible in the original program. Current research adopts Information Flow Control (IFC) to ensure that the code within an enclave does not leak sensitive information to the non-enclave environment [15], [16], [17]. This research, however, either takes a limited view of a passive attacker that only observes the data leaving the enclave, or it incorporates the effects of an active attacker into the execution semantics and the security condition, thus requiring additional verification effort to secure enclave programs.

These challenges lead us to the following key research questions addressed by the paper: (a) How to enable seamless integration of enclaves and managed languages like Java? (b) What is the right security model for realistic enclave attacks and how to statically verify the security of enclave programs with respect to these attacks? (c) How to harden state-of-the-art IFC tools to verify security in the TEE context? (d) How to demonstrate feasibility via realistic use cases?

Accessible and secure confidential computing To address the questions above, we present J_E , a programming model that offers language-level integration of enclave technology. We leverage IFC to secure applications running within TEEs and propose a security condition and an enforcement mechanism

targeting realistic attackers in the context of TEEs.

To support **seamless integration** of enclave programming, J_E defines a programming model for developers of enclave applications as a Java extension based on three annotations and two operators. Programmers can execute computations securely just by adding the `@Enclave` annotation to classes to be placed inside the enclave and the `@Gateway` annotation to the methods accessible from the non-enclave environment. Also, programmers can label secret data with `@Secret` annotations and control the release of secret information to the non-enclave environment via the `declassify` operator as well as the influence of untrusted information from the non-enclave environment via `endorse` operator. J_E builds on the Java information flow (Jif) security-typed language [18] to statically check confidentiality and integrity policies of code running within an enclave. A J_E program is automatically translated into an equivalent Jif program.

To provide security against **realistic attackers**, our key observation is that a program is secure in the presence of enclave active attackers iff the program does not leak additional sensitive information as compared to executing the program in the presence of a passive attacker. Inspired by the line of works on IFC for distributed applications [19], [20], we propose *robustness*, a semantic characterization of security for enclaves for two realistic active attackers. Robustness captures the interplay between the integrity of untrusted data coming from the non-enclave environment and the confidentiality of secret data within the enclave, ensuring that an active attacker does not learn additional information. We enforce robustness with a security type system for a core of J_E to check that the active attackers' control over inputs at the enclave interface and over the execution order of interface components does not enable them to learn more sensitive information than a passive attacker who merely observes the outputs at the interface. Importantly, our security type system can be used to check robustness for the partitioned programs leveraging existing verification efforts for the original program before partitioning. In contrast to existing type systems for robust declassification [19], [20], [21], our security type system is flow sensitive [22] which poses additional challenges with declassification policies and enclave attackers.

To validate the design of J_E , we verify several case studies from the literature. To validate our approach to security for enclaves, we prove our security type system sound with respect to robustness, showing that it accepts only secure programs. The implementation of J_E and the case studies discussed in the paper are publicly available ¹. In summary, this paper makes the following contributions:

- We present J_E , a programming model that seamlessly supports computations inside an enclave. In this model, programmers use annotations to identify the computations to be executed inside the enclave.
- We provide a core calculus for J_E and propose a semantic security model to capture the essence of information flows

in the presence of active enclave attackers.

- We enforce robustness statically via a security type system, which we prove sound for programs implemented in J_E .
- We evaluate the applicability of J_E using different case studies.

The paper is structured as follows. Section II introduces enclave technology and the different approaches for enclave software development. Section III overviews the J_E design. In Section IV, we present a security framework for J_E . Section VI describes the implementation. The evaluation is in Section VII. Section VIII presents related work and Section IX concludes.

II. TRUSTED EXECUTION ENVIRONMENTS IN A NUTSHELL

In this section, we introduce TEEs and we refer to Intel SGX, as a concrete implementation. TEEs make use of dedicated processor instructions and modified memory access mechanisms to enable private computations. The principle behind a TEE is to provide applications with memory isolation. With the help of the dedicated instructions, applications can create private memory regions known as *enclaves*. An enclave is essentially a reserved area in the system memory protected by the CPU. Data inside an enclave is protected from high privilege software such as OS, VMM and BIOS. This design leads to two distinct execution environments – the enclave environment (trusted environment) and the non-enclave environment (regular system memory).

Intel Software Guard Extensions (SGX) [1], [2] is an enclave implementation from Intel introduced with the sixth generation Intel core and Intel Xeon E3 v6 server processors.

A. The SGX Mechanism

In SGX enabled systems, the BIOS reserves a contiguous part of the DRAM as Processor Reserved Memory (PRM). Currently, the size of the PRM is limited to 128 MB. Out of this PRM, a memory region of about 90 MB is used as Enclave Page Cache (EPC) which stores the enclave code and data. The content of the enclave memory is encrypted using the Memory Encryption Engine (MEE) and is decrypted only when inside the CPU. This solution protects the enclave data from an attacker having physical access to DRAM. An application creates a new enclave using the `ECREATE` instruction. The enclave memory can only be accessed from inside the enclave. The CPU rejects any attempt to access the enclave memory from the non trusted environment. SGX provides a Remote Attestation (RA) service to verify the integrity of the data loaded inside the enclave. We refer to Costan and Devadas [23] for a detailed description of SGX.

B. Software Development with SGX

There are two main development methodologies used for programming with SGX enclaves. One is to use the C/C++ interface provided by Intel along with the Microsoft Visual Studio IDE. In this approach, a programmer writes their application in C/C++, programs to be executed inside and outside the enclave are written in two separate projects and

¹<https://prg-grp.github.io/jc-lang>

together they form a complete application. The application is compiled using the Intel compiler integrated into the IDE. The benefit of this approach is that the developer can keep the trusted code base (TCB) minimal. Unfortunately, no such support is available for applications written in managed languages. The other approach is to use systems based on library OSes [24], [25], [26], [27], [28]. They allow running unmodified applications inside the enclave. The application and the library OS are compiled together into an image file and the host OS runs the image inside the enclave.

C. Attacker models with enclaves

In this work, we consider two attacker models: a **passive** attacker and (two variants of) an **active** attacker. A passive attacker cannot modify the execution state of the program (inside or outside the enclave) and can only observe the entire execution state of the non-enclave environment. The rationale is that a developer can partition the original program into an enclave component and a non-enclave component, assuming that the partitioned program will be as secure as the original program, so long as the enclave component contains the secret data as well as any code that accesses these data.

The active attacker model extends the threat model considered in [15], [16] and considers that an active attacker can arbitrarily alter the state of the non-enclave environment and hence it can modify the data exchanged between the two environments. An active attacker can influence the execution of the enclave component either by compromising the integrity of inputs at the interface or by controlling the execution order of the interface components. Both cases may lead to triggering execution paths and side effects that were not possible in the original program, thus enabling learning more sensitive information than a passive attacker.

For both attacker models, we only assume that the enclave hardware and software, except for the J_E source program, are trusted. The J_E uses static analysis to check that a passive attacker learns no more information about secret data than allowed by the original program, and to check that an active attacker learns no more information than a passive attacker. Denial-of-service attacks, arbitrary destruction of enclaves by the host OS, hardware side-channel attacks, and power analysis attacks are out of the scope of this paper. We also ignore covert channels including termination and timing. While termination channels have well-understood information leakage bounds [29], orthogonal techniques such as constant-time programming [30] and predictive mitigation mechanisms [31] can help mitigating some of these concerns.

III. J_E DESIGN

We propose a language design that supports computations inside an SGX enclave. Our approach, J_E , relieves the programmer from dealing with the lower-level details of SGX, such as enclave creation, initialization, and destruction. Moreover, J_E supports security annotations to define application-level policies that are subsequently interpreted as information flow policies and are verified via a security type system.

Listing 1: Language constructs

```

1 @Enclave
2 class Encrypter {
3
4     @Secret static String key;
5
6     @Gateway
7     public static String encrypt(String plaintext) {
8         String cipher = encode(plaintext, key);
9         return declassify(cipher);
10 } }
```

Listing 2: The Endorse operator

```

1 @Enclave
2 class Encrypter {
3
4     @Secret static String key;
5
6     @Gateway
7     public static String encrypt(String plaintext) {
8         String cipher = "";
9         String plaintextE = endorse(plaintext);
10        if (plaintextE.length() == 8) {
11            String var1 = encode(plaintextE, key);
12            cipher = declassify(var1);
13        }
14        return cipher;
15 } }
```

A. J_E Annotations

J_E extends the Java language with annotations to specify the parts of the program that run inside SGX as well as the sensitive data. We demonstrate these features using a routine to securely encrypt data (Listing 1). Encryption is based on a key that should be kept secret; hence the whole encryption operation should run within the enclave. J_E supports the following annotations to define security policies.

a) *Class annotation @Enclave*: A programmer can annotate any class with the @Enclave annotation. Hereafter, we refer to such classes as enclave classes. The code and the data that belong to enclave classes are placed inside the enclave. To ensure that data and computations related to encryption take place within the enclave, the class Encrypter in Listing 1 is annotated with the @Enclave annotation.

b) *Field annotation @Secret*: A programmer can annotate fields of an enclave class with the @Secret annotation – we refer to such fields as *secret* fields. The annotation denotes that the field holds sensitive information. Any program construct influenced by a secret field will also be considered as secret in order to prevent flows of any sensitive data from the enclave to the non-enclave environment. In Listing 1, the key field is used as a key for encryption. The key field is annotated with the @Secret annotation to denote that its value must not be leaked to the non-enclave environment.

c) *Method annotation @Gateway*: Static methods defined inside an enclave class can be annotated with the @Gateway annotation. We refer to these methods as *gateway* methods. These methods act as an interface between the enclave and the non-enclave environment: the execution switches from the non-enclave environment to the enclave when a gateway method is called from the non-enclave environment. In Listing 1, the encrypt method is annotated with the @Gateway annotation to ensure that encryption is performed within the enclave. The encrypt method encrypts the plaintext argument with the

secret key and returns the corresponding cipher text to the external environment. The return value of a gateway method should not be influenced by secret information.

d) Operator declassify: The declassify unary operator downgrades a secret value into a public value. In Listing 1, the value in cipher (Line 8) is influenced by the secret field key. Hence, cipher is considered sensitive and cannot be returned to the non-enclave environment. However, the encryption can be considered secure as an observer cannot learn useful information by observing only the ciphertext. The declassify operator is used to control the release of sensitive information by explicitly declassifying a secret value provided it can be trusted (see the following paragraph).

e) Operator endorse: The endorse operator endorses an untrusted value into a trusted one. A declassification can get triggered based on the specific value of a variable. As a result, a malicious user can influence the value of a variable to trigger the declassification. J_E ensures that only trusted values are declassified and application of the declassify operator does not depend on untrusted values. By default, the arguments to the gateway methods are considered untrusted as they are received from the non-enclave environment. In Listing 2, the gateway method encrypt accepts the plaintext argument from the non-enclave environment. We explicitly trust the identifier plaintext (Line 9). As a result, plaintextE is a trusted version of plaintext and the declassification no longer depends on an untrusted variable. The program in Listing 2 is valid because the var1 identifier is declassified (Line 12) and it is trusted. For this reason, the application of the declassify operator (Line 12) does not depend on an untrusted value. Also, the code in Listing 1 is invalid as the identifier cipher being declassified (Line 8) is untrusted.

B. Compilation phases

The first step of the compilation process involves automatically translating a J_E program into an equivalent Jif program. The generated program is verified using the Jif compiler. If the compilation succeeds, the code for data exchange between the enclave and non-enclave environments along with the initialization code is added. Finally, two JARs corresponding to each environment are generated. A detailed description of the compilation steps is provided in Appendix A

C. Execution Model

We execute J_E programs in two separate JVMs, one running in the external environment and one within the enclave. By default, a J_E program starts its execution inside the non-enclave environment. When a program running in the non-enclave environment calls a gateway method, the corresponding parameters are passed to the enclave-environment and the called gateway method executes inside the enclave.

When a gateway method is called from the non-enclave environment, the corresponding arguments are serialized and copied (deep copy) into the enclave. Hence, this semantics results in a distinct copy of each argument inside the enclave. To avoid inconsistent copies of an object in the two

```

CL ::= class C { $\bar{f}$ ,  $\bar{m}$ } | classEnclave C { $\bar{f}$ ,  $\bar{m}$ }
m   ::= method $_{\phi}(\bar{p})\{S; \text{return}(e)\}$ 
S   ::= skip | if e then S1 else S2 | S1; S2 | while e do S
      x := declassify(e) | x := eG | C.f := eG
e   ::= x | v | e.f | e1  $\oplus$  e2
eG ::= e | e.m( $\bar{p}$ )
v   ::= n | C | unit

```

Figure 1: Syntax of J_E

environments, once a gateway method is called, we prevent any further usage of all its arguments in the non-enclave environment.

IV. SECURITY FRAMEWORK

This section presents a security framework for reasoning about the confidentiality and integrity properties of enclave programs. We formalize a core of J_E by defining the syntax and semantics (Section IV-A and IV-B). Drawing on IFC, we present a security model (Section IV-C) with the following ingredients: a security policy specifying the parts of the program that contain secret/trusted information, and the parts that contain public/untrusted information (Section IV-C1); an attacker model specifying the capabilities of active and passive attackers (Section IV-C2); a security condition capturing a semantic characterization of security with respect to the program semantics, the attacker model, and the security policy. Finally, we present a security type system that enforces the security condition in a sound manner (Section IV-D).

A. J_E Syntax

J_E is an imperative language extended with constructs for static classes, methods and fields as shown in Figure 1. CL represents the list of class definitions in a program. We write \bar{l} for a finite list of elements l_1, \dots, l_n . We distinguish two types of classes: *normal* classes (*class*) which are executed in the non-enclave environment and *enclave* classes (*classEnclave*) which are executed in the enclave environment. Each class is defined with a list of fields \bar{f} and methods \bar{m} . We assume classes $C \in \text{Class}$, methods $m \in \text{Method}$, and fields $f \in \text{Field}$ are uniquely identified.

We define a method by the list of formal parameters \bar{p} and method body S , and use the annotation $\phi \in \{G, NG\}$ to indicate whether (G) or not (NG) a method is a gateway. The method body S is a sequence of commands that is executed when the method is called, followed by $\text{return}(e)$ for the return value of the method. J_E distinguishes two types of expressions, side-effect free expressions e including variables $x \in \text{Vars}$, values $v \in \text{Val}$, fields accesses $e.f$, and binary operations \oplus , and side-effectful expressions e_G which extend e with method calls $e.m(\bar{p})$. Commands S include standard features such as assignment to class fields and variables, conditionals, loops, and sequencing. Command $x := \text{declassify}(e)$ is semantically equivalent to an assignment and it is used for declassification. Command

$x := e_G$, assigns the result of evaluating e_G to the variable x . If e_G is a method call, e.g., the assignment $x := C.m(\bar{p})$ denotes a call to method m of class C with actual parameters \bar{p} , and stores the return value in variable x . Similarly $C.f := e_G$ assigns the result of evaluating e_G to the field f of class C . Values are integers $n \in \mathbb{Z}$, class identifiers or unit.

B. J_E Operational Semantics

We define the (big-step) operational semantics of J_E . A configuration $\langle S, \mathcal{M}, \mathcal{H} \rangle$ consists of a command S , a memory $\mathcal{M} = \text{Vars} \rightarrow \text{Val}$ mapping variables to values, and a heap $\mathcal{H} = \text{Class} \times \text{Field} \rightarrow \text{Val}$ mapping class and field identifiers to values. A state $\sigma = \langle \mathcal{M}, \mathcal{H} \rangle$ is a pair of a memory \mathcal{M} and a heap \mathcal{H} . An observation trace t is a (possibly empty) sequence of events $\beta \in \text{Val}$, and $t_1.t_2$ denotes trace concatenation.

We use judgments of the form $\alpha \vdash_\delta \langle S, \mathcal{M}, \mathcal{H} \rangle \Downarrow_t \mathcal{M}', \mathcal{H}'$ to denote that a configuration $\langle S, \mathcal{M}, \mathcal{H} \rangle$ evaluates to memory \mathcal{M}' and heap \mathcal{H}' in *execution mode* $\alpha \in \{N, E\}$ with *static mode* $\delta = \text{Class} \cup \text{Var} \rightarrow \{N, E\}$, and produces an observation trace t . We write $\alpha \vdash_\delta \langle S, \mathcal{M}, \mathcal{H} \rangle \Downarrow \mathcal{M}', \mathcal{H}'$ if the observation trace is empty and $\alpha \vdash_\delta \langle S, \mathcal{M}, \mathcal{H} \rangle \Downarrow_t _$ to ignore the final state. We use modes to distinguish between the enclave environment (E) and the non-enclave environment (N). An execution mode α indicates the current execution environment of a configuration, while a static mode δ associates a class identifier or a variable with the execution environment it was assigned to, statically. Abusing notation, we write $\alpha \vdash_\delta \langle S; \text{return}(e), \mathcal{M}, \mathcal{H} \rangle \Downarrow \mathcal{M}, \mathcal{H} \triangleright v$ for commands that yield a value v , via the *return* command. To simplify the presentation, we assume that the sets of variable, field, and class identifiers of the enclave environment and non-enclave environment are disjoint. An initial configuration starts in the non-enclave execution mode and it executes a sequence of commands S_0 from an initial state $\langle \mathcal{M}_0, \mathcal{H}_0 \rangle$, i.e., $N \vdash_\delta \langle S_0, \mathcal{M}_0, \mathcal{H}_0 \rangle$.

The semantics of expressions, Figure 2, is mostly standard. We use judgments of the form $\alpha \vdash_\delta \langle e_G, \mathcal{M}, \mathcal{H} \rangle \Downarrow_\beta \langle v, \mathcal{M}', \mathcal{H}' \rangle$ to denote that an expression e_G evaluates to value v , memory \mathcal{M}' , and heap \mathcal{H}' in state $\langle \mathcal{M}, \mathcal{H} \rangle$, execution mode α , static mode δ , and it emits the event β . Rule METHOD interacts with the execution semantics of commands to execute the method body, hence it can potentially alter the execution state. We use the auxiliary functions $\text{getMethod} : \text{Class} \times \text{Method} \rightarrow m$ to extract a method definition, and $\text{fields} : \text{Class} \rightarrow \wp(\text{Field})$ extract the set of fields of class C . We write $\mathcal{M}[x \mapsto v]$ to denote a memory \mathcal{M} with variable x assigned the value v . Similarly, $\mathcal{H}[(C, f) \mapsto v]$ denotes a heap \mathcal{H} with field f of class C assigned the value v . The full semantics of expressions is reported in Figure 8.

Rule FIELD ACCESS evaluates expression e to a class identifier C , ensures that the execution mode and static mode are the same, and performs a lookup of its field f in the heap \mathcal{H} . It then returns the resulting value v as well as the memory and heap, which are unchanged.

Rule METHOD evaluates expression e to a class identifier C and uses the getMethod function to get the definition of

its method m . Next, it checks that either the execution mode and static mode are the same ($\alpha = \delta(C)$) or that the execution mode is non-enclave (N), the static mode is enclave (E), and m is a gateway method ($\alpha = N \wedge \delta(C) = E \wedge \phi = G$). This condition ensures that only methods from the classes that have the same static mode as the current execution mode can be called, with the exception that the non-enclave execution mode can call the gateway methods of enclave classes. For the latter, our rules enforce the copy semantics in a call-by-value fashion. The rule substitutes the actual parameters for the formal parameters and executes the method body, returning a value. Finally, if the method is a gateway call, the emitted event is the return value v , otherwise the empty event ϵ .

Figure 3 provides an excerpt of evaluation rules for J_E commands. The full set of rules can be found in Figure 10. Rule STORE updates the field in the heap with the associated value and checks that $\alpha = \delta(C)$ to ensure that execution mode, and the static mode of class C are the same. The event that STORE emits is equal to the event emitted during the evaluation of expression e_G . Rule RETURN evaluates the expression in the current state and returns the associated value.

C. Security Model

1) *Security Policy*: J_E adopts security labels to specify application-level policies as information flow policies. A security label ℓ is a tuple $\langle \ell_C, \ell_I \rangle$ of a confidentiality label ℓ_C and an integrity label ℓ_I . We use two levels, *Public* (\mathbb{P}) and *Secret* (\mathbb{S}), for confidentiality, and two levels, *Trusted* (\mathbb{T}) and *Untrusted* (\mathbb{U}), for integrity. Intuitively, for confidentiality, data from *Secret* sources should not flow to *Public* sinks unless it is explicitly declassified by the developer, and, for integrity, data from *Untrusted* sources should not flow to *Trusted* sinks unless it is explicitly endorsed by the developer. These requirements are reflected by the ordering relation between security labels, namely $\mathbb{P} \sqsubseteq \mathbb{S}$ and $\mathbb{T} \sqsubseteq \mathbb{U}$. The product lattice \mathcal{L} lifts the constraints to an ordering relation over label pairs such that $\ell_1 \sqsubseteq \ell_2$ iff $\ell_{1C} \sqsubseteq \ell_{2C}$ and $\ell_{1I} \sqsubseteq \ell_{2I}$. The lattice \mathcal{L} defines the *join* $\ell_1 \sqcup \ell_2$ and *meet* $\ell_1 \sqcap \ell_2$ operators to compute *least upper bound* and *greatest lower bound* of two labels, respectively. A security policy is then defined by an assignment of security labels to variables, classes and fields of an J_E program.

2) *Attacker Model*: We use the security labels to define the view of the memory and heap from the attacker's perspective. In our setting, the attacker is at level $\langle \mathbb{P}, \mathbb{U} \rangle$ and it can observe all of initial program state (i.e., the initial memory and the initial heap) that is labeled as public \mathbb{P} . This is because an attacker has full control of the non-enclave state and it is allowed to learn any public data of the enclave state. We assume that each variable, class, and field has an associated security label from the lattice \mathcal{L} as defined by a security mapping $\Gamma : (\text{Var} \cup \text{Class} \cup \text{Field}) \rightarrow \mathcal{L}$. We then define indistinguishability over pairs of program states for a security mapping Γ and an attacker at security level $A = \langle \mathbb{P}, \mathbb{U} \rangle$.

Definition 1 (State Indistinguishability). *Two memories \mathcal{M}_1 and \mathcal{M}_2 are indistinguishable for the attacker A (written*

$$\begin{array}{c}
\text{FIELD ACCESS} \frac{\alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle C, \mathcal{M}, \mathcal{H} \rangle \quad \alpha = \delta(C) \quad v = \mathcal{H}(C, f)}{\alpha \vdash_{\delta} \langle e.f, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle v, \mathcal{M}, \mathcal{H} \rangle} \\
\\
\text{METHOD} \frac{\begin{array}{l} \alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle C, \mathcal{M}, \mathcal{H} \rangle \quad \text{method}_{\phi}(\bar{p})\{S; \text{return}(e)\} = \text{getMethod}(C, m) \\ \left((\alpha = \delta(C)) \vee (\alpha = N \wedge \delta(C) = E \wedge \phi = G) \right) \quad \mathcal{M}^* = \mathcal{M}[p_i \mapsto \sigma(q_i)] \quad i = 1, \dots, |\bar{p}| \quad \delta(C) \vdash_{\delta} \langle S; \text{return}(e), \mathcal{M}^*, \mathcal{H} \rangle \Downarrow \mathcal{M}^{*'}, \mathcal{H}' \triangleright v \\ \mathcal{M}' = \mathcal{M}^{*'} \setminus [p_i] \quad i = 1, \dots, |\bar{p}| \quad (\alpha = \delta(C) \Rightarrow \beta = \epsilon) \quad (\alpha = N \wedge \delta(C) = E \wedge \phi = G \Rightarrow \beta = v) \end{array}}{\alpha \vdash_{\delta} \langle e.m(\bar{q}), \mathcal{M}, \mathcal{H} \rangle \Downarrow_{\beta} \langle v, \mathcal{M}', \mathcal{H}' \rangle}
\end{array}$$

Figure 2: Excerpt of J_E expression semantics

$$\begin{array}{c}
\text{STORE} \frac{\alpha \vdash_{\delta} \langle e_G, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{\beta} \langle v, \mathcal{M}', \mathcal{H}' \rangle \quad \alpha = \delta(C) \quad \mathcal{H}'' = \mathcal{H}'[(C, f) \mapsto v]}{\alpha \vdash_{\delta} \langle C.f := e_G, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{\beta} \mathcal{M}', \mathcal{H}''} \quad \text{RETURN} \frac{\alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle v, \mathcal{M}, \mathcal{H} \rangle}{\alpha \vdash_{\delta} \langle \text{return}(e), \mathcal{M}, \mathcal{H} \rangle \Downarrow \mathcal{M}, \mathcal{H} \triangleright v}
\end{array}$$

Figure 3: Excerpt of J_E command semantics

$\mathcal{M}_1 =_A \mathcal{M}_2$ iff $\forall x. \delta(x) = N, \mathcal{M}_1(x) = \mathcal{M}_2(x)$, and $\forall x. \delta(x) = E$ such that $\Gamma(x) = \langle \mathbb{P}, - \rangle, \mathcal{M}_1(x) = \mathcal{M}_2(x)$.

Two heaps \mathcal{H}_1 and \mathcal{H}_2 are indistinguishable for the attacker A (written $\mathcal{H}_1 =_A \mathcal{H}_2$) iff $\forall C. \delta(C) = N, \forall f \in \text{fields}(C), \mathcal{H}_1(C, f) = \mathcal{H}_2(C, f)$ and $\forall C. \delta(C) = E, \forall f \in \text{fields}(C)$ such that $\Gamma(C.f) = \langle \mathbb{P}, - \rangle, \mathcal{H}_1(C, f) = \mathcal{H}_2(C, f)$.

Two program states $\sigma_1 = \langle \mathcal{M}_1, \mathcal{H}_1 \rangle$ and $\sigma_2 = \langle \mathcal{M}_2, \mathcal{H}_2 \rangle$ are indistinguishable for the attacker A (written $\sigma_1 =_A \sigma_2$) iff $\mathcal{M}_1 =_A \mathcal{M}_2$ and $\mathcal{H}_1 =_A \mathcal{H}_2$.

Intuitively, two indistinguishable memories assign the same value to the public variables inside and outside of enclave. Similarly, two indistinguishable heaps assign the same value to all of the fields of *non-enclave classes* and the public fields of *enclave classes*.

As secret values are stored inside the enclave, the only way for an attacker to learn secret information is by observing the return values of gateway methods. In fact, the attacker observations are captured by our semantics in Figure 3 via traces t and events β . Therefore, we can define indistinguishability for program executions by requiring that any two indistinguishable initial states produce the same observation traces. This implies that an attacker cannot discriminate the two initial states, thus it cannot learn secret information from the enclave environment.

Definition 2 (Execution Indistinguishability). *Let S be a J_E program and σ_1 and σ_2 be two initial states such that $\sigma_1 =_A \sigma_2$. Two executions are indistinguishable (written $N \vdash_{\delta} \langle S, \sigma_1 \rangle \approx_A N \vdash_{\delta} \langle S, \sigma_2 \rangle$) if $N \vdash_{\delta} \langle S, \sigma_1 \rangle \Downarrow_{t_1} -$ and $N \vdash_{\delta} \langle S, \sigma_2 \rangle \Downarrow_{t_2} -$ then $t_1 = t_2$.*

In line with existing works [29], execution indistinguishability ignores information leaks that are due to program (non) termination. The definition ensures security w.r.t. a passive non-enclave attacker that observes only the results of gateway method calls. In particular, it rejects all programs that leak secret information to the non-enclave environment. As such condition can be restrictive for most practical scenarios (see Section III for examples), developers resort to various forms of declassification operations to release secret information in

a controlled manner. We refer to prior works for an overview of various dimensions of declassification [32]. In our setting, declassification can be dangerous as it can be abused by an active attacker to release secret information in a way that it was not intended by the developer [20]. Next, we define the attackers that are relevant in the context of enclave programs.

3) *Passive and Active Attackers*: We consider three types of attackers: the passive attacker (PA), the *havoc* active attacker (HAA) limited to modifying parameters passed to gateway calls, and the *havoc reordering* active attacker (HRAA) capable of controlling the order and frequency in which gateway methods are called. Both active attackers are reasonable in our context as the attacker fully controls the non-enclave environment. We use the PA attacker as a reference model to show that the HAA attacker and the HRAA attacker do not learn more secret information than the PA attacker.

The PA does not intervene in the execution of the program, it just observes observation traces (as in Figure 3) to learn secret information from the enclave. The active attackers can influence the execution. The HAA attacker can modify the non-enclave state and hence the parameters passed to gateway methods, thus modifying the behavior of the program executing inside the enclave. This may change the behavior of the program in a way that leaks information about the enclave secrets, e.g., by abusing declassification operations that were intended in the context of a PA attacker. We illustrate the issue via an example inspired by Askarov and Myers [21].

The program in Listing 3 contains a gateway method *foo* whose return value depends on the *time* parameter. The intention of the developer, who assumes a PA attacker, is to declassify *secretVal* only after the *releaseTime* has elapsed. This is implemented by comparing the *time* in which method call was issued with the predefined *releaseTime* (Line 10). Yet, an HAA attacker can arbitrarily change the value of *time* and control the release of information via declassification. This example motivates the need for a security condition that rejects scenarios where an active attacker learns more secret information than a passive attacker.

To model the active capability of the HAA attacker, we

Listing 3: HAA Attacker

```

1 @Enclave
2 class FooClass {
3
4     @Secret static int secretVal;
5     static int releaseTime = 2025
6
7     @Gateway
8     public static int foo(int time) {
9         int res = 0;
10        if (time >= releaseTime)
11            res = declassify(secretVal);
12        else
13            res = 0;
14
15        return res;
16    } }

```

Listing 4: HRAA Attacker

```

1 @Enclave
2 class FooClass {
3
4     @Secret static int secret1, secret2;
5     static boolean releaseTrigger = false;
6
7     @Gateway
8     public static void bar() {
9         releaseTrigger = true;
10    }
11
12    @Gateway
13    public static int foo() {
14        int res = 0;
15        if (releaseTrigger) {
16            releaseTrigger = !releaseTrigger;
17            res = declassify(secret1); }
18        else {
19            releaseTrigger = !releaseTrigger;
20            res = declassify(secret2); }
21
22        return res;
23    } }

```

Listing 5: Delayed Declassification

```

1 @Enclave
2 class FooClass {
3
4     @Secret static int secretVal;
5
6     @Gateway
7     public static int foo(int input) {
8         int x = declassify(secretVal);
9         int l = 0;
10        if (input > 0)
11            l = x;
12        else
13            l = 7
14
15        return l;
16    } }

```

introduce program holes $[\bullet]$ [20]. Holes represent program contexts where an HAA attacker can insert an untrusted code a to modify the program's state. Because the only way an attacker can affect the execution of enclave code is via gateways, we extend our program syntax by adding $[\bullet]; x := e.m(\bar{p})$ and $[\bullet]; C.f := e.m(\bar{p})$, and define the active attacks.

Definition 3 (J_E program with holes). *A program with holes $S[\vec{\bullet}]$ is defined by extending the syntax in Figure 1 as follows:*

$$S[\vec{\bullet}] ::= \dots \mid [\bullet]; x := e.m(\bar{p}) \mid [\bullet]; C.f := e.m(\bar{p})$$

where m is a gateway method.

An HAA attacker can execute any untrusted code a in a hole before method calls. However, this code can only contain variables, classes, and fields in the non-enclave environment which are by definition public and untrusted. We define the attacker's code as follows:

$$a ::= \text{skip} \mid a_1; a_2 \mid q := e \text{ (where } q \in \bar{p}) \quad (1)$$

While an HAA attacker can inject arbitrary untrusted code in the non-enclave environment, we argue that the definition above captures the most powerful attack strategies of HAA attacker.

Lemma 1. *The attack code definition presented in 1, captures the most powerful attack strategies available to an HAA attacker, who controls the non-enclave environment.*

Proof. The proof of this lemma is presented in Appendix C. \square

We write $S[\vec{a}]$ for a program under an HAA attack \vec{a} . In this setting, a passive attack can be modeled as $S[\vec{\text{skip}}]$.

An HRAA attacker controls both the code and data memory outside the enclave. Hence in addition to the HAA attack capability, an HRAA attacker can change the *order* and *frequency* of gateway method calls issued from the non-enclave environment. This can cause information leaks as the (attacker-controlled) order and frequency of gateway calls may influence

the values returned to the non-enclave environment. Listing 4 illustrates the problem. Consider the non-enclave program `FooClass.bar(); FooClass.foo()`. The intended order of issuing gateway methods is `bar();foo()`, hence this program is secure with respect to an HAA attacker. `secret1` is always going to be declassified, and the resulting trace depends on its value. However, an HRAA attacker that controls the code memory outside the enclave can change the order of gateway calls to `foo();bar()` and learn the declassified value of `secret2`.

A similar argument applies to calling gateway methods multiple times. For instance, if a gateway was intended to be called only once, calling it more than once might leak sensitive information. We revisit Listing 4 to illustrate the problem. Consider the non-enclave program `FooClass.foo()` revealing the value of `secret1`. An HRAA attacker can instead call `FooClass.foo(); FooClass.foo()` and learn the declassified value of `secret2`.

Since the HRAA attacker has full control over the non-enclave code and memory, and the secrets reside only inside the enclave, we model them as sequences of gateway calls.

Definition 4 (Program under HRAA control). *We define the program under HRAA control as a sequence of gateway calls:*

$$S'[\vec{\bullet}] ::= S'_1[\vec{\bullet}]; S'_2[\vec{\bullet}] \mid [\bullet]; x := C.m(\bar{p})$$

where m is a gateway method defined in $S[\vec{\bullet}]$.

In this model, the attack definition of 1 will remain unchanged, indicating that HRAA attacker subsumes the power of HAA attacker.

Lemma 2. *The attacker code defined in 1, captures the most powerful attack strategies available to the HRAA attacker.*

Proof. The proof is similar to Lemma 1 and is presented in Appendix C. \square

4) Security Condition: In our setting, programs use declassification to release secret information in a controlled manner. Intuitively, a program is secure if an active attacker cannot learn more secret information than a passive attacker.

Drawing on the idea of robust declassification [20], we present robustness, a security condition that formalizes this intuition.

Definition 5 (Robustness under HAA). *Program $S[\vec{\bullet}]$ is robust w.r.t an HAA attacker A if for all $\sigma_1, \sigma_2, \vec{a}_1, \vec{a}_2$*

$$\begin{aligned} N \vdash_\delta \langle S[\vec{a}_1], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}_1], \sigma_2 \rangle \Rightarrow \\ N \vdash_\delta \langle S[\vec{a}_2], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}_2], \sigma_2 \rangle \end{aligned}$$

Robustness holds whenever for an attack vector \vec{a} and two indistinguishable initial states σ_1, σ_2 , if the program $S[\vec{a}_1]$ satisfies execution indistinguishability (definition 2), then for another attack vector \vec{a}_2 , the program $S[\vec{a}_2]$ also satisfies execution indistinguishability. In other words, the attacker observations of $S[\vec{a}_2]$'s observation traces do not reveal any secrets apart from what the attacker already knows by the observation traces of the program $S[\vec{a}_1]$. The PA attacker is captured by executions of the program $S[\vec{skip}]$.

We extend robustness to capture HRAA attackers.

Definition 6 (Robustness under HRAA). *Program $S[\vec{\bullet}]$ is robust w.r.t an HRAA attacker A if for all $\sigma_1, \sigma_2, \vec{a}_1, \vec{a}_2$ and for all $S'[\vec{\bullet}]$:*

$$\begin{aligned} N \vdash_\delta \langle S[\vec{a}_1], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}_1], \sigma_2 \rangle \Rightarrow \\ N \vdash_\delta \langle S'[\vec{a}_2], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a}_2], \sigma_2 \rangle \end{aligned}$$

This definition ensures that the HRAA attacker does not learn more information by changing the order and frequency of gateway calls. Observe that if $S'[\vec{a}']$ triggers an execution that was not possible $S[\vec{a}]$, and that execution's return value depended on some declassified secret, the active power may enable the HRAA attacker to learn information that they would not have learned originally. In Listing 3, the program $S[\vec{a}] ::= [a_1]; FooClass.bar(); [a_2]; FooClass.foo()$ is not robust under HRAA because the program $S'[\vec{a}'] ::= [a'_1]; FooClass.foo(); [a'_2]; FooClass.bar()$ enables the HRAA attacker to reveal the declassified value of `secret2`. A similar argument applies to calling gateway methods that were defined in $S[\vec{\bullet}]$ but were never called. As expected, this definition is stronger than Definition 5.

5) *Delayed Declassification*: While our security condition extends the definition of robust declassification [20] to the setting of realistic enclave attackers, there are some key differences pertaining to traces and attacker observations. Robust declassification considers every assignment to public variables as immediately visible to the attacker, because it defines the observations as projection over the public part of the memory. This definition does not reflect the enclave attacker model, because TEEs encrypt the entire enclave memory, so even if a public variable is modified inside the enclave, it is not visible to the attacker unless it is written to the non-enclave memory. This motivates our use of a trace-based observation model as generated by the return values of gateway method calls.

This model poses additional challenges with handling of declassification policies. For example, the program in Listing 5

satisfies robust declassification because the attacker *input* neither affects the decision to declassify nor the declassified value itself. This is achieved by making the value in variable x observable to the attacker immediately. However, in our model, the declassified value in x will not be visible to the attacker until it is returned by the gateway method. Because the attacker controls the input and therefore the assignment in line 11, this results in controlling the decision to declassify the secret value. Hence, the program is not robust. In fact, for attack vectors $a ::= input := 0$ and $a' ::= input := 1$, our definition of robustness will correctly reject the program. We dub this concept *delayed declassification*. Observe that delayed declassification is orthogonal to the well-known *Where* and *What* dimensions of declassification [32] and it appears as result of the trace-based observation model.

D. Security Type System

This section presents a security type system to enforce robustness. In line with the enclave attacker model, the program from the non-enclave environment is public and untrusted, while program from the enclave environment is trusted and its input data can be labeled by the developer as either public or secret. We label secret fields with $\langle \mathbb{S}, \mathbb{T} \rangle$ as they contain sensitive information and are protected by the enclave. The security label of the arguments of a gateway method is defaulted to $\langle \mathbb{P}, \mathbb{U} \rangle$ and while the return parameter is labeled as $\langle \mathbb{P}, \mathbb{T} \rangle$ or $\langle \mathbb{P}, \mathbb{U} \rangle$.

The goal of our type system is to ensure robustness against the active attackers. To achieve this, the type system checks that the decision to declassify a secret, or to return it, is not influenced by untrusted non-enclave data, thus ensuring that the attacker cannot control the decision to release secret data. The security type system enforces the *Where* dimension of declassification for a PA attacker [32] and it has been proved sound with respect to the security condition of *gradual release* [33]. We extend the type system to additionally account for active attacks in our setting and prove it sound for robustness. This has the advantage of reusing existing verification efforts via security type systems, which assume a PA attacker, and verifying only on the effect of an active attacker whenever these programs have been verified for the PA attacker. In our setting, this may happen whenever a developer partitions an existing (secure) program to execute with enclaves.

Our security type system uses a typing environment $\Gamma : (Var \cup Class \cup Field) \rightarrow \mathcal{L}$ mapping variables, classes, and fields to security labels from the security lattice \mathcal{L} . We also use another environment $\Pi : (Var \cup Class \cup Field) \rightarrow \mathcal{B}$ that maps variables, fields, and classes to boolean flags. A flag $d \in \mathcal{B}$ can be true (T) or false (F), and is used to track the propagation of declassified values. Initially, every variable and field has the flag initialized to false and the flag is set to true whenever they store a value that may be affected by declassification. We define the ordering relation $F \sqsubseteq T$ on flags, which will be useful in the sub-typing and method rules.

The label of every variable, field, and class in our setting, is a tuple consisting of a security label and a flag (ℓ, d) . We use

indexes ℓ and d to access elements of this tuple. e.g., pc_ℓ will show the security label of pc . Methods are typed in isolation using type signatures of the form $(\Gamma^-, \Pi^- \xrightarrow{pc'} \Gamma^+, \Pi^+)_{(\ell, d)}$ which require an environments Γ^- and Π^- before the method is invoked, environments Γ^+ and Π^+ after the method invocation, the label of its return value (ℓ, d) , and the program counter label pc' capturing the lower bound on method's side-effects [16]. The typing judgments for expressions are of the form $\Gamma, \Pi \vdash_\delta e_G : \tau$, meaning that in mode δ , and environments Γ and Π , an expression e_G has the type τ . If the expression e_G is a method, then τ is a method type, otherwise τ is a type (ℓ, d) representing the security label of e . Similarly, the typing judgments for commands have the form $pc, \Gamma, \Pi \vdash_\delta S : \Gamma', \Pi'$ where Γ and Π are the environments before, and Γ' and Π' are the environments after the execution of command S , δ is the static, and pc is the program counter label used to prevent implicit flows. The judgment for the return command is of the form $pc, \Gamma, \Pi \vdash_\delta \text{return}(e) : \Gamma, \Pi \triangleright (\ell, d)$ to capture the security label ℓ and the flag d of the returned value.

Figure 4 depicts an excerpt of the typing rules for expressions. The full list of rules is reported in Figure 9. The only non-standard rule here is T-METHOD. Rule T-METHOD checks that method's body $S; \text{return}(e)$ is well-typed under $pc', \Gamma^-, \Gamma^+, \Pi^-$, and Π^+ and it returns the label (ℓ, d) of the result.

Figure 5 presents a few interesting rules for commands; we refer to Figure 11 for the full list. In our type system, variables are flow-sensitive (see rule T-ASSIGN in Figure 11), while fields are flow-insensitive, i.e., their security label is defined via annotations. In fact, rule T-STORE ensures that the join of the security labels of pc and expression e is at least as restrictive as the label of field $C.f$. Moreover, the rule ensures that if the security label of $\ell_1 \sqcup \ell_2 \sqcup pc_\ell$ is secret $\langle \mathbb{S}, - \rangle$, the store is defined in the enclave, and it ensures that if the flag of e is true, this command can only be executed in a trusted context. This is to prevent untrusted input from controlling the propagation and release of declassified values. Finally, even though $C.f$'s security labels are flow-insensitive, its flag is not, and it is updated to the disjunction of the flags of e and pc . This rule also updates the flag of class C , so if a class has a field affected by a declassified value, the whole class is going to be flagged true. The rule T-STORE is only for side-effect free expressions (i.e., e). There is also $e.m(\bar{p})$ that combined with Store can act as a method call or a gateway method call. Rule T-STORE-CALL and T-STORE-GATEWAY-CALL type check method calls and gateway method calls, respectively. Similarly, type checking assignments is broken into three rules, T-ASSIGN, T-ASSIGN-CALL, and T-ASSIGN-GATEWAY-CALL.

Rule T-DECLASSIFY ensures that only trusted data is allowed to be declassified $\ell \sqsubseteq \langle \mathbb{S}, \mathbb{T} \rangle$ and declassification can only happen in public and trusted context $pc_\ell \sqsubseteq \langle \mathbb{P}, \mathbb{T} \rangle$. This prevents attacker-controlled untrusted data from influencing the decision to declassify secret information. The security label of variable x will be $\langle \mathbb{P}, \mathbb{T} \rangle$ and its flag will be true T .

Rule T-RETURN sets the security label of the returned value to the join of the security labels of program context and the expression. This is to prevent the implicit flows that may happen when returning in a secret context. Additionally, if the expression e 's flag is true, return can only be executed in a trusted context. This is to prevent attacker from affecting the decision to return a declassified value.

Rule T-STORE-GATEWAY-CALL handles the type checking of gateway methods call from a store. The security labels of gateway parameters are explicitly defined in Γ^- and must have the $\langle \mathbb{P}, \mathbb{U} \rangle$ security label ($\forall p \in \bar{p}. \Gamma^-(p) = \langle \mathbb{P}, \mathbb{U} \rangle$), and their flag (defined in Π^-) should be false $\forall p \in \bar{p}. \Pi^-(p) = F$. The rule ensures that the gateway can only return public values ($\ell_2 \sqsubseteq \langle \mathbb{P}, \mathbb{U} \rangle$), the labels of actual parameters are less restrictive than the predefined labels of formal parameters $\Gamma(q_i) \sqsubseteq \Gamma^-(p_i)$, and the method's typing environment is satisfied ($\forall y \in \text{dom}(\Gamma^-). \Gamma(y) \sqsubseteq \Gamma^-(y)$). Additionally, the typing environment after type checking the method body should respect the method's predefined post typing environment ($\forall y \in \text{dom}(\Gamma^+). \Gamma^+(y) \sqsubseteq \Gamma_{out}(y)$), while ensuring that the type of identifiers other than those used by the method remains unchanged ($\forall y \in \text{dom}(\Gamma) \setminus \text{dom}(\Gamma^+). \Gamma(y) = \Gamma_{out}(y)$). Similar conditions apply to Π environment. At last, the rule updates the flag of field $C.f$ to false, so even a declassified value, after returning from a gateway method has a false flag.

To illustrate our type system, we revisit the example of Listing 3. The type system rejects this program because the declassify operator can only be used in a trusted program context pc . Rules T-OP and T-IF-ELSE set the pc security label to the join of the security labels of $time$ and $releaseTime$. Since $time$ has security label $\langle \mathbb{P}, \mathbb{U} \rangle$ (as it comes from outside of the enclave), and $releaseTime$ has security label $\langle \mathbb{P}, \mathbb{T} \rangle$ (as it comes from the enclave), we have that $pc_\ell = \langle \mathbb{P}, \mathbb{U} \rangle \sqcup \langle \mathbb{P}, \mathbb{T} \rangle = \langle \mathbb{P}, \mathbb{U} \rangle$. Next, rule T-DECLASSIFY does not allow the declassification since $pc_\ell \not\sqsubseteq \langle \mathbb{P}, \mathbb{T} \rangle$.

We prove that our security type system enforces robustness for the HAA attacker.

Theorem 1. *If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$ then $S[\vec{\bullet}]$ satisfies robustness under HAA.*

We now show how it can be extended with minimal changes to enforce robustness for HRAA attackers. The additional power of HRAA attacks comes from the ability to control the order and frequency of gateway calls. Whenever a gateway call modifies trusted identifiers inside the enclave, this can be used to influence declassification operations that are performed by another method. This can be prevented by computing the set of *shared identifiers* Σ containing all global variables and fields that are *assigned* to in at least one method and *accessed* in at least one method. (i.e., Σ contains all global variables and fields that are used and modified in one or more gateways methods) We compute Σ in a preprocessing step up to reaching a fixed point. We can then use the typing rules of Figures 4 and 5 under the constraint that the initial environment considers the integrity label of every variable and field in Σ as *untrusted*. i.e., $\forall x \in \Sigma. \Gamma_0[x \mapsto \langle \Gamma_0(x)_C, \mathbb{U} \rangle]$

$$\begin{array}{c}
\text{T-INT} \frac{}{\Gamma, \Pi \vdash_\delta n : (\langle \mathbb{P}, \mathbb{T} \rangle, F)} \quad \text{T-OP} \frac{\Gamma, \Pi \vdash_\delta e_1 : (\ell_1, d_1) \quad \Gamma, \Pi \vdash_\delta e_2 : (\ell_2, d_2)}{\Gamma, \Pi \vdash_\delta e_1 \oplus e_2 : (\ell_1 \sqcup \ell_2, d_1 \vee d_2)} \\
\text{T-METHOD} \frac{\text{method}_\phi(\bar{p})\{S; \text{return}(e)\} = \text{getMethod}(C, m) \quad pc', \Gamma^-, \Pi^- \vdash_\delta S; \text{return}(e) : \Gamma^+, \Pi^+ \triangleright (\ell, d)}{\Gamma, \Pi \vdash_\delta C.m(\bar{p}) : \left(\Gamma^-, \Pi^- \xrightarrow{pc'} \Gamma^+, \Pi^+ \right)_{(\ell, d)}}
\end{array}$$

Figure 4: Excerpt of typing rules for expressions

$$\begin{array}{c}
\text{T-STORE} \frac{\Gamma, \Pi \vdash_\delta e : (\ell_2, d_2) \quad \ell_2 \sqcup pc_\ell \sqsubseteq \ell_1 \quad \ell_1 \sqcup \ell_2 \sqcup pc_\ell = \langle \mathbb{S}, - \rangle \Rightarrow \delta(C) = E \quad pc_d \vee d_2 = T \Rightarrow pc_\ell \sqcup \ell_1 = \langle -, \mathbb{T} \rangle \quad d' = d_2 \vee pc_d}{pc, \Gamma, \Pi \vdash_\delta C.f := e : \Gamma, \Pi[C.f \mapsto d', C \mapsto \Pi(C) \vee d']} \\
\text{T-DECLASSIFY} \frac{\Gamma, \Pi \vdash_\delta e : (\ell, d) \quad \ell \sqsubseteq \langle \mathbb{S}, \mathbb{T} \rangle \quad pc_\ell \sqsubseteq \langle \mathbb{P}, \mathbb{T} \rangle \quad \delta(x) = E}{pc, \Gamma, \Pi \vdash_\delta x := \text{declassify}(e) : \Gamma[x \mapsto \ell \sqcap \langle \mathbb{P}, \mathbb{T} \rangle], \Pi[x \mapsto T]} \quad \text{T-RETURN} \frac{\Gamma, \Pi \vdash_\delta e : (\ell, d) \quad d = T \Rightarrow pc_\ell = \langle -, \mathbb{T} \rangle}{pc, \Gamma, \Pi \vdash_\delta \text{return}(e) : \Gamma, \Pi \triangleright (pc_\ell \sqcup \ell, d)} \\
\text{T-STORE-GATEWAY-CALL} \frac{\Gamma, \Pi \vdash_\delta C.f : (\ell_1, d_1) \quad \Gamma, \Pi \vdash_\delta C'.m(\bar{p}) : \left(\Gamma^-, \Pi^- \xrightarrow{pc'} \Gamma^+, \Pi^+ \right)_{(\ell_2, d_2)} \quad \delta(C) = N \quad \delta(C') = E}{\ell_2 \sqsubseteq \langle \mathbb{P}, \mathbb{U} \rangle \quad \forall p \in \bar{p}. \Gamma^-(p) = \langle \mathbb{P}, \mathbb{U} \rangle \quad \forall p \in \bar{p}. \Pi^-(p) = F \quad \ell_2 \sqcup pc_\ell \sqsubseteq \ell_1 \quad \Gamma(q_i) \sqsubseteq \Gamma^-(p_i) \quad i = 1 \dots |\bar{p}| \\ \Pi(q_i) = \Pi^-(p_i) \quad i = 1 \dots |\bar{p}| \quad \forall y \in \text{dom}(\Gamma^-). \Gamma(y) \sqsubseteq \Gamma^-(y) \quad \forall y \in \text{dom}(\Gamma^+). \Gamma^+(y) \sqsubseteq \Gamma_{out}(y) \\ \forall y \in (\text{dom}(\Gamma) \setminus \text{dom}(\Gamma^+)). \Gamma(y) = \Gamma_{out}(y) \quad \forall y \in \text{dom}(\Pi^+). \Pi(y) \sqsubseteq \Pi^-(y) \\ \forall y \in \text{dom}(\Pi^+). \Pi^+(y) \sqsubseteq \Pi_{out}(y) \quad \forall y \in (\text{dom}(\Pi) \setminus \text{dom}(\Pi^+)). \Pi(y) = \Pi_{out}(y)}{pc, \Gamma, \Pi \vdash_\delta C.f := C'.m(\bar{q}) : \Gamma_{out}, \Pi_{out}[C.f \mapsto F]}
\end{array}$$

Figure 5: Excerpt of typing rules for commands

and $\forall C.f \in \Sigma. \Gamma_0[C.f \mapsto \langle \Gamma_0(C.f)_C, \mathbb{U} \rangle]$.

The intuition is that by considering shared identifiers as untrusted, we enable the security type system to reject programs that use these identifiers to declassify information (explicitly or implicitly) and return it via a gateway method. For example, the program in Listing 4 is rejected by our type system. After the preprocessing phase, $\Sigma = \{\text{releaseTrigger}\}$ since *releaseTrigger* is assigned to in one method call and accessed in another. By assigning the security label $\langle \mathbb{P}, \mathbb{U} \rangle$ to *releaseTrigger*, rule T-DECLASSIFY (Line 16) will fail since declassification is not allowed in an untrusted context.

Recall from last section that if different program executions call different gateways and those gateways return *declassified values*, the HRAA attacker can learn more by calling the gateways that were not called in the original program. In order to prevent this, we add another step to the type checking process to ensure that all gateways which declassify secret values are called in all possible executions of the (non-enclave) program.

This process is performed in several steps:

- 1) Identify the set of all the *gateway* methods in $S[\vec{\bullet}]$ denoted by G^D such that their return value's flag is true (i.e., $\Gamma, \Pi \vdash_\delta C.m(\bar{p}) : (-)_{(-, T)}$).
- 2) Enumerate the paths of the program and extract the set of gateway calls along those paths. This is achieved by a depth-first traversal of the program's graph. We use $\text{pathsGW}(S[\vec{\bullet}])$ to denote the set of all possible paths of $S[\vec{\bullet}]$, and pathGW_i for the set of gateway calls in path i .
- 3) Check that G^D is a subset of all of the possible paths

of $S[\vec{\bullet}]$. In other words:

$$G^D \subseteq \text{pathGW}_i \quad \forall i \in \text{pathsGW}(S[\vec{\bullet}])$$

This process is performed after calculating Σ and type checking methods in isolation, but before type checking the program itself. If the above process fails, we reject the program as not robust against HRAA attacker. This requirement reflects the power of the HRAA attacker outside the enclave, rejecting programs that do not call all declassifying gateways. To improve permissiveness and security, such programs can be moved to the enclave and exposed to the non-enclave environment as a single gateway method.

We remark that our enforcement accepts programs with noninterfering runs that do not always declassify information. Listing 6 presents a program with both interfering and noninterfering runs. The program satisfies robustness (Definition 6) and is also accepted by the type system. In fact, if variable *trustedLow* (Line 4, a public and trusted variable) is set to true, only noninterfering computations will be executed and an HRAA attacker program such as `[choice=v]; ComputeArray.compute(choice);` will observe the result "Done with computation" + *choice*. Otherwise, if *trustedLow* is set to false, an interfering run will be executed and the attacker will observe the declassified average value.

This example demonstrates that even though our type system requires calling all of the declassifying gateways along all program's executions, it does not mean that all of the runs of the program are necessarily interfering.

We prove soundness of the security type system with respect to the HRAA attacker.

Listing 6: Noninterfering Runs under HRAA

```

1 @Enclave
2 class ComputeArray {
3
4     static boolean trustedLow;
5     @Secret static int[] array;
6
7     @Gateway
8     public static String compute(int choice) {
9         int avg = computeAvg(array);
10        String res = 0;
11        if (trustedLow) {
12            if (choice == 1) {
13                computation1(array); //noninterfering
14            } else if (choice == 2) {
15                computation2(array); //noninterfering
16            } else if (choice == 3) {
17                computation3(array); //noninterfering
18            }
19            res = "Done with computation" + choice;
20        }
21    }
22    else {
23        res = String.valueOf(declassify(avg));
24    }
25    return res;
26 } }

```

Theorem 2. *If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$ with regard to Σ and G^D , then $S[\vec{\bullet}]$ satisfies robustness under HRAA.*

We refer to Appendix D for the proofs of theorems.

Use case for the HRAA attacker Our enforcement mechanism for the HRAA attacker requires a program to call all gateways that declassify secret information in any execution of a program. While this condition may seem restrictive, it is necessary in order to ensure that an attacker as powerful as HRAA cannot manipulate the program to tamper with declassification in unintended ways.

We identify the setting of IoT app platforms as a promising use case for enforcing security against the HRAA attacker [34]. IoT apps allow users to run simple trigger-action apps in cloud-based IoT platforms to seamlessly connect their IoT services and devices. Upon the triggering of an event, e.g., "EZVIZ camera senses motion at home", the app executes code to perform an action, e.g., "Send an email with the camera stream". Currently, the users have to trust the cloud provider with the sensitive information of their services and devices to run apps on their behalf. TEEs can help executing user apps securely in an untrusted IoT cloud platform and protect against the HRAA attacker. Specifically, the user can leverage enclaves to securely connect their services and smart devices, e.g., Email and EZVIZ camera, via authentication tokens (using code patterns similar to Listing 9) and use these tokens to execute trigger-action automations via gateway methods that simply transfer data between services and devices, and do not declassify sensitive information. In this setting, only the authentication gateway declassifies sensitive information, thus making these programs amenable to verification by our type system.

V. ENDORSEMENT AND NONMALLEABLE ATTACKS

A. Endorsement

We extend our security framework to accommodate explicit endorsement of untrusted information coming from the

non-enclave environment via gateway calls. This enables a developer to mark untrusted expressions as trusted, explicitly, to indicate that the security policy should be insensitive to their value. Following the approach of Askarov and Myers [21], we extend the syntax and semantics of J_E with command $x := \text{endorse}_\eta(e)$. Each endorsement command has a unique label η , and produces an endorsement event $\text{endorse}(\eta, v)$, which records the label η along with the endorsed value v .

$$\text{ENDORSE} \frac{E \vdash_\delta \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle v, \mathcal{M}, \mathcal{H} \rangle \quad \delta(x) = E \quad \mathcal{M}' = \mathcal{M}[x \mapsto v]}{E \vdash_\delta \langle x := \text{endorse}_\eta(e), \mathcal{M}, \mathcal{H} \rangle \Downarrow_{\text{endorse}(\eta, v)} \mathcal{M}', \mathcal{H}}$$

We define *irrelevant* attacks as the set of attacks that are endorsed, and therefore can be excluded from the set of attacks used in definition 5. Using this concept, we argue that definition 5 should only hold for relevant attacks. Given a program $S[\vec{\bullet}]$, starting state σ , and attacker vector \vec{a} which produces trace t (i.e., $\langle S[\vec{a}], \sigma \rangle \Downarrow_t$), relevant attacks, denoted by $\Omega(S[\vec{a}], \sigma)$, are the attacks that lead to the same sequence of endorsement events as in t .

We define relevant attacks by using irrelevant traces, which given trace t , are the set of all traces that agree with t on some prefix of *endorsement events* until they necessarily disagree on some endorsement. Formally:

Definition 7 (Irrelevant Traces). *Given a trace t , where endorsements are marked as $\text{endorse}(\eta_j, v_j)$, define a set of irrelevant traces based on the number of endorsements in t*

$$\psi_i(t) = \{t' \mid t' = k.\text{endorse}(\eta_i, v'_i).k'\}$$

where k is a prefix of t' with $i - 1$ events all of which agree with *endorse events* in t , and $v_i \neq v'_i$.

We define $\psi(t) \triangleq \bigcup_i \psi_i(t)$ as a set of irrelevant traces w.r.t. trace t .

Now, we can define the relevant attacks as the set of attacks that *do not* lead to irrelevant traces.

Definition 8 (Relevant Attacks). *Given a program $S[\vec{\bullet}]$, starting state σ , and attacker vector \vec{a} such that $\langle S[\vec{a}], \sigma \rangle \Downarrow_t$, relevant attacks $\Omega(S[\vec{a}], \sigma)$ are defined as:*

$$\Omega(S[\vec{a}], \sigma) = \{a' \mid \langle S[\vec{a'}], \sigma \rangle \Downarrow_{t'} \wedge t' \notin \psi(t)\}$$

Using the definition relevant attacks, we can redefine the robustness property. This new security condition accounts for the fact that the active effect on endorsed expressions does not matter.

Definition 9 (Robustness with Endorsement). *Program $S[\vec{\bullet}]$ is robust w.r.t an HAA attacker A if for all $\sigma_1, \sigma_2, \vec{a}_1$ and for all $\vec{a}_2 \in \Omega(S[\vec{a}_1], \sigma_1)$*

$$N \vdash_\delta \langle S[\vec{a}_1], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}_1], \sigma_2 \rangle \Rightarrow N \vdash_\delta \langle S[\vec{a}_2], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}_2], \sigma_2 \rangle$$

This definition is similar to the original robustness property, except that instead of ensuring indistinguishability for all

$$\text{T-ENDORSE} \frac{\Gamma, \Pi \vdash_{\delta} e : (\ell, d) \quad pc_{\ell} \sqsubseteq \langle \mathbb{P}, \mathbb{T} \rangle \quad \delta(x) = E}{pc, \Gamma, \Pi \vdash_{\delta} x := \text{endorse}_{\eta}(e) : \Gamma[x \mapsto \ell \sqcap \langle \mathbb{S}, \mathbb{T} \rangle], \Pi[x \mapsto d]}$$

Figure 6: Typing rule for endorse command

possible attacks, it only ensures indistinguishability for the relevant attacks, effectively ignoring the influence of irrelevant attacks.

Figure 6 presents the typing rule for endorsement. Similar to the T-DECLASSIFY, T-ENDORSE ensures that $\text{endorse}_{\eta}(e)$ can only be used inside the enclave environment, and endorsement can only occur in a public and trusted context $pc_{\ell} \sqsubseteq \langle \mathbb{P}, \mathbb{T} \rangle$. This is to prevent attacker-controlled untrusted data to influence the decision to endorse untrusted information.

Now, we present type soundness to prove that any well-typed program under the extended type system satisfies the new robustness property of Definition 9.

Theorem 3. *If $pc, \Gamma, \Pi \vdash_{\delta} S[\vec{\bullet}] : \Gamma', \Pi'$ then $S[\vec{\bullet}]$ satisfies robustness with endorsement.*

We can use a similar approach to model endorsement of HRAA attacker. However, we have to modify the definition of relevant attacks to account for the difference between the order and frequency of endorsement events in $S'[\vec{\bullet}]$ and $S[\vec{\bullet}]$.

In the new definition of relevant attacks, we use the unique endorsement label η to ensure that the value of an endorsement used in $S[\vec{\bullet}]$ is equal to the value of that same endorsement in $S'[\vec{\bullet}]$ independently of the order and frequency of that endorsement's event. We lift the definition of irrelevant traces and relevant attacks to this new setting.

Definition 10 (HRAA Irrelevant Traces). *Given a trace t , where endorsements are marked as $\text{endorse}(\eta_j, v_j)$, define the set of irrelevant traces as:*

$$\psi_R(t) = \{t' \mid \exists \text{endorse}(\eta_j, v_j) \in t, \exists \text{endorse}(\eta_i, v_i) \in t', \eta_j = \eta_i \text{ and } v_j \neq v_i\}$$

Observe that we need Definition 10 because, if we used Definition 7, there might be cases where there is an irrelevant trace t' , but there exist no k with $i - 1$ events such that all of its endorse events agree with t , thus, t' cannot be marked as irrelevant.

Definition 11 (HRAA Relevant Attacks). *Given programs $S[\vec{\bullet}]$ and $S'[\vec{\bullet}]$, starting state σ , and attacker vector \vec{a} such that $\langle S[\vec{a}], \sigma \rangle \Downarrow_t$, the set of relevant attacks w.r.t HRAA attacker $\Omega_R(S[\vec{a}], S'[\vec{\bullet}], \sigma)$ are defined as:*

$$\Omega_R(S[\vec{a}], S'[\vec{\bullet}], \sigma) = \{a' \mid \langle S'[\vec{a'}], \sigma \rangle \Downarrow_{t'} \wedge t' \notin \psi(t)\}$$

Now, using this new definition of relevant attacks, we can redefine the robustness property w.r.t HRAA attacker.

Definition 12 (Robustness under HRAA with Endorsement). *Program $S[\vec{\bullet}]$ is robust w.r.t an HAA attacker A if for all $\sigma_1, \sigma_2, \vec{a}_1$, and for all $S'[\vec{\bullet}]$ such that $\forall \vec{a}_2 \in$*

$$\Omega_R(S[\vec{a}_1], S'[\vec{\bullet}], \sigma_1)$$

$$N \vdash_{\delta} \langle S[\vec{a}_1], \sigma_1 \rangle \approx_A N \vdash_{\delta} \langle S[\vec{a}_1], \sigma_2 \rangle \Rightarrow$$

$$N \vdash_{\delta} \langle S'[\vec{a}_2], \sigma_1 \rangle \approx_A N \vdash_{\delta} \langle S'[\vec{a}_2], \sigma_2 \rangle$$

This definition ensures that in a robust program w.r.t HRAA attacker, as long as the endorsed values are equal, the result of execution for all possible orders and frequencies of gateway calls is indistinguishable.

We can augment the security type system of HRAA attacker with rule of Figure 6 to enforce robustness. The following theorem proves that any well-typed program under the extended type system w.r.t HRAA attacker satisfies the robustness property of Definition 12.

Theorem 4. *If $pc, \Gamma, \Pi \vdash_{\delta} S[\vec{\bullet}] : \Gamma', \Pi'$ with regard to Σ and G^D , then $S[\vec{\bullet}]$ satisfies robustness under HRAA with endorsement.*

We refer the readers to Appendix E for the proofs of Theorems 3 and 4.

B. Nonmalleable Attacks

Transparent endorsement [35] was introduced as a dual to robust declassification to prevent attacks arising from trusted agents endorsing information that their provider could not have seen. It is a common practice in TEE to have encrypted data passed to the enclave from outside, if these encrypted data are labeled *secret*, then the type system will reject the program and prevent the malleable attacks. Otherwise, if these encrypted ciphertexts are labeled *public* even though they contain secret plaintexts, then it is possible to receive an input which its provider cannot read, thus opening up the system to malleable attacks. While from a technical perspective transparent endorsement can be accommodated into our framework along the lines of Cecchetti et al. [35], we postpone it to future work.

VI. IMPLEMENTATION

In this section, we describe the implementation of J_E for Intel SGX and discuss the gap with our formal model.

The J_E runtime operates with two separate execution environments, namely the non-enclave environment and enclave environment, which correspond to two separate Java Virtual Machines. The communication between the two environments is achieved using Java RMI [36].

The J_E compilation process involves multiple steps. A J_E program is first partitioned, then translated to Jif to check security, and finally transformed to use RMI communication. A detailed description of the code transformations implemented by the J_E compiler with a complete step-by-step example is in Appendix A.

In addition to the abstractions supported in the formal model, i.e., static classes and static fields, the implementation supports additional Java features, including objects and generics. We convert J_E programs that include objects and generics into equivalent Jif programs. Because the implementation

Listing 7: Password checker

```

1 @Enclave
2 class PasswordChecker {
3     @Secret static String password;
4
5     @Gateway
6     public static boolean checkPassword(String guess) {
7         return (getHash(guess) == declassify(getHash(password)));
8     } }

```

leverages Jif for label propagation, J_E security analysis is flow-insensitive, in contrast to the security type system, where the analysis is flow-sensitive on program variables. In the implementation, we do not assign any default Jif label to the local variables, hence, local variables get the security label of the expression they are initialized with. We remark that this decision is sound with respect to the language subset considered in our formal model. In fact, if a (declassification-free) program is deemed secure by Jif’s flow-insensitive analysis, then it is also secure with respect to our flow-sensitive analysis. However, soundness may come at the expense of additional manual annotations due to the flow insensitivity of Jif’s analysis.

Another difference with our formal model is the enforcement of delayed declassification. We not do implement the propagation of the declassify flag (d) used to track delayed declassification in section IV-C5. Instead, we require that every declassify operator is used only as a parameter of a return statement. Since any declassifications are performed within the scope of the return statement, it simply disallows the programmer to write code subject to delayed declassification leaks. We implement this check during the static analysis.

We use the SGX-LKL [37] framework to run the JAR corresponding to the enclave partition. SGX-LKL uses Linux Kernel Library [38] to handle system calls from the application within the enclave. In SGX-LKL, the application JAR, a JVM, and the necessary LKL binaries are compiled to a single image file. SGX-LKL loads the image and runs it inside the enclave.

VII. EVALUATION

In this section, we present the evaluation of J_E . We implemented case studies to demonstrate how the design of J_E can address the security requirements of distributed applications. The goal of the case studies is to exhibit the features of the language. Additional case studies are provided in Appendix B.

A. Password Checker

Listing 7 shows a password checker. The class PasswordChecker is annotated with the @Enclave annotation. As a result, during run time, the PasswordChecker class is placed inside the enclave. The field password is annotated with the @Secret annotation. The method checkPassword returns a boolean which is the result of the comparison of hashes of the field password and parameter guess. The code snippet without any annotations is a valid Java code. The case study demonstrates the HAA attacker scenario. The parameter

guess 6 to the gateway method checkPassword is controlled by the attacker and is considered untrusted.

B. Apache Spark

Spark [39] is a popular processing framework for big data. It is widely used for machine learning and streaming jobs within clusters. We show a J_E implementation of secure medical data processing using the enclave and Spark. The goal is to process encrypted medical records inside the enclave to extract statistics about the records. Listing 8 represents the code running outside the enclave. We create a Spark context object sc (Line 5) and store the encrypted records to be processed in a JavaRDD (resilient distributed dataset) object recList (Line 7) which is a distributed data structure to store the data to be processed. EncRec is an encrypted medical record. The StatUtil class (Line 12) contains a gateway method process (Line 17) that accepts an encrypted record, it decrypts it with the secret key key, and returns the corresponding statistical information StatRecord. Hence, a medical record is only decrypted inside the enclave, protecting the medical data. The code without any annotations is a valid Java program, and the programmer only needs to add the annotations without any extra modifications to the original program.

C. Updatable Password Checker

We consider a password checker with a provision for updating the password via a gateway method. Listing 9 shows the class UpdatablePasswordChecker annotated with the @Enclave annotation. It contains a secret field password and two gateway methods, namely, checkPassword and updatePassword. The checkPassword method (Line 12) compares the argument guess with the field password and returns the result of the comparison as a boolean. Method updatePassword (Line 7) assigns the value of the argument newPass to the secret field password (Line 9). This case study illustrates the HRAA attack scenario. The attacker can reorder the calling sequence of gateway methods. Since the gateway method updatePassword assigns an untrusted value to the secret field password (Line 9), the secret field can no longer be trusted. Therefore, we infer the security level of the password secret field as $\langle S, U \rangle$. The field password needs to be endorsed (Line 15) to allow the declassification (Line 18). J_E detects secret fields that can be modified by an attacker and the programmer needs to explicitly trust such fields to allow any dependent declassification.

VIII. RELATED WORK

This section compares J_E with closely related work on securing distributed applications with TEEs, focusing on information flow control, frameworks for TEEs, and secure program partitioning.

1) *Information flow control and enclaves*: Recent works leverage IFC and enclaves to build applications with strong security guarantees. These works target specific challenges in the domain space, including security foundations of enclave applications, secure code partitioning for enclaves

Listing 8: Medical data processing using Spark

```

1 class MainClass {
2
3     public static void main(String[] args) {
4         JavaSparkContext sc = new JavaSparkContext(
5             new SparkConf().setAppName("Foo"));
6         JavaRDD<EncRec> recList =
7             sc.parallelize(getEncryptedRecords());
8         List<StatRecord> recList = recList.map(x ->
9             StatUtil.process(x));
10    } }

```

```

11 @Enclave
12 class StatUtil {
13
14     @Secret static String key;
15
16     @Gateway
17     public static StatRecord process(EncRec rec) {
18         EncRec recE = endorse(rec);
19         Record rec = decrypt(key, recE);
20         // process the decrypted record
21         return declassify(rec);
22    } }

```

Listing 9: Updatable password checker

```

1 @Enclave
2 class UpdatablePasswordChecker {
3
4     @Secret static String password;
5
6     @Gateway
7     public static void updatePassword(String currPass, String newPass) {
8         if (checkPassword(currPass) == true) {
9             password = newPass;
10        } }
11
12     @Gateway
13     public static boolean checkPassword(String guess) {
14         boolean result = false;
15         String guessE = endorse(guess);
16         if (endorse(password).equals(guessE)) { // endorsing password
17             result = true; // result of type <S,T>
18         }
19         return declassify(result); // declassifying result
20    } }

```

and language support for programming in TEEs. J_E pushes the boundary in several directions providing: (a) language support for Java applications in Intel SGX; (b) a security type system enabling verification of secure partitioning against realistic attacker models; (c) formalization and soundness proofs with respect to semantic characterization of security; (d) a prototype implementation supported by case studies.

Sinha et al. [15] present Moat, a system for statically verifying confidentiality properties of Intel SGX programs in the presence of passive and active attackers. They combine a flow sensitive type system with automated theorem proving to check confidentiality policies with declassification. Their active attacker is similar to our HAA attacker, assuming arbitrary modification of non-SGX code. By contrast, we focus on verifying robustness of an application against active attacker which has the advantage of using a more lightweight analysis whenever a program is already proved secure against passive attackers. Moreover, we consider more powerful attackers, and target languages with managed runtimes, thus shielding developers from SGX-specific details. Follow-up work by Sinha et al. [40] enforces Information Release Confinement, an access control policy which allows SGX code to perform arbitrary computations ensuring that it can only generate output to the non-SGX memory through encrypted channels. While this approach cannot enforce expressive information flow policies, it implements Control Flow Integrity to ensure that an active attacker does not compromise the control flow of an executions. Our approach would require similar techniques to enforce security against the HRAA attacker.

Inspired by Moat, Gollamudi and Chong [16] propose a security type system to enforce flexible information flow policies

for an SGX-enabled imperative language. They abstract away the details of enclave management and develop a translation tool to automatically infer the parts of the program to be executed inside the enclaves. They consider active attackers and formalize security using knowledge-based conditions. By contrast, our approach enforces robustness for a mainstream language like Java. Recent work by Gollamudi et al. [41] uses enclaves to enforce more expressive confidentiality and integrity against passive attackers in a distributed setting. Like us, both works come with soundness proofs of security, with the key difference that our work targets robustness instead of variants of noninterference. Liu et al. [42], [43] study automated program partitioning with SGX for passive attackers. We argue that passive attackers are too weak in the context of SGX, hence our work provides the formal grounds for extending and proving their techniques in a realistic setting.

A large array of works study IFC in distributed settings [19], [44], [45], [46], [47], [48]. These works address the problem of secure program partitioning across nodes in a distributed system under the assumption that low integrity nodes may be controlled by an active attacker. Our work draws inspiration from these approaches and extends them to capture attacker models arising in TEEs.

Existing works target the foundations of IFC in presence of active attackers [49], [21], [50], [51], [52], [20] and enforce robustness via security type systems.

2) *Programming frameworks for TEEs*: To reduce the trusted computing base, in contrast to running unmodified applications inside SGX, some approaches focus on partitioning the application into components that execute within the enclave and the rest, which executes in the untrusted environment. Glamdring [53] propose the first source-level partitioning framework for securing C applications with Intel SGX. Developers annotate sensitive data and Glamdring automatically partitions the application into untrusted and enclave parts. Panoply [54] generates application binaries from the annotated source code where annotations specify the parts of the application to be run inside separate enclaves. DynSGX [55] provides tools to dynamically load, unload and execute compiled functions inside SGX enclaves efficiently. While these works leverage enclaves to enforce strong isolation properties, they do not enforce application-level policies and lack formal security guarantees.

Several works aim to ease the programming of applications that (partially) execute within an enclave. Coppolino et al. [56]

present a comparative analysis of the existing approaches for securing Java applications with Intel SGX. RUST-SGX [57] provides memory-safe SGX support for Rust through a memory management scheme to control the interface between Rust and Intel’s C/C++ APIs. Civet [17] is a programming framework for Java using an XML file to specify the classes that are executed inside the enclave. Civet leverages dynamic information flow control to track insecure flows within the enclave. Uranus [58] supports executing Java functions inside SGX enclaves. It provides two method-level annotations `JECall` and `JOCall` to indicate methods to be executed inside and outside the enclave respectively. Secure Routines [59] is a programming framework for SGX in the Go language. Programmers can execute Go functions (goroutine) inside an enclave, and use low-overhead channels to communicate with the untrusted environments. Like us, these works aim at developing programming models with enclaves. However they lack provable security guarantees and require security analysis of the partitioned application from scratch.

Other works propose processor model calculi to capture necessary conditions for safe remote execution of enclave programs. Subramanyan et al. [60] introduce an abstract processor model to verify the security guarantees of Intel SGX under specific adversary capabilities. Sinha [61] studies confidentiality risks that can be exploited by application and infrastructure attacks in SGX applications. Autarky [62] is a controlled-channel attack resistant framework based on a modified SGX ISA.

3) *Running unmodified applications inside SGX:* Haven [24] was the first system built to run unmodified Windows legacy applications inside Intel SGX. SCONE [25] is a Docker extension that uses SGX to protect container processes. SGX-LKL [28], SGXKernel [26], Graphene-SGX [27], and Occlum [63] are library OS based frameworks designed to run unmodified Linux applications inside the SGX enclave.

4) *Multitier programming and secure program partitioning:* The J_E programming model is inspired by the multitier programming paradigm [64], [65], [66] – for a comprehensive overview of multitier programming, we refer to the survey by Weisenburger et al. [67]. In multitier programming, the code for different tiers is written as a single compilation unit and the compiler automatically splits it into the components associated to each individual tier. Different works extend the multitier programming model with information flow policies to build secure web applications via security type systems [68], [69], [70], [71] and symbolic execution [72]. In contrast to J_E , none of these approaches enforce robustness properties.

IX. CONCLUSION

In this paper we present J_E , a language for confidential computing which supports enclave-enabled applications. First, J_E seamlessly integrates with a high-level, managed language, and enables programmers to develop secure enclave-enabled applications by adding annotations to Java programs. Second, J_E comes with a security model that accounts for realistic attackers, that, in the case of enclave programming, can tamper

with the code and the data of the non-enclave environment. We define the notion of *robustness* of enclave-enabled programs and prove that it is correctly enforced by the J_E type system.

We evaluate our approach on several use cases from the literature, including a battleship game, a secure event processing system, and a popular processing framework for big data, showing that it can correctly handle complex cases of partitioning, information flow, declassification and trust.

We envision different avenues for future work including a generalization of our framework to multiple enclaves and nonmalleable information flow. On the practical side, we plan to extend our automated partitioning and compilation algorithms to handle Java programs beyond the core J_E fragment.

ACKNOWLEDGMENTS

Thanks are due to Owen Arden and anonymous reviewers for their helpful feedback on this paper. This work is partially supported by the Deutsche Forschungsgemeinschaft (DFG) – SFB 1119 – 236615297, the BRF Project 1025524 from the University of St.Gallen, the Swedish Foundation for Strategic Research (SSF), the Swedish Research Council (VR), and Digital Futures.

REFERENCES

- [1] F. McKeen, I. Alexandrovich, A. Berenzon, C. V. Rozas, H. Shafi, V. Shanbhogue, and U. R. Savagaonkar, “Innovative instructions and software model for isolated execution,” ser. HASP ’13, 2013.
- [2] Intel, “Intel® software guard extensions developer guide,” 2014, accessed 2021-05-20. [Online]. Available: <https://software.intel.com/content/dam/develop/public/us/en/documents/intel-sgx-developer-guide.pdf>
- [3] ARM, “Building a secure system using trustzone® technology,” 2009.
- [4] S. Pinto and C. Garlati, “Multi zone security for arm cortex-m devices,” 2020, <https://hex-five.com/wp-content/uploads/2020/02/Multi-Zone-Security-White-Paper-20200224.pdf>.
- [5] V. Costan, I. Lebedev, and S. Devadas, “Sanctum: Minimal hardware extensions for strong software isolation,” 2016, pp. 857–874.
- [6] D. Kaplan, J. Powell, and T. Woller, “Amd memory encryption,” 2016.
- [7] T. Bourgeat, I. Lebedev, A. Wright, S. Zhang, Arvind, and S. Devadas, “Mi6: Secure enclaves in a speculative out-of-order processor,” in *MICRO* ’52, 2019, p. 42–56.
- [8] D. Lee, D. Kohlbrenner, S. Shinde, K. Asanović, and D. Song, “Keystone: An open framework for architecting trusted execution environments,” in *EuroSys* ’20, 2020.
- [9] Apple, “Apple platform security,” https://manuals.info.apple.com/MANUALS/1000/MA1902/en_US/apple-platform-security-guide.pdf, 2020, accessed 2021-05-20.
- [10] Fortanix, “The fortanix runtime encryption,” <https://fortanix.com/products/runtime-encryption>, accessed 2021-05-20.
- [11] Anjuna, “Anjuna enterprise enclaves,” <https://www.anjuna.io/enterprise-enclaves>, accessed 2021-05-20.
- [12] Oracle, “Java native interface,” <https://docs.oracle.com/javase/8/docs/technotes/guides/jni>, accessed 2021-05-20.
- [13] GraalVM, “Graalvm native image,” <https://www.graalvm.org/docs/reference-manual/native-image>, accessed 2021-05-20.
- [14] A. Sabelfeld and A. C. Myers, “Language-based information-flow security,” *IEEE Journal on Selected Areas in Communications*, 2003.
- [15] R. Sinha, S. K. Rajamani, S. A. Seshia, and K. Vaswani, “Moat: Verifying confidentiality of enclave programs,” in *CCS* ’15, 2015, pp. 1169–1184.
- [16] A. Gollamudi and S. Chong, “Automatic enforcement of expressive security policies using enclaves,” in *OOPSLA* ’16, 2016, p. 494–513.
- [17] C.-C. Tsai, J. Son, B. Jain, J. McAvey, R. A. Popa, and D. E. Porter, “Civet: An efficient java partitioning framework for hardware enclaves,” in *USENIX Security* ’20, 2020.

- [18] A. C. Myers, "JFlow: Practical Mostly-static Information Flow Control," in *POPL'99*, 1999, pp. 228–241.
- [19] S. Zdancewic, L. Zheng, N. Nystrom, and A. C. Myers, "Untrusted hosts and confidentiality: Secure program partitioning," in *SOSP'01*, 2001, p. 1–14.
- [20] A. C. Myers, A. Sabelfeld, and S. Zdancewic, "Enforcing robust declassification," in *CSFW'04*, 2004, pp. 172–186.
- [21] A. Askarov and A. C. Myers, "Attacker control and impact for confidentiality and integrity," *LMCS'11*, 2011.
- [22] S. Hunt and D. Sands, "On flow-sensitive security types," in *POPL'06*, 2006, pp. 79–90.
- [23] V. Costan and S. Devadas, "Intel sgx explained," Cryptology ePrint Archive, Report 2016/086, 2016, <https://eprint.iacr.org/2016/086>.
- [24] A. Baumann, M. Peinado, and G. Hunt, "Shielding applications from an untrusted cloud with haven," *ACM Transactions on Computer Systems (TOCS)*, p. 8, 2015.
- [25] S. Arnaudov, B. Trach, F. Gregor, T. Knauth, A. Martin, C. Priebe, J. Lind, D. Muthukumaran, D. O'Keeffe, M. L. Stillwell, D. Goltzsche, D. Eysers, R. Kapitza, P. Pietzuch, and C. Fetzer, "Scone: Secure linux containers with intel sgx," in *OSDI'16*, 2016, pp. 689–703.
- [26] H. Tian, Y. Zhang, C. Xing, and S. Yan, "Sgxkernel: A library operating system optimized for intel sgx," in *Proceedings of the Computing Frontiers Conference*, 2017, pp. 35–44.
- [27] C.-C. Tsai, D. E. Porter, and M. Vij, "Graphene-sgx: A practical library os for unmodified applications on sgx," in *USENIX ATC'17*, 2017, pp. 645–658.
- [28] C. Priebe, D. Muthukumaran, J. Lind, H. Zhu, S. Cui, V. A. Sartakov, and P. Pietzuch, "Sgx-lkl: Securing the host os interface for trusted execution," 2019.
- [29] A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands, "Termination-insensitive noninterference leaks more than just a bit," in *ESORICS*, 2008.
- [30] J. B. Almeida, M. Barbosa, G. Barthe, F. Dupressoir, and M. Emmi, "Verifying constant-time implementations," in *USENIX Security*, 2016.
- [31] A. Askarov, D. Zhang, and A. C. Myers, "Predictive black-box mitigation of timing channels," in *ACM CCS*, 2010.
- [32] A. Sabelfeld and D. Sands, "Declassification: Dimensions and principles," *Journal of Computer Security*, 2009.
- [33] A. Askarov and A. Sabelfeld, "Gradual release: Unifying declassification, encryption and key release policies," in *S&P'07*, 2007, pp. 207–221.
- [34] M. Balliu, I. Bastys, and A. Sabelfeld, "Securing IoT Apps," *IEEE Security & Privacy Magazine*, 2019.
- [35] E. Cecchetti, A. C. Myers, and O. Arden, "Nonmalleable information flow control," in *CCS'17*, 2017, pp. 1875–1891.
- [36] Oracle, "Java remote method invocation - distributed computing for java," <https://www.oracle.com/technetwork/java/javase/tech/index-jsp-138781.html>, accessed 2021-05-20.
- [37] SGX-LKL, "Sgx-lkl library os for running linux applications inside of intel sgx enclaves," <https://github.com/llds/sgx-lkl>, accessed 2021-05-20.
- [38] O. Purdila, L. A. Grijincu, and N. Tapus, "Lkl: The linux kernel library," in *9th RoEduNet IEEE International Conference*, 2010, pp. 328–333.
- [39] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica, "Spark: Cluster computing with working sets," in *Hot Topics in Cloud Computing*, 2010, p. 10.
- [40] R. Sinha, M. Costa, A. Lal, N. P. Lopes, S. K. Rajamani, S. A. Seshia, and K. Vaswani, "A design and verification methodology for secure isolated regions," in *PLDI'16*, 2016, pp. 665–681.
- [41] A. Gollamudi, S. Chong, and O. Arden, "Information flow control for distributed trusted execution environments," in *CSF'19*, 2019, pp. 304–30414.
- [42] S. Liu, G. Tan, and T. Jaeger, "Ptrsplit: Supporting general pointers in automatic program partitioning," in *CCS'17*, 2017, pp. 2359–2371.
- [43] S. Liu, D. Zeng, Y. Huang, F. Capobianco, S. McCamant, T. Jaeger, and G. Tan, "Program-mandering: Quantitative privilege separation," in *CCS'19*, 2019, pp. 1023–1040.
- [44] S. Chong, J. Liu, A. C. Myers, X. Qi, K. Vikram, L. Zheng, and X. Zheng, "Secure web applications via automatic partitioning," in *SOSP'07*, 2007, pp. 31–44.
- [45] J. Liu, O. Arden, M. D. George, and A. C. Myers, "Fabric: Building open distributed systems securely by construction," *Journal of Computer Security*, pp. 367–426, 2017.
- [46] L. Zheng, S. Chong, A. C. Myers, and S. Zdancewic, "Using replication and partitioning to build secure distributed systems," in *S&P'03*, 2003, p. 236.
- [47] S. Zdancewic, L. Zheng, N. Nystrom, and A. C. Myers, "Secure program partitioning," *ACM Trans. Comput. Syst.*, p. 283–328, Aug. 2002.
- [48] S. Chong, K. Vikram, and A. C. Myers, "SIF: Enforcing confidentiality and integrity in web applications," in *USENIX Security 07*. Boston, MA: USENIX Association, Aug. 2007.
- [49] E. Cecchetti, A. C. Myers, and O. Arden, "Nonmalleable information flow control," in *CCS'17*, 2017, pp. 1875–1891.
- [50] M. Balliu and I. Mastroeni, "A Weakest Precondition Approach to Active Attacks Analysis," ser. PLAS '09, 2009, pp. 59–71.
- [51] C. Fournet, G. L. Guernic, and T. Rezk, "A security-preserving compiler for distributed programs: From information-flow policies to cryptographic mechanisms," in *CCS'09*, 2009, p. 432–441.
- [52] M. Balliu and I. Mastroeni, "A Weakest Precondition Approach to Robustness," *Transactions on Computational Science X*, vol. 10, pp. 261–297, 2010.
- [53] J. Lind, C. Priebe, D. Muthukumaran, D. O'Keeffe, P.-L. Aublin, F. Kelbert, T. Reiher, D. Goltzsche, D. Eysers, R. Kapitza, C. Fetzer, and P. Pietzuch, "Glamdring: Automatic application partitioning for intel sgx," in *USENIX ATC'17*, 2017, p. 285–298.
- [54] S. Shinde, D. L. Tien, S. Tople, and P. Saxena, "Panoply: Low-tcb linux applications with SGX enclaves," in *NDSS'17*, 2017.
- [55] R. Silva, P. Barbosa, and A. Brito, "Dynsgx: A privacy preserving toolset for dynamically loading functions into intel(r) sgx enclaves," in *CloudCom'17*, 2017, pp. 314–321.
- [56] L. Coppolino, S. D'Antonio, G. Mazzeo, and L. Romano, "A comparative analysis of emerging approaches for securing java software with intel sgx," *Future Generation Computer Systems*, pp. 620 – 633, 2019.
- [57] H. Wang, P. Wang, Y. Ding, M. Sun, Y. Jing, R. Duan, L. Li, Y. Zhang, T. Wei, and Z. Lin, "Towards memory safe enclave programming with rust-sgx," in *CCS'19*, 2019, p. 2333–2350.
- [58] J. Jiang, X. Chen, T. Li, C. Wang, T. Shen, S. Zhao, H. Cui, C. Wang, and F. Zhang, "Uranus: Simple, Efficient SGX Programming and its Applications," in *ASIA CCS'20*, 2020.
- [59] A. Ghosn, J. R. Larus, and E. Bugnion, "Secured routines: Language-based construction of trusted execution environments," in *USENIX ATC'19*, 2019, p. 571–585.
- [60] P. Subramanyan, R. Sinha, I. Lebedev, S. Devadas, and S. A. Seshia, "A formal foundation for secure remote execution of enclaves," in *CCS'17*, 2017, p. 2435–2450.
- [61] R. Sinha, "Secure computing using certified software and trusted hardware," Ph.D. dissertation, 2017.
- [62] M. Orenbach, A. Baumann, and M. Silberstein, "Autarky: Closing controlled channels with self-paging enclaves," in *EuroSys'20*, 2020.
- [63] Y. Shen, H. Tian, Y. Chen, K. Chen, R. Wang, Y. Xu, Y. Xia, and S. Yan, "Occlum: Secure and efficient multitasking inside a single enclave of intel sgx," in *ASPLOS'20*, 2020, p. 955–970.
- [64] M. Serrano, E. Gallezio, and F. Loitsch, "Hop, a language for programming the web 2.0," ser. OOPSLA Companion '06. ACM, 2006.
- [65] E. Cooper, S. Lindley, P. Wadler, and J. Yallop, "Links: Web programming without tiers," ser. FMCO '06. Springer-Verlag, 2006, pp. 266–296.
- [66] P. Weisenburger, M. Köhler, and G. Salvaneschi, "Distributed system development with scalaloci," *Proc. ACM Program. Lang.*, vol. 2, no. OOPSLA, Oct. 2018. [Online]. Available: <https://doi.org/10.1145/3276499>
- [67] P. Weisenburger, J. Wirth, and G. Salvaneschi, "A survey of multitier programming," *ACM Comput. Surv.*, Sep. 2020.
- [68] B. J. Corcoran, N. Swamy, and M. W. Hicks, "Cross-tier, label-based security enforcement for web applications," in *SIGMOD*, 2009.
- [69] D. Schoepe, D. Hedin, and A. Sabelfeld, "SeLINQ: tracking information across application-database boundaries," in *ICFP*, 2014.
- [70] M. Balliu, B. Liebe, D. Schoepe, and A. Sabelfeld, "JSLINQ: building secure applications across tiers," in *CODASPY'16*, 2016, pp. 307–318.
- [71] J. Parker, N. Vazou, and M. Hicks, "Lweb: information flow security for multi-tier web applications," *Proc. ACM Program. Lang.*, vol. 3, no. POPL'19, pp. 75:1–75:30, 2019.
- [72] A. Chlipala, "Ur/web: A simple model for programming the web," in *POPL '15*. Association for Computing Machinery, 2015, p. 153–165.
- [73] P. Weisenburger, M. Luthra, B. Koldehofe, and G. Salvaneschi, "Quality-aware runtime adaptation in complex event processing," in *Proceedings of the 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2017, p. 140–151.
- [74] J. Daemen and V. Rijmen, "Aes proposal: Rijndael," 1999.

A. Compilation Steps

The J_E compilation process involves multiple steps. Figure 7 shows how a J_E program (Listing 10) is partitioned (Listing 11 and 12), translated to Jif to check security (Listing 13 and 14), and augmented with RMI communication (Listing 15 and 16). The following are the individual steps in details.

Static Analysis: The first step of the compilation involves performing static analysis to collect the secret variables that are modified by multiple gateway methods (shared variables) and detect locations of the gateway calls. This information is used to decide security labels of shared variables and to warn the programmer about the possible state changes inside the enclave (see A).

Code Partitioning: A J_E program is first analyzed and, based on the annotations, it is split into two partition – the enclave and the non-enclave partition. All the classes annotated with the @Enclave annotation and all their required dependencies belong to the enclave partition. All the remaining classes belong to the non-enclave partition. Listing 10 shows a complete J_E program using the routine to securely encrypt data from Listing 1 (Section III). The complete program includes a Main class (Line 1) and an Encrypter class (Line 8). The Encrypter class is annotated with the @Enclave annotation hence it belongs to the enclave partition. Listing 12 and 11 show the partitioned J_E programs. In this phase, the J_E compiler also performs some correctness checks. For example, a class inheriting a class annotated with the @Enclave annotation also needs to be annotated with @Enclave. and the arguments of the gateway methods should implement the java.io.Serializable interface. During execution, arguments to the gateway methods are serialized and passed to the enclave environment. This results into two distinct copies of the parameters and any state changes of the arguments inside the enclave environment are not visible to into the non-enclave environment. We provide a warning at gateway call locations mentioning that changes to the gateway parameters inside the enclave environment are not visible to the non-enclave environment.

Conversion to Jif: Next, the partition to be run inside the enclave is converted into an equivalent Jif [18] program. Jif extends Java with security labels to statically enforce information flow control. A Jif security label is a pair consisting of a confidentiality level and an integrity level.

In the example, Listing 14 is the equivalent Jif program of Listing 12. The non-enclave partition remains unchanged (Listing 13 and Listing 11). Conversion of the J_E program into Jif involves the following steps. (1) J_E secret fields are converted into Jif fields with {Enclave->*;Enclave<-*} label (Listing 12, line 3 and Listing 14, line 2). Such label represents values that are *secret* and *trusted*. (2) For gateway methods, the arguments and the return value are labeled with the Jif's {} and {Enclave->_ ;Enclave<-*} labels respectively. The {} label denotes *public* and *untrusted* information while the {Enclave->_ ;Enclave<-*} label represents the *public* and

trusted information (Listing 14, line 4). (3) The declassify operator in J_E corresponds to the declassify operator in Jif (Listing 14, line 7). The security label propagation mechanism in Jif is identical to the propagation mechanism in J_E . The obtained Jif program is compiled using the Jif compiler to ensure proper label propagation and checking.

Remote Communication: The next step introduces the code required for the communication via Java RMI [36] between the enclave and the non-enclave partition. As in RMI, remote objects are reachable only through an interface, for each class annotated with the @Enclave annotation, the J_E compiler generates an interface implementing java.rmi.Remote with all the the gateway methods. Next, the J_E compiler creates a wrapper class implementing the interface above for each enclave class. This way, all the gateway methods of an enclave class are exposed remotely to the non-enclave environment through the interface. Finally, the method calls to the enclave class from the non-enclave environment are replaced with an RMI lookup that returns a remote reference to the interface of the wrapper class and to a method call to a method of the interface.

For example, in Listing 16, the J_E compiler generates the RemoteEncr interface (Line 2) and the EncrypterWrapper class, which acts as a wrapper for the class Encrypter, and implements the RemoteEncr interface (Line 5). The EncrypterWrapper class defines a method wrapEncrypt which calls the static method Encrypter.encrypt (Line 9). In the non-enclave environment (Listing 15), the J_E compiler transforms the direct callss Encrypter.encrypt(plaintext) to a lookup of enclave remote class (Line 6) and to a call the wrapEncrypt method of the remote interface EncrypterWrapper (Line 8).

Packaging: All the classes to be placed inside and outside the enclave are packaged into two separate executable JAR files. Both the JAR file contain an executable class, which includes code for initialization to set up the RMI registry and to publish remote objects required for communication. The user code executes after the initialization phase is complete.

B. Case Studies

1) Secure Complex Event Processing: We implement a secure distributed event processing (CEP) system based on AdaptiveCEP [73]. In CEP, event sources produce events that can carry data and are processed in a cluster of distributed nodes. AdaptiveCEP adapts the placement of the event processing operators to maximize throughput. In the secure CEP case study, every event is encrypted with the AES [74] before it is sent to the cluster. Event processing nodes decrypt the events inside the enclave, perform the processing, and emit the corresponding output events. Listing 17 shows the code for secure CEP using a *filter* event processor node. The EventProcessor class implements a filter node that can be placed on a remote machine. It contains a gateway method filter (Line 20) that accepts an encrypted event and an encrypted predicate. The event is decrypted inside the enclave (Line 27), the predicate is applied to the event payload, and

Figure 7: Compilation steps

Listing 10: J_E code

```

1 class Main {
2   public static void main(String[] args) {
3     String plaintext = "message";
4     String cipher = Encrypter.encrypt(plaintext);
5   } }
6
7 @Enclave
8 class Encrypter {
9   @Secret private static String key;
10
11   @Gateway
12   public static String encrypt(String plaintext) {
13     String plaintextE = endorse(plaintext);
14     String cipher = encode(plaintext, key);
15     return declassify(cipher);
16   } }

```

Listing 11: Partition outside the enclave

```

1 class Main {
2   public static void main(String[] args) {
3     String plaintext = "message";
4     String cipher = Encrypter.encrypt(plaintext);
5   } }

```

Listing 12: Partition inside the enclave

```

1 @Enclave
2 final class Encrypter {
3   @Secret private static String key;
4
5   @Gateway
6   public static String encrypt(String plaintext) {
7     String plaintextE = endorse(plaintext);
8     String cipher = encode(plaintext, key);
9     return declassify(cipher);
10  } }

```

Listing 13: Jif code outside the enclave

```

1 class Main {
2   public static void main(String[] args) {
3     String plaintext = "message";
4     String cipher = Encrypter.encrypt(plaintext);
5   } }

```

Listing 14: Jif code inside the enclave

```

1 final class Encrypter [principal SGX] authority(SGX) {
2   private String{Enclave->*; Enclave<-*} key;
3
4   public String{Enclave->_; Enclave<-*} encrypt{Enclave<-*}(String{
    plaintext) where authority(Enclave) {
5     String plaintextE = endorse(plaintext, {} to {Enclave<-*});
6     String{Enclave->*; Enclave<-*} cipher = encode(plaintext, key);
7     return declassify (cipher, {Enclave->*; Enclave<-*} to {Enclave->_;
        Enclave<-*});
8   } }

```

Listing 15: Non-enclave partition – communication

```

1
2 class Main{
3   public static void main(String[] args) {
4     String plaintext = "message";
5
6     Remote remoteServer = Naming.lookup("rmi://localhost/RemoteEncr");
7     RemoteEncr remSrvStub = (RemoteEncr) remoteServer;
8     String cipher = remSrvStub.wrapEncrypt(plaintext);
9   } }

```

Listing 16: Enclave partition – communication

```

1
2 interface RemoteEncr extends Remote {
3   public String wrapEncrypt(String plaintext);
4 }
5 class EncrypterWrapper extends UnicastRemoteObject implements
    RemoteEncr {
6
7   @Override
8   public String wrapEncrypt(String plaintext) {
9     return Encrypter.encrypt(plaintext);
10  } }
11 final class Encrypter { ... }

```

the result is returned (Line 28). This way, the event data are protected from an attacker that controls the OS. The attacker can only observe the encrypted events `EncEvent` passed as an argument to the process method.

2) *Secure Calculator*: We show a secure calculator as described in [16]. In this test-case, confidential data is placed inside the enclave and during run time, a user provides a set of operations to be performed on the confidential data in the form tasks from the non-enclave environment. The operations are executed inside the enclave and the result is returned to the non-enclave environment. Listing 18 illustrates tax computation on salary information using the secure calculator.

The `TaskProcessor` class is annotated with the `@Enclave` annotation because it contains a secret field `salary` (Line 14) and a gateway method `process` (Line 17) that must be protected within the enclave. A user submits a list of tax-related computation tasks (`taskList`) to the `process` method (Line 4) where the tasks are executed sequentially inside the enclave. The untrusted input `taskList` is first sanitized (Line 18) to verify that it holds certain integrity properties such as it is not null and not empty. If the sanitization check succeeds, then the `taskList` field is endorsed (21); otherwise, null is returned. When translating to Jif, we use Jif's *checked*

Listing 17: Secure complex event processing

```

1 class Main {
2   public static void main(String[] args) {
3     List<IntEvent> events = Generator.getIntEvents(100);
4     List<EncIntEvent> encEvents = encrypt(events);
5     List<EncIntEvent> result = encEvents.stream().
6       filter(FilterNode.filter).
7       collect(Collectors.toList());
8   } }
9 class EncIntEvent extends EncEvent {
10   private EncInt val; // encrypted integer
11   private String origin;
12 }

13 @Enclave
14 class FilterNode {
15
16   @Secret private static String key;
17   static private Predicate predicate;
18
19   @Gateway
20   public static boolean filter(EncIntEvent event) {
21     if(!sanitize(event)) {
22       return null;
23     }
24     boolean result;
25     try {
26       EncInt encInt = endorse(event).getVal();
27       Integer val = decrypt(encInt, key);
28       result = apply(predicate, val);
29     } catch (Exception e) {
30       return null;
31     }
32     return declassify(result);
33 } }

```

endorsement construct to implement input sanitization. The input sanitization reduces the information leakage by only allowing the verified tasks to interact with the secret fields and influence the returned value. Finally, the result is returned to the non-enclave environment. The case study demonstrates that the code can be easily partitioned using annotations such that the sensitive information is kept inside the enclave, and a user from the non-enclave environment can only observe the result of predefined sensitive operations that are declassified explicitly by the developer.

3) *Battleship Game*: Listing 19 shows a J_E implementation of the Battleship game as in [20]. In the battleship game, each of the two players owns a secret grid. Initially, players position their battleships randomly on the grid. The game then proceeds in rounds and in each round, players guess a position on the opponent's grid. At the end of a round, players are told if the guessed location contains a battleship. The goal is to successfully guess all the battleship locations on the opponent's grid. The game ends when a player guesses all the battleship locations. This case study demonstrates a scenario where we need to declassify some secret information depending on the external inputs. Initially, both grids are secret and in each round, a grid location needs to be made public via declassification. The location to declassify depends on the

guess of the opponent; thus, we need to endorse the opponent's guess to declassify the location. Listing 19 shows the code run on every player's machine. The Main class is placed outside the enclave and it stores the status of the opponent's grid (Line 2). The main method consists of a while loop in which a player sends and receives guesses. The Grid class is placed inside the enclave. The secret variable grid (Line 17) stores the status of a player's grid. The gateway method applyGuess (Line 20) accepts the opponent's guess as an argument and checks if the guessed location contains a battleship. The argument guess has the security label $\langle \mathbb{P}, \mathbb{U} \rangle$ and we use the endorse operator (Line 21) to raise its integrity level to \mathbb{T} . The apply method (Line 25) extracts the array indices from the argument guess and checks if a battleship is present at the array location. The variable result has the security label $\langle \mathbb{S}, \mathbb{T} \rangle$ as it is implicitly influenced by the secret field grid (inside the apply method, when checking for the battleship location). The subsequent declassification operation downgrades the security label of the variable result to $\langle \mathbb{P}, \mathbb{T} \rangle$. Checking the presence of battleship is only allowed if argument guess to the apply method has the integrity level \mathbb{T} , thus to meet this condition, we endorse the argument guess (Line 21). The case study demonstrates the use of the endorse and declassify operators to declassify secret information securely by endorsing the untrusted values explicitly.

Listing 18: Secure calculator

```

1 class Main {
2     public static void main(String[] args) {
3         List<Task> taskList = getTaskSeq("TAX");
4         Double tax = TaskProcessor.process(taskList);
5     }
6 }
7 class Task {
8     public Double run (Double input) {
9         // Task computation
10    }
11 }

12 @Enclave
13 class TaskProcessor {
14     @Secret static Double salary;
15
16     @Gateway
17     public static Double process(List<Task> taskList) {
18         if !(sanitize(taskList)) {
19             return null;
20         }
21         List<Task> taskListE = endorse(taskList);
22         try {
23             Double result = salary;
24             for (Task task : taskList) {
25                 result = task.run(result);
26             }
27         } catch (Exception e) {
28             return null;
29         }
30         return declassify(result);
31     } }

```

Listing 19: Battleship game

```

1 class Main {
2     static boolean[][] gridOpp; // opponent's grid status
3
4     public static void main(String[] args) {
5         boolean gameOver = false;
6         while( !gameOver) {
7             Guess g2 = getGuess(); // opponent's guess
8             updateOppGrid(g2);
9             int result = Grid.applyGuess(g2);
10            // generate and send the guess to the opponent
11            // along with the result of the previous guess
12        }
13    } }

14 @Enclave
15 class Grid {
16
17     @Secret private static boolean[][] grid;
18
19     @Gateway
20     public static int applyGuess(Guess guess) {
21         Guess guessE = endorse(guess);
22         int result = apply(guessE); // ship here?
23         return declassify(result);
24     }
25     private static int apply(Guess guess) {
26         // check for the presence of battleship
27         return result;
28     } }

```

C. Attacker Code

In this section, we present proofs for Lemma 1 and Lemma 2.

Lemma 1. *The attack code definition presented in 1, captures the most powerful attack strategies available to an HAA attacker, who controls the non-enclave environment.*

Proof. Because HAA attacker is in control of the non-enclave environment, he can potentially insert his code anywhere in the program. But the effects of his code is limited to what is presented in 1. Because there are no secret values in non-enclave environment, all of the active attacker's commands executed outside of enclave can eventually have one of the following effects: modifying the parameters passed to a gateway during gateway method calls, or making the program diverge. We don't consider the case with divergence, because in this case the active attacker makes even less observations than the passive attacker.

We justify this argument by structural induction on the possible commands available to HAA attacker:

skip This case is straightforward because it has no effect on the return value of gateways.

$x := e_G$ If x is not one of the parameters passed to a gateway (i.e., \bar{p}), then this assignment has no effect on the values inside of enclave. The case that x is one of gateway parameters is captured by attack $q := e$. We also have to consider the cases when e_G is a method call, either to a gateway or a non-enclave method. *We limit HAA attacker from calling additional gateway methods, and this ability will be captured by HRAA attacker.* On the other hand if e_G is a non-enclave method call, we can simply replace $x := e_G$ with $x := e'$, where e' is the value returned from the method.

$C.f := e_G$ Similar to the assignment case.

if e then S_1 else S_2 Because this command is outside of enclave, e is either public value from σ , or a value set by the attacker. In both cases, the attacker knows its value, thus instead of writing an if statement, he can directly write the command S_1 or S_2 . Moreover, the attacker can only have an effect on the values inside enclave if S_1 or S_2 update the values of gateway parameters. Thus, this case can be reduced to a command of form $q := e$ which captures the intended effect of *if e then S_1 else S_2* .

while e do S This case is similar to the conditional, except that by using a non-terminating loop, the attacker can make the program diverge and limit his own observations. We *ignore* this kind of active attacker because it makes even less observations than the passive attacker.

$S_1; S_2$ we have three cases:

- Neither S_1 nor S_2 modifies the gateway parameter, then this command will have no effect on the enclave, hence the attack code can be ignored.
- Only one of S_1 or S_2 modify the parameters of a gateway method. Without loss of generality, assume S_2 is the command that makes this modification. Using the above cases, we can combine $S_1; S_2$ to a command S of the form $q := e$ that has the same effect as $S_1; S_2$.

- Both of S_1 and S_2 modify gateway parameters. In this case $S_1; S_2$ can be reduced to $q_1 := e; q_2 := e'$, and this structure is presentable by the syntax of a in 1.

□

Lemma 2. *The attacker code defined in 1, captures the most powerful attack strategies available to the HRAA attacker.*

Proof. The HRAA attacker has complete control over the non-enclave environment including code and memory, and he can potentially modify the code outside of the enclave in anyway he wants. Since sensitive data is stored inside the enclave and is only accessible via gateways, the HRAA attacker can only leverage his power by calling the gateways in any order and with any frequency, and the effects of his code is limited to the attack code presented in 1.

The HRAA attacker can execute any J_E program that uses the gateways, but this will not make them more powerful than the attacks presented in 1. The HRAA attacker's commands executed outside of enclave can eventually have one of the following effects: modifying the parameters passed to a gateway during gateway method calls, or making the program diverge.

Because the proof of this lemma is very similar to the proof of Lemma 1, we only investigate some of the most interesting commands available to HRAA attacker, as the rest of them can be reconstructed straightforwardly from the proof of Lemma 1.

$x := e_G$ If x is not one of the parameters passed to a gateway (i.e., \bar{p}), then this assignment has no effect on the values inside of enclave. The case that x is one of gateway parameters is captured by attack $q := e$. We also have to consider the cases when e_G is a method call, either to a gateway or a non-enclave method. HRAA attacker can call additional gateway methods, and the case that e_G is a gateway method call can be captured by Definition 4 and the ability of HRAA attacker to call any gateway methods. On the other hand if e_G is a non-enclave method call and it does not call any gateways, we can argue that since everything outside of enclave is public, the attacker can simply replace $x := e_G$ with $x := e'$, where e' is the value returned from the method. If e_G is a non-enclave method call and it does call some other gateways, the attacker can capture the effect of running $x := e_G$ by simply calling the gateway which was inside e_G .

if e then S_1 else S_2 Since this command is outside of the enclave, e is either a public value, or a value set by the attacker. In both cases, the attacker knows its value, therefore instead of writing the if statement, he can directly write the command S_1 or S_2 , and since S_1 or S_2 can have one or a sequence of gateway calls, The behavior of the if statement can be captured by Definition 4 as one or a sequence of gateway calls.

while e do S In this case, similar to the if statement, the attacker knows the value of e . We should consider 4 cases:

- The loop is finite and S has no gateway calls: The attacker can only have an effect on the values inside enclave if S updates the values of gateway parameters. Thus, this case can be reduced to a command of form $q := e$ which captures the intended effect of *while* e *do* S .
- The loop is finite and S has gateway calls: Which means that it can be represented by Definition 4 as a sequence of gateway calls.
- The loop is infinite: It does not matter whether S has gateway calls or not. Since the execution never terminates and Definition 2 is only defined on terminating programs, we can *ignore* this case.

□

D. Type System Soundness

In this section, we present the proof of Theorem 1. We have to show that if a program is well-typed, then changing the attack vector \vec{a} will not change the attacker's observations. We use events and traces to keep track of these observations. In our setting, only gateway method expressions ($e.m(\vec{p})$) emit an event, and this event is the return value of the gateway method. So, before proving Theorem 1, we present two lemmas about gateway method calls.

Based on the syntax of our language, a gateway method can only be executed in an Assignment or a Store command. First we present a lemma for the assignment:

Lemma 3 (Assignment Gateway Method Calls). *Suppose that $S[\vec{\bullet}]$ is $[\bullet]; x := C.m(\vec{p})$. If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$, Class C is inside enclave, and m is a gateway method, then for all possible attacks \vec{a} and \vec{a}' , if $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$ then $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$.*

Proof. In our system, the body of a method is type-checked in isolation using the rule T-METHOD. Because we are in a Non-Enclave-Enclave context, method m is definitely a gateway, the body of the method no longer has any holes, and all of the parameters of the method (\vec{p}) are now variables with the label $(\langle \mathbb{P}, \mathbb{U} \rangle, F)$.

The event that a method call emits is determined by its return value. In other words, the last command of a method's body is *return*(e) and the value of the expression e is the event. By induction on the structure of S :

skip This case is trivial, as it does not affect the emitted event.

$x := e_G$ By rule T-ASSIGN, the label of variable x is updated according to the label of e_G and current context pc . Thus, if e_G or the context are affected by attacker's inputs, the resulting label will be untrusted \mathbb{U} . Note that e_G can be a method expression, this case will be an enclave-enclave method call, so it will not produce any events, and even though there are no explicit holes in this command, the attacker's inputs can still be passed as one of the method's parameters, so it is possible to have a return value that is untrusted. Additionally, if e_G is affected by some declassified value, it is going to have a True flag, and

by rule T-ASSIGN, this assignment can only happen in trusted contexts.

$C.f := e_G$ This case is very similar to the Assign case, but more limited because of the flow insensitivity of T-STORE. C is a class inside of enclave, so f 's security label is either $\langle \mathbb{P}, \mathbb{T} \rangle$ or $\langle \mathbb{S}, \mathbb{T} \rangle$. In any case, its value can only be updated in trusted context and by trusted values. So, changing the attack vector cannot have any effect on a trusted field's value.

$x := \text{declassify}(e)$ Using T-DECLASSIFY rule, we ensure that expression e should be trusted and it can only be declassified in a trusted context. Thus, the resulting variable x always has $\langle \mathbb{P}, \mathbb{T} \rangle$ label. However, because e is secret, its value can be different in σ_1 and σ_2 . To keep track of the propagation of this declassified value, and control its delayed release, in rule T-DECLASSIFY we make sure x 's flag is True.

return(e) For gateway calls, the rule T-ASSIGN-GATEWAY-CALL uses the precondition $\ell \sqsubseteq \langle \mathbb{P}, \mathbb{U} \rangle$ to ensure the return value's label is public, but based on whether it was affected by the attacker or not, it can be trusted or untrusted. Since the assumption $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma$ implies that $S[\vec{\bullet}]$ is well-typed then e 's label is indeed $\sqsubseteq \langle \mathbb{P}, \mathbb{U} \rangle$. Having this assumption, we continue by induction on the structure of e :

- v : This case is trivial. The return value is independent of the attacks. In other words

$$\begin{aligned} \text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle) &= \text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle) = \\ \text{Trace}(N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle) &= \text{Trace}(N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle) \\ &= v \end{aligned}$$

- x : Following the semantics of ASSIGN ($x := e_G$), the value of x was determined by some expression e_G . We have to consider four cases:
 - e_G security label is $\langle \mathbb{P}, \mathbb{T} \rangle$ and its flag is False: Because e_G value was not affected by any declassified value, and because the security label is public and trusted, following the initial indistinguishability of states $\sigma_1 =_A \sigma_2$, and the fact that trusted values are not effected by attacker inputs, we can easily conclude e_G evaluates to the same value in $\langle S[\vec{a}], \sigma_1 \rangle$, $\langle S[\vec{a}'], \sigma_1 \rangle$, $\langle S[\vec{a}], \sigma_2 \rangle$, and $\langle S[\vec{a}'], \sigma_2 \rangle$.
 - e_G security label is $\langle \mathbb{P}, \mathbb{T} \rangle$ and its flag is True: e_G is affected by declassified values, and because its security label is trusted, we know that it was not affected by the attacker's input. Obviously, the declassified values might be different in σ_1 and σ_2 , however, by the initial indistinguishability of states $\sigma_1 =_A \sigma_2$ and the assumption $\text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle)$, we can conclude that the declassified values were equal in σ_1 and σ_2 . Thus it is straightforward to deduce

that e_G will evaluate to equal values in $\langle S[\vec{a}], \sigma_1 \rangle$, $\langle S[\vec{a'}], \sigma_1 \rangle$, $\langle S[\vec{a}], \sigma_2 \rangle$, and $\langle S[\vec{a'}], \sigma_2 \rangle$.

- e_G security label is $\langle \mathbb{P}, \mathbb{U} \rangle$ and its flag is False: e_G value was not affected by any declassified value, it was just affected by the public values and the attacker's inputs. So following the initial indistinguishability of states $\sigma_1 =_A \sigma_2$ and the assumption $\text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle)$, we can show that $\text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_2 \rangle)$ will also hold. This is because the public values are unchanged, and the only difference is in the attacker's input, which is the same in both $\langle S[\vec{a'}], \sigma_1 \rangle$ and $\langle S[\vec{a'}], \sigma_2 \rangle$. Thus, even though the return value is *possibly* different, it is still going to be the same in $\langle S[\vec{a'}], \sigma_1 \rangle$ and $\langle S[\vec{a'}], \sigma_2 \rangle$, hence $\text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_2 \rangle)$.
- e_G security label is $\langle \mathbb{P}, \mathbb{U} \rangle$ and its flag is True: This means that e_G was effected by both attacker input and declassified values. However, by rules T-ASSIGN and T-STORE we know that after getting affected by declassified values, it was only affected by untrusted values directly (*not* through a condition or loop). Thus by assumption $\text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle)$ we can learn that the declassified value is the same in σ_1 and σ_2 . The rest is similar to the previous case, when the attacker's input changes, $\langle S[\vec{a'}], \sigma_1 \rangle$ and $\langle S[\vec{a'}], \sigma_2 \rangle$ might produce *possibly* different yet still equal events. Thus $\text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_2 \rangle)$ will still hold.

- $C.f$: Following the semantics of STORE ($C.f := e_G$), the value of $C.f$ was determined by some expression e_G . This case is easier than the assignment because C is a class inside of enclave, so the security label of its field f can be either $\langle \mathbb{P}, \mathbb{T} \rangle$ or $\langle \mathbb{S}, \mathbb{T} \rangle$. By rule T-ASSIGN-GATEWAY-CALL we already know that the return value has a label $\sqsubseteq \langle \mathbb{P}, \mathbb{U} \rangle$, thus $C.f$'s label is definitely $\langle \mathbb{P}, \mathbb{T} \rangle$. This label means that the value stored in $C.f$ is not affected by the attacker's input, so following the initial indistinguishability of states $\sigma_1 =_A \sigma_2$, and the assumption $\text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle)$, we can conclude that $\text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_2 \rangle)$ will also hold.

- $e_1 \oplus e_2$: By induction on the structure of e_1 and e_2 , we see that they can be one of the above cases. Thus it is straightforward to show that $\text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_2 \rangle)$ will hold.

if e then S_1 else S_2 We have four cases:

- e is *public and trusted*: Because e is public and trusted it evaluates to equal values in $\langle S[\vec{a}], \sigma_1 \rangle$ and $\langle S[\vec{a'}], \sigma_1 \rangle$ and (in case of declassification) to *possibly*

different yet still equal values in $\langle S[\vec{a}], \sigma_2 \rangle$ and $\langle S[\vec{a'}], \sigma_2 \rangle$. Therefore, by assumption $\text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle)$, it's easy to conclude that the if statement takes the same branch and $\text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_2 \rangle)$ also holds.

- e is *public and untrusted*: In this case neither S_1 nor S_2 can have any declassification. A public value evaluates to the same values in σ_1 and σ_2 , however because e is also untrusted, it can also be affected by attacker's input. However because attacker's input only changes when the attack vector changes, we can easily conclude that by having the initial indistinguishability of states $\sigma_1 =_A \sigma_2$, e will evaluate to equal values in $\langle S[\vec{a}], \sigma_1 \rangle$ and $\langle S[\vec{a'}], \sigma_2 \rangle$ and (based on the attacker's input) to *possibly different yet still equal* values in $\langle S[\vec{a'}], \sigma_1 \rangle$ and $\langle S[\vec{a'}], \sigma_2 \rangle$. So it is possible for the if condition to take different branches when the attacker vector changes. Now because e is untrusted, pc 's security label also becomes untrusted, rules T-ASSIGN and T-STORE, prevent the assignment or store of declassified values in the body of the if statement. So the values that depend on declassified values cannot be affected by this condition. (ultimately this is what we wanted to achieve by using flags, to disallow attacker input to control the branch that can indirectly affect the final declassification value). From here on it is straightforward to show that by the initial indistinguishability of states $\sigma_1 =_A \sigma_2$ and the assumption $\text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle)$, $\text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_2 \rangle)$ will also hold.
- e is *secret and trusted*: By rule T-DECLASSIFY neither S_1 nor S_2 can have any declassification. Any assignment or store inside S_1 and S_2 is going to be secret, and have to be declassified in order to be able to affect the return value. We can use this intuition, the initial indistinguishability of states $\sigma_1 =_A \sigma_2$, and the assumption $\text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle)$ to conclude that $\text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_1 \rangle) = \text{Trace}(N \vdash_\delta \langle S[\vec{a'}], \sigma_2 \rangle)$ holds.
- e is *secret and untrusted*: This case is similar to the last one, but by rule T-STORE any store command is not possible in S_1 and S_2 and because the result of any assignment inside S_1 and S_2 is going to be secret and untrusted, it cannot be declassified later on, so it can never effect the return value.

while e do S This case is similar to the case for conditionals. The only difference is that in this case, untrusted attacker inputs can cause the program to diverge, and by definition of trace indistinguishability, we consider non-terminating traces indistinguishable.

$S_1; S_2$ We have two cases:

- $S_1; \text{return}(e)$: S_1 can be any of the above commands except $\text{return}(e)$. Additionally, we can use an argument similar to the one we used for $\text{return}(e)$ case to show that the return_value will be the same for programs $S[\vec{a}]$ and $S[\vec{a}']$.
- $S_1; S_2$: This case is fairly straightforward, because neither S_1 nor S_2 have a return command, they do not affect the trace.

□

We also need another Lemma for calling a gateway method through the STORE command:

Lemma 4 (Store Gateway Method Calls). *Suppose that $S[\vec{\bullet}]$ is $[\bullet]; C'.f := C.m(\vec{p})$. If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$, Class C' is outside the enclave, C is inside enclave, and m is a gateway method, then for all possible attacks \vec{a} and \vec{a}' , if $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$ then $N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$.*

Proof. Similar to Lemma 3. □

Now, using Lemma 3 and Lemma 4, we can prove Theorem 1.

Theorem 1. *If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$ then $S[\vec{\bullet}]$ satisfies robustness under HAA.*

Proof of Theorem 1. Following the approach of [20], suppose that for the well-typed program $S[\vec{a}]$, and states $\sigma_1 =_A \sigma_2$, we have $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$. We need to show that for all possible attacks \vec{a}' , the execution indistinguishability holds. Formally:

$$\forall \vec{a}'. N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$$

Execution is in non-enclave environment, and the security label of all variables, fields and classes outside of enclave is $\langle \mathbb{P}, \mathbb{U} \rangle$. It is worth noting that the values returned from gateways can be $\langle \mathbb{P}, \mathbb{T} \rangle$, but because the security label of pc outside of enclave is $\langle \mathbb{P}, \mathbb{U} \rangle$, the final label of the variable holding this returned value will be $\langle \mathbb{P}, \mathbb{T} \rangle \sqcup \langle \mathbb{P}, \mathbb{U} \rangle = \langle \mathbb{P}, \mathbb{U} \rangle$, furthermore, since all of the fields outside of enclave are $\langle \mathbb{P}, \mathbb{U} \rangle$, by sub-typing rule we can store a $\langle \mathbb{P}, \mathbb{T} \rangle$ inside it, but it will be treated as a $\langle \mathbb{P}, \mathbb{U} \rangle$ after that.

We present the proof by structural induction on $S[\vec{\bullet}]$. It is worth noting that we will not consider $x := \text{declassiy}(e)$ command, because the semantics of DECLASSIFY ensures that this command is only executed inside of enclave.

If $S[\vec{\bullet}]$ is Skip then we are done. Because it does not have a hole. So we continue by structural induction for the other cases of $S[\vec{a}]$:

$[\vec{a}]; x := e_G$: If e_G is not a method expression, then it will not have any effect on the trace, and we are done. If e_G is a method expression $e.m(\vec{p})$, we have to consider two cases: If m is a method outside of enclave it does not produce any events, so the trace remains unchanged. If m is a gateway method, by Lemma 3, we already know that its returned value (event) will be indistinguishable

for all attacks, thus $\forall \vec{a}'. N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$.

$[\vec{a}]; C'.f := e_G$: Similar to the previous case. If e_G is not a method expression or m is outside of enclave, we are done. Otherwise, if e_G is a method expression $e.m(\vec{p})$ and m is a gateway method, by Lemma 4, we know that the return values (events) are indistinguishable for all attacks, thus $\forall \vec{a}'. N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$ will hold.

if e then $S_1[\vec{a}_1]$ else $S_2[\vec{a}_2]$: If $S_1[\vec{a}_1]$ and $S_2[\vec{a}_2]$ does not have gateway calls, then executing this command will have no effect on the trace. So we only consider cases in which at least one of $S_1[\vec{a}_1]$ or $S_2[\vec{a}_2]$ has a gateway method call.

By the fact that the execution is outside of enclave we know that e is public. By using the initial indistinguishability of states $\sigma_1 =_A \sigma_2$, and the results of Lemmas 3 and 4, we can infer that e will evaluate to the same value v in both configurations $\langle e, \sigma_1 \rangle \Downarrow \langle v, \sigma_1 \rangle$ and $\langle e, \sigma_2 \rangle \Downarrow \langle v, \sigma_2 \rangle$. Therefore, the conditional takes the same branch in both configurations.

Without loss of generality, assume that v is non-zero, then configurations $\langle S[\vec{a}], \sigma_1 \rangle$ and $\langle S[\vec{a}], \sigma_2 \rangle$ will be reduced to $\langle S_1[\vec{a}_1], \sigma_1 \rangle$ and $\langle S_1[\vec{a}_1], \sigma_2 \rangle$, respectively. Similarly, $\langle S[\vec{a}'], \sigma_1 \rangle$ will reduce to $\langle S_1[\vec{a}_1], \sigma_1 \rangle$ and $\langle S[\vec{a}'], \sigma_2 \rangle$ will reduce to $\langle S_1[\vec{a}_1], \sigma_2 \rangle$. By assumption $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$ we have $N \vdash_\delta \langle S_1[\vec{a}_1], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S_1[\vec{a}_1], \sigma_2 \rangle$. So, by induction hypothesis we can obtain $N \vdash_\delta \langle S_1[\vec{a}_1], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S_1[\vec{a}_1], \sigma_2 \rangle$. Which in turn implies that $N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$.

while e do $S_1[\vec{a}]$: We only consider the case in which $S_1[\vec{a}]$ has a gateway method call. By an argument similar to the conditional, we can show e will evaluate to the same value v in both σ_1 and σ_2 . If v is zero, then both configurations $\langle S_1[\vec{a}], \sigma_1 \rangle$ and $\langle S_1[\vec{a}], \sigma_2 \rangle$ will terminate without making any changes to the trace. Thus $N \vdash_\delta \langle S_1[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S_2[\vec{a}], \sigma_2 \rangle$ will hold. If v is non-zero, then by $N \vdash_\delta \langle \text{while } e \text{ do } S_1[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle \text{while } e \text{ do } S_1[\vec{a}], \sigma_2 \rangle$, it is necessary that $N \vdash_\delta \langle S_1[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S_1[\vec{a}], \sigma_2 \rangle$ holds. Then by induction hypothesis we have $N \vdash_\delta \langle S_1[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S_1[\vec{a}], \sigma_2 \rangle$. If either $\langle S_1[\vec{a}], \sigma_1 \rangle$ or $\langle S_1[\vec{a}], \sigma_2 \rangle$ diverges, then its corresponding loop $\langle \text{while } e \text{ do } S_1[\vec{a}], \sigma_1 \rangle$ or $\langle \text{while } e \text{ do } S_1[\vec{a}], \sigma_2 \rangle$ also diverges. Then by definition of execution indistinguishability, $N \vdash_\delta \langle S_1[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S_1[\vec{a}], \sigma_2 \rangle$ holds. On the other hand, if both $\langle S_1[\vec{a}], \sigma_1 \rangle$ and $\langle S_1[\vec{a}], \sigma_2 \rangle$ terminate, it is straightforward to show that the gateway calls inside of $S_1[\vec{a}']$ are issued the same number of times under σ_1 and σ_2 . By Lemmas 3 and 4, we already know that events resulted from these gateway calls are indistinguishable, thus the traces of $N \vdash_\delta \langle S_1[\vec{a}'], \sigma_1 \rangle$

and $N \vdash_\delta \langle S_1[\vec{a}_1], \sigma_2 \rangle$ will be indistinguishable, meaning that $N \vdash_\delta \langle S_1[\vec{a}_1], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S_1[\vec{a}_1], \sigma_2 \rangle$ holds. $S_1[\vec{a}_1]; S_2[\vec{a}_2]$: We only consider the cases in which $S_1[\vec{a}_1]$ and $S_2[\vec{a}_2]$ have gateway method calls. By induction on structure of $S_1[\vec{a}_1]$ we can show that it either diverges or terminates. Similar to the case of while loop, if it diverges in at least one of the configurations, by definition $N \vdash_\delta \langle S[\vec{a}_1], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}_1], \sigma_2 \rangle$ holds. If it terminates in both configurations, we can use the result of Lemmas 3 and 4 to infer that these execution will result in indistinguishable traces. Thus $N \vdash_\delta \langle S_1[\vec{a}_1], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S_1[\vec{a}_1], \sigma_2 \rangle$ will hold. If $S_1[\vec{a}_1]$ terminates in both configurations, we can consider $S_2[\vec{a}_2]$. Imagine $\langle S_1[\vec{a}_1], \sigma_1 \rangle \Downarrow_t \mathcal{M}'_1, \mathcal{H}'_1$ and $\langle S_1[\vec{a}_1], \sigma_2 \rangle \Downarrow_t \mathcal{M}'_2, \mathcal{H}'_2$. By the result of Lemmas 3 and 4 and initial state indistinguishability, it is fairly straightforward to show that $\mathcal{M}'_1 =_A \mathcal{M}'_2$ and $\mathcal{H}'_1 =_A \mathcal{H}'_2$. Now, a similar argument can be used for $S_2[\vec{a}_2]$ to infer that $N \vdash_\delta \langle S_2[\vec{a}_2], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S_2[\vec{a}_2], \sigma_2 \rangle$ also holds. \square

We can present the proof of Theorem 2 similar to the proof of Theorem 1. First, we change Lemma 3 to account for the *untrusted* label of the values in shared state Σ .

Lemma 5 (Gateway Method Calls). *Suppose that $S[\vec{\bullet}]$ is $[\bullet]; x := C.m(\vec{p})$. If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$ w.r.t HRAA attacker, Class C is inside enclave, and m is a gateway method, then for all possible attacks \vec{a} and \vec{a}' , if $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$ then $N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$.*

Notice that the definition of Lemma 5 is similar to Lemma 3. This is because we are not considering any ordering of gateway calls yet, we are just using this Lemma to show that the indistinguishability of gateways' return values is independent of the *attacker input* and the *shared state*.

Only Fields can act as shared states, since variables are local. In rule T-STORE, $pc_d \vee d_2 = T \Rightarrow pc_\ell \sqcup \ell_1 = \langle -, \mathbb{T} \rangle$ ensures that the value of shared state (which is untrusted) cannot directly or indirectly be affected by declassified values.

To prove this Lemma we just have to extend the definition of untrusted value to attacker inputs and all of the fields in Σ . Gateway return value's independence from attacker input can be shown from the change in attack vector a' , and its independence of shared state can be proved from the public indistinguishability of states σ_1 and σ_2 , and the untrustedness of shared state. We omit this proof because its structure is similar to the proof presented for Lemma 3.

Now, using this lemma we can prove Theorem 2.

Theorem 2. *If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$ with regard to Σ and G^D , then $S[\vec{\bullet}]$ satisfies robustness under HRAA.*

Proof of Theorem 2. Suppose that for the well-typed program with fixed order of execution $S[\vec{a}]$, and states $\sigma_1 =_A \sigma_2$, we

have $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$. We need to show that in all possible ordering and frequency of gateways calls, and for all possible values for the parameters of gateways, the execution indistinguishability holds. Formally:

$$\forall S', a'. N \vdash_\delta \langle S'[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a}'], \sigma_2 \rangle$$

Using Lemma 5, by induction on the structure of S' :

$[\vec{a}']; x := C.m(\vec{p})$: If $C.m(\vec{p})$ was a gateway method originally called in $S[\vec{\bullet}]$, $N \vdash_\delta \langle S'[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a}'], \sigma_2 \rangle$ is going to be a direct result of Lemma 7, which shows that the indistinguishability of gateways' return values (and its resulting event) is independent of the value of parameters passed to them by the attacker and the shared state.

However, it is also possible that $C.m(\vec{p})$ is a gateway method, not originally called in $S[\vec{\bullet}]$. Because of the preprocessing step of G^D we can conclude that $C.m(\vec{p})$ does not have any *new* declassifications, by initial indistinguishability of states σ_1 and σ_2 , and equality of attack vector and shared state on both sides, we can conclude that $N \vdash_\delta \langle S'[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a}'], \sigma_2 \rangle$ holds.

$S'_1[\vec{a}_1]; S'_2[\vec{a}_2]$: We know that both $S'_1[\vec{a}_1]$ and $S'_2[\vec{a}_2]$ have gateway calls. By Rule T-DECLASSIFY we know that untrusted values cannot have any effect on the decision to declassify. Therefore, we can conclude that changing the order or frequency of gateway calls cannot lead to any *new* declassifications. By the assumption $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$, and the fact that any of the gateway calls in $S'_1[\vec{a}_1]$ and $S'_2[\vec{a}_2]$ that had a declassification were also in $S[\vec{a}]$ (because G^D was a subset of it), we conclude that declassified secrets had the same value in σ_1 and σ_2 .

We know that the return value of gateway calls can depend on declassified secrets (which are equal in σ_1 and σ_2), attacker input (which is the same (\vec{a}') for both programs), public values (that are equal in both states), and the values in shared state. Of these values, only shared state can be different between $S[\vec{\bullet}]$ and $S'[\vec{\bullet}]$. We know that they cannot effect declassifications, we also know by the definition of shared state that their value might change by the order of execution, resulting in different traces for $S'_1[\vec{a}_1]; S'_2[\vec{a}_2]$ compared to any other order of calling these gateways. However, since their value is determined by the order of execution, we are going to get similar traces for each possible order of execution. Thus, we can conclude that $N \vdash_\delta \langle S'_1[\vec{a}_1]; S'_2[\vec{a}_2], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'_1[\vec{a}_1]; S'_2[\vec{a}_2], \sigma_2 \rangle$ will always hold no matter the order of gateway method calls in $S'_1[\vec{a}_1]$ and $S'_2[\vec{a}_2]$.

Using these cases it is straightforward to conclude that the indistinguishability of traces will be preserved $\forall S', a'. N \vdash_\delta \langle S'[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a}'], \sigma_2 \rangle$. \square

E. Type Soundness for Endorsement

In this section, we present the proof of Theorem 3 which follows a similar structure as the proof of Theorem 1. We just have to extend Lemma 3 and Lemma 4 to account for relevant attacks. We revisit Lemma 3:

Lemma 6 (Assignment Gateway Method Calls). *Suppose that $S[\vec{\bullet}]$ is $[\bullet]; x := C.m(\vec{p})$. If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$, Class C is inside enclave, and m is a gateway method, then for all possible attacks \vec{a} and \vec{a}' such that $\vec{a}' \in \Omega(S[\vec{a}], \sigma_1)$, if $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$ then $N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$.*

All of the parameters of the method (\vec{p}) are now variables in the body of method m with the label $(\mathbb{P}, \mathbb{U}, F)$. Let's continue by induction on the structure of S for some of the interesting cases:

$x := e_G$ Following the rule T-ASSIGN, the label of variable x is updated according to the label of e_G and current context pc . Thus, if e_G or the context are affected by attacker's inputs, the resulting label will be untrusted \mathbb{U} . However, now it is possible that e_G or context were endorsed. In this case we can use the fact that a' is a relevant attacks to infer that the endorsed values will be the same in $S[\vec{a}]$ and $S[\vec{a}']$. Additionally, if e_G was affected by some declassified value it is going to have a True flag, and if it eventually affects the return value of the method, we can use $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$ to conclude the declassified secrets were equal in states σ_1 and σ_2 .

$x := \text{endorse}_\eta(e)$ Rule T-ENDORSE ensures that e can only be endorsed in a trusted context. Using this rule, combined with assumption $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$ and the fact the a' is an attack relevant to a , allows us to conclude that the result of endorsement will be equal under a and a' .

$x := \text{declassify}(e)$ Rule T-DECLASSIFY ensures that expression e should be trusted and it can only be declassified in a trusted context. But now that we have endorsement, it is possible that e or the context were affected by attacker's input. However, because a' is a relevant attack, we can be sure that endorsed values are the same under a and a' . Moreover, because e can be secret, its value can be different in σ_1 and σ_2 . To keep track of the propagation of this declassified value, and control its delayed release, rule T-DECLASSIFY makes x 's flag True. Now if e is eventually returned, we can conclude the result of declassification was equal under a and a' , so the indistinguishability of traces $N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$ will also hold.

We should also revisit Lemma 4, but its proof is going to be similar to that of Lemma 6. Utilizing these Lemmas, the proof of Theorem 3 will be completely similar to the proof of Theorem 1.

The proof of Theorem 4 is similar to the proof of Theorem 2, with two main differences:

- 1) Variables and fields in shared state are untrusted, and therefore can be endorsed.
- 2) It is possible that $S'[\vec{\bullet}]$ lead to additional endorsement.

Similar to the proof of Theorem 2, we propose a lemma to check the indistinguishability of single gateway calls.

Lemma 7 (Gateway Method Calls). *Suppose that $S[\vec{\bullet}]$ is $[\bullet]; x := C.m(\vec{p})$. If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$ w.r.t HRAA attacker, Class C is inside enclave, and m is a gateway method, then for all possible attacks \vec{a} and \vec{a}' such that $\vec{a}' \in \Omega_R(S[\vec{a}], S[\vec{\bullet}], \sigma_1)$, if $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$ then $N \vdash_\delta \langle S[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}'], \sigma_2 \rangle$.*

Notice that in Lemma 7, in the definition of relevant attacks $S'[\vec{\bullet}]$ is equal to $S[\vec{\bullet}]$. This is because we are not considering any ordering of gateway calls yet. However, the values of attacker input and *untrusted* shared state can be endorsed, and affect the return value. nevertheless, because a' is a relevant attack, we know that the endorsed value are the same in $S[\vec{a}]$ and $S[\vec{a}']$, regardless of whether they originated from the attacker inputs or the shared state. We omit its proof because it is similar to the proof presented for Lemma 6.

Now, using this lemma we can prove Theorem 4.

Theorem 4. *If $pc, \Gamma, \Pi \vdash_\delta S[\vec{\bullet}] : \Gamma', \Pi'$ with regard to Σ and G^D , then $S[\vec{\bullet}]$ satisfies robustness under HRAA with endorsement.*

Proof. Suppose that for the well-typed program with fixed order of execution $S[\vec{a}]$, and states $\sigma_1 =_A \sigma_2$, we have $N \vdash_\delta \langle S[\vec{a}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a}], \sigma_2 \rangle$. We need to show that in all possible ordering and frequency of gateways calls, and for all relevant attacks a' , the execution indistinguishability holds. Formally:

$$\forall S', a' \in \Omega_R(S[\vec{a}], S'[\vec{\bullet}], \sigma_1).$$

$$N \vdash_\delta \langle S'[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a}'], \sigma_2 \rangle$$

Using Lemma 7, by induction on the structure of S' :

$[\vec{a}']; x := C.m(\vec{p})$: If $C.m(\vec{p})$ was a gateway method originally called in $S[\vec{\bullet}]$, $N \vdash_\delta \langle S'[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a}'], \sigma_2 \rangle$ is going to be a direct result of Lemma 7. However, it is also possible that $C.m(\vec{p})$ is a gateway method, not originally called in $S[\vec{\bullet}]$. Because of the preprocessing step of G^D we can conclude that $C.m(\vec{p})$ does not have any *new* declassifications, but it is still possible for it to have some *new* endorsements. We consider two cases:

- If $C.m(\vec{p})$ did not had new endorsement, by initial public indistinguishability of states σ_1 and σ_2 , and equality of attack vector and shared state on both sides, we can conclude that $N \vdash_\delta \langle S'[\vec{a}'], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a}'], \sigma_2 \rangle$ holds.
- If it had any *new* endorsement, this new endorsed value is either from the shared state or attacker input. Nevertheless, shared state is equal in both sides

because we do not have any reordering of gateways, attacker input is also equal because both sides use a' . These facts along with the public indistinguishability of states σ_1 and σ_2 , allows us to show that $N \vdash_\delta \langle S'[\vec{a'}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a'}], \sigma_2 \rangle$ also holds in this case.

$S'_1[\vec{a'_1}]; S'_2[\vec{a'_2}]$: We have to consider two cases:

- $S'_1[\vec{a'_1}]; S'_2[\vec{a'_2}]$ calls *no new gateways*: In this case different order and frequency of execution can change the shared state. We know that values of the shared state are untrusted, but they can be endorsed. However, because a' is a relevant attack w.r.t $S'[\vec{\bullet}]$ and $S[\vec{\bullet}]$, any endorsed shared state is going to have the same value in $S[\vec{\bullet}]$ and $S'[\vec{\bullet}]$. This allows us to conclude that even though the values of shared state can be different between $S[\vec{\bullet}]$ and $S'[\vec{\bullet}]$, the shared state values that are endorsed and can affect declassifications are equal. The rest of the proof for this case is similar to theorem 2.

By Rule T-DECLASSIFY and the preprocessing step of G^D we know that changing the order or frequency of gateway calls cannot lead to any *new* declassifications, and by the assumption $N \vdash_\delta \langle S[\vec{a'}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S[\vec{a'}], \sigma_2 \rangle$, and the fact that any of the gateway calls in $S'_1[\vec{a'_1}]$ and $S'_2[\vec{a'_2}]$ that had a declassification were also in $S[\vec{a'}]$ (because G^D was a subset of it), we can conclude that declassified secrets had the same value in σ_1 and σ_2 .

The return value of gateway calls can depend on declassified secrets (which we shown are equal in σ_1 and σ_2), attacker input (which is the same (a') for both programs), public values (that are equal in both states), and the values in shared state. Of these values, only shared state can be different between $S[\vec{\bullet}]$ and $S'[\vec{\bullet}]$, resulting in different traces for $S'_1[\vec{a'_1}]; S'_2[\vec{a'_2}]$ compared to any other order of calling these gateways.

However, since the value of shared state is determined by the order of execution, we are going to get similar traces for each possible order of execution.

- $S'_1[\vec{a'_1}]; S'_2[\vec{a'_2}]$ calls a new gateway not originally called in $S[\vec{\bullet}]$. We have to consider two cases:
 - They have *no* new endorsements: The return value of gateway call can depend on attacker input (which is the same (a') for both programs), public values (that are equal in both states), and the values in shared state which can be different between $S[\vec{\bullet}]$ and $S'[\vec{\bullet}]$. However, since the value of shared state is determined by the order of execution, we are going to get similar traces for each possible order of execution. Thus $N \vdash_\delta \langle S'[\vec{a'}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a'}], \sigma_2 \rangle$ holds.
 - $S'[\vec{\bullet}]$ leads to new endorsements: By G^D we know that these new endorsements cannot lead to any new declassifications. Additionally, these new endorsed values can either be from the shared state or attacker input. Shared state is equal in both sides because we have similar ordering of gateways in $S'_1[\vec{a'_1}]; S'_2[\vec{a'_2}]$, attacker input is also equal because both sides use a' . These facts along with the public indistinguishability of states σ_1 and σ_2 , allows us to show that $N \vdash_\delta \langle S'[\vec{a'}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a'}], \sigma_2 \rangle$ also holds in this case.

Therefore, we can conclude that $N \vdash_\delta \langle S'_1[\vec{a'_1}]; S'_2[\vec{a'_2}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'_1[\vec{a'_1}]; S'_2[\vec{a'_2}], \sigma_2 \rangle$ will always hold no matter the order of gateway method calls in $S'_1[\vec{a'_1}]$ and $S'_2[\vec{a'_2}]$.

Using these cases it is straightforward to conclude that for all S' and relevant attacks a' the indistinguishability of traces will be preserved $N \vdash_\delta \langle S'[\vec{a'}], \sigma_1 \rangle \approx_A N \vdash_\delta \langle S'[\vec{a'}], \sigma_2 \rangle$ \square

$$\begin{array}{c}
\text{INT} \frac{}{\alpha \vdash_{\delta} \langle n, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle n, \mathcal{M}, \mathcal{H} \rangle} \quad \text{UNIT} \frac{}{\alpha \vdash_{\delta} \langle \text{unit}, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle \text{unit}, \mathcal{M}, \mathcal{H} \rangle} \quad \text{CLASS} \frac{}{\alpha \vdash_{\delta} \langle C, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle C, \mathcal{M}, \mathcal{H} \rangle} \\
\\
\text{VAR} \frac{\alpha = \delta(x) \quad v = \mathcal{M}(x)}{\alpha \vdash_{\delta} \langle x, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle v, \mathcal{M}, \mathcal{H} \rangle} \quad \text{FIELD ACCESS} \frac{\alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle C, \mathcal{M}, \mathcal{H} \rangle \quad \alpha = \delta(C) \quad v = \mathcal{H}(C, f)}{\alpha \vdash_{\delta} \langle e.f, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle v, \mathcal{M}, \mathcal{H} \rangle} \\
\\
\text{OP} \frac{\alpha \vdash_{\delta} \langle e_1, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle n_1, \mathcal{M}, \mathcal{H} \rangle \quad \alpha \vdash_{\delta} \langle e_2, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle n_2, \mathcal{M}, \mathcal{H} \rangle \quad n' = n_1 \oplus n_2}{\alpha \vdash_{\delta} \langle e_1 \oplus e_2, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle n', \mathcal{M}, \mathcal{H} \rangle} \\
\\
\text{METHOD} \frac{\begin{array}{c} \alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle C, \mathcal{M}, \mathcal{H} \rangle \quad \text{method}_{\phi}(\bar{p})\{S; \text{return}(e)\} = \text{getMethod}(C, m) \\ \left((\alpha = \delta(C)) \vee (\alpha = N \wedge \delta(C) = E \wedge \phi = G) \right) \quad \mathcal{M}^* = \mathcal{M}[p_i \mapsto \sigma(q_i)] \quad i = 1, \dots, |\bar{p}| \quad \delta(C) \vdash_{\delta} \langle S; \text{return}(e), \mathcal{M}^*, \mathcal{H} \rangle \Downarrow \mathcal{M}^{*'}, \mathcal{H}' \triangleright v \\ \mathcal{M}' = \mathcal{M}^{*'} \setminus [p_i] \quad i = 1, \dots, |\bar{p}| \quad (\alpha = \delta(C) \Rightarrow \beta = \epsilon) \quad (\alpha = N \wedge \delta(C) = E \wedge \phi = G \Rightarrow \beta = v) \end{array}}{\alpha \vdash_{\delta} \langle e.m(\bar{q}), \mathcal{M}, \mathcal{H} \rangle \Downarrow_{\beta} \langle v, \mathcal{M}', \mathcal{H}' \rangle}
\end{array}$$

Figure 8: Large-step semantics for J_E expressions

$$\begin{array}{c}
\text{T-UNIT} \frac{}{\Gamma, \Pi \vdash_{\delta} \text{unit} : (\langle \mathbb{P}, \mathbb{T} \rangle, F)} \quad \text{T-INT} \frac{}{\Gamma, \Pi \vdash_{\delta} n : (\langle \mathbb{P}, \mathbb{T} \rangle, F)} \quad \text{SUB TYPING} \frac{\Gamma, \Pi \vdash_{\delta} e : (\ell, d) \quad \ell \sqsubseteq \ell' \quad d \sqsubseteq d'}{\Gamma, \Pi \vdash_{\delta} e : (\ell', d')} \\
\\
\text{T-CLASS} \frac{\Gamma(C) = \ell \quad \Pi(C) = d}{\Gamma, \Pi \vdash_{\delta} C : (\ell, F)} \quad \text{T-VAR} \frac{\Gamma(x) = \ell \quad \Pi(x) = d}{\Gamma, \Pi \vdash_{\delta} x : (\ell, d)} \quad \text{T-FIELD ACCESS} \frac{\Gamma(C.f) = \ell \quad \Pi(C.f) = d}{\Gamma, \Pi \vdash_{\delta} C.f : (\ell, d)} \\
\\
\text{T-OP} \frac{\Gamma, \Pi \vdash_{\delta} e_1 : (\ell_1, d_1) \quad \Gamma, \Pi \vdash_{\delta} e_2 : (\ell_2, d_2)}{\Gamma, \Pi \vdash_{\delta} e_1 \oplus e_2 : (\ell_1 \sqcup \ell_2, d_1 \vee d_2)} \\
\\
\text{T-METHOD} \frac{\text{method}_{\phi}(\bar{p})\{S; \text{return}(e)\} = \text{getMethod}(C, m) \quad pc', \Gamma^-, \Pi^- \vdash_{\delta} S; \text{return}(e) : \Gamma^+, \Pi^+ \triangleright (\ell, d)}{\Gamma, \Pi \vdash_{\delta} C.m(\bar{p}) : \left(\Gamma^-, \Pi^- \xrightarrow{pc'} \Gamma^+, \Pi^+ \right)_{(\ell, d)}}
\end{array}$$

Figure 9: Typing rules for expressions

$$\begin{array}{c}
\text{SKIP} \frac{}{\alpha \vdash_{\delta} \langle \text{skip}, \mathcal{M}, \mathcal{H} \rangle \Downarrow \mathcal{M}, \mathcal{H}} \\
\text{ASSIGN} \frac{\alpha \vdash_{\delta} \langle e_G, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{\beta} \langle v, \mathcal{M}', \mathcal{H}' \rangle \quad \alpha = \delta(x) \quad \mathcal{M}'' = \mathcal{M}'[x \mapsto v]}{\alpha \vdash_{\delta} \langle x := e_G, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{\beta} \mathcal{M}'', \mathcal{H}'} \\
\text{STORE} \frac{\alpha \vdash_{\delta} \langle e_G, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{\beta} \langle v, \mathcal{M}', \mathcal{H}' \rangle \quad \alpha = \delta(C) \quad \mathcal{H}'' = \mathcal{H}'[(C, f) \mapsto v]}{\alpha \vdash_{\delta} \langle C.f := e_G, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{\beta} \mathcal{M}', \mathcal{H}''} \\
\text{SEQ} \frac{\alpha \vdash_{\delta} \langle S_1, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{t_1} \mathcal{M}_1, \mathcal{H}_1 \quad \alpha \vdash_{\delta} \langle S_2, \mathcal{M}_1, \mathcal{H}_1 \rangle \Downarrow_{t_2} \mathcal{M}_2, \mathcal{H}_2}{\alpha \vdash_{\delta} \langle S_1; S_2, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{t_1.t_2} \mathcal{M}_2, \mathcal{H}_2} \\
\text{SEQ-RETURN} \frac{\alpha \vdash_{\delta} \langle S, \mathcal{M}, \mathcal{H} \rangle \Downarrow_t \mathcal{M}_1, \mathcal{H}_1 \quad \alpha \vdash_{\delta} \langle \text{return}(e), \mathcal{M}_1, \mathcal{H}_1 \rangle \Downarrow \mathcal{M}_1, \mathcal{H}_1 \triangleright v}{\alpha \vdash_{\delta} \langle S; \text{return}(e), \mathcal{M}, \mathcal{H} \rangle \Downarrow_t \mathcal{M}_1, \mathcal{H}_1 \triangleright v} \\
\text{IF-ELSE-T} \frac{\alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle n, \mathcal{M}, \mathcal{H} \rangle \quad n \neq 0 \quad \alpha \vdash_{\delta} \langle S_1, \mathcal{M}, \mathcal{H} \rangle \Downarrow_t \mathcal{M}', \mathcal{H}'}{\alpha \vdash_{\delta} \langle \text{if } e \text{ then } S_1 \text{ else } S_2, \mathcal{M}, \mathcal{H} \rangle \Downarrow_t \mathcal{M}', \mathcal{H}'} \\
\text{IF-ELSE-F} \frac{\alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle 0, \mathcal{M}, \mathcal{H} \rangle \quad \alpha \vdash_{\delta} \langle S_2, \mathcal{M}, \mathcal{H} \rangle \Downarrow_t \mathcal{M}', \mathcal{H}'}{\alpha \vdash_{\delta} \langle \text{if } e \text{ then } S_1 \text{ else } S_2, \mathcal{M}, \mathcal{H} \rangle \Downarrow_t \mathcal{M}', \mathcal{H}'} \\
\text{WHILE-F} \frac{\alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle 0, \mathcal{M}, \mathcal{H} \rangle}{\alpha \vdash_{\delta} \langle \text{while } e \text{ then } S, \mathcal{M}, \mathcal{H} \rangle \Downarrow \mathcal{M}, \mathcal{H}} \\
\text{WHILE-T} \frac{\alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle n, \mathcal{M}, \mathcal{H} \rangle \quad n \neq 0 \quad \alpha \vdash_{\delta} \langle S, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{t_1} \mathcal{M}', \mathcal{H}' \quad \alpha \vdash_{\delta} \langle \text{while } e \text{ then } S, \mathcal{M}', \mathcal{H}' \rangle \Downarrow_{t_2} \mathcal{M}'', \mathcal{H}''}{\alpha \vdash_{\delta} \langle \text{while } e \text{ then } S, \mathcal{M}, \mathcal{H} \rangle \Downarrow_{t_1.t_2} \mathcal{M}'', \mathcal{H}''} \\
\text{DECLASSIFY} \frac{E \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle v, \mathcal{M}, \mathcal{H} \rangle \quad \delta(x) = E \quad \mathcal{M}' = \mathcal{M}[x \mapsto v]}{E \vdash_{\delta} \langle x := \text{declassify}(e), \mathcal{M}, \mathcal{H} \rangle \Downarrow \mathcal{M}', \mathcal{H}} \\
\text{RETURN} \frac{\alpha \vdash_{\delta} \langle e, \mathcal{M}, \mathcal{H} \rangle \Downarrow \langle v, \mathcal{M}, \mathcal{H} \rangle}{\alpha \vdash_{\delta} \langle \text{return}(e), \mathcal{M}, \mathcal{H} \rangle \Downarrow \mathcal{M}, \mathcal{H} \triangleright v}
\end{array}$$

Figure 10: Large-step semantics of J_E commands

$$\begin{array}{c}
\text{T-SKIP} \frac{}{pc, \Gamma, \Pi \vdash_{\delta} \text{skip} : \Gamma, \Pi} \quad \text{T-SUB TYPING} \frac{pc_1, \Gamma_1, \Pi_1 \vdash_{\delta} S : \Gamma'_1, \Pi'_1 \quad pc_1 \sqsubseteq pc_2 \quad \Gamma_1 \sqsubseteq \Gamma_2 \quad \Pi_1 \sqsubseteq \Pi_2}{pc_2, \Gamma_2, \Pi_2 \vdash_{\delta} S : \Gamma'_2, \Pi'_2} \\
\\
\text{T-ASSIGN} \frac{\Gamma, \Pi \vdash_{\delta} e : (\ell, d) \quad \ell \sqcup pc_{\ell} = \langle \mathbb{S}, - \rangle \Rightarrow \delta(x) = E \quad d = T \Rightarrow pc_{\ell} = \langle -, \mathbb{T} \rangle}{pc, \Gamma, \Pi \vdash_{\delta} x := e : \Gamma[x \mapsto pc_{\ell} \sqcup \ell], \Pi[x \mapsto d \vee pc_d]} \\
\\
\text{T-STORE} \frac{\Gamma, \Pi \vdash_{\delta} e : (\ell_2, d_2) \quad \ell_2 \sqcup pc_{\ell} \sqsubseteq \ell_1 \quad \ell_1 \sqcup \ell_2 \sqcup pc_{\ell} = \langle \mathbb{S}, - \rangle \Rightarrow \delta(C) = E \quad pc_d \vee d_2 = T \Rightarrow pc_{\ell} \sqcup \ell_1 = \langle -, \mathbb{T} \rangle \quad d' = d_2 \vee pc_d}{pc, \Gamma, \Pi \vdash_{\delta} C.f := e : \Gamma, \Pi[C.f \mapsto d', C \mapsto \Pi(C) \vee d']} \\
\\
\text{T-SEQ} \frac{pc, \Gamma, \Pi \vdash_{\delta} S_1 : \Gamma', \Pi' \quad pc, \Gamma', \Pi' \vdash_{\delta} S_2 : \Gamma'', \Pi''}{pc, \Gamma, \Pi \vdash_{\delta} S_1; S_2 : \Gamma'', \Pi''} \\
\\
\text{T-SEQ-RETURN} \frac{pc, \Gamma, \Pi \vdash_{\delta} S : \Gamma', \Pi' \quad pc, \Gamma', \Pi' \vdash_{\delta} \text{return}(e) : \Gamma', \Pi' \triangleright (\ell, d)}{pc, \Gamma, \Pi \vdash_{\delta} S; \text{return}(e) : \Gamma', \Pi' \triangleright (\ell, d)} \\
\\
\text{T-IF-ELSE} \frac{\Gamma, \Pi \vdash_{\delta} e : (\ell, d) \quad (pc_{\ell} \sqcup \ell, pc_d \vee d), \Gamma, \Pi \vdash_{\delta} S_1 : \Gamma', \Pi' \quad (pc_{\ell} \sqcup \ell, pc_d \vee d), \Gamma, \Pi \vdash_{\delta} S_2 : \Gamma', \Pi'}{pc, \Gamma, \Pi \vdash_{\delta} \text{if } e \text{ then } S_1 \text{ else } S_2 : \Gamma', \Pi'} \\
\\
\text{T-WHILE} \frac{\Gamma, \Pi \vdash_{\delta} e : (\ell, d) \quad (pc_{\ell} \sqcup \ell, pc_d \vee d), \Gamma, \Pi \vdash_{\delta} S : \Gamma', \Pi'}{pc, \Gamma, \Pi \vdash_{\delta} \text{while } e \text{ then } S : \Gamma', \Pi'} \\
\\
\text{T-DECLASSIFY} \frac{\Gamma, \Pi \vdash_{\delta} e : (\ell, d) \quad \ell \sqsubseteq \langle \mathbb{S}, \mathbb{T} \rangle \quad pc_{\ell} \sqsubseteq \langle \mathbb{P}, \mathbb{T} \rangle \quad \delta(x) = E}{pc, \Gamma, \Pi \vdash_{\delta} x := \text{declassify}(e) : \Gamma[x \mapsto \ell \sqcap \langle \mathbb{P}, \mathbb{T} \rangle], \Pi[x \mapsto T]} \\
\\
\text{T-RETURN} \frac{\Gamma, \Pi \vdash_{\delta} e : (\ell, d) \quad d = T \Rightarrow pc_{\ell} = \langle -, \mathbb{T} \rangle}{pc, \Gamma, \Pi \vdash_{\delta} \text{return}(e) : \Gamma, \Pi \triangleright (pc_{\ell} \sqcup \ell, d)}
\end{array}$$

Figure 11: Typing rules for commands

$$\begin{array}{c}
\Gamma, \Pi \vdash_{\delta} C.m(\bar{p}) : \left(\Gamma^-, \Pi^- \xrightarrow{pc'} \Gamma^+, \Pi^+ \right)_{(\ell, d)} \\
pc_{\ell} \sqsubseteq pc'_{\ell} \quad \delta(x) = \delta(C) \quad d = T \Rightarrow pc_{\ell} = \langle -, \mathbb{T} \rangle \quad \Gamma(q_i) \sqsubseteq \Gamma^-(p_i) \quad i = 1 \dots |\bar{p}| \quad \Pi(q_i) = \Pi^-(p_i) \quad i = 1 \dots |\bar{p}| \\
\forall y \in \text{dom}(\Gamma^-). \Gamma(y) \sqsubseteq \Gamma^-(y) \quad \forall y \in \text{dom}(\Gamma^+). \Gamma^+(y) \sqsubseteq \Gamma_{out}(y) \quad \forall y \in (\text{dom}(\Gamma) \setminus \text{dom}(\Gamma^+)). \Gamma(y) = \Gamma_{out}(y) \\
\forall y \in \text{dom}(\Pi^-). \Pi(y) \sqsubseteq \Pi^-(y) \quad \forall y \in \text{dom}(\Pi^+). \Pi^+(y) \sqsubseteq \Pi_{out}(y) \quad \forall y \in (\text{dom}(\Pi) \setminus \text{dom}(\Pi^+)). \Pi(y) = \Pi_{out}(y) \\
\text{T-ASSIGN-CALL} \frac{}{pc, \Gamma, \Pi \vdash_{\delta} x := C.m(\bar{q}) : \Gamma_{out}[x \mapsto pc_{\ell} \sqcup \ell], \Pi_{out}[x \mapsto d \vee pc_d]}
\end{array}$$

$$\begin{array}{c}
\Gamma, \Pi \vdash_{\delta} C.m(\bar{p}) : \left(\Gamma^-, \Pi^- \xrightarrow{pc'} \Gamma^+, \Pi^+ \right)_{(\ell, d)} \quad \delta(x) = N \\
\delta(C) = E \quad \ell \sqsubseteq \langle \mathbb{P}, \mathbb{U} \rangle \quad \forall p \in \bar{p}. \Gamma^-(p) = \langle \mathbb{P}, \mathbb{U} \rangle \quad \forall p \in \bar{p}. \Pi^-(p) = F \quad \Gamma(q_i) \sqsubseteq \Gamma^-(p_i) \quad i = 1 \dots |\bar{p}| \\
\Pi(q_i) = \Pi^-(p_i) \quad i = 1 \dots |\bar{p}| \quad \forall y \in \text{dom}(\Gamma^-). \Gamma(y) \sqsubseteq \Gamma^-(y) \quad \forall y \in \text{dom}(\Gamma^+). \Gamma^+(y) \sqsubseteq \Gamma_{out}(y) \\
\forall y \in (\text{dom}(\Gamma) \setminus \text{dom}(\Gamma^+)). \Gamma(y) = \Gamma_{out}(y) \quad \forall y \in \text{dom}(\Pi^-). \Pi(y) \sqsubseteq \Pi^-(y) \\
\forall y \in \text{dom}(\Pi^+). \Pi^+(y) \sqsubseteq \Pi_{out}(y) \quad \forall y \in (\text{dom}(\Pi) \setminus \text{dom}(\Pi^+)). \Pi(y) = \Pi_{out}(y) \\
\text{T-ASSIGN-GATEWAY-CALL} \frac{}{pc, \Gamma, \Pi \vdash_{\delta} x := C.m(\bar{q}) : \Gamma_{out}[x \mapsto pc_{\ell} \sqcup \ell], \Pi_{out}[x \mapsto F]}
\end{array}$$

$$\begin{array}{c}
\Gamma, \Pi \vdash_{\delta} C.f : (\ell_1, d_1) \quad \Gamma, \Pi \vdash_{\delta} C'.m(\bar{p}) : \left(\Gamma^-, \Pi^- \xrightarrow{pc'} \Gamma^+, \Pi^+ \right)_{(\ell_2, d_2)} \\
\delta(C) = \delta(C') \quad d_2 = T \Rightarrow pc_{\ell} = \langle -, \mathbb{T} \rangle \quad \ell_1 \sqcup \ell_2 \sqcup pc_{\ell} = \langle \mathbb{S}, - \rangle \Rightarrow \delta(C) = E \\
pc_{\ell} \sqsubseteq pc'_{\ell} \quad \ell_2 \sqcup pc_{\ell} \sqsubseteq \ell_1 \quad \Gamma(q_i) \sqsubseteq \Gamma^-(p_i) \quad i = 1 \dots |\bar{p}| \quad \Pi(q_i) = \Pi^-(p_i) \quad i = 1 \dots |\bar{p}| \\
\forall y \in \text{dom}(\Gamma^-). \Gamma(y) \sqsubseteq \Gamma^-(y) \quad \forall y \in \text{dom}(\Gamma^+). \Gamma^+(y) \sqsubseteq \Gamma_{out}(y) \quad \forall y \in (\text{dom}(\Gamma) \setminus \text{dom}(\Gamma^+)). \Gamma(y) = \Gamma_{out}(y) \\
\forall y \in \text{dom}(\Pi^-). \Pi(y) \sqsubseteq \Pi^-(y) \quad \forall y \in \text{dom}(\Pi^+). \Pi^+(y) \sqsubseteq \Pi_{out}(y) \quad \forall y \in (\text{dom}(\Pi) \setminus \text{dom}(\Pi^+)). \Pi(y) = \Pi_{out}(y) \\
\text{T-STORE-CALL} \frac{}{pc, \Gamma, \Pi \vdash_{\delta} C.f := C'.m(\bar{q}) : \Gamma_{out}, \Pi_{out}[C.f \mapsto d_2 \vee pc_d]}
\end{array}$$

$$\begin{array}{c}
\Gamma, \Pi \vdash_{\delta} C.f : (\ell_1, d_1) \quad \Gamma, \Pi \vdash_{\delta} C'.m(\bar{p}) : \left(\Gamma^-, \Pi^- \xrightarrow{pc'} \Gamma^+, \Pi^+ \right)_{(\ell_2, d_2)} \quad \delta(C) = N \quad \delta(C') = E \\
\ell_2 \sqsubseteq \langle \mathbb{P}, \mathbb{U} \rangle \quad \forall p \in \bar{p}. \Gamma^-(p) = \langle \mathbb{P}, \mathbb{U} \rangle \quad \forall p \in \bar{p}. \Pi^-(p) = F \quad \ell_2 \sqcup pc_{\ell} \sqsubseteq \ell_1 \quad \Gamma(q_i) \sqsubseteq \Gamma^-(p_i) \quad i = 1 \dots |\bar{p}| \\
\Pi(q_i) = \Pi^-(p_i) \quad i = 1 \dots |\bar{p}| \quad \forall y \in \text{dom}(\Gamma^-). \Gamma(y) \sqsubseteq \Gamma^-(y) \quad \forall y \in \text{dom}(\Gamma^+). \Gamma^+(y) \sqsubseteq \Gamma_{out}(y) \\
\forall y \in (\text{dom}(\Gamma) \setminus \text{dom}(\Gamma^+)). \Gamma(y) = \Gamma_{out}(y) \quad \forall y \in \text{dom}(\Pi^-). \Pi(y) \sqsubseteq \Pi^-(y) \\
\forall y \in \text{dom}(\Pi^+). \Pi^+(y) \sqsubseteq \Pi_{out}(y) \quad \forall y \in (\text{dom}(\Pi) \setminus \text{dom}(\Pi^+)). \Pi(y) = \Pi_{out}(y) \\
\text{T-STORE-GATEWAY-CALL} \frac{}{pc, \Gamma, \Pi \vdash_{\delta} C.f := C'.m(\bar{q}) : \Gamma_{out}, \Pi_{out}[C.f \mapsto F]}
\end{array}$$

Figure 11: Typing rules for commands