



RIPencapsulation: Defeating IP Encapsulation on TI MSP Devices

Prakhar Sah and Matthew Hicks, *Virginia Tech*

<https://www.usenix.org/conference/woot24/presentation/sah>

This paper is included in the Proceedings of the
18th USENIX WOOT Conference on Offensive Technologies.

August 12-13, 2024 • Philadelphia, PA, USA

ISBN 978-1-939133-43-4

Open access to the
Proceedings of the 18th USENIX WOOT
Conference on Offensive Technologies
is sponsored by USENIX.



RIPencapsulation: Defeating IP Encapsulation on TI MSP Devices

Prakhar Sah
Virginia Tech
sprakhar@vt.edu

Matthew Hicks
Virginia Tech
mdhicks2@vt.edu

Abstract

Internet of Things (IoT) devices sit at the intersection of unwieldy software complexity and unprecedented attacker access. This unique position comes with a daunting security challenge: how can we protect both proprietary code and confidential data on a device that the attacker has unfettered access to? Trusted Execution Environments (TEEs) promise to solve this challenge through hardware-based separation of trusted and untrusted computation and data. While TEEs do an adequate job of protecting secrets on desktop-class devices, we reveal that trade-offs made in two of the most widely-used commercial IoT devices undermine their security.

This paper uncovers two fundamental weaknesses in IP Encapsulation (IPE), the TEE deployed by Texas Instruments for MSP430 and MSP432 devices. We observe that lack of call site enforcement and residual state after unexpected TEE exits enable an attacker to reveal all proprietary code and secret data within the IPE. We design and implement an attack called RIPencapsulation, which systematically executes portions of code within the IPE and uses the partial state revealed through the register file to exfiltrate secret data and to identify gadget instructions. The attack then uses gadget instructions to reveal all proprietary code within the IPE. Experiments with commodity devices and a production compiler show that—even after following all manufacturer secure coding practices—RIPencapsulation reveals, within minutes, both the code and keys from third-party cryptographic software, as well as allowing unrestricted writes to TEE memory.

1 Introduction

The global IoT industry is projected to become a trillion-dollar industry by 2027 [2]. IoT devices are widely deployed in both safety- and mission-critical roles in government, healthcare, transportation, manufacturing, defense, and telecommunications industries. Thus, these devices are a treasure trove of sensitive information. Data security concerns are a major obstacle to the growth of the IoT sensor market as data breaches

continue to rise with the advancement of technology. In addition to data security, the growing software complexity of IoT devices and proliferation of Artificial Intelligence mandate code security, i.e., the protection of proprietary algorithms and models. Failing to protect proprietary code and secret data puts both consumers and companies at risk.

“Trying to design information security solutions without due consideration of the complex human nature may prove to be an Achilles heel” [9]. Cryptographic algorithms like AES and RSA provide confidentiality of data, but the key still ends up in device memory. This leaves keys vulnerable to exfiltration by an attacker with physical access. Unfortunately, physical access is the common case for IoT devices. To address the threat of co-resident attackers, manufacturers provide a Trusted Execution Environment (TEE). TEEs bifurcate hardware (either physically or virtually) into security domains, where code and data in the high-security domain are protected from the low-security domain. For IoT-class devices, Texas Instruments provides a TEE called IP Encapsulation (IPE), which physically partitions device memory into a protected region and an unprotected region.

Texas Instruments is the world’s second-largest manufacturer of microcontrollers [7]. Their MSP family of devices consists of over 2000 unique devices [3] making them one of the most widely deployed microcontrollers [4, 5]. MSP430s are 16-bit industrial-grade microcontrollers with low power consumption at a low cost. The MSP432 line of microcontrollers extends the capabilities of the MSP430 with a 32-bit, ARM-based architecture. Both devices have TEE support in the form of IPE, making IPE the most widely deployed TEE.

While the community continues to probe the security of TEEs provided by higher-end devices, the security of IoT-class TEEs remains under-explored. This paper fills that gap by analyzing the security of TI’s IPE. IPE protects code and data within the IPE zone from all non-IPE zone read and write accesses [37, 45, 48]. IPE is enforced by the Memory Protection Unit (MPU) for MSP430 and the System Controller (SYSCTL) module for MSP432, which restricts direct external IPE zone accesses by checking the origin of mem-

ory accesses. In addition, the MPU/SYSCTL also restricts all JTAG/DMA accesses to the IPE zone. Thus, only memory accesses from within the IPE are allowed to the IPE zone, making it ideal for storing secret data (e.g., keys) and proprietary code (e.g., AI models). Our exploration of IPE security reveals that—despite following the TI-recommended secure programming practices—the state-of-art compilers produce assemblies that leak information from and allow unrestricted writes to the IPE zone via the unprotected register file. These weaknesses render IP encapsulation insecure.

IPE has two weaknesses that undermine security by enabling **arbitrary read and write access to secure memory**:

- IPE allows all code outside the IPE zone to branch to arbitrary instructions within the IPE zone.
- When execution leaves the IPE zone unexpectedly (e.g., via an interrupt) the contents of the register file remain.

We construct an attack leveraging these weakness called RIPencapsulation, which exfiltrates all code and data protected by the IPE—within minutes. Inspired by SGX-Step [49] and Interrupt-oriented Bugdoor Programming [43], but extending them to the IoT domain, RIPencapsulation combines interrupt-based control flow attack patterns, with data-oriented attack patterns, and side-channel attack patterns to break IPE. **The result of RIPencapsulation is unrestricted read, write, and execute access to IPE-protected memory.**

We implement and evaluate RIPencapsulation on TI MSP430 and MSP432 devices. Our evaluation shows that RIPencapsulation reveals the entire contents of the IPE zone with third-party implementations of AES, SHA256, and RSA, using a variety of optimization settings on a production compiler. RIPencapsulation exfiltrates all keys and code from both devices, automatically, within minutes—even when following all of TI’s secure coding best-practices.

This paper makes the following technical contributions:

- **Create a side channel:** We design and implement an interrupt-based side channel that reveals the state within the IPE zone of TI MSP devices (§4, §5, §6).
- **Reconstruct firmware:** We use partial state, timing, and size information to reconstruct the IP-encapsulated assembly (§4.2, §5.2).
- **Demonstrate generality:** We show the problem is pervasive across TI MSP devices, memory types, cryptographic implementations, and compiler optimizations (§7, §7.6).
- **Modify firmware:** We uncover interrupt-based data-oriented gadgets to modify the IP-encapsulated code and data (§7.7).
- **Discuss mitigation:** We qualitatively analyze various mitigation techniques with respect to their impact on an IoT device (§8).
- **Open-source RIPencapsulation.** We open-source all attack and evaluation code [8].




Figure 1: Block diagram illustrating how the MPU enforces IP Encapsulation on MSP430 devices.

Responsible disclosure. We have disclosed RIPencapsulation and the threat it poses to the security guarantees of IP Encapsulation on MSP-class devices to Texas Instruments.

2 Background

This section covers topics underpinning RIPencapsulation.

2.1 Trusted Execution Environment

A Trusted Execution Environment (TEE) is a secure processing area that provides isolated execution and secure storage of code and data inside a tamper-resistant module. TEEs include process isolation, the integrity of applications running inside the TEE, and the confidentiality of data associated with it. TEEs achieve this by using a memory protection mechanism restricting access to the security-critical software module [39].

Intel Software Guard Extensions (SGX) and AMD Secure Encrypted Virtualization (SEV) are multi-tenant systems running with operating systems and full-fledged commodity desktop/server-class systems. They use hardware memory-based encryption to isolate sensitive code and data. This type of deployment scenario is not available in low-end microcontroller devices like the MSPs due to the hardware overhead of the encryption engine and the run time overhead of encryption in an environment with single-cycle main memory access.

ARM TrustZone is a microcontroller-focused TEE that splits the physical memory into secure and normal regions, with specifically assigned computational units like Secure Attribution Unit (SAU), ensuring safe context switches between secure and normal processes, or lightweight remote attestation schemes like SMART [21] and VRASED [35], which provide user trust with minimal HW/SW modifications. SANCUS [34] is another secure architecture that provides

isolated execution and privacy of code and data in low-level networked embedded systems.

The only TEE for the low-end devices common to large-scale IoT deployments is TI’s Intellectual Property Encapsulation (IPE). IPE protects the encapsulated memory region from all direct non-IPE accesses, whether it be from on-chip execution or off-chip via the debugger, i.e., only program code executed from the IPE region itself has access privileges to IPE code and data. Figure 1 illustrates the IPE implementation for MSP430 devices: the Memory Protection Unit (MPU) verifies IPE region accesses by snooping the Memory Address Bus (MAB) and the program counter to check whether the access request is made by code in the IPE zone. Any unauthorized access to the IPE zone causes the MPU to drive the Memory Data Bus (MDB) with $0x3FFF$. To execute code stored inside the IPE segment, the program must call functions within the IPE zone. The MSP432’s IPE mechanism is slightly different technically, but works in a similar manner and has the same flaws as the MSP430 IPE (§4).

Existing attacks demonstrate ways to undermine deployed TEEs. CLKscrew [44] exploits software-exposed energy management mechanisms to introduce faults in the secure part of the memory, exfiltrating IP from TrustZone. Volt Boot [32] is another attack that demonstrates the vulnerability of on-chip volatile memories due to the physical separation common to modern system-on-chip power distribution networks. Recent research reveals vulnerabilities in TEEs like ARM TrustZone [15] and AMD SEV [30, 31], calling into question the reliability of these protection mechanisms. **Our work adds to the TEE attack literature with an attack on TI’s IPE.**

2.2 Code-reuse Attacks

Attackers turn to code reuse attacks when injecting their own code is prevented. Instead of injecting new code, the attacker constructs malicious functionality by chaining existing code snippets, called “gadgets” found in the target program. These gadgets are short sequences of instructions, typically ending with a “return” instruction. By crafting a chain of gadgets and manipulating the program’s control flow, the attacker is able to execute arbitrary commands. Many techniques exist in the literature that utilize various aspects of the memory control plane [11–13, 16, 40] and data plane [18, 25] to create Turing-complete gadgets for malicious purposes. We take inspiration from such techniques, replacing a `ret` with a carefully crafted interrupt and leveraging data-oriented attack principles.

2.3 Context Switches and Call Site Verification

Context switching is the process of transitioning the execution from the untrusted zone to the secure/trusted zone and vice versa. When a context switch occurs, the current state of the regular execution environment, including registers, memory contents, and program counter is saved, and the system

transitions between zones. Secure context switching in TEEs is crucial for ensuring the isolation of trusted code and data from the rest of the system as any data that remains serves as a side channel to untrusted software.

In addition to context clearing, a TEE must enforce call site verification to ensure that untrusted code interfaces with trusted code in acceptable ways. When a program interacts with the TEE, it makes calls to procedures within the enclave. These calls are known as “call sites”. Call site verification verifies that the callee is a function allowed to be called from untrusted code. This prevents attackers from hijacking the control flow and executing arbitrary gadgets within the enclave. We show that call site verification is a necessary component of any secure TEE.

3 Threat Model

The defender’s code is bug-free. The defender follows TI’s secure coding practices: (1) they clear IPE state on exit and (2) disable interrupts upon IPE entry. The defender uses commodity cryptographic algorithms and compilers.

We assume the attacker has the capability to run any code in the untrusted world, this can either come from co-resident software or from having physical access [9]. This means the attacker can configure the timer and interrupt service routines. We assume the attacker’s software has a way to communicate with the attacker; The attacker wishes to have arbitrary read, write, and execute access to IPE-protected code and data. Note that this level of access is in-keeping with how TI expects users to interface with the IPE, i.e., the IPE is a place for protected, third-party code that untrusted user software can use as a library.

4 RIPencapsulation Design

RIPencapsulation is an interrupt-based side-channel memory exfiltration attack on TI MSP IP-Encapsulated (IPE) memory. As Figure 2 illustrates, the attack consists of three phases. After the first two phases of the attack, we are able to reverse engineer 80% of the IPE-protected code. If the attacker’s goal is only to get the keys from commodity cryptographic implementations, this generally suffices, as we prove in §7. However, completing all three phases enables the attacker to exfiltrate 100% of IPE-protected memory.

4.1 Creating a Side Channel (Phase 1)

The fundamental weakness in the design of IPE is that it does not protect the register file when context switches outside the IPE zone. This first phase exploits this weakness as a side channel that reveals the internal state of the IPE zone. The naive approach is to single-step IPE code using a debugger, but the MPU/SYSCTL monitors the memory address bus and

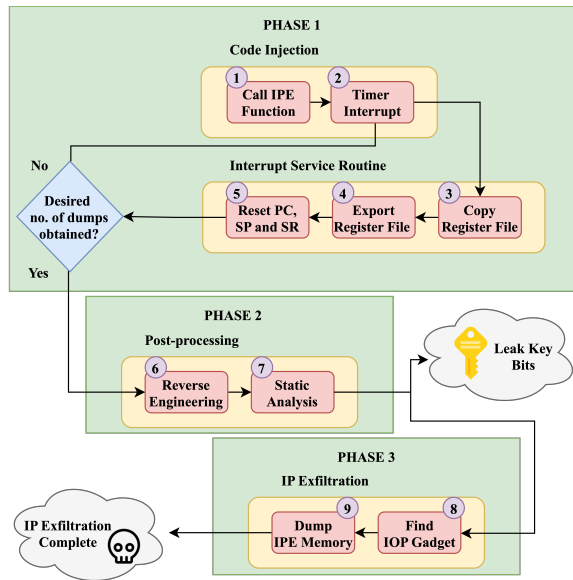


Figure 2: High-level IPE attack flow.

prevents JTAG or DMA from accessing the IPE-protected memory zone. This prevents breakpoints from working inside the IPE zone. While we observe that it is still possible to send HALT signals while in the IPE zone, it only provides timing granularity of 250 ms—not single-instruction precision.

The ability to process signals indicates that, in general, interrupts are processed while executing in the IPE zone. Thus, we can leverage a more precise interrupt source for single-instruction precision: the timer. Timer module interrupts have the capability to interrupt a program at single-cycle granularity (i.e., sub-instruction level).

Figure 2 shows the flow of the RIPencapsulation attack. In ①, attacker-controlled code sets up the timer module to interrupt the processor at a counter value of one clock cycle into the IPE code. Then the attacker code jumps directly to the instruction after the instructions which disables external interrupts.¹ In ②, the victim IPE code executes until the timer module interrupts it and passes control to the attacker’s ISR. Note that on an interrupt, any currently executing instruction is completed, and the PC of the next instruction, along with the SR, is pushed onto the stack. The context of the IPE code at the time of the interrupt remains intact.

Inside the ISR, ③ we copy the register states to attacker-controlled memory. This includes all the general-purpose registers as well as the stack pointer. In ④, the memory space available on MSP430/MSP432 is limited, with memory sizes around 256KB. So saving the register states of the entire

¹TI’s secure coding practice recommends disabling interrupts while executing in the IPE zone. We leverage the lack of call site verification to bypass any interrupt disabling code. It is also possible to use external interrupts or even a power reset in place of timer interrupts.

IPE-analysis process on the target MSP microcontroller is not feasible and we use its UART interface to transmit the exfiltrated register state to a workstation for analysis.² In ⑤, the attacker’s ISR resets the PC, SP, and SR registers to restore the program state to attacker-controlled code. The attacker code increments the timer counter value by one cycle and calls the victim code again. We repeat Phase 1 until the state-change due to all IPE code is captured. The end result is similar to single-stepping through the IPE code.

4.2 IPE State Dump Post-processing (Phase 2)

Phase 2 consists of post-processing the register state dumps and does not require access to the device. The goal is to decode the IPE assembly, extracting as much information as possible from the state dumps. In ⑥ we Look at the changes in PC deduce instruction boundaries. Comparing its PC with the PC of the next instruction, we are able to infer the size of the completed instruction unless it is a control-flow discontinuity. The frequency of successively repeating PCs in the state dumps provides instruction cycle count.

In the MSP430 instruction set architecture, the instruction size and cycles depend only on the operands and not the operation. An MSP430 instruction has broadly two types of operands, register or memory, with the possibility of combining these two for the source and destination operands. Identifying the operation is straightforward when both operands are direct register values. In such cases, the critical piece of information is the change in values of the general-purpose registers (GPRs) and how those changes correlate to the previous values of the registers. For instance, an ADD instruction with two register operands causes a difference in the value of the destination operand. By doing so, we decode the operation as well as the operands. However, we need to probe memory-to-register or register-to-memory instructions further in order to decipher them, as we detail in §5.2. We also observe that memory-to-memory operations are indecipherable just by looking at the register file, as at least one of the GPRs must change for identification.³

Status bit changes are also helpful in resolving the instruction guesses. A carry bit set implies that the result of the operation produces a carry. Zero or negative bit changes help us infer whether the result of an operation is zero or negative. The overflow bit signifies an overflow in the signed-variable range [46]. It also helps resolve certain operations, e.g., an AND instruction resets the overflow bit. These distinctions are most beneficial when the destination is not a register.

After reverse engineering, we are able to decode all the register and register-indirect addressing mode operations, when

²We also experimentally verify that it is possible to use the debugger to transmit the exfiltrated register state—albeit 22x slower than the UART.

³The ARM instruction set of the MSP432 only allows for register-to-register state-changing operations, which increases the power of Phase 1.

both operands are in one of those addressing modes.⁴ ⑦ The next step in post-processing depends on whether the attacker aims to obtain the key or exfiltrate the entire IPE memory.

4.3 IPE Memory Access Gadgets (Phase 3)

A ROP gadget is a security exploit that allows an attacker to execute code in the presence of code protection mechanisms [16, 38]. For TI’s IPE, there is no context switch gateway between calls from untrusted code to trusted code. We observe that an attacker has the capability to branch/jump to any instruction inside the IPE code without breaching the IPE security. Taking inspiration from ROP, we devise an interrupt- and data-oriented exploit for memory access purposes, which we call *Interrupt-Oriented Programming* (IOP). Interrupt-oriented Bugdoor Programming (IOBP) [43] first introduced the idea of IOP in the context of complete read/write access locked microcontrollers. However, these IOBP gadgets are rare in real-world firmware. The fundamental idea of our IOP is that if we setup the registers for a specific IPE instruction, branch directly to that IPE instruction, wait for it to complete execution, interrupt execution before the following instruction executes, and look at the results in the attacker’s ISR, we get a Turing-complete set required for return-oriented programming [16, 38]. Register-indirect addressing mode instructions in the IPE code that directly modify the value of a destination register are ideal candidates for memory exfiltration as these instructions enable the attacker to read an IPE memory location by modifying the memory address value in the source register, followed by a readout of the value returned in the destination register. The same is true for register-based memory store instructions. Henceforth, we refer to this class of instructions as *IOP gadgets*.

⑧ The first step of arbitrary IPE memory access is to find an IOP gadget. ⑨ Next, the attacker injects code into the non-IPE memory that sets up the timer module with a counter value, allowing the execution of the targeted instruction to complete. The malicious function also writes the desired IPE segment base address to the source register of the instruction. Then the program jumps to the IOP gadget inside the IPE zone. After execution of the gadget instruction(s), the timer module interrupts IPE execution, invoking the attacker’s ISR. In the case of exfiltration, the ISR sends the value of the destination register to the attacker workstation and restores the PC, SP, and SR of the program to the attacker’s control code. For memory writes, the host uploads a new malicious value. Repeating this process, the attacker is able read and write the entire IPE-protected memory.

⁴An attacker may leverage the lack of call-site verification to control register values and memory addresses to further disambiguate instructions as part of Phase 1, but none of our attacks required it.

Platform	MSP430FR5994	MSP432P401R
ISA	MSP430 CPUX	ARMv7-M
NV Memory Type	FRAM	Flash
NV Memory Size	256 KB	256 KB
Max Clock Frequency	16 MHz	48 MHz
IP Encapsulation	Yes	Yes
UART Interface	Yes	Yes
Timer Module	Yes	Yes

Table 1: RIPencapsulation target platforms.

Algorithm 1 Interrupt-based side-channel routine for exfiltrating IPE process register states

```

1: ...Set timer ISR routine...
2:  $i \leftarrow 0$ 
3:  $TIMER\_COUNT \leftarrow 0$ 
4: while  $i \neq DESIRED\_DUMPS$  do
5:    $TIMER\_COUNT \leftarrow TIMER\_COUNT + 1$ 
6:   ...Set timer and UART parameters...
7:   ...Start timer counter...
8:   ...Enable interrupts...
9:   CALL IPE_Function()
10:  // Secure process gets interrupted
11:  ...Copy all core registers to unprotected memory...
12:  ...Copy saved state to UART TX Buffer...
13:   $i \leftarrow i + 1$ 
14:  ...Reset SP, SR and PC registers...
15: end while

```

5 RIPencapsulation Implementation

We implement and evaluate RIPencapsulation on Texas Instruments MSP430FR5994 and MSP432P401R launchpads. Table 1 details the devices’ relevant specifications. Since the MSP430 and MSP432 are based on different instruction sets, there are some implementation differences in RIPencapsulation on the two microcontrollers; §6 covers these differences.

5.1 Capturing IPE Register State

RIPencapsulation’s interrupt-based side-channel procedure transfers the single-cycle register state dumps to the attacker’s machine. Algorithm 1 provides the details of RIPencapsulation’s side channel routine. MSP430FR5994 comes with two asynchronous general-purpose timer modules, each with four operating modes, and supports multiple captures/compares.⁵ The three operating modes besides Stop mode are – Up, Continuous, and Up/down; they have similar behavior: the timer counts to a value before overflowing and restarting the

⁵We use the timer module instead of the watchdog timer because it has finer-grain control over timer interval ranges and supports interrupt intervals as low as one clock cycle.

Source Operand	Destination Operand				
	Rm	PC	x(Rm)	TONI	&TONI
Rn	✓	✓	✓	×	×
@Rn	✓	✓	✓	×	×
@Rn+	✓	✓	✓	×	×
#N	-	-	-	×	×
x(Rn)	✓	✓	✓	×	×
EDE	×	×	×	×	×
&EDE	×	×	×	×	×

Table 2: MSP430 decoded instructions. EDE and TONI are generic labels and have no special meaning. Certain instructions cannot be guessed accurately as there is ambiguity when the source operand is an immediate value or label (-).

count. It generates interrupts when the timer counter overflows or reaches the `capture/compare` register value, depending on the operating mode.

After setting up the timer parameters, we start the timer and enable interrupts. Then we branch to the IPE code, where the timer interrupts it, invoking the interrupt service routine (ISR). Inside the ISR, we copy the latent IPE state from the register file and sent it to the workstation via the UART. Next, the ISR restores PC, SP, and SR to the attacker’s state machine code, where we reset the timer parameters and restart the counter with the incremented `TIMER_COUNT` value so that the device goes one cycle deeper into the IPE code.

5.2 Reverse Engineering and Heuristics

MSP430 consist of the 27 implemented instructions of the MSP430 CPU [46]. Besides this, some MSP430 devices also support additional `CPUX` instructions that are used when a greater than 16-bit address space is required. The instructions are categorized into three types based on the operand sizes: byte, word, and 20-bit. The MSP430 instructions support seven different addressing modes. Table 2 summarizes the various source and destination addressing modes that the MSP430 instruction set architecture supports and highlights the instructions RIPencapsulation is able to guess correctly. RIPencapsulation successfully decodes the instructions that use register/indexed/indirect register mode source and destination operands. It can also decipher the indirect autoincrement mode instructions, where the source operand is `@Rn+`. In contrast to the indirect register mode instructions, these instructions increment the source register value by 1, 2, or 4 depending on whether the operand size is byte, word, or 20-bit respectively. It is important to note that these four types of instructions constitute the bulk of the assembly in our benchmarks. We are unable to decipher instructions with symbolic or absolute addressing mode source or destination operands

as they are IP encapsulated and cannot be directly modified to attacker-controlled non-IPE address locations for disambiguation. Also, we do not consider guesses for instructions with immediate mode source operand and register/indexed/indirect register mode destination operand in our evaluation of decoded assembly instructions, as it is ambiguous whether the source operand is an immediate value or a value stored in some IPE symbol or absolute address.

Algorithm 2 Sorting algorithm for register mode instructions

```

1:  $i \leftarrow 0$ 
2: while  $i \neq TOTAL\_DUMPS - 1$  do
3:   if  $PC[i + 1] - PC[i] \neq 0$  then
4:     if  $Rm[i + 1] - Rm[i] \neq 0$  then
5:        $R_{dest} \leftarrow Rm$ 
6:     end if
7:     if  $R_{dest}[i + 1] \leftarrow f(R_{dest}[i], Rn[i])$  then
8:        $R_{source} \leftarrow Rn$ 
9:        $OPERATION = f$ 
10:    end if
11:  end if
12:   $i \leftarrow i + 1$ 
13: end while

```

Algorithm 2 details the sorting algorithm for register mode instructions where the source and destination operands are register contents. For indexed or indirect register mode instructions, the source or destination operand contains a register value used as a pointer to the actual operand value. Algorithm 3 highlights the procedure for identifying these instructions. Before branching to the suspected instruction, the attacker sets all the general-purpose registers (GPRs) to controlled non-IPE memory addresses. Based on the addressing mode, the suspected instruction will modify either the corresponding destination register or the memory address value it points to. Observing which register or memory address value changes, we guess the destination operand and its addressing mode (register/indexed/indirect). Analyzing the new value at the destination location and comparing it with the older GPR values, we are able to figure out the source operand, its addressing mode, and the operation performed. For instance, a "MOV @R10, R15" loads the value at the memory address pointed by R10 to R15. For indexed instructions, we perform multiple tests with different memory address inputs to get a complete resolution, as we also need to identify the offset to the address value stored in the register operand. Besides this, we also perform multiple tests to disambiguate unintended clashes between operations. For instance, XOR between `0xFE` and `0x00` produces the same output as MOV with `0xFE`.

Single operand instructions such as `PUSH`, `POP`, `CALL`, and `RETI` modify the stack and have unique effects on the stack pointer, giving away their identity. We also perform heuristic analysis on the obtained register dumps to clarify the victim program assembly state better. Figure 3 illustrates the register




Figure 3: Register state plot of a decrementing loop-counter showing the change of R12 register with the PC.

Algorithm 3 Routine for identifying indexed and indirect mode instructions

```

1: ...Set timer ISR routine...
2:  $i \leftarrow 0$ 
3:  $TIMER\_COUNT \leftarrow CYCLES\_REQUIRED$ 
4:  $ADDR \leftarrow SUSPECTED\_INSTRUCTION$ 
5: while  $i \neq NUMER\_OF\_TESTS$  do
6:   ...Set timer and UART parameters...
7:   ...Start timer counter...
8:   ...Enable interrupts...
9:   ...Set all GPRs to controlled memory addresses...
10:   $JUMP\ ADDR$ 
11:  // Secure process gets interrupted
12:  ...Copy all core registers to unprotected memory...
13:  ...Copy saved state to UART TX Buffer...
14:   $i \leftarrow i + 1$ 
15:  ...Reset SP, SR and PC registers...
16: end while

```

state plot of a decrementing loop-counter program. Here the R12 register holds the counter value decremented at the end of each loop. The return of the program counter to the lower memory address signifies JUMPing back to the start of the loop. We interpret PC value discontinuities which do not modify the SP as JUMP instructions.

5.3 IPE Memory Exfiltration

Instructions which use indexed/indirect register mode source operand are ideal for exfiltrating the IPE memory. However, any indexed/indirect register value also works as we modify the memory address pointed to by that register

Algorithm 4 MSP430 IP exfiltration routine using IOP gadget

```

1: ...Set timer ISR routine...
2:  $i \leftarrow 0$ 
3:  $TIMER\_COUNT \leftarrow CYCLES\_REQUIRED$ 
4:  $ADDR \leftarrow IPE\_START - OPERAND\_SIZE$ 
5: while  $i \neq DESIRED\_DUMPS$  do
6:   ...Set timer and UART parameters...
7:   ...Start timer counter...
8:   ...Enable interrupts...
9:    $ADDR \leftarrow ADDR + OPERAND\_SIZE$ 
10:   $R_{source} \leftarrow ADDR$ 
11:   $JUMP\ IOP\_Gadget$ 
12:  // Secure process gets interrupted
13:  ...Copy  $R_{dest}$  value to unprotected memory...
14:  ...Copy saved state to UART TX Buffer...
15:   $i \leftarrow i + 1$ 
16:  ...Reset SP, SR and PC registers...
17: end while

```

to some attacker-controlled memory location. After the attacker identifies an IOP gadget, they inject malicious code that writes the IPE segment base address to the source register of the victim instruction and branches to it. Following the timer interrupt, the ISR transmits the saved destination operand containing the IPE memory value over the UART. The malicious code then increments the victim address before the ISR transmits the next destination operand value to the attacker's machine. Algorithm 4 details RIPencapsulation's MSP430 IP exfiltration routine using the IOP gadget.

6 RIPencapsulation on MSP432

We also test the efficacy of RIPencapsulation on MSP432 with the Cortex-M4 Flash-based microcontroller MSP432P401R. This device has a different ISA, different system-on-chip, and different non-volatile memory, so by showing that RIPencapsulation works for the MSP432, we show that the flaws in TI's IPE are fundamental.

6.1 CPU Halts Inside IPE Zone

MSP432 IP Encapsulation has some key differences in its implementation, besides the fact that it is only configurable using the bootloader (due to Flash vs. FRAM differences). MSP432 IPE is more secure in that it disables all kinds of CPU halts inside the IPE zone. The SYSCTL security control monitors all debug accesses from the DAP (debug) port inside the IPE zone. In fact, SYSCTL does not permit any CPU halts inside the IPE zone, including breakpoint addresses pointing inside the IPE memory (unlike MSP430 IPE). Even with this precaution, we find that we are able to trigger timer interrupts inside the IPE zone. Thus, like for the MSP430, we use timer interrupts to build our attack, however, there are some differences in the interrupt handler routine.

6.2 Return from Interrupt Handler

The setup procedure for timer interrupts is the same as on the MSP430. However, the working of interrupts is slightly different. On the MSP432, all interrupts and exceptions are handled by the Nested Vectored Interrupt Controller (NVIC). When the NVIC detects that an interrupt signal is HIGH, it changes the state of the interrupt to pending. Interrupts remain pending until the processor enters the interrupt service routine (ISR), upon which the NVIC changes the interrupt status to active. When the interrupt is serviced by the ISR, the processor loads an EXC_RETURN value stored in the link register (LR) to the program counter (PC). In our test code using timer interrupts, an EXC_RETURN value of 0xFFFFF9E9 is loaded into PC, which returns to the original thread using the state from the main stack pointer (MSP). On return from the ISR, the NVIC changes the interrupt status from active to inactive. As such, when we try to replicate RIPencapsulation on the MSP432 by jumping to the malicious code, like we do on the MSP430, the interrupt status remains active, and the NVIC does not allow re-entry into the ISR even when the timer interrupt becomes pending again. To resolve this, we exit the ISR using the EXC_RETURN value. So, instead of jumping directly to the malicious code, we write the malicious return state on the MSP. This way, when the EXC_RETURN value is loaded into the PC, the NVIC changes the interrupt status to inactive and the processor returns to the malicious code instead of the original thread, where we reset the timer interrupts and

Algorithm 5 MSP432 IP exfiltration routine using IOP gadget

```
1: ...Set timer ISR routine...
2:  $i \leftarrow 0$ 
3:  $TIMER\_COUNT \leftarrow CYCLES\_REQUIRED$ 
4:  $ADDR \leftarrow IPE\_START - OPERAND\_SIZE$ 
5: while  $i \neq DESIRED\_DUMPS$  do
6:   ...Set timer and UART parameters...
7:   ...Start timer counter...
8:   ...Enable interrupts...
9:    $R_{source\_STORE} \leftarrow 0X695A$ 
10:   $R_{dest\_STORE} \leftarrow SYS\_SECDATA\_UNLOCK$ 
11:   $R_{dest\_BRANCH} \leftarrow LOAD\_INSTRUCTION\_ADDR$ 
12:   $ADDR \leftarrow ADDR + OPERAND\_SIZE$ 
13:   $R_{source\_LOAD} \leftarrow ADDR$ 
14:  JUMP IOP_Gadget
15:  // Secure process gets interrupted
16:  ...Stop timer and clear interrupt flags...
17:  ...Copy  $R_{dest\_LOAD}$  value to unprotected memory...
18:  ...Copy saved state to UART TX Buffer...
19:   $i \leftarrow i + 1$ 
20:  ...Write malicious return address to MSP...
21: end while
```

increment the timer counter for single-step execution of the victim IPE code.

6.3 Clearing Interrupt Flags

Even though the program control returns to the malicious code upon completion of the ISR, as we desire, we observe that the program immediately re-enters the ISR, not executing any victim code. We discover that this is happening because the interrupt signal stays asserted even after completion of the ISR. The NVIC detects the asserted signal and immediately changes the interrupt status to pending, leading to re-entry inside the ISR. Clearing the interrupt flags de-asserts the interrupt signal and the NVIC waits for the next timer interrupt upon completion of the ISR.

6.4 Unlocking Read Access

Exfiltrating the data from the MSP432 IPE region is not so straightforward. The attacker must unlock data access each time the control goes inside the IPE zone. If we try to jump to an arbitrary load instruction inside the IPE zone without unlocking the data access first, the processor throws an exception. Writing 0x695A to the memory-mapped SYS_SECDATA_UNLOCK register unlocks the data access for the IPE region writing to that register (since MSP432 allows the creation of up to four isolated IPE regions). To bypass this, we need to construct a more sophisticated read IOP gadget. Algorithm 5 details RIPencapsulation's MSP432 IP exfiltration routine using the IOP gadget. In our implementation, we

Optimization Level	Instructions decoded	Reveal key bits?	Contains IOP Gadget?
-O0	60.5%	✓	✓
-Og	66.6%	✓	✓
-O1	68.4%	✓	✓
-O2	68.4%	✓	✓
-O3	68.4%	✓	✓
-Os	67.9%	✓	✓
-Ofast	68.4%	✓	✓

Table 3: Test cases with different compiler optimizations for AES (`tiny AES`). The second column depicts the percentage of assembly instructions we are able to reverse engineer after Phase 2.

jump inside the IPE zone to a store instruction to write the unlock value to the memory-mapped `SYS_SECDATA_UNLOCK` register followed by an indirect register branch instruction to branch to the actual load instruction which we then use to exfiltrate the IPE firmware. The attacker is capable of leveraging other sophisticated IOP gadgets to perform this read exploit.

7 Evaluation

We evaluate RIPencapsulation against four commodity cryptographic algorithms, namely AES (`tiny AES`), SHA256 (`saddi`), SHA256 (`gladman`), and RSA (`codebase`). We select these specific benchmarks in order to evaluate RIPencapsulation across different forms of cryptographic algorithms like symmetric-key and public-key cryptography and cryptographic hashing. We compile all benchmarks using the open-source MSP430 GCC toolchain developed by Texas Instruments [47], testing them against a range of optimization levels. Our evaluation answers the following questions:

- Are cryptographic implementations generally susceptible to the RIPencapsulation attack?
- What effect does the compiler optimization level have on the vulnerability of these cryptographic implementations?
- How long does it take to carry out the RIPencapsulation attack?
- Is IPE vulnerable across device types?

7.1 AES

Our static analysis results show that the AES implementation leaks key bits to the registers in its `KeyExpansion` function. At optimization levels `-O1` and above, the `KeyExpansion` function is embedded inside the `AES_init_ctx` function. All the instructions leaking the key bits are of the form `MOV.B @Rn, Rm`. The destination registers in these instructions leak

Optimization Level	Instructions decoded	Reveal 'pt' bits?	Contains IOP Gadget?
-O0	59.4%	✓	✓
-Og	56.7%	✓	✓
-O1	53.9%	✓	✓
-O2	49.7%	✓	✓
-O3	50.1%	✓	✓
-Os	53%	✓	✓
-Ofast	50.1%	✓	✓

Table 4: Test cases with different compiler optimizations for SHA256 (`saddi`). The second column depicts the percentage of assembly instructions we are able to reverse engineer after Phase 2.

the last four bytes of the secret key and all the remaining round keys. Reverse engineering the AES secret key using the round keys is a deterministic process, and literature exists describing the same [19]. We take the second-round key and the last four bytes of the secret key to obtain the original 128-bit key using the following formula:

If $i\%4 \neq 0$,

$$\omega_{i-4} = \omega_i \oplus (\omega_{i-1})$$

Else,

$$\omega_{i-4} = \omega_i \oplus \text{sbox}(\text{shift}(\omega_{i-1})),$$

followed by XOR of ω_{i-4} 's 1st byte with $Rcon[j]$

Here ω_i is the i th word in the complete AES key. $\text{sbox}(x)$ represents the byte substitution using the S-Box lookup table, $\text{shift}(x)$ is a cyclical shift of the bytes of ω_i , and $Rcon[j]$ is the round constant for round j , whose key we reverse engineer here. This would be the first round in our case. It is also worth mentioning that the `-O0` optimization level assembly reveals the secret key location in the stack. Assemblies produced at all optimization levels contain an instruction of the form `MOV.B @Rn, Rm`. We use them as IOP gadgets to exfiltrate the AES IPE code and data. Table 3 summarizes the results for our AES-128 implementation.

7.2 SHA256

Table 4 highlights the results of the RIPencapsulation attack on `saddi SHA256` implementation. Not only do all the assemblies leak the plaintext location to the registers, but also all the plaintext bits. The plaintext bits are leaked in the `SHA256Guts` function, which is embedded inside the `sha256_update` function at optimization levels `-O1` and above. The instructions containing the leaking registers are word instructions, and

Optimization Level	Instructions decoded	Reveal 'pt' bits?	Contains IOP Gadget?
-O0	59.2%	✓	✓
-Og	60.9%	✓	✓
-O1	58.3%	✓	✓
-O2	60.4%	✓	✓
-O3	60.7%	✓	✓
-Os	57.7%	✓	✓
-Ofast	60.7%	✓	✓

Table 5: Test cases with different compiler optimizations for SHA256 (*gladman*). The second column depicts the percentage of assembly instructions we are able to reverse engineer after Phase 2.

Optimization Level	Instructions decoded	Reveal key bits?	Contains IOP Gadget?
-O0	58.3%	✓	✓
-Og	58.4%	✓	✓
-O1	57.3%	✓	✓
-O2	54.3%	✓	✓
-O3	54.3%	✓	✓
-Os	60.8%	✓*	✓
-Ofast	54.3%	✓	✓

Table 6: Test cases with different compiler optimizations for RSA (*codebase*). The second column depicts the percentage of assembly instructions we are able to reverse engineer after Phase 2. *Readout from the stack is required for some bits of the private key.

the plaintext bytes are present in reverse order within each word as the memory model that MSP430FR5994 uses is little endian. All our SHA256 assemblies contain instructions of the form `MOV.W x(Rn), Rm`. We use them as IOP gadgets to exfiltrate our SHA256 IPE code and data.

We also evaluate RIPencapsulation on the *gladman* SHA256 implementation. Table 5 summarizes its evaluation results. All assemblies of this SHA256 implementation leak the plaintext bits to the registers inside the `sha_end1` function. The compiled code also contains a prologue instruction inside the calling function for SHA256, which performs a SUBA operation on the stack pointer. We bypass this instruction and, as such, include it in our calling routine for the IP-encapsulated SHA256 function. Assemblies produced at all optimization levels contain IOP gadgets that look like `MOV.W @Rn, Rm`.

7.3 RSA

The *codebase* RSA implementation is the smallest code of all the benchmarks evaluated in this paper. Table 6 summa-

rizes the evaluation results for our RSA implementation. Assemblies produced at optimization levels -O0, -Og, and -O1 store the private key location in the stack and leak the private key bits to the registers inside the `modexp` function, which is called by the `rsaDecrypt` function. Level -O2, -O3, and -Ofast leak both the private key location and the key itself to the general-purpose registers inside the `rsaDecrypt` function. The -Os assembly stores the private key location in the stack and only partially leaks the private key bits to the registers inside the `modexp` function. The assembly directly leaks the first two words to the registers but accesses the next two words in the `BIS.W x(SP), R12` instruction, which performs an OR between the two operands. Since OR is lossy, we cannot reverse-engineer the operands from the result of the operation. So, we need to read out the last two words of the private key from the stack. The -Os optimization assembly contains an IOP gadget of the form `MOV.W @SP, R12`, that is used to exfiltrate the private key bits from the stack. The -O0 level produces assembly with a prologue consisting of a PUSH operation on the stack followed by a SUBA operation on the stack pointer. All other assemblies of RSA contain only the SUBA operation on the stack pointer in the calling function.

7.4 Attacking Real Firmware

-Opt	O0	Og	O1	O2	O3	Os	Ofast
IOP Gadget?	✓	✓	✓	✓	✓	✓	✓

Table 7: Test cases with different compiler optimizations for a flight controller GPS/GPRS tracker.

While cryptographic algorithms are a natural target in TEEs, we demonstrate the generalizability of our exploit on real-world firmware. To this end, we use an existing GPS/GPRS tracker application from a drone flight controller [6]. We verify that this application is vulnerable to RIPencapsulation attacks as it contains several indirect load and store instructions which serve as our IOP gadgets. Table 7 shows that assemblies produced at all optimization levels contain IOP gadgets. Indirect load instructions serve as read IOP gadgets which enable the exfiltration of all the IPE protected code and data. We also verify the presence of indirect store instructions in the IPE code which serve as write IOP gadgets. Such write IOP gadgets give the attacker the ability to plant bugs in the GPS/GPRS tracker code.

7.5 Attack Time

IPE exfiltration rate is characterized by the delay in transmitting the data packets to the external device over the UART communication channel. At a baud rate of 115200 and clock

frequency of 1 MHz, it takes 10 seconds on average to obtain the register states leaking the second round key of the tiny AES encryption process. `saddi` SHA256 directly leaks the plaintext bits to the register states in under 75 seconds for the `-O0` optimization level, whereas `-Ofast` optimized code leaks the plaintext in 23 seconds. Although the entire `gladman` SHA256 register state exfiltration process takes the longest time due to its large code size (2x in size compared to `saddi` SHA256 and 4x compared to `codebase` RSA), it leaks the plaintext bits to the registers in roughly the same amount of time as `saddi`, ranging from 30 to 70 seconds, with the `-O0` level assembly being the most time-consuming. `Codebase` RSA is the simplest benchmark of all and leaks the private key to the registers within 5 seconds. However, since `-Os` optimized RSA code requires exfiltration of some key bits from the stack, `Phase 3` is necessary and the attacker must use an IOP gadget for exfiltrating the private key. In general, we are able to receive 10,000 dumps over the UART channel in under 2 minutes. The time taken to execute `Phase 2` of the attack depends on the attacker's approach to analyzing the exfiltrated register state. To accelerate `Phase 2`, we automate the reverse engineering process on our workstation.

7.6 MSP432 Results

We were able to replicate all of the attacks from the MSP430 on the MSP432. Even though the MSP432 is a newer ISA with a more security-oriented system-on-chip, the RIPencapsulation is easier on the MSP432 due to its Thumb architecture set, which predominantly uses register mode operations. Thus, the reverse engineering capability of RIPencapsulation on MSP432 increases to 80% without `Phase 3`.

7.7 Write Exploit

Besides leaking the IPE memory to the outside world, specific IOP gadgets are also able to write to the IPE region, giving the attacker the power to modify the IPE code and possibly get access control over it that way. This breaks the integrity and authenticity guarantees of code and data inside the IPE region. Register direct/indirect addressing mode instructions in the IPE code with register indirect destination addressing are ideal candidates for such an exploit as they allow deploying a payload to any desired IPE memory location. For the MSP430 instruction set these instructions look like `MOV Rn, x(Rm)`, `MOV @Rn, x(Rm)`, `MOV x(Rn), x(Rm)`, `MOV Rn, x(SP)` and many more variations of the same. We find that all our MSP430 benchmarks contain one of these instructions across all compiler optimizations. Meanwhile, the MSP432 uses Flash for its non-volatile memory hence we need the Flash Controller for writes. However, MSP432 IPE only allows storing code and constants in its IPE region and the Flash Controller is disabled inside the IPE zone.

8 Mitigation Strategies

We find two fundamental shortcomings in Texas Instruments's (TI's) IP Encapsulation (IPE) design, which make it vulnerable to RIPencapsulation. First, it does not clear the IPE state on context switches to non-IPE execution. This enables the attacker to invoke an Interrupt Service Routine (ISR) outside the IPE zone but retain access to the latent state of IPE execution through the register file. Secondly, IPE has no call site verification. This allows the attacker to jump anywhere inside the IPE zone, opening the door to control- and data-oriented attacks. Even though TI recommends the IP author to disable interrupts from non-IPE code as a secure coding practice, we show that this is futile as attackers can bypass any protections in the IPE zone by jumping to the instruction right after. Thus, a defense must address both the fundamental IPE flaws.

"The privacy of register and on-chip caches should be protected by the trusted computing base from software attacks" [42]. In the AEGIS Architecture [42], a secure context manager (SCM) stores all the process's register values in the SCM table on interrupt and clears the register states before invoking the ISR so that the ISR cannot access the internal state of the secure process. The SCM restores the register states from the SCM table on return from the ISR. Clercq et al. [20] provide a hybrid implementation of AEGIS for MSP430 devices that improves SANCUS [34]. When extended to include clearing the residual state present in the shared memory region of SRAM (e.g., the stack), AEGIS-based system should address the problem of latent IPE state upon unexpected IPE exits—at the cost of hardware modification.

The lack of valid call site verification allows the attacker to orchestrate data-oriented control flow attacks. Almost 90% of exploit-based software attacks use some form of Return-Oriented Programming (ROP) [27]. Address Space Layout Randomization (ASLR) is a well-known class of code security techniques [22, 26, 28, 52, 53] that randomize the memory address of a program's sections in order to reduce the chances of code reuse exploits that rely on knowing the exact location of process objects. Not only does ASLR entail a heavy overhead, but it also supposes a higher level of ability to intrude into the code and introspect its insecure uses. The facilities available on MSP430 and MSP432 devices do not lend themselves to an efficient implementation of such a high overhead approach.

A better defense for control- and data-oriented attacks is a call gateway veneer. In the embedded space, ARM TrustZone enforces this gateway veneer by introducing a Non-Secure Callable (NSC) memory region. All calls from a normal program to a secure function must go through the gateway veneer residing in the NSC memory. Calls to invalid entry points inside the secure code cause a hardware exception which always traps into a secure state. Fault injection [44] and short-term data remanence [32] attacks break the security guarantees of ARM TrustZone by pausing the trusted execution in a

controlled manner to reveal its internal state.

We believe that ARM TrustZone, combined with fault attack protection, in conjunction with AEGIS-style secure context-switching represents the best solution to defend IoT-class devices against RIPencapsulation. Given the many trade-offs at play for ultra-constrained devices like MSPs, we believe that the design, implementation, and evaluation of such a defense is important future work that this paper motivates.

9 Related Work

Trusted Execution Environments (TEEs) are process isolation and secure storage solutions that are finding their way into IoT-class embedded devices. Recent work indicates a rise in the trend of secure process preemption or exception-based exploits, which infer the program's internal state by studying their effects on the unprotected areas of the compromised device. In light of this trend and given RIPencapsulation is a TEE attack, we cover attacks against other TEEs on devices ranging from desktop-class to embedded devices.

9.1 Interrupt-based Attacks

SGX-Step [49] is an interrupt-based side-channel attack on Intel Software Guard eXtensions (SGX), which builds on previous kernel-level SGX exploits that preempt the enclave execution to leak information from page tables (PTE) [51, 54] or branch prediction units [29]. SGX-Step exploits the Advanced Programmable Interrupt Controller (APIC) timer to interrupt the secure process at several-instruction granularity in order to gain fine-grained control of the side channels, improving the temporal resolution of previous enclave preemption attacks [24, 29, 51, 54]. Nemesis [50] extends the SGX-Step by using an interrupt-based side channel to leak instruction-level information from TEE execution. Nemesis requires precisely timed interrupts to capture per-instruction latency differences.

RIPencapsulation extends the idea of SGX-Step and Nemesis to IoT class devices. RIPencapsulation provides single-cycle granularity and uncovers 100% of TEE-protected memory by combining interrupt-based attacks with data-oriented attack patterns. RIPencapsulation also demonstrates an exploit to modify the TEE-protected code and data for its target device class. Note that, the compiler-based defenses that reduce the effectiveness of SGX-Step, apply to heavyweight TEEs and rely on detecting high rates of page faults or interrupts by leveraging the x86 Transitional Synchronization eXtensions (TSX) and as such do not apply to the IoT-class devices targeted by RIPencapsulation [17, 41].

The Page-Fault Weird Machine [10] has a similar high-level idea to Interrupt-Oriented Programming (IOP) for using interrupted execution with pre-existing gadgets in the device to get useful computation. However, RIPencapsulation's IOP differs in the context of gadget aims, target device class, and security

violation. We use gadgets on the firmware level (instruction set architecture) as opposed to microarchitecture gadgets. Our target class devices are bare-metal resource-constrained microcontrollers as opposed to x86 full-fledged desktop class devices (for example our device class does not have virtual memory). Finally, we use our gadgets to create read/write exploits and break the security of TEEs on our target devices.

Interrupt-oriented Bugdoor Programming (IOBP) [43] is conceptually the same as RIPencapsulation's IOP. The difference lies in the adversarial setting. IOBP exploits require some a priori knowledge to find useful IOBP gadgets in read/write access blocked microcontrollers. We too consider our protected firmware (IPE protected library) to be read/write access blocked, however a part of the memory is unlocked to read/write accesses. This is a realistic threat model for IP theft/imitation in cases such as flight controllers, where a competitor (potential adversary) uses the licensed security critical library (flight controller) APIs in their development code. When IOP is used in such a context, a) we do not require a priori knowledge of the IPE firmware, as we can use IOP gadgets to dump CPU states and reverse engineer partial information about the underlying code, enough to detect read/write IOP gadgets and b) find more usefully exploitable IOP gadgets for protected state extraction/modification.

9.2 Debugger-based Attacks

Shedding too much Light on a Microcontroller's Firmware Protection [36] analyzes the security of protected Flash memory on STM32 microcontrollers. The paper uses fine-grain CPU resets and a vulnerability in the Flash protection logic protocol to extract the entire firmware by accessing iterative addresses of the firmware via the debugger and capturing latent Flash data in unprotected SRAM. While this attack requires optical fault injection and chemical etching to manipulate particular Flash bits, RIPencapsulation can leverage their controlled CPU resets in the event that timer interrupts are not available.

Brosch [14] presents a firmware dumping technique for an ARM Cortex-M0 SoC that uses the debugger to manipulate the CPU register values at single-step intervals. By single-stepping through the program code and observing the CPU register changes, they find load instructions to exfiltrate the protected memory. Their attack is based on the debugger's single-stepping capability inside protected memory, which is prevented by TI's IPE implementation. RIPencapsulation side-steps TI's protection through a combination of interrupts and data-oriented attacks, achieving single-step IPE execution. Additionally, RIPencapsulation presents a write exploit, breaking the integrity and authenticity of firmware inside the protected memory.

9.3 Blind Attacks

Code reuse attacks require varying degrees of information on the target. Blind attacks aim to create exploits when neither the source nor binary code is available. Half-Blind attacks [23] presents a stack overflow/ROP gadget attack to gain privileged access in the bootstrap loader (BSL) of an MSP430 device, which can then be used to extract the firmware image. Hacking Blind [11] is a more generic and advanced version of half-blind attack which is fully blind and presents techniques to find and chain multiple, different gadgets. While these attacks do not work with TI's IPE protection, they serve as inspiration for attacking unknown binaries with RIPencapsulation.

9.4 Other attacks on TEEs

CipherLeaks [31] is a ciphertext side-channel attack on AMD's Secure Encrypted Virtualization (SEV) TEE. SEV protects Virtual Machines (VM) from an untrusted hypervisor by using hardware-enforced memory encryption. On domain switch between the guest and host VM, SEV stores the encrypted register states in the Virtual Machine Save Area (VMSA). The CipherLeaks attack model assumes that the attacker has read access to the VMSA but no write access. CipherLeaks uses the hypervisor to monitor specific offsets of the VMSA to infer changes of any 16-bit plaintext. Non-Automatic VM Exits (NAE) expose some plaintext register values to the hypervisor. In essence, the attacker triggers an NAE to collect a dictionary of plaintext-ciphertext pairs for these registers stored in the VMSA. Cipherleaks then uses this plaintext-ciphertext dictionary to crack the entire OpenSSL RSA key in 410 rounds with 100% accuracy. In order to patch this vulnerability, AMD added randomization in stored register values when encrypting and saving them into the VMSA during VMEXITs [1]. This fix is available in the AMD SEV-SNP TEE. Unfortunately, encrypting the register file is not a complete defense against RIPencapsulation, as the ability to enter the IPE zone at arbitrary points allows the attack to create IPE zone altering and exfiltration gadgets.

CLKscrew [44] is an energy management-based exploit that manipulates the voltage and frequency of the processor to induce faults. Dynamic Voltage and Frequency Scaling (DVFS) is an energy management scheme, ubiquitous on commodity devices, that trades off processing speed for energy savings. CLKscrew is able to break the confidentiality and integrity of ARM TrustZone using software-only control of the regulators. In essence, the attacker increases the frequency of the processor beyond the limits dictated by the operating voltage to induce instability and halt the TrustZone process from the normal world. Performing Differential Fault Analysis on the correct and faulty decrypted plaintext pair, the attacker is able to infer the AES key. We envision an exploit that combines the Interrupt-Oriented Programming described

in this paper with CLKscrew to get very fine-grained control over fault injections inside the TEE.

During preparation of the camera-ready version of this paper, we discovered a paper citing a previous version of this paper that leverages our arbitrary read/write/execute access to IPE-protected memory to strengthen their attack on IPE [33]. While the practical variant of their attack relies heavily on RIPencapsulation, they provide their own unique capabilities worth noting. Their main contribution is an attack primitive that they refer to as *controlled call corruption* which exploits a microarchitecture bug in the IPE access control mechanism. If an adversary sets the `SP` to some protected memory address, followed by a `call` to the IPE zone, the return address is pushed to the stack pointer location, subverting the IPE access control check. Their defensive landscape analysis agrees with ours that a hybrid solution is required for a comprehensive defense.

10 Conclusion

Texas Instruments MSP IP Encapsulation (IPE) aims to provide confidentiality of data stored inside the IP-encapsulated memory zone; this includes proprietary code and keys. RIPencapsulation breaks this guarantee by leveraging two fundamental drawbacks in MSP430 IPE design: residual state on context switches and lack of call site verification. We exploit these flaws to create an interrupt-based side channel to gain cycle-accurate control IPE execution and exfiltrate all IPE secrets. The evaluation shows that this attack works using production tools and settings of popular open-source cryptographic implementations.

This paper shows that Trusted Execution Environment (TEE) designers must pay careful attention to unexpected TEE entries and exits. Without guarding the entries to IPE code, attackers can bypass defenses and create gadget-like instruction sequences. Without cleaning up residual state shared across security domains on every possible exit, attackers have access to secret-revealing side-channel information. These requirements extend beyond any single TEE implementation, serving as necessary conditions for ensuring code and data confidentiality by all TEEs.

Acknowledgments

The project depicted is sponsored by the Defense Advanced Research Projects Agency. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. Approved for public release; distribution is unlimited. This material is based upon work supported by the National Science Foundation under Grant No. 2240744.

References

- [1] AMD Secure Encryption Virtualization (SEV) Information Disclosure. <https://www.amd.com/en/resources/product-security/bulletin/amd-sb-1013.html>. Accessed: 2023-4-14.
- [2] IoT market size, trends and global forecast to 2032. <https://www.thebusinessresearchcompany.com/report/iot-global-market-report>. Accessed: 2023-4-13.
- [3] MSP430 microcontrollers. <https://www.ti.com/microcontrollers-mcus-processors/msp430-microcontrollers/overview.html>. Accessed: 2023-4-13.
- [4] Embedded Survey: The Current State of Embedded Development. <https://www.embedded.com/embedded-survey/>, September 2016. Accessed: 2023-7-18.
- [5] 2019 Embedded Markets Study, 2019.
- [6] GPS/GPRS tracker. <https://github.com/wudidashao/tracker>, 2019.
- [7] Microcontroller (MCU) Market Analysis. <https://www.utmel.com/blog/categories/microcontrollers/microcontroller-mcu-market-analysis>, February 2022. Accessed: 2023-2-3.
- [8] RIPencapsulation Code. <https://github.com/ForTE-Research/RIPencapsulation>, 2023.
- [9] Adnan Akhuzada, Mehdi Sookhak, Nor Badrul Anuar, Abdullah Gani, Ejaz Ahmed, Muhammad Shiraz, Steven Furnell, Amir Hayat, and Muhammad Khurram Khan. Man-at-the-end attacks: Analysis, taxonomy, human aspects, motivation and future directions. *Journal of Network and Computer Applications*, 48:44–57, 2015.
- [10] Julian Bangert, Sergey Bratus, Rebecca Shapiro, and Sean W Smith. The Page-Fault Weird Machine: Lessons in Instruction-less Computation. In *7th USENIX Workshop on Offensive Technologies (WOOT 13)*, 2013.
- [11] Andrea Bittau, Adam Belay, Ali Mashtizadeh, David Mazières, and Dan Boneh. Hacking blind. In *2014 IEEE Symposium on Security and Privacy*, pages 227–242. IEEE, 2014.
- [12] Tyler Bletsch, Xuxian Jiang, Vince W Freeh, and Zhenkai Liang. Jump-oriented programming: a new class of code-reuse attack. In *Proceedings of the 6th ACM symposium on information, computer and communications security*, pages 30–40, 2011.
- [13] Erik Bosman and Herbert Bos. Framing signals—a return to portable shellcode. In *2014 IEEE Symposium on Security and Privacy*, pages 243–258. IEEE, 2014.
- [14] Kris Brosch. Firmware dumping technique for an ARM Cortex-M0 SoC. <https://blog.includesecurity.com/2015/11/>, November 2015. Accessed: 2023-7-27.
- [15] David Cerdeira, Nuno Santos, Pedro Fonseca, and Sandro Pinto. Sok: Understanding the prevailing security vulnerabilities in trustzone-assisted tee systems. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 1416–1432. IEEE, 2020.
- [16] Stephen Checkoway, Lucas Davi, Alexandra Dmitrienko, Ahmad-Reza Sadeghi, Hovav Shacham, and Marcel Winandy. Return-oriented programming without returns. In *Proceedings of the 17th ACM conference on Computer and communications security*, pages 559–572, 2010.
- [17] Sanchuan Chen, Xiaokuan Zhang, Michael K Reiter, and Yinqian Zhang. Detecting privileged side-channel attacks in shielded execution with Déjà Vu. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, pages 7–18, 2017.
- [18] Shuo Chen, Jun Xu, Emre Can Sezer, Prachi Gauriar, and Ravishankar K Iyer. Non-control-data attacks are realistic threats. In *USENIX security symposium*, volume 5, page 146, 2005.
- [19] William R Cordwell. Aes key recovery from round keys. Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 2008.
- [20] Ruan De Clercq, Frank Piessens, Dries Schellekens, and Ingrid Verbauwhede. Secure interrupts on low-end microcontrollers. In *2014 IEEE 25th International Conference on Application-Specific Systems, Architectures and Processors*, pages 147–152. IEEE, 2014.
- [21] Karim Eldefrawy, Gene Tsudik, Aurélien Francillon, and Daniele Perito. Smart: secure and minimal architecture for (establishing dynamic) root of trust. In *Ndss*, volume 12, pages 1–15, 2012.
- [22] Stephanie Forrest, Anil Somayaji, and David H Ackley. Building diverse computer systems. In *Proceedings. The Sixth Workshop on Hot Topics in Operating Systems (Cat. No. 97TB100133)*, pages 67–72. IEEE, 1997.
- [23] Travis Goodspeed and Aurélien Francillon. Half-blind attacks: mask ROM bootloaders are dangerous. In *Proceedings of the 3rd USENIX conference on Offensive technologies. USENIX Association*, pages 6–6, 2009.

- [24] Marcus Hähnel, Weidong Cui, and Marcus Peinado. High-resolution side channels for untrusted operating systems. In *USENIX Annual Technical Conference*, volume 17, 2017.
- [25] Hong Hu, Shweta Shinde, Sendriou Adrian, Zheng Leong Chua, Prateek Saxena, and Zhenkai Liang. Data-oriented programming: On the expressiveness of non-control data attacks. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 969–986. IEEE, 2016.
- [26] Vivek Iyer, Amit Kanitkar, Partha Dasgupta, and Raghunathan Srinivasan. Preventing overflow attacks by memory randomization. In *2010 IEEE 21st International Symposium on Software Reliability Engineering*, pages 339–347. IEEE, 2010.
- [27] Ayman Jarrous. Anti-ROP: A moving target defense. <https://securityintelligence.com/anti-rop-a-moving-target-defense/>, September 2016. Accessed: 2023-4-14.
- [28] Chongkyung Kil, Jinsuk Jun, Christopher Bookholt, Jun Xu, and Peng Ning. Address space layout permutation (aslp): Towards fine-grained randomization of commodity software. In *2006 22nd Annual Computer Security Applications Conference (ACSAC'06)*, pages 339–348. IEEE, 2006.
- [29] Sangho Lee, Ming-Wei Shih, Prasun Gera, Taesoo Kim, Hyesoon Kim, and Marcus Peinado. Inferring fine-grained control flow inside SGX enclaves with branch shadowing. In *USENIX Security Symposium*, volume 19, pages 16–18, 2017.
- [30] Mengyuan Li, Luca Wilke, Jan Wichelmann, Thomas Eisenbarth, Radu Teodorescu, and Yinqian Zhang. A systematic look at ciphertext side channels on amd sevsnp. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 337–351. IEEE, 2022.
- [31] Mengyuan Li, Yinqian Zhang, Huibo Wang, Kang Li, and Yueqiang Cheng. Cipherleaks: Breaking constant-time cryptography on amd sev via the ciphertext side channel. In *USENIX Security Symposium*, pages 717–732, 2021.
- [32] Jubayer Mahmud and Matthew Hicks. Sram has no chill: exploiting power domain separation to steal on-chip secrets. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 1043–1055, 2022.
- [33] Cas Magnus Marton Bognar, Frank Piessens, and Jo Van Bulck. Intellectual Property Exposure: Subverting and Securing Intellectual Property Encapsulation in Texas Instruments Microcontrollers. In *33rd USENIX Security Symposium*, 2024.
- [34] Job Noorman, Pieter Agten, Wilfried Daniels, Raoul Strackx, Anthony Van Herrewege, Christophe Huygens, Bart Preneel, Ingrid Verbauwhede, and Frank Piessens. Sancus: Low-cost trustworthy extensible networked devices with a zero-software trusted computing base. In *22nd USENIX Security Symposium (USENIX Security 13)*, pages 479–498, 2013.
- [35] Ivan De Oliveira Nunes, Karim Eldefrawy, Norrathep Rattanavipanon, Michael Steiner, and Gene Tsudik. {VRASED}: A verified {Hardware/Software}{Co-Design} for remote attestation. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 1429–1446, 2019.
- [36] Johannes Obermaier and Stefan Tatschner. Shedding too much light on a microcontroller’s firmware protection. In *11th USENIX Workshop on Offensive Technologies (WOOT 17)*, 2017.
- [37] Katie Pier. *MSP Code Protection Features*. Texas Instruments, December 2015.
- [38] Ryan Roemer, Erik Buchanan, Hovav Shacham, and Stefan Savage. Return-oriented programming: Systems, languages, and applications. *ACM Transactions on Information and System Security (TISSEC)*, 15(1):1–34, 2012.
- [39] Mohamed Sabt, Mohammed Achemlal, and Abdelmadjid Bouabdallah. Trusted execution environment: what it is, and what it is not. In *2015 IEEE Trust-com/BigDataSE/Ispa*, volume 1, pages 57–64. IEEE, 2015.
- [40] Hovav Shacham. The geometry of innocent flesh on the bone: Return-into-libc without function calls (on the x86). In *Proceedings of the 14th ACM conference on Computer and communications security*, pages 552–561, 2007.
- [41] Ming-Wei Shih, Sangho Lee, Taesoo Kim, and Marcus Peinado. T-SGX: Eradicating controlled-channel attacks against enclave programs. In *NDSS*, 2017.
- [42] G Edward Suh, Dwaine Clarke, Blaise Gassend, Marten Van Dijk, and Srinivas Devadas. Aegis: Architecture for tamper-evident and tamper-resistant processing. In *ACM International Conference on Supercomputing 25th Anniversary Volume*, pages 357–368, 2003.
- [43] Samuel Junjie Tan, Sergey Bratus, and Travis Goodspeed. Interrupt-oriented bugdoor programming: A minimalist approach to bugdooring embedded systems firmware. In *Proceedings of the 30th Annual Computer Security Applications Conference*, pages 116–125, 2014.

- [44] Adrian Tang, Simha Sethumadhavan, and Salvatore J Stolfo. Clkscrew: Exposing the perils of security-oblivious energy management. In *USENIX Security Symposium*, volume 2, pages 1057–1074, 2017.
- [45] Texas Instruments. *MSP432P4xx Family Technical Reference Manual*, December 2016. Rev. E.
- [46] Texas Instruments. *CPUX*, March 2018. Rev. F.
- [47] Texas Instruments. *MSP430 GCC Toolchain*, June 2020. Rev. F.
- [48] Texas Instruments. *MSP430FR58xx, MSP430FR59xx, and MSP430FR6xx Family User’s Guide*, April 2020. Rev. P.
- [49] Jo Van Bulck, Frank Piessens, and Raoul Strackx. SGX-Step: A practical attack framework for precise enclave execution control. In *Proceedings of the 2nd Workshop on System Software for Trusted Execution*, pages 1–6, 2017.
- [50] Jo Van Bulck, Frank Piessens, and Raoul Strackx. Nemesis: Studying microarchitectural timing leaks in rudimentary cpu interrupt logic. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 178–195, 2018.
- [51] Jo Van Bulck, Nico Weichbrodt, Rüdiger Kapitza, Frank Piessens, and Raoul Strackx. Telling your secrets without page faults: Stealthy page table-based attacks on enclaved execution. In *Proceedings of the 26th USENIX Security Symposium*, pages 1041–1056. USENIX Association, 2017.
- [52] David Williams-King, Graham Gobieski, Kent Williams-King, James P Blake, Xinhao Yuan, Patrick Colp, Michelle Zheng, Vasileios P Kemerlis, Junfeng Yang, and William Aiello. Shuffler: Fast and deployable continuous code re-randomization. In *OSDI*, volume 10, pages 3026877–3026906, 2016.
- [53] Jun Xu, Zbigniew Kalbarczyk, and Ravishankar K Iyer. Transparent runtime randomization for security. In *22nd International Symposium on Reliable Distributed Systems, 2003. Proceedings.*, pages 260–269. IEEE, 2003.
- [54] Yuanzhong Xu, Weidong Cui, and Marcus Peinado. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *2015 IEEE Symposium on Security and Privacy*, pages 640–656. IEEE, 2015.