# SIGY: Breaking Intel SGX Enclaves with Malicious Exceptions & Signals

Supraja Sridhara, Andrin Bertschi, Benedict Schlüter, Shweta Shinde

ETH Zürich

*Abstract*—**User programs recover from hardware exceptions and respond to signals by executing custom handlers that they register specifically for such events. We present SIGY attack, which abuses this programming model on Intel SGX to break the confidentiality and integrity guarantees of enclaves. SIGY uses the untrusted OS to deliver fake hardware events and injects fake signals in an enclave at any point. Such unintended execution of benign program-defined handlers in an enclave corrupts its state and violates execution integrity. 7 runtimes and library OSes (OpenEnclave, Gramine, Scone, Asylo, Teaclave, Occlum, EnclaveOS) are vulnerable to SIGY. 8 languages supported in Intel SGX have programming constructs that are vulnerable to SIGY. We use SIGY to demonstrate 4 proof of concept exploits on webservers (Nginx, Node.js) to leak secrets and data analytics workloads in different languages (C and Java) to break execution integrity.**

## I. Introduction

User programs rely on exceptions and signals to manage unexpected events or errors that may occur during execution. Programming languages allow developers to express event and error-specific logic in the form of *handlers*, such that when the program is notified of a particular event, the handler is executed automatically. These rich abstractions are facilitated by the cooperation of the hardware and the operating system (OS). When the hardware encounters runtime errors during the program execution (e.g., page faults, segmentation faults, timers, divide by zero), it notifies the OS via interrupts. The OS either handles the faults itself (e.g., load the page into memory) or forwards it to the user program's signal handler (e.g., DivideByZeroException). The program can also request the OS for notifications about events of its interest that either emanate from the system (e.g., Ctrl+C) or other processes on the system (e.g., synchronization between parent and child processes) and register handlers that should execute on such events [19]. Thus, the OS not only monitors for such events on behalf of the application and notifies it, but also diverts the control of the application to the event handlers in the program. Both these mechanisms facilitate rich functionality in the programs, while the hardware and the OS provide efficient notification and handler invocation.

Intel SGX provides a user-level abstraction called *enclaves* that protects sensitive data and code execution [40], [33]. The hardware protects enclave confidentiality and integrity even when the OS and other user processes are compromised. With such a strong threat model, Intel SGX limits the attack surface to the critical code running in an enclave. Since enclave memory is rendered inaccessible to the OS, traditional programs that were written with the assumption of a trusted OS simply cannot execute inside enclaves (e.g., `syscall` instruction is illegal inside an enclave).

Due to its unique programming model, existing programs do not execute out of the box on Intel SGX [27], [8], [9], [17], [53], [18], [11], [35], [12], [3], [20], [10], [23], [13]. As a solution, programmers can use SGX runtimes that provide a small trusted runtime that interfaces with the SGX hardware to expose a new high-level interface to the programmer. Alternatively, programmers can use a trusted library OS inside an enclave that can execute unmodified applications that were not programmed for enclaves. Runtimes and Library OSes for Intel SGX support exception and signal delivery to enclaves since it is a much-required feature for programs. The hardware or the OS can inform the enclave about an exception or a signal by inducing an asynchronous exit. The enclave safely stores its current execution state and exits to untrusted code. The enclave can then be re-entered from another fixed entry point to execute corresponding pre-registered handlers for the exceptions or signals. During this flow, the hardware and the trusted software ensure that the OS cannot subvert the execution of the enclave—it executes the handler and then resumes the execution at the point where it was interrupted. This mechanism allows enclaves to handle runtime events even when the OS is untrusted.

Heckler and WeSee introduce a new class of attacks called Ahoi attacks where an attacker uses notifications to compromise Confidential VMs (CVMs) enabled by Intel TDX and AMD SEV-SNP [52], [51]. Ahoi attacks use interrupts under the control of a malicious hypervisor that can trigger interrupt handlers in the CVMs. These interrupt handlers alter the global execution state of the CVMs and compromise them. In light of these findings, we revisit Intel SGX and analyze if an attacker can use notifications to compromise the security of enclaves. We investigate two lines of inquiry: (i) what events can the OS fake to trigger handler execution inside the enclave? and (ii) can such handler execution bring about direct changes to the enclave's global execution state (e.g., variables)?

In this paper, we introduce a new attack called SIGY where the OS compromises the enclave execution by inducing fake events and signals to execute benign handlers registered by the enclave. Intuitively, enclaves want to recover from a divide-by-zero and expect signals from another enclave. To handle such events, enclaves will register handlers that explicitly update the enclave state, say by changing the denominator to a non-zero value or invoking an event handler to respond to another enclave's request. If the OS convincingly tricks the enclave into falsely believing that such an event occurred, the enclave will stop its current execution and execute the handler that will explicitly update the enclave state (e.g., change a variable to a non-zero value or execute a function). In the least, this will result in corruption of the enclave's state resulting in a crash. If the OS injects the event at an opportune moment, it

can use the effects of the handler to compromise the enclave. We demonstrate this phenomenon by introducing a new attack called SIGY, which exploits the OS's ability to fake signals to execute enclave handlers and subvert SGX guarantees.

We show that existing runtimes, library OSes, and programming language constructs are vulnerable to SIGY. We first analyze existing support to execute SGX applications: 8 runtimes (Intel SGX SDK, Open Enclave, Teaclave SGX-SDK, Asylo, Rust EDP, GoTEE, Enarx, and EGo) and 6 library OSes (Gramine, Scone, EnclaveOS, EdgelessRT, MystikOS, and Occlum). Then, we analyze the signal delivery mechanism and handlers for programs written in 9 languages (C, C++, Java, Python, Go, JavaScript, Rust, Julia and Wasm) to observe their behavior in enclaves. We find that 3/8 runtimes (Open Enclave, Teaclave SGX-SDK, and Asylo) and 4/6 library OSes (Gramine, Scone, EnclaveOS, and Occlum) are susceptible to SIGY because they do not detect the fake signals injected by the OS. Of the 9 languages we study, 8 (C, C++, Java, Python, Go, JavaScript, Rust, and Julia) offer language constructs for programs to register custom handlers. We hand-code applications in each of these languages to register handlers and execute them in vulnerable SGX runtimes and library OSes to confirm that they are indeed vulnerable to SIGY. Next, we demonstrate that SIGY breaks the confidentiality and integrity of 4 open-source applications (Nginx, Node.js, machine learning) that have been ported to Intel SGX by prior works. Our proof of concept exploits on these enclaves leak secrets and change outputs. Depending on the victim enclave, SIGY may require injecting signals at a particular window of execution. We construct a proof of concept exploit against a worst-case application, a multi-layer perceptron, that requires 186 billion injections to bias the output.

Our proposed software defenses serve as point-wise solutions against SIGY. The vulnerable runtimes and library OSes have to make a design choice between either disabling functionality for security or leaving the onus on the developer to reason about the security. While the latter can be a pragmatic solution, new attacks like SIGY serve as an example that programmers using runtimes and library OSes for lift-and-shift should not be burdened with this decision. We conclude that some programs simply cannot be protected without limiting functionality. Our conclusion encourages runtime and library OS maintainers to disable vulnerable exception and signal delivery interfaces. Our detailed analysis of existing enclave ecosystems spanning from runtimes, library OSes, programming language support, and existing enclave-bound applications provides an in-depth exploration to help future lift-and-shift solutions in making judicious choices. In summary, SIGY brings attention to a new attack surface that requires a re-examination of the enclave ecosystem.

**Contributions.** The paper makes three main contributions:

1) SIGY *Attack.* We present a novel attack on Intel SGX where a malicious OS can send fake signals enclave and trick them into executing enclave-registered handlers that change the enclave state.

2) *Analysis.* Of the 14 popular frameworks for running applications in Intel SGX, 7 are vulnerable to SIGY because they forward fake signals to the enclave whereas 8/9 popular languages used for enclave programming support custom handlers.



```
1.  int old_mean;
2.  int mean, n;
3.  void add(int data){
4.    try{
5.      old_mean = mean;
6.      //this can cause overflow
7.      mean = addExact(old_mean,data)/n;
8.      n += 1;
9.    }catch(ArithmeticException e){
10.         mean = old_mean;
11.   }
12. }
```

| | before execution | after |
|---|---|---|
| | | normal execution |
| | old_mean=0 | old_mean=15 |
| | mean=0 | mean=20 |
| | n=0 | n=3 |
| execute | | attack execution |
| add(10) | | old_mean=0 |
| add(20) | | mean=0 |
| add(30) | | n=0 |

Fig. 1. SIGY on Java applications. Attacker injects `sigfpe` 3 times to change the execution and data integrity.

3) *Exploits.* We build exploits on 4 open-source enclave applications to demonstrate SIGY.

**Responsible Disclosure.** As per the responsible disclosure guidelines, we have informed all the 7 impacted runtimes and library OSes. We are coordinating patching with the maintainers and will update the paper once the process concludes.

## II. SIGY OVERVIEW

For functionality, applications register custom exception and signal handlers that alter the global state of the program. To preserve this functionality on Intel SGX, runtimes and library OSes provide mechanisms to send signals to applications that execute in enclaves. SIGY tricks the benign runtime and library OS signal handling mechanisms which results in the enclave executing the handlers. Therefore, the custom handlers in applications put together with the signal propagation infrastructure of runtimes and library OSes render the enclaves vulnerable to SIGY.

**Threat Model.** We trust the hardware in the Intel CPU package and assume that it is free from bugs. The enclaves are launched and attested according to the Intel SGX specification, and all software that executes inside the enclave is assumed to be bug-free. This includes the enclave application code, trusted runtime, and library OSes. We assume that all software that executes outside the enclave, including the OS, untrusted runtime, and other processes are untrusted and can be malicious.

### A. Motivating Examples

**Faking hardware exceptions.** Consider a Java function in Figure 1 that executes in an enclave whose `add` function (Line 3) is called 3 times with data as $10, 20, 30$. The function computes and stores the new mean on every invocation. Now, consider a malicious operating system (OS) that wants to compromise the data integrity of this enclave. In this example, the attacker's goal is to ensure that the application never updates the `mean`, i.e., it remains $0$ despite the 3 invocations of `add`. Because this code executes in an enclave, the OS cannot change the values in memory directly. However, observe that this enclave catches and handles `ArithmeticExceptions` to deal with bad data. The handler reverts the value of `mean`, discarding the effects of the bad data. If the OS can trigger this handler every time the `add` function is invoked, then the `mean` will not be updated.

The OS can use SIGY to achieve its goal. Specifically, the Java runtime converts the signal for floating point exceptions (sigfpe) that it gets from the OS to an ArithmeticException. This exception is then caught and handled by the enclave. Therefore, using SIGY the OS injects sigfpe to the Java code in the enclave when it executes Lines 4-9 in Figure 1. This will result in the enclave always executing the exception handler. With 3 such signal injections, the OS ensures that the mean does not change, thus breaking integrity.

**Faking user-defined signals.** Several library OSes (e.g., Gramine, Scone) allow lifting and shifting unmodified applications like Nginx to execute in enclaves. Nginx is a web server that is highly optimized to provide maximum uptime, configured using a config file. Consider an unmodified Nginx server that executes in an enclave to serve http data with config in Listing 1. At a later point, an Nginx administrator introduces authentication tokens (JSON Web Tokens (jwt) [14]) in the config file ( Listing 2). To enable the administrator to refresh the config file without downtime, Nginx performs the configuration refresh on receiving sighup. When the administrator sends sighup to the Nginx process, the Nginx process reads the new config file and starts using it.

Listing 1.  Old config

```
1   ...
2   server {
3     listen 80;
4       location /products/
5         ...
6
7
8
9     }
10  }
```

Listing 2.  New config

```
1   ...
2   server {
3     listen 80;
4       location /products/{
5     auth_jwt   "API";
6     auth_jwt_type
7                 encrypted
8         ...
9     }
10  }
```

Now, consider a malicious OS that aims to disable this authentication mechanism in the Nginx server. To do this, the OS should be able to force the server to use the old configuration file without the authentication enabled. Library OSes protect enclave files by encrypting them. Therefore, the malicious OS cannot directly edit the config file. However, the OS can capture the old config file as an encrypted blob from the file-system before it is replaced by the administrator. Then, once the Nginx configuration upgrade is complete, and checked by the administrator, the OS writes the encrypted blob back to replace the new config file. Note that, this is not sufficient to trick the Nginx server into using the compromised configuration without restarting the enclave. The OS's goal is to ensure that when users connect to the Nginx server, they are served using the older configuration instead of the configuration that the administrator upgraded and checked, thus mounting a time-of-check time-of-use (TOCTOU) attack. For this, the OS uses SIGY to force the Nginx server to use this compromised configuration after the administrator checks that the configuration reload was successful by injecting sighup to the Nginx process.

### B. SIGY *Attacks on Real-world Enclaves*

SIGY uses asynchronous signal injection to compromise enclaves by triggering expressive signal handlers. This requires
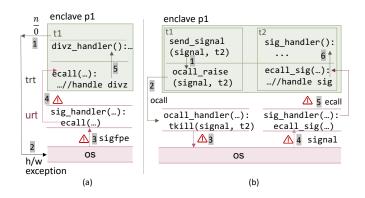


Fig. 2.  SDK interfaces. (a) Handling hardware exceptions (b) Handling intra-enclave signals.

a runtime or library OS that propagates hardware exceptions and signals to the enclave application. Further, SIGY uses handlers in the applications that perform computations that alter the enclave's global state. The enclave handlers depend on programming language constructs used in the application (e.g., signal registration constructs, and signal handling constructs). Therefore, we first evaluate 14 runtimes and library OSes and check if they propagate signals to the enclaves (Section III and Section IV). Next, we examine 9 programming languages and systematically analyze the constructs they provide for programs to register and execute custom signal handlers (Section V-B). Finally, we use the insights from our runtime and programming language analysis to demonstrate SIGY on 4 publicly available enclave applications (Section VI). We extract secrets to break confidentiality and change programming execution to break integrity.

## III.   FAKING SIGNALS IN SDKS

**Background: Sending signals to threads.** SDKs are typically used to execute a single process in an enclave. So, they may not enable mechanisms for the enclave to send signals to other processes. However, they do support multi-threading inside the same enclave. For sending signals between threads, the OS exposes the tkill system call. Along similar lines, some SDKs add mechanisms to enable enclave threads to send signals to each other. To support sending signals to other threads, the SDKs add a new ocall interface to send a signal to the target enclave thread via the OS (Steps 1-4 in Figure 2(b)). When the OS sends the signal to the target enclave thread, the untrusted runtime's handler propagates it to the trusted runtime using an ecall (Step 5) which in turn invokes the target thread's signal handler in the enclave.

### A. Intel SGX SDK

The Intel SGX SDK allows enclaves to register handlers for hardware exceptions (e.g., divide-by-zero, overflow) to perform custom handling when hardware exceptions occur. During normal operation, if the enclave executes an instruction that triggers a hardware exception (Figure 2(a)), the SGX hardware triggers an asynchronous exit. On an asynchronous exit, the trusted hardware stores the current execution state of the enclave into the protected state save area (SSA). Then, it raises an exception that the untrusted OS traps on where

the hardware stores information about the exit into the SSA as shown in Listing 3. Crucially, in SGX 2, the hardware stores information about the exit into the SSA as shown in Listing 3. Note that, SGX 1 does not store this hardware information [34].

Listing 3. Exit information stored in the SSA.
```
1  struct _exit_info_t {
2   uint32_t  vector;    // exception vector number
3   uint32_t  exit_type; // HW or SW exceptions
4   uint32_t  valid;     // supported/unsupported
5  };
```

This includes the validity, type, and reason (i.e., exception vector) for the asynchronous exit. Enclave software can access this information while untrusted software (e.g., untrusted runtime, OS) cannot. The Intel SGX SDK's trusted runtime uses the exit information from the SSA to deduce the validity (see Line 1 in Listing 4) and reason for any asynchronous exit (Line $4-5$).

Listing 4. Psuedocode of exception handling for Intel SGX SDK.
```
1  if (ssa_gpr->exit_info.valid == 1) {
2    // info used to forward exception to enclave app
3    info->exception_valid = ssa_gpr->exit_info.valid;
4    info->exception_vector = ssa_gpr->exit_info.vector;
5    info->exception_type = ssa_gpr->exit_info.exit_type;
6    ...
7  }
```

When a hardware exception causes an exit from the enclave (Step 1 in Figure 2(a)), the trusted hardware saves the exit information in the SSA and raises a hardware exception to the OS (Step 2). The OS converts the hardware exception to a signal.[1] Then, it identifies the enclave process that caused the exception and sends it a signal (Step 3). This signal is caught and handled by the untrusted runtime of the Intel SGX SDK. The signal handler converts the signal back to the corresponding hardware exception. Then, it notifies the enclave with the hardware exception information by performing an enclave enter (Step 4).[2]

**SGX 1.** The trusted runtime uses the hardware exception vector from the untrusted runtime to call the enclave application's exception handler. Therefore, SGX 1 enclaves that do not have hardware support to store the exit information are vulnerable to SIGY. Specifically, the OS can arbitrarily inject a signal to the enclave to cause an asynchronous exit. Then, the untrusted runtime enters the enclave with the hardware exception vector corresponding to the signal. The exception handling in the trusted runtime executes the enclave's exception handler (Step 5) without any filtering. Therefore, an attacker can use SIGY to asynchronously send signals to the enclave and trigger its exception handlers.

**SGX 2.** Enclaves that execute in SGX 2 with the Intel SGX SDK are not vulnerable to SIGY. In SGX 2, the trusted runtime uses the exit information from the SSA to determine the validity and exception vector when performing exception handling. If the OS maliciously injects signals to the enclave, it causes an invalid exit (i.e., the hardware stores 0 in Line 4 in Listing 3). When the trusted runtime checks the exit information, the guard check on Line 1 of Listing 4 will fail.

---

[1]OS does not raise a signal for page faults. It simply resumes operation in the untrusted runtime

[2]The trusted runtime implements 2-level exception handling for asynchronous exits. We abstract this detail in our discussion.
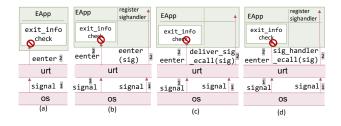


Fig. 3. (a) Intel SGX SDK (b) Open enclave (c) Asylo (d) Teaclave SGX-SDK.

TABLE I. LIBRARY OS AND RUNTIME ANALYSIS FOR SIGY. ✓: INTERFACE SUPPORTED ✗: INTERFACE NOT SUPPORTED ✳: CANNOT BE ANALYZED AS THEY ARE CLOSED SOURCE.

| Type | Name | Interfaces | | Can inject with SIGY | |
|---|---|---|---|---|---|
| | | HW exception interface | Other signal interface | HW exception signals | Process signals |
| SDK | Intel SGX SDK [13] | ✓ | ✗ | no | no |
| | Open Enclave [18] | ✓ | ✓ | yes | yes |
| | Teaclave [23] | ✓ | ✓ | yes | yes |
| | Asylo [3] | ✓ | ✓ | yes | yes |
| | Rust EDP [20] | ✗ | ✗ | no | no |
| libos: 1-process per enclave | Gramine [12] | ✓ | ✓ | yes | no |
| | Scone [21] | ✳ | ✳ | yes | yes |
| | EnclaveOS [11] | ✳ | ✳ | yes | yes |
| | EdgelessRT [8] | ✗ | ✗ | no | no |
| libos: n-process per enclave | Mystikos [17] | ✗ | ✓ | no | no |
| | Occlum [53] | ✓ | ✓ | yes | yes |
| language runtime | GoTEE [35] | ✗ | ✗ | no | no |
| | Enarx [10] | ✗ | ✗ | no | no |
| | EGo [9] | ✗ | ✗ | no | no |

So the trusted runtime will discard the exception and will not execute the handler in the enclave (see Figure 3(a)).

### B. Open Enclave

Open Enclave is a widely used open-source SDK that allows developers to write custom applications to execute in SGX enclaves.

**Filtering using exit information.** To support exception handling, Open Enclave allows applications to register handlers for hardware exceptions. An instruction that raises a hardware exception in the enclave causes an asynchronous exit and traps into the OS. The OS and the untrusted runtime then propagate this exception to the enclave's trusted runtime. Open Enclave's exception handling uses the exit information that the hardware stores in the SSA like in Listing 4. If the OS attacks an enclave by maliciously injecting a signal, this will cause an asynchronous exit from the enclave. However, the hardware will indicate that this is an invalid exit in the exit information stored in the SSA. When the trusted runtime handles this exception, it will detect the attack and not execute the exception handler (LHS of Figure 3(b)).

**Supporting inter-thread signals.** The hardware exception interface allows applications to only register handlers for hardware exception events. Open Enclave does not allow enclaves to register handlers for other signals (e.g., `sigusr1`, `sighup`) which limits the expressiveness of applications that can execute with Open Enclave. To improve expressiveness, Open Enclave introduces a separate mechanism to allow enclave threads to send signals to each other. This allows enclave threads to explicitly enable signals from the host and register signal handlers (Line 2 and Line 3 in Listing 5). With this, the threads can send signals (Line 4) to other enclave threads that are routed through the trusted runtime.

Listing 5. Enable, register, and send a signal in Open Enclave.

```
1  ...
2  oe_add_vectored_exception_handler(false, sigusr_handle)
3  //enable signal
4  oe_sgx_td_register_host_signal(thread, SIGUSR1)
5  //do ocall
6  host_send_interrupt(target_thread, SIGUSR1)
7  ...
```



Fig. 4. $T_1$ and $T_2$ are threads of the same enclave process. Black: Normal operation with ocall and ecall interfaces in Teaclave. Pink: Interfaces that can be maliciously invoked for SIGY.

When one enclave thread wants to send a signal (e.g., `sigusr1`) to the target enclave thread, it invokes the trusted runtime which performs an ocall (Step 2 in Figure 2(b)). In the ocall context, the untrusted runtime raises a signal to the target enclave thread using the `tkill` system call. When the target enclave thread is resumed through `eenter` (Step 5 in Figure 2(b)), the trusted runtime calls the target enclave thread's signal handler (Step 6 in Figure 2(b)).

**Attacking Open Enclave.** When threads send signals to each other, the hardware exit information cannot be used to determine if the signals are legitimate. Specifically, the exit information stored in the SSA is only useful to determine the legitimacy of hardware exceptions and not explicit exits caused by sending signals. Therefore, the enclave cannot validate if the signal was legitimately raised by one of its threads or if it was maliciously injected by untrusted software. We can use the signal injection interface to compromise the security of Open Enclave using SIGY. Concretely, untrusted software like the OS can directly send signals to enclave threads. This causes the untrusted runtime to enter the enclave with the signal. Because the trusted runtime does not validate the source of this signal, it will invoke the enclaves signal handler (RHS of Figure 3). Therefore, an attacker can arbitrarily execute the enclave's signal handler using SIGY.

*C. Teaclave SGX-SDK*

Teaclave SGX-SDK (Teaclave for short) enables developers to write programs in Rust and execute them in enclaves. It is based on Intel SGX SDK and implements different Rust libraries to ease enclave application development (i.e., standard library functionality).

**Hardware exceptions.** For hardware exceptions, it relies on the Intel SGX SDK's exception handling mechanism. The trusted runtime of Intel SGX SDK detects and discards all maliciously injected exceptions by checking the exit information in the SSA as discussed in Section III-A. Hence, this interface cannot be used to compromise Teaclave enclaves using SIGY.

**Signal support.** Next, we analyze Teaclave's support that enables enclave threads to send signals to each other. Teaclave includes wrappers for Rust signal libraries. This allows enclave threads to register handlers and send signals to each other. To enable this functionality, Teaclave introduces a new ocall that its Rust library wrapper invokes as shown in Figure 4. The ocall function in Teaclave's untrusted runtime performs `sigaction` and `raise` libc calls Teaclave also adds a public ecall (`t_signal_handler_ecall`) to be invoked by the untrusted runtime to forward signals from one enclave thread to another.

During normal operation, an enclave thread can raise a signal to another thread using Teaclave's signal library. This is translated i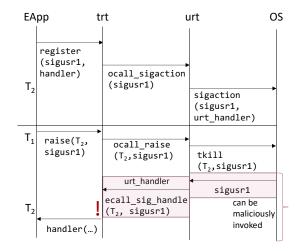nto an ocall (`u_raise_ocall`) by the library which transfers control to Teaclave's untrusted runtime. The untrusted runtime sends a signal to the target thread using `tkill`. When the OS sends the signal to the target thread, Teaclave's untrusted runtime invokes the ecall (`t_signal_handler_ecall`) to handle the signal in the enclave. The ecall implementation in Teaclave's trusted runtime then triggers the signal handler registered in the target thread.

**Attacking Teaclave.** SIGY needs the ability to arbitrarily inject signals that trigger the signal handler in the enclave. To gain this capability, SIGY can abuse the `t_signal_handler_ecall` interface. If an enclave registers signal handlers, SIGY can arbitrarily inject signals into the untrusted runtime and trigger this ecall. Alternatively, because the untrusted runtime is attacker-controlled, SIGY could directly invoke this ecall without the need for signal injection. In both cases, the enclave application will always execute the signal handler (Figure 4). Therefore, an attacker can use SIGY to trigger computation in the enclave. Note that the trusted runtime does not have any mechanism to distinguish legitimate signals (e.g., from one enclave thread to another) from those that are maliciously injected by the OS. Further, the `t_signal_handler_ecall` is a public root ecall (i.e., invoked from any untrusted software) for functionality such that it can be invoked by the untrusted runtime to forward signals to the enclave. Therefore, Teaclave's signal handling design makes it vulnerable to SIGY (RHS in Figure 3(d)).

*D. Asylo*

Asylo is an open-source framework that provides a POSIX interface to enable enclave application development. It implements wrappers for POSIX functions that invoke ocalls to interact with the untrusted OS. Asylo uses the Intel SGX SDK and preserves the hardware exception handling interface from the SGX SDK. This interface checks the exit information in the SSA and discards maliciously injected exceptions. So, this interface is not vulnerable to SIGY.

**Signal support.** Asylo introduces a new signal handling mechanism to allow enclave applications to register and
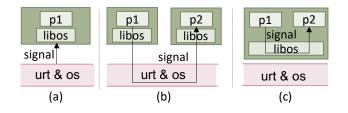
5

Fig. 5. Signal propagation with library OSes. (a) OS or untrusted runtime sends signal to enclave process. (b) One enclave process sends signal to another enclave process through the untrusted runtime and OS. (c) LibOS creates a process abstraction such that 2 processes run in the same enclave. These processes can send signals to each other via the LibOS.



Fig. 6. Hardware exception and signal handling in Gramine.

handle signals. In Asylo, this signal handling interface can be used by the enclave application to register all POSIX signals. To support signals, Asylo introduces a new ocall (`ocall_enc_untrusted_register_signal _handler`) to register a signal handler and an ecall (`ecall_deliver_signal`) to propagate the signal from the OS to the enclave. When the OS sends a signal, the untrusted runtime invokes the ecall and transfers the execution to the enclave. The enclaves trusted runtime uses the signal from the ecall's parameters, looks up the corresponding handler that was registered, and invokes it.

**Attacking Asylo.** The ecall used to deliver the signal to the enclave is a public root ecall. Therefore, SIGY can use this interface to attack enclaves in Asylo. Concretely, when the OS sends a signal to the enclave, the untrusted runtime invokes the ecall to enter the enclave's trusted runtime (RHS Figure 3(c)) The trusted runtime executes the enclave's signal handler without any additional checks.

In summary, our analysis shows that all SDKs use the exit information that the hardware stores to handle hardware exceptions. However, they introduce new mechanisms to support signal handling between threads in the enclaves which render 3 out of the 4 SDKs vulnerable to SIGY.

## IV. FAKING SIGNALS IN LIBRARY OSES

**Background.** Library OSes support rich exception handling and signal interfaces for enclave applications. Unlike the SDKs, they also implement mechanisms to execute multi-process applications by adding support for calls like `fork`, `vfork`, and `execv`. In some library OSes (e.g., Gramine, Scone) calls to these functions spawn new enclave processes Figure 5(b). Like the OS, the library OSes also add support to send signals from one enclave process to another. For this, they route the signals through the untrusted operating system and runtime (Figure 5(b)). Other library OSes support multi-processing, by implementing process abstractions inside the enclave (e.g., Occlum). Such library OSes route inter-process signals inside the enclave itself and do not have to exit or use the untrusted runtime and OS to propagate signals (Figure 5(c)).

### A. Gramine

Gramine is an open-source library OS that enables executing unmodified multi-process applications in SGX enclaves. It has 3 parts—a l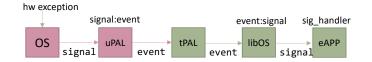ibrary OS, a Platform Adaptation Layer (PAL), and a patched C standard library. For simplicity, we refer to the part of the PAL that runs outside the enclave as untrusted PAL (uPAL) and the part that executes inside the enclave as (tPAL). Figure 6 shows exception and signal handling in Gramine.

**Gramine's signal support.** Crucially for SIGY, Gramine's C standard library allows support for signal handling in enclave applications. Specifically, the wrappers for C's `sigaction` calls allow the application to register handlers for signals. When the enclave triggers a hardware exception, the untrusted OS traps on it and raises a signal that is caught by the uPAL. The uPAL's signal handling function (`handle_sync_signal`) converts the signal into a corresponding Gramine-specific PAL event. For example, the uPAL maps `sigfpe` to `pal_event_arithmetic_error`. Then, it invokes tPAL using the `sgx_raise` (Figure 6) who forwards the event to the library OS. The library OS converts the event to a signal, creates the signal struct (`siginfo_t`) like in Linux, and raises the signal to the enclave application. This finally results in the enclave executing its signal handler.

**Attacking Gramine.** The trusted PAL does not check if the event was raised because of a real hardware exception in the enclave or by a signal injected by the untrusted OS, thus making Gramine vulnerable to SIGY. Specifically, the untrusted OS can inject a signal arbitrarily to a Gramine application (see Figure 5(a)). This signal is converted to an event in the uPAL and forwarded to the tPAL which eventually executes the enclave application's handler.

Although Gramine supports executing multi-process applications, it doesn't explicitly enable support for enclave processes to send signals to each other. Instead, its signal handling mechanism has limited expressiveness and is built to only support hardware exceptions (as shown in Table I). Therefore, with SIGY we can only inject signals that map to hardware exceptions to compromise enclaves that execute with Gramine (c.f. Section V-A).

### B. Scone and EnclaveOS

They are closed-source library OSes that support executing unmodified applications in SGX based on Intel SGX SDK. As we don't have the source code, we perform a black-box analysis of these library OSes. Specifically, we write a program that registers handlers for all signals (c.f. Section V-A) and observe which handlers are executed. For both Scone [27] and EnclaveOS [11], our analysis shows that we can arbitrarily inject most signals from the untrusted OS to the enclaves (c.f. Section V-A). This successfully triggers signal handlers in the enclave making them vulnerable to SIGY. Scone forwards all signals from the OS to the enclaves. For EnclaveOS, our analysis shows that `sigusr2` is reserved for library OS-specific operations, and enclave applications cannot register signal handlers for these. Besides this, we observe that all other

TABLE II.    SIGNAL SUPPORT IN SDKs & LIBRARY OSes. ✓EXECUTES SIGNAL HANDLER, ✗ CRASHES THE ENCLAVE, ✽ NO OBSERVABLE BEHAVIOR. (C.F. TABLE III FOR MAPPINGS FROM SIGNAL NUMBER TO SIGNAL NAME).

| Runtime/LibOS | 1 | 2 | 3 | 4* | 5* | 6 | 7* | 8* | 9 | 10 | 11* | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Open Enclave | ✓ | ✗ | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✗ | ✓ | ✗ | ✓ | ✓ | ✓ | ✗ | ✗ | ✽ | ✽ | ✗ | ✽ | ✽ | ✽ | ✽ | ✗ | ✗ | ✗ | ✗ | ✗ | ✽ | ✽ | ✗ |
| Teaclave | ✓ | ✓ | ✓ | ✗ | ✗ | ✓ | ✗ | ✗ | ✗ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Asylo | ✓ | ✓ | ✓ | ✽ | ✓ | ✓ | ✽ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✽ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✽ | ✓ | ✓ |
| Gramine | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✓ | ✓ | ✗ | ✗ | ✓ | ✽ | ✽ | ✗ | ✽ | ✗ | ✽ | ✽ | ✗ | ✽ | ✽ | ✽ | ✽ | ✽ | ✗ | ✗ | ✗ | ✗ | ✽ | ✗ | ✗ |
| Scone | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| EnclaveOS | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Occlum | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |

signals from the OS execute enclave handlers. Because these are closed-source we cannot comprehensively analyze their behavior. We suspect that they allow all signals because they introduce a new ecall which enables the OS to send signals to the enclaves.

### C. Occlum

Occlum is an open-source library OS that enables multi-process applications by executing them in a single enclave [53]. As shown in Figure 5(c), Occlum operates under the threat model of untrusted co-resident processes inside an enclave [54]. It performs inter-process isolation using Intel MPX for software-fault isolation. Crucially, it allows application processes in the enclave to send signals to each other. In this threat model that assumes untrusted processes inside the enclave, SIGY can inject signals from one process to another. For Occlum, we do not need the untrusted OS to inject signals into the enclave. Instead, we can use malicious attacker-controlled processes to send signals to the victim process in the same enclave. The library OS in Occlum does not filter such injections and simply forwards the signals from the attacker-controlled process to the victim process. For completeness, we checked if the OS can inject signals into the enclave to trigger signal handlers. We report that while we can inject signals, these signals do not result in the enclave executing its signal handlers. Instead, the library OS invokes the kernel's default signal handling which always crashes the process.

### D. Other Runtimes

**Runtimes without signal support.** Of the runtimes and library OSes that we surveyed, the language runtimes GoTEE, EGo, and Enarx are not vulnerable to SIGY because they do not support signal handling [9], [35], [10]. Similarly, EdgelessRT (a library OS) and RustEDP (an SDK) also do not support signal handling in enclaves [8], [20]. Therefore, these also cannot be attacked using SIGY.

**Runtimes with limited signal support.** MystikOS is a library OS that allows multiple application processes to execute in a single enclave process [17]. Similar to Occlum, MystikOS implements mechanisms to allow application processes inside the enclave to send signals to each other. All signals are routed through the library OS and do not leave the enclave. Therefore, it doesn't expose interfaces for the untrusted software OS to inject signals. Unlike Occlum, it assumes that processes in the enclave mutually trust each other. So, while processes can send signals to each other, this cannot be abused by an attacker to compromise the enclave's execution. Therefore, MystikOS is not vulnerable to SIGY.

TABLE III.    SIGNAL NUMBER TO SIGNAL NAME MAPPINGS. * SIGNALS FROM HARDWARE EXCEPTIONS

| # | Name | # | Name |
|---|---|---|---|
| 1 | SIGHUP | 2 | SIGINT |
| 3 | SIGQUIT | 4* | SIGILL |
| 5* | SIGTRAP | 6 | SIGABRT/SIGIOT |
| 7* | SIGBUS | 8* | SIGFPE |
| 9 | SIGKILL | 10 | SIGUSR1 |
| 11* | SIGSEGV | 12 | SIGUSR2 |
| 13 | SIGPIPE | 14 | SIGALRM |
| 15 | SIGTERM | 16 | SIGSTKFLT |
| 17 | SIGCHLD | 18 | SIGCONT |
| 19 | SIGSTOP | 20 | SIGTSTP |
| 21 | SIGTTIN | 22 | SIGTTOU |
| 23 | SIGURG | 24 | SIGXCPU |
| 25 | SIGXFSZ | 26 | SIGVTALRM |
| 27 | SIGPROF | 28 | SIGWINCH |
| 29 | SIGIO | 30 | SIGPWR |
| 31* | SIGSYS/SIGUNUSED | | |

## V.    WHICH SIGNALS CAN WE INJECT?

In Section III and Section IV, we discussed SDKs and library OSes that are vulnerable to SIGY. Next, we systematically analyze each of the vulnerable SDKs and library OSes to determine which signal handlers are potentially of interest to the attacker. Then, we examine programming language support for signals to determine if applications written in them might be vulnerable to SIGY.

### A. SDK and Library OS

To determine which signals are of interest to SIGY in vulnerable SDKs and library OSes, we write test applications that register handlers for all signals in the range 1-31. We execute our test applications in each of the 7 vulnerable SDKs and library OSes. Then, for each run of the test application we inject a signal from the OS and check if the application: (a) executes the registered handler, (b) crashes, or (c) has no effect as shown in Table II.

**Our findings.** Our experiments show that when our test applications invoke the `rt_sigaction` system call and try to register handlers for `sigkill` and `sigstp`, the system call fails and the application always crashes for all SDKs and library OSes. This is because these signals are reserved by the operating system for process management (e.g., `sigkill` is used to force kill the process). Occlum and Scone allow handlers for all other signals except `sigkill` and `sigstp` to be registered and we report that their signal handlers are executed. EnclaveOS reserves `sigusr2` for library OS operations but forwards all other signals except `sigkill` and `sigstp` to the application and executes its registered handlers. Notably, Teaclave does not allow injecting the signals that it expects to get from hardware exceptions (`sigill`, `sigtrap`, `sigbus`, `sigfpe`, and `sigsegv`) through the signal interface used for inter-thread communication. In contrast, Gramine only

| Language | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| C | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| C++ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Java | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ |
| Python | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Go | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✶ | ✓ | ✓ | ✓ | ✓ |
| JS | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Rust | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Wasm | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ |
| Julia | ✗ | ✓ | ✗ | ✗ | ✶ | ✗ | ✗ | ✓ | ✗ | ✓ | ✗ | ✶ | ✶ | ✗ | ✗ | ✗ | ✶ | ✶ | ✗ | ✓ | ✓ | ✓ | ✓ | ✗ | ✗ | ✗ | ✗ | ✶ | ✗ | ✗ | ✗ |

| Language | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| C | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ |
| C ++ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ |
| Java | ✓ | ✓ | ✓ | ✓ | ✗ | ✗ | ✓ | ✓ | ✗ | ✓ | ✗ | ✓ | ✓ | ✓ | ✗ | ✓ | ✗ | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✓ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ |
| Python | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ |
| Go | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✗ | ✓ | ✗ | ✓ | ✓ | ✓ | ✗ | ✓ | ✗ | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| JS | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✗ | ✗ |
| Rust | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ |
| Wasm | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✓ | ✓ | ✗ | ✗ | ✓ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ |
| Julia | ✗ | ✓ | ✗ | ✓ | ✗ | ✓ | ✓ | ✓ | ✗ | ✓ | ✓ | ✓ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✓ |

TABLE VI.    COMPILER/INTERPRETER VERSION.

| Language | Version |
|---|---|
| C | gcc (GCC) 13.2.1 |
| C++ | gcc (GCC) 13.2.1 |
| Java | OpenJDK 17.0.10 |
| Python | Python 3.11.6 |
| Go | go 1.21.6 |
| Node.js | v21.6.1 |
| Rust | cargo 1.75.0 |
| Wasm | wasmtime-cli 17.0.0 |
| Julia | julia version 1.10.2 |

TABLE VII.    IMPLICIT: RUNTIME CONVERTS SIGNALS INTO S/W EXCEPTIONS. EXPLICIT: ALLOW PROCESSES TO REGISTER SIGNAL-SPECIFIC HANDLERS.

| Language | Implicit | Explicit | Language | Implicit | Explicit |
|---|---|---|---|---|---|
| C | ✗ | ✓ | Go | ✗ | ✓ |
| C++ | ✗ | ✓ | JS | ✗ | ✓ |
| Java | ✓ | ✗ | Rust | ✗ | ✓ |
| Python | ✗ | ✓ | Wasm | ✗ | ✗ |
| Julia | ✓ | ✗ | | | |

allows applications to use signals that map to a limited set of hardware exceptions (e.g., sigill, sigbus, sigfpe, and sigsegv). Openenclave only allows a small subset of signals (sigill, sigbus, sigfpe, and sigsegv) to be forwarded to the enclave through its signal handling interface.

### B. Programming Languages

The SDKs and library OSes that we analyzed allow developers to execute programs written in different languages in the enclaves. While some of them provide support for specific languages (e.g., Teaclave is used to develop and run Rust programs), others support a wide range of languages (e.g., Scone supports executing programs written in C, C++, Java, Python, Rust, Go, JavaScript).[3] Therefore, we analyzed 9 popular programming languages (see Table VI) to check if they allow applications to register and execute signal handlers. To do this, we wrote programs in each of the languages to register signal handlers for all signals from 1 to 31. We report if the corresponding signal handlers were executed for each

---

[3]We analyze server-side JS, specifically Node.js

programming language in Table IV. The language and signal pairs that execute the registered signal handlers are of specific interest for SIGY. An attacker can use these signals to compromise applications written in the corresponding language. For completeness, we also analyzed which signal handlers the programming language standard libraries or interpreters register by default i.e., when the program does not register any signal. We report our findings in Table V. Our experiments show that signal number 10 (sigusr1) and signal number 3 (sigquit) in Julia and Java respectively write debug logs to stderr and stdout. In SIGY, stdout and stderr are not accessible to the attacker and are therefore not interesting. Further, Go starts profiling on signal 27 (sigprof), however this does not lead to any changes in the program's global state. Interestingly, NodeJS starts a debug server on signal 10 (sigusr1). We exploit this behavior to demonstrate SIGY on a NodeJS application in Section VI-B. Finally, we find that all other signal handlers that programming language runtimes register do not perform any computation that changes the program's global state and simply crash the application. These handlers are therefore not interesting to SIGY.

**Our findings.** WebAssembly system interface (WASI), the standard interface definition for WebAssembly, does not yet support signals [24]. So, programs compiled to Wasm binaries which use WASI cannot register or execute signal handlers. For the other languages we broadly classify the signal support into 2 categories: explicit and implicit support (Figure 7). Programming languages that offer explicit signal support allow applications to register signal handlers for specific signals (e.g., using the signal(signum, handler) libc function) which are executed when the OS sends a signal to the enclave. Therefore, SIGY can asynchronously trigger these handlers when enclave programs are written in these languages.

Languages which provide implicit support for signals (e.g., Java, Julia) register signal handlers directly with the OS. They do not provide any constructs for the programs to register signal-specific handlers. Instead, the language runtime converts signals into software exceptions. These software exceptions are caught and handled by the applications. For example, Java converts sigfpe to ArithmeticException and forwards it to the application. The application can execute
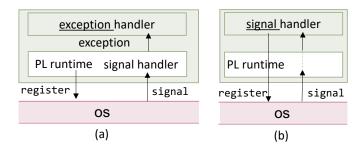
Fig. 7. (a) Languages with implicit support for signals. The PL runtime registers signal handlers with the OS and converts signals from the OS to exceptions for the applications. (b) Languages with explicit support for signals. The PL runtime allows the applications to directly register handlers with the OS which are invoked when OS sends signals.



Fig. 8. Maliciously injecting signals to Nginx to trigger insecure states (pink). C: Configuration, B: Binary.

custom handling for the exception in its `catch` block (e.g., Line 9 in Figure 1). Similarly, Julia converts `sigfpe` to `DivideError` which the application can catch and handle. SIGY can trigger these `catch` blocks to change the execution and data integrity of applications written in these programming languages.

Go uses `sigprof` for internal CPU profiling and so allows programs to register handlers for all other signals except `sigprof`. All other languages that provide explicit signal support, allow programs to register handlers for all signals between 1-31 except 9 and 19 (i.e, `sigkill` and `sigstp`). Note that, this is a direct consequence of the kernel blocking all requests to register handlers for `sigkill` and `sigstp`. On the other hand, Java only converts `sigfpe` to `ArithmeticException`. Therefore, SIGY can only use `sigfpe` to compromise programs written in Java.

In summary, of the 9 languages that we analyzed, we found that 1 does not support any signal handling, 2 only supports implicit signal handling, and the 6 others support explicit signal handling. Programs written in any of the 8 languages that provide signal support may be vulnerable to SIGY and should be reanalyzed in light of our findings.

## VI. CASE STUDIES

We confirmed our findings from Section III-Section V using hand-coded enclaves. In this section, we first discuss end-to-end case studies and explain how SIGY can be used to compromise the security when executing in enclaves. Then, we build a framework and demonstrate the feasibility of SIGY's signal injection architecture to show that it can be used to build exploits that require a large number of signal injections.

### A. Nginx

We surveyed widely used open-source webservers optimized for high uptime and found that many of them use signals to upgrade configuration without have to restart the serve. Specifically, httpd and Nginx use `sigusr1`, and squid proxy server uses `sighup` to upgrade the server's configuration. We choose to demonstrate SIGY on Nginx as it is ported by library OSes (Gramine and Scone) to run in SGX enclaves.

By default, Nginx allows a system administrator to upgrade its configuration and binaries using signals (`sighup`
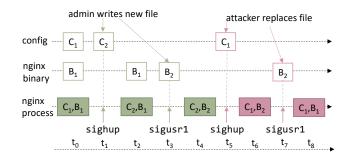
and `sigusr1`) without degrading the uptime of the server. For SIGY, this gives the attacker the ability to change the configuration and binary of the server by injecting `sighup` and `sigusr1`. Gramine, Scone, and EnclaveOS port Nginx to run in SGX enclaves. Our analysis from Section V-A shows that SIGY cannot inject `sigusr1` into enclaves running with Gramine. Further, the port of Nginx by EnclaveOS reduces the functionality of Nginx and does not allow administrators to refresh configuration files without restarting the server. Therefore, we demonstrate SIGY on Nginx using Scone.

**Benign Nginx configuration and binary upgrade.** Consider an Nginx webserver executing inside an enclave ($t_0$ in Figure 8). At time $t_1$, the Nginx administrator replaces the configuration file with an updated file and sends `sighup` to the Nginx process. On receiving this, the Nginx process (at time $t_2$) reads the new configuration file and begins to use it. Similarly, the administrator upgrades the binary of the server by writing a new file and sending `sigusr1` at time $t_3$ which is read and used by the server at $t_4$. Note that, in an enclave setting, the administrator encrypts the configuration and binary files before writing them to the OS accessible file-system. These files are only decrypted inside the enclave.

**Attacking Nginx with** SIGY**.** Because the administrator encrypts the configuration and binary files before writing them into the file-system, the OS cannot directly manipulate them to compromise Nginx. However, a malicious OS can record the encrypted blobs of configuration and binary files when the administrator writes them. The OS can also observe the new configuration and binary files when the administrator sends the signals to upgrade them in the enclave. Once the administrator has finished the configuration and binary upgrade, the OS uses SIGY to manipulate the state of the Nginx server.

Specifically, at $t_5$, the OS replaces the new configuration file with the old file it captured and sends malicious `sighup` to the enclave. The enclave reads in this configuration and updates the Nginx process ($t_6$). This ensures that the Nginx process uses an old configuration with a new binary. Furthermore, the OS can also replace the new binary file with an old one that it captured and send `sigusr1` ($t_7$). With this, the OS has successfully used SIGY to restore the Nginx process to the state that it was in at time $t_0$ before the administrator performed the upgrades. This exploit will undo any performance and security improvements provided by the new binary and configuration and can be used by the OS to bring back old security bugs (e.g., forcing the enclave to use an Nginx version before v1.23.2

which patched critical issues [6]).

### B. Node.js Server

Node.js is an open-source JavaScript runtime used for server-side scripting. By default, a Node.js server starts a debug web socket when it receives `sigusr1`, even if the server was started without debugging enabled. To demonstrate SIGY, we use Scone's Node.js port that executes an HTTPs server in an enclave. When the server is up and running, the malicious OS sends `sigusr1` to the Node.js server which opens a debug web socket. Using this web socket, the attacker can leak the server's memory and inject arbitrary code. In our exploit, we use the web socket to leak the RSA private keys used for TLS by the server using 1 injection of `sigusr1`. Note that, a malicious network adversary cannot do this exploit (e.g., by sending a socket request for the debug socket to the Node.js server).

### C. Multi-Normal Distribution

We demonstrate SIGY on a Java application, inspired by Heckler, to show how implicit signal support in programming languages make applications vulnerable [52]. JSAT is a statistical analysis tool for machine learning applications in Java. JSAT implements a `MultiVariateNormal` class that can be used to create a Multivariate Gaussian distribution. The class implements a function that updates the mean and covariance of the distribution using the `setUsingData` function as shown in Listing 6. If data added to the dataset causes an `ArithmeticException`, the `setUsingData` function discards the data and reverts the mean to the original value.

Listing 6. JSAT mean and covariance.
```
1  public boolean setUsingData(List dataSet) {
2   Vec origMean = this.mean;
3   try { //can overflow
4    Vec newMean = meanVector(dataSet);
5    Matrix covarianc = covarianceMatrix(newMean,dataSet);
6    this.mean = newMean;
7    setCovariance(covarianc);
8   } catch(ArithmeticException ex) {
9    this.mean = origMean;
10  }
11  }
```

As noted in Section V-A and Section V-B: (a) the Java runtime converts `sigfpe` to ArithmeticException that can be caught and handled by Java applications, and (b) Gramine forwards `sigfpe` to the enclaves. Therefore, we use SIGY to trigger expressive application logic in JSAT executing in Gramine by injecting `sigfpe`. We use one of the tests in JSAT to train a Learning Vector Quantization (LVQ) with Multivariate Gaussian distribution as a local classifier. This setup calls the `setUsingData` function during the training process of the classifier. We use SIGY to inject `sigfpe` every time the program executes lines $4 - 8$ in Listing 6. In our attack, we inject `sigfpe` 240 times and drop the error rate of the classifier from nearly $0\%$ to $66\%$. Our attack adds an additional overhead of overhead of $3.4$ seconds ($3.34\times$) to the training. Therefore, SIGY can be used to bias the classifier and consequently any inference that a user might execute on the model that it builds. We choose Java instead of Julia for our case study because library OSes support running Java applications in enclaves. However, if library OSes support Julia

execution, a similar attack would be possible on Julia's text analysis framework [52].

### D. Multi-Layer Perceptron

To demonstrate the feasibility of SIGY when an application requires a large number of signal injections we build a signal injection framework with sgx-step and Gramine. To test the robustness of our framework, we identify an application that requires a large number of signal injections (on the order of $10^8$). To start with, we used an open-source implementation of a multi-layer perceptron written in C [16]. This library uses `tanh` as an activation function which is called over $10^8$ times during the training process. Mathematically, $tanh(x)$ tends to 1 when $x$ tends to $\infty$. Pragmatically, this occurs in the `tanh` function if the input overflows. To handle this, we introduce a signal handler to the tanh function as shown in Listing 7. We train the MLP model with 3 hidden layers and up to 6 units in each layer (tanh activation function in hidden layers, sigmoid in output layer). Our training consists of 2000 epochs and 1096 training samples from the Banknote data set [5].

Listing 7. Tanh activation function for MLP.
```
1  jmp_buf buf;
2  void sigfpe_handler(int signum)
3    longjmp(buf, 1);
4  void tanh(int input[...])
5    ...
6    for (i = 0; i < n; i++){
7     if (sigsetjmp(buf, 1))
8      output[i+1] = 1;            // on SIGFPE
9     else
10     output[i+1] = tanh(input[i]); // no overflow
```

Using SIGY, if we inject `sigfpe` when execution is in between lines 8-10, it will trigger the signal handler and set `output[i+1]` to 1 (Line 8). We inject `sigfpe` into the `tanh` function $1,8636 \times 10^7$ times during the training of the neural network. Without SIGY, the training achieves an accuracy of $97.09\%$. With SIGY, the accuracy drops to $59.27\%$. Further, SIGY adds 970 seconds ($53.7\times$) overhead to the training process. This large overhead is primarily because of the asynchronous exits that sgx-step triggers and the signal handling in the enclave on every execution of the `tanh` function which the training invokes $\approx 10^8$ times in total.

## VII. PROOF OF CONCEPT EXPLOITS

We perform our experiments on machine setups as shown in Table VIII. We set up the laptop with sgx-step compatible configurations. The server has a newer kernel with an in-tree driver. Both machines have hardware support to store exit information in the SSA.

**Sending Signals.** We send signals to enclaves with the `util-linux` program `kill` [15]. The utility program is a wrapper around the `kill` syscall in Linux. The `kill` syscall populates the `siginfo_t` struct to indicate that the signal originated from user-space. Some of our experiments (e.g.,

injecting to Java applications) require the struct to indicate that the signal is the result of an integer division by zero error (`fpe_intdiv`) or floating point overflow (`fpe_fltovf`). To do this, we implement a kernel module that correctly populates this information and sends it to the user-space process. (98 LoC).

**Nginx.** Scone is closed-source and only supports running Nginx servers in the paid versions. Therefore, to perform our experiments we ported Nginx versions 1.22.1 and 1.24.0 to execute in a Scone v5.8.0 enclave [27]. To do this, we remove some system calls (e.g., `fcntl`) that Scone does not support. Further, we adapted Nginx to correctly propagate Scone configurations when spawning new processes (e.g., by invoking the `execve` system call). Note that our modifications do not change the functional behavior of Nginx. We build Nginx with the `select` event method and a minimal configuration to serve HTTP websites.

**Node.js.** We run a Node.js v10.14.1 web server (28 LoC) with standard TLS libraries such as express and https using Scone's v5.8.0 publicly supported Node.js port and configuration [22]. After sending `sigusr1` to Node.js, we use Chrome v120.0.6099.224 Developer tools to connect to the server, dump its memory, and extract the RSA private key used for TLS.

**MLP and JSAT.** To demonstrate SIGY on the MLP training (Section VI-D) we build a framework using sgx-step v1.5.0 [58] and Gramine commit `211ec447e` [12]. First, as a preparatory step, we run the training of the model in Gramine in debug mode, to identify the instruction pages of the `tanh` function. To bias the training process, we must inject `sigfpe` every time the `tanh` function is invoked. When the `tanh` function starts executing, sgx-step generates an asynchronous exit using timer interrupts and page faults. We should inject `sigfpe` on this event but this is not straightforward. On an asynchronous exit, control switches form the enclave to Gramine's uPAL. If we inject `sigfpe` when uPAL is executing, Gramine will crash. First, we change Gramine's uPAL (with 212 LoC) to ignore our signals to avoid it from crashing. Next, we use our framework to ensure that we inject `sigfpe` only when the enclave has resumed executing the `tanh` function. For this, we use 2 threads. The main thread executes the enclave and handles the asynchronous exit on each `tanh` invocation. Then, a worker thread injects `sigfpe` to the enclave after it is resumed. Java applications require more profiling than ahead-of-time compiled languages (e.g., C). For simplicity, we perform our experiments on JSAT by ensuring that our target function (Listing 6) waits for an `ArithmeticException` on every invocation.

## VIII. POTENTIAL DEFENSES

Current signal handling mechanisms render SDKs and library OSes vulnerable to SIGY. We propose techniques to potentially address the root cause of these issues. Then, we consider orthogonal protection techniques that can be used to diminish SIGY impact on a per-application basis.

### A. Detecting Fake Signals

**Hardware exceptions in SDKs.** All the SDKs that we analyzed implement a hardware exception handling interface that checks the exit information before invoking application-registered exception handlers as shown in Figure 2(a). Except for Intel SGX SDK, all of them introduce additional mechanisms to support enclave threads to send and receive signals. This makes the SDKs vulnerable to SIGY (Figure 2(b)) and should be hardened. First, they should not accept signals from the OS that usually are a result of hardware exceptions (e.g., `sigfpe`, `sigsegv`) as these signals should be sent through the hardware exception handling interface and never through the signal handling interface for custom signals.

**Protecting inter-thread signals in SDKs.** To prevent SIGY, the SDKs should introduce a mechanism to check if a signal that a thread receives was in fact raised by an enclave thread. For example, the trusted runtime can set up a protected shared memory between the enclave threads. Then, to indicate that a signal was raised by an enclave thread, the trusted runtime can write the signal number and target threadID in the shared memory region and only then exit the enclave with an ocall. Further, when the untrusted runtime enters the enclave with the signal (Step 5), the trusted runtime in the target thread can look up the shared memory region to check the legitimacy of the signal. This defense breaks functionality of applications that require signal injection from outside the enclave (e.g., upgrading Nginx configuration). If enclaves need signals from outside that are not raised for hardware exceptions for functionality, there is no mechanism to defend them against SIGY.

**Hardware exceptions in library OSes.** Gramine only supports signals that are raised because of hardware exceptions (c.f. Section IV-A). To stop SIGY, Gramine's exception handling can be enhanced to use the exit information that the hardware stores in the SSA. Gramine has used this technique in a patch which stops SIGY [2]. Unlike Gramine, other library OSes support signals that are both raised because of hardware exceptions (e.g., `sigfpe`) and explicitly by the processes (e.g., `sigusr1`, `sighup`). Currently, library OSes do not distinguish between signals from hardware exceptions and signals that originate from within the application in the enclave. First, their signal handling logic should be separated for these two cases. Then, for the first case of hardware exceptions, the library OSes can filter illegitimate injections using the exit information stored by the hardware. Next, we discuss mechanisms to protect against injections of other signals that processes originate.

**Interprocess-signals in library OSes.** Our analysis in Section IV shows that library OSes that create distinct enclaves for application processes route the signal through the untrusted runtime and OS which makes them susceptible to SIGY. Protecting this interface is challenging as enclave processes need a mechanism to check if the signal is from another trusted enclave process. SGX does not support setting up shared memory regions between multiple enclave processes. Therefore, a method similar to the one that we outlined for signals between threads is not straightforward to implement between enclave processes. To protect against SIGY, the library OSes should implement a message passing framework to communicate the signal information to the target trusted processes which their library OS looks up to check the signal's legitimacy.

Occlum is vulnerable to SIGY because it assumes a threat model where untrusted processes reside within the same en-

clave. While setting up a shared memory region to communicate the legitimacy of signals between processes in Occlum is easier, establishing the trust relationship between the processes is more challenging. To this end, Occlum should introduce a notion of attestation across process groups inside the same enclave. Currently, all processes inside the enclave are mutually distrusting so this would be a fundamental architectural change to Occlum. With the trusted process groups notion, Occlum can use a message-passing mechanism to protect against SIGY. Finally, library OSes should not accept signals from untrusted software. If such signals that are not a result of hardware exceptions are required for functionality then there is no defense against SIGY.

### B. Limiting SIGY Impact

SIGY relies on the fact that the signal is propagated to the enclave application by the runtime or library OS. Once the signal is sent to the application, it will execute the signal handler in the enclave which compromises the enclave's execution integrity. To amplify the effect of the signal handler, the attacker might need other capabilities. For example, to compromise Nginx (Section VI-A), the attacker needs the ability to replace encrypted blobs of files to change the configuration or binary that the Nginx server uses. Similarly, for Node.js (Section VI-B) the attacker needs to be able to connect to the port. If there are orthogonal protections like file-system rollback prevention, network firewalls, or dynamic attestation in place, these end-to-end attacks will be stopped. Of the library OSes we surveyed, only Gramine and Scone provide protections for filesystems and network. Gramine has filesystem protections to add trusted (i.e., hashed), and encrypted files to the enclave, but does not provide rollback protection. Scone paper does not discuss any port or protocol filtering in network protection for enclaves; when filesystem protection is enabled it ensures that the files are encrypted, rollback, and integrity protected. Next, we checked publicly available Scone versions—its network shields can be configured during enclave creation with a whitelist of allowed ports. However, this network shield is not activated by default allowing SIGY to compromise Node.js. Thus, developers should consider the new attack surface introduced by SIGY when configuring enclaves. Lastly, orthogonal protections cannot stop SIGY attacks which exploit `sigfpe` and compromise enclaves as these attacks directly affect enclave memory and do not need any external subsystem (e.g., network, files). Therefore, library OSes should implement comprehensive defenses for this new attack surface.

### C. Detecting SIGY

SIGY injects signals into enclaves which cause asynchronous exits. One way to protect an enclave against SIGY is to detect if an asynchronous exit was caused by fake signal injections. Intel introduced a hardware extension called AEX-notify to perform checks on reentry after an interrupt or exception [4]. AEX-notify proposes a defense to prevent precise single-stepping of enclaves using the new hardware [32]. For this, they implement handlers for timer-interrupts that speed up the execution of the successive instructions. AEX-notify cannot mitigate SIGY when the applications register handlers that only need to be invoked once (Section VI-B and Section VI-A)

to compromise the enclave. Specifically, the OS can send the signal when the enclave has legitimately exited (e.g., timer interrupt, page-fault) without AEX-notify detecting it. This is because SIGY will not cause any asynchronous exits, let alone require any single-stepping.

On the other hand, our attacks that need to send signals when the enclave application is executing a specific set of instructions can be harder, but not impossible, to perform in the presence of AEX-notify. However, this does not completely stop SIGY. Note that, for SIGY the attacker does not need to single-step the enclave. Instead, the attacker needs a mechanism to determine when the enclave is executing a set of instructions. Further, if the size of this instruction set is larger than the number of instructions executed between timer interrupts, there will always be a legitimate enclave exit for the OS to then inject the signal. Then, when a genuine timer interrupt occurs in between these instructions, the attacker can inject the signal. For applications where this is not the case, the frequency of timer interrupts is controlled by the untrusted OS. Therefore, if needed, the attacker can tune the frequency to any value to send signals to enclaves when the enclave executes the target instruction set.

## IX. RELATED WORK

**Malicious Synchronous Interfaces.** Ports et al. comprehensively analyze the threats to applications from an untrusted OS [48]. They perform this analysis in the context of Overshadow [29], a framework that protects the confidentiality and integrity of applications in the presence of a malicious operating system and a trusted hypervisor. They note that because the untrusted OS manages the signals it can maliciously alter them. However, the analysis focuses on synchronous attacks from the OS that redirect signals or send bad return values. Iago builds on this, where the attacker *synchronously* manipulates the return values for system calls when the enclave makes *explicit* requests [28]. These bad return values can be used to trick the enclave into performing unintended computations. Iago demonstrates how the limited sanitization in the system call interface between an enclave and OS can be used to synchronously alter enclave execution. This synchronous interface has been comprehensively studied and analyzed in the context of SGX enclaves [39], [43], [56], [25].

**Malicious Asynchronous Interfaces.** Previous works have demonstrated using *asynchronous* timer interrupts and page faults at arbitrary points of the enclave's execution [58]. The attacker manipulates the software to trigger these hardware events, which then trigger handlers that have *fixed* effects—timer interrupts and page faults cause an exit from the enclave, but do not execute any handlers in the enclave that change the program state. In contrast, to the best of our knowledge, SIGY is the first work on Intel SGX that injects signals and exceptions *asynchronously* at any point during the enclave's execution. Depending on enclave-specific logic, such arbitrary signal injection can induce *varied effects* depending on the signal that is injected and the logic that the enclave has specified for it. This allows SIGY to execute such signal handlers at any point during enclave execution to bring about programmatic changes to the enclave's state.

**Bugs in Enclave runtimes and Library OSes.** Previous works have studied and demonstrated attacks on buggy enclave-

OS interface implementations (e.g., Application Binary Interface, Application Programming Interface) [56], [43]. Async-Shock and Game of Threads exploit synchronization bugs in multi-threaded enclave applications by interrupting the threads at specific points in their execution [62] or using race-conditions [50]. Along similar lines, SmashEx demonstrates attacks that use the lack of atomicity in signal handlers to compromise enclave execution [34]. Unlike these previous works, SIGY does not rely on bugs in the runtimes, library OSes, or signal handlers.

**Bugs in Enclave applications.** Prior works have used vulnerabilities in enclave application code to compromise it using code-reuse attacks [44]. SIGY does not rely on bugs in enclave application implementations. However, SIGY can be used to bring back old bugs. For example, by forcing Nginx to use old binaries, SIGY brings back vulnerabilities in older Nginx versions.

**Detection Tools.** Several works have used fuzzing [31], [42] and symbolic execution [25], [30], [43], [61] to detect vulnerabilities in enclaves. Notably, SEnFuzzer uses fuzzing to detect ocall and ecall interface bugs [65]. We manually analyzed the runtimes and library OSes to find vulnerable runtimes, library OSes, and signal handlers that SIGY can exploit. These detection tools can be used to further analyze applications with the intention of finding other effects of signal handling in enclaves. Lefeuvre et al., investigate the fragility of interfaces for software compartmentalization and build a fuzzer to detect interface vulnerabilities but do not consider signals [46]. Beyond Intel SGX, SIGY's observations can be applied to other such avenues of research.

**Side-channels.** Prior works have used timer interrupts and page faults to leak side-channel information [60], [64], [38], [49], [59]. Similarly, others have leveraged other side-channels to compromise enclaves [36], [47]. Sgx-step is a framework that enables attackers to precisely single-step enclaves [57]. While SIGY is not a side-channel attack, side-channel information can be used to amplify the effects of SIGY.

**Kernel bugs due to interrupts and signals.** Buggy implementation of signal handlers in the kernel can compromise the security of applications by inducing memory-corruption from race conditions in the handlers [1]. ExpRace exploits kernel races by using interrupts in a non-TEE setting [45]. Prior works have analyzed the possibility of using buggy signal handlers and interrupt remapping to corrupt applications executing with trusted OSes [7], [63]. Signal handlers have been known to have race condition bugs that can be exploited to compromise applications [1]. SIGY does not assume any buggy or racy signal handlers.

**Malicious Interrupts.** There are three ways of delivering notifications to a user-level computation: interrupts, exceptions, and signals. Since Intel TDX and AMD SEV offer a VM abstraction, the attacker can only use interrupts as a notification mechanism. Specifically, when the hypervisor delivers an interrupt, the guest OS can either handle it in its kernel with an interrupt handler or convert it to an exception or signal for the user process currently executing on the vCPU. A recent series of works dubbed Ahoi attacks, exploit CVMs provided by AMD-SEV and Intel TDX by injecting malicious interrupts from hypervisors [52], [51]. In particular, Heckler injects

interrupts which are then delivered as signals to the victim user-level program. SIGY is inspired by Ahoi attacks, but investigates Intel SGX which is a user-level abstraction. SIGY shows that a malicious OS can deliver interrupts, exceptions, and signals to an enclave. When a CPU executing an enclave triggers a hardware interrupt (e.g., INT0 for divide by zero faults), Intel SGX hardware sets the SSA to indicate that the interrupt did indeed originate on the core. This way, the enclave software can check the SSA before executing the handler. Fortunately, our investigation shows that all current enclave SDKs (Intel SGX SDK, Open Enclave, Teaclave SGX-SDK) do indeed perform the interrupt authenticity check before executing the handler. However, this covers only a fraction of the attack surface. The OS can still deliver software-based exceptions and signals directly to the enclave. Since these notifications are not based on hardware-events (e.g., executing an illegal instruction) triggered by the enclave, the CPU cannot provide any information about their authenticity. Thus, SIGY is an instance of Ahoi attack that uses exceptions and signals as the notification mechanism to compromise Intel SGX enclaves.

**SIGY on Arm TEEs.** Besides SGX, several interface attacks have been demonstrated on Arm TrustZone [41], [55]. Further, prior works have used fuzzing to test TrustZone's trusted software [37]. Unlike SGX, TrustZone does not expose signal interfaces to the malicious OS and filters interrupts from the untrusted OS to the trusted applications [26]. Therefore SIGY cannot be used to attack TrustZone. Heckler shows that Arm CCA's protected VMs are not vulnerable to attacks using malicious interrupts because of Arm's interrupt architecture [52]. Specifically, Arm's interrupt architecture does not map interrupts to exceptions. Therefore, a malicious hypervisor cannot abuse interrupts to induce exceptions or trigger signal handlers in Arm CCA's VM. Furthermore, in Arm CCA's protected VM setting, the guest OS is trusted and the VM does not accept exceptions and signals from the untrusted hypervisor. Arm CCA is not vulnerable to SIGY.

## X. CONCLUSION

SIGY demonstrates that a malicious OS can exploit Intel SGX enclaves by delivering malicious exceptions and signals to trick the enclave into executing handlers. Our analysis of various runtimes and library OSes shows that they are vulnerable to SIGY. Programming languages as well as native and ported enclave-bound programs that need exceptions and signal handling should consciously to choose between functionality and security.

## REFERENCES

[1] "Unsafe function call from a signal handler ," https://owasp.org/www-community/vulnerabilities/Unsafe_function_call_from_a_signal_handler, accessed 17.04.2024.

[2] "[PAL/Linux-SGX] Cross-verify SW signals vs HW exceptions," https://github.com/gramineproject/gramine/commit/a390e33e16ed374a40de2344562a937f289be2e1, accessed 22.04.2024.

[3] "Asylo Github," https://github.com/google/asylo, accessed 28.01.2024.

[4] "Asynchronous Enclave Exit Notify and the EDECCSSA User Leaf Function," https://www.intel.com/content/www/us/en/content-details/736463/white-paper-asynchronous-enclave-exit-notify-and-the-edeccssa-user-leaf-function.html, accessed 28.01.2024.

[5] "Banknote Database," https://banknotedb.com, accessed 28.01.2024.

[6] "CVE-2022-41741," https://nvd.nist.gov/vuln/detail/CVE-2022-41741, accessed 28.01.2024.

[7] "Delivering Signals for Fun and Profit," https://lcamtuf.coredump.cx/signals.txt, accessed 28.01.2024.

[8] "Edgelessrt Github," https://github.com/edgelesssys/edgelessrt, accessed 28.01.2024.

[9] "Ego Github," https://github.com/edgelesssys/ego, accessed 28.01.2024.

[10] "Enarx Github," https://github.com/enarx/enarx, accessed 28.01.2024.

[11] "Fortanix Runtime Encryption® Platform," https://resources.fortanix.com/hubfs/Fortanix\_RTE\_Platform\_Whitepaper.pdf, accessed 28.01.2024.

[12] "Gramine Github," https://github.com/gramineproject/gramine/tree/211ec447ee69f16139520fc3a17c561c36a00943, accessed 28.01.2024.

[13] "Intel SDK," https://github.com/intel/linux-sgx, accessed 28.01.2024.

[14] "JSON Web Token (JWT)," https://datatracker.ietf.org/doc/html/rfc7519, accessed 28.01.2024.

[15] "kill(1) — Linux manual page," https://man7.org/linux/man-pages/man1/kill.1.html, accessed 28.01.2024.

[16] "Multi Layer Perceptron in C," https://github.com/manoharmukku/multilayer-perceptron-in-c, accessed 28.01.2024.

[17] "MystikOS Github," https://github.com/deislabs/mystikos, accessed 28.01.2024.

[18] "Openenclave Github," https://github.com/openenclave/openenclave, accessed 28.01.2024.

[19] "POSIX.1-2017," https://pubs.opengroup.org/onlinepubs/9699919799.2018edition/, accessed 28.01.2024.

[20] "Rust EDP Github," https://github.com/fortanix/rust-sgx/tree/master, accessed 28.01.2024.

[21] "Scone," https://sconedocs.github.io, accessed 28.01.2024.

[22] "Scone Nodejs," https://sconedocs.github.io/Nodejs, accessed 28.01.2024.

[23] "Teaclave Github," https://github.com/apache/incubator-teaclave-sgx-sdk, accessed 28.01.2024.

[24] "WASI: signal handling," https://github.com/WebAssembly/WASI/issues/166, accessed 28.01.2024.

[25] F. Alder, L.-A. Daniel, D. Oswald, F. Piessens, and J. Van Bulck, "Pandora: Principled symbolic validation of intel sgx enclave runtimes."

[26] ARM, "Learn the Architecture: TrustZone for AArch64," https://developer.arm.com/architectures/learn-the-architecture/trustzone-for-aarch64/trustzone-in-the-processor, 2021.

[27] S. Arnautov, B. Trach, F. Gregor, T. Knauth, A. Martin, C. Priebe, J. Lind, D. Muthukumaran, D. O'Keeffe, M. L. Stillwell, D. Goltzsche, D. Eyers, R. Kapitza, P. Pietzuch, and C. Fetzer, "SCONE: Secure linux containers with intel SGX," in *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. Savannah, GA: USENIX Association, Nov. 2016, pp. 689–703. [Online]. Available: https://www.usenix.org/conference/osdi16/technical-sessions/presentation/arnautov

[28] S. Checkoway and H. Shacham, "Iago attacks: why the system call api is a bad untrusted rpc interface," *ACM SIGARCH Computer Architecture News*, vol. 41, no. 1, pp. 253–264, 2013.

[29] X. Chen, T. Garfinkel, E. C. Lewis, P. Subrahmanyam, C. A. Waldspurger, D. Boneh, J. Dwoskin, and D. R. Ports, "Overshadow: A virtualization-based approach to retrofitting protection in commodity operating systems," *SIGOPS Oper. Syst. Rev.*, vol. 42, no. 2, p. 2–13, Mar. 2008.

[30] T. Cloosters, M. Rodler, and L. Davi, "TeeRex: Discovery and exploitation of memory corruption vulnerabilities in SGX enclaves," in *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020, pp. 841–858. [Online]. Available: https://www.usenix.org/conference/usenixsecurity20/presentation/cloosters

[31] T. Cloosters, J. Willbold, T. Holz, and L. Davi, "SGXFuzz: Efficiently synthesizing nested structures for SGX enclave fuzzing," in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 3147–3164. [Online]. Available: https://www.usenix.org/conference/usenixsecurity22/presentation/cloosters

[32] S. Constable, J. V. Bulck, X. Cheng, Y. Xiao, C. Xing, I. Alexandrovich, T. Kim, F. Piessens, M. Vij, and M. Silberstein, "AEX-Notify: Thwarting precise Single-Stepping attacks through interrupt awareness for intel SGX enclaves," in *32nd USENIX Security Symposium (USENIX Security 23)*. Anaheim, CA: USENIX Association, Aug. 2023, pp. 4051–4068. [Online]. Available: https://www.usenix.org/conference/usenixsecurity23/presentation/constable

[33] V. Costan and S. Devadas, "Intel sgx explained." *IACR Cryptology ePrint Archive*, vol. 2016, no. 086, pp. 1–118, 2016.

[34] J. Cui, J. Z. Yu, S. Shinde, P. Saxena, and Z. Cai, "Smashex: Smashing sgx enclaves using exceptions," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '21, 2021.

[35] A. Ghosn, J. R. Larus, and E. Bugnion, "Secured routines: Language-based construction of trusted execution environments," in *2019 USENIX Annual Technical Conference (USENIX ATC 19)*. Renton, WA: USENIX Association, Jul. 2019, pp. 571–586. [Online]. Available: http://www.usenix.org/conference/atc19/presentation/ghosn

[36] J. Götzfried, M. Eckert, S. Schinzel, and T. Müller, "Cache attacks on intel sgx," in *Proceedings of the 10th European Workshop on Systems Security*, ser. EuroSec'17. New York, NY, USA: Association for Computing Machinery, 2017. [Online]. Available: https://doi.org/10.1145/3065913.3065915

[37] L. Harrison, H. Vijayakumar, R. Padhye, K. Sen, and M. Grace, "PARTEMU: Enabling dynamic analysis of Real-World TrustZone software using emulation," in *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020, pp. 789–806. [Online]. Available: https://www.usenix.org/conference/usenixsecurity20/presentation/harrison

[38] W. He, W. Zhang, S. Das, and Y. Liu, "Sgxlinger: A new side-channel attack vector based on interrupt latency against enclave execution," in *2018 IEEE 36th International Conference on Computer Design (ICCD)*. IEEE, 2018, pp. 108–114.

[39] O. S. Hofmann, S. Kim, A. M. Dunn, M. Z. Lee, and E. Witchel, "Inktag: Secure applications on an untrusted operating system," *SIGPLAN Not.*, vol. 48, no. 4, p. 265–278, Mar. 2013.

[40] Intel, "Intel software guard extensions," https://software.intel.com/content/www/us/en/develop/topics/software-guard-extensions.html, accessed 28.01.2024.

[41] P. Jiang, Q. Wang, J. Cheng, C. Wang, L. Xu, X. Wang, Y. Wu, X. Li, and K. Ren, "Boomerang: Metadata-Private messaging under hardware trust," in *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. Boston, MA: USENIX Association, Apr. 2023, pp. 877–899. [Online]. Available: https://www.usenix.org/conference/nsdi23/presentation/jiang

[42] A. Khan, M. Zou, K. Kim, D. Xu, A. Bianchi, and D. J. Tian, "Fuzzing sgx enclaves via host program mutations," in *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*, 2023, pp. 472–488.

[43] M. R. Khandaker, Y. Cheng, Z. Wang, and T. Wei, "Coin attacks: On insecurity of enclave untrusted interfaces in sgx," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 971–985. [Online]. Available: https://doi.org/10.1145/3373376.3378486

[44] J. Lee, J. Jang, Y. Jang, N. Kwak, Y. Choi, C. Choi, T. Kim, M. Peinado, and B. B. Kang, "Hacking in darkness: Return-oriented programming against secure enclaves," in *26th USENIX Security Symposium (USENIX Security 17)*. Vancouver, BC: USENIX Association, Aug. 2017, pp. 523–539. [Online]. Available: https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/lee-jaehyuk

[45] Y. Lee, C. Min, and B. Lee, "ExpRace: Exploiting kernel races through raising interrupts," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 2363–2380. [Online]. Available: https://www.usenix.org/conference/usenixsecurity21/presentation/lee-yoochan

[46] H. Lefeuvre, V.-A. Bădoiu, Y. Chien, F. Huici, N. Dautenhahn, and P. Olivier, "Assessing the impact of interface vulnerabilities in compartmentalized software," *arXiv preprint arXiv:2212.12904*, 2022.

[47] O. Oleksenko, B. Trach, R. Krahn, M. Silberstein, and C. Fetzer, "Varys: Protecting SGX enclaves from practical Side-Channel attacks," in *2018 USENIX Annual Technical Conference (USENIX ATC 18)*. Boston,

MA: USENIX Association, Jul. 2018, pp. 227–240. [Online]. Available: https://www.usenix.org/conference/atc18/presentation/oleksenko

[48] D. R. K. Ports and T. Garfinkel, "Towards application security on untrusted operating systems," in *Proceedings of the 3rd Conference on Hot Topics in Security*, ser. HOTSEC'08. USA: USENIX Association, 2008.

[49] I. Puddu, M. Schneider, M. Haller, and S. Capkun, "Frontal attack: Leaking Control-Flow in SGX via the CPU frontend," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 663–680. [Online]. Available: https://www.usenix.org/conference/usenixsecurity21/presentation/puddu

[50] J. R. Sanchez Vicarte, B. Schreiber, R. Paccagnella, and C. W. Fletcher, "Game of threads: Enabling asynchronous poisoning attacks," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 35–52. [Online]. Available: https://doi.org/10.1145/3373376.3378462

[51] B. Schlüter, S. Sridhara, A. Bertschi, and S. Shinde, "WeSee: Using Malicious #VC Interrupts to Break AMD SEV-SNP," in *IEEE S&P*, 2024.

[52] B. Schlüter, S. Sridhara, M. Kuhne, A. Bertschi, and S. Shinde, "Heckler: Breaking Confidential VMs with Malicious Interrupts," in *USENIX Security*, 2024.

[53] 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 *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 955–970. [Online]. Available: https://doi.org/10.1145/3373376.3378469

[54] ——, "Occlum: Secure and efficient multitasking inside a single enclave of intel sgx," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020, pp. 955–970.

[55] D. Suciu, S. McLaughlin, L. Simon, and R. Sion, "Horizontal privilege escalation in trusted applications," in *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020. [Online]. Available: https://www.usenix.org/conference/usenixsecurity20/presentation/suciu

[56] J. Van Bulck, D. Oswald, E. Marin, A. Aldoseri, F. D. Garcia, and F. Piessens, "A tale of two worlds: Assessing the vulnerability of enclave shielding runtimes," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1741–1758. [Online]. Available: https://doi.org/10.1145/3319535.3363206

[57] J. Van Bulck and F. Piessens, "Sgx-step: An open-source framework for precise dissection and practical exploitation of intel sgx enclaves."

[58] J. Van Bulck, F. Piessens, and R. Strackx, "Sgx-step: A practical attack framework for precise enclave execution control," in *Proceedings of the 2nd Workshop on System Software for Trusted Execution*, ser. SysTEX'17. New York, NY, USA: Association for Computing Machinery, 2017. [Online]. Available: https://doi.org/10.1145/3152701.3152706

[59] ——, "Nemesis: Studying microarchitectural timing leaks in rudimentary cpu interrupt logic," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 178–195. [Online]. Available: https://doi.org/10.1145/3243734.3243822

[60] J. Van Bulck, N. Weichbrodt, R. Kapitza, F. Piessens, and R. Strackx, "Telling your secrets without page faults: Stealthy page {Table-Based} attacks on enclaved execution," in *26th USENIX Security Symposium (USENIX Security 17)*, 2017, pp. 1041–1056.

[61] Y. Wang, Z. Zhang, N. He, Z. Zhong, S. Guo, Q. Bao, D. Li, Y. Guo, and X. Chen, "Symgx: Detecting cross-boundary pointer vulnerabilities of sgx applications via static symbolic execution," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '23. New York, NY, USA:

Association for Computing Machinery, 2023, p. 2710–2724. [Online]. Available: https://doi.org/10.1145/3576915.3623213

[62] N. Weichbrodt, A. Kurmus, P. Pietzuch, and R. Kapitza, "Asyncshock: Exploiting synchronisation bugs in intel sgx enclaves," in *Computer Security–ESORICS 2016: 21st European Symposium on Research in Computer Security, Heraklion, Greece, September 26-30, 2016, Proceedings, Part I 21*. Springer, 2016, pp. 440–457.

[63] R. Wojtczuk and J. Rutkowska, "Following the White Rabbit: Software attacks against Intel (R) VT-d technology," https://invisiblethingslab.com/resources/2011/Software\%20Attacks\%20on\%20Intel\%20VT-d.pdf, 2011.

[64] Y. Xu, W. Cui, and M. Peinado, "Controlled-channel attacks: Deterministic side channels for untrusted operating systems," in *IEEE S&P*, 2015.

[65] D. Yu, J. Wang, H. Fang, Y. Fang, and Y. Zhang, "Senfuzzer: Detecting sgx memory corruption via information feedback and tailored interface analysis," in *Proceedings of the 26th International Symposium on Research in Attacks, Intrusions and Defenses*, ser. RAID '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 485–498. [Online]. Available: https://doi.org/10.1145/3607199.3607215