



OPEC: Operation-based Security Isolation for Bare-metal Embedded Systems

Xia Zhou, Jiaqi Li, Wenlong Zhang, Yajin Zhou, Wenbo Shen, Kui Ren
Zhejiang University
Hangzhou, Zhejiang, China
{zhouxia_icsr,lijiaqi93,wl_zhang,yajin_zhou,shenwenbo,kuiren}@zju.edu.cn

Abstract

Bare-metal embedded systems usually lack security isolation. Attackers can subvert the whole system with a single vulnerability. Previous research intends to enforce both privilege isolation (to run application code at the unprivileged level) and resource isolation for global variables and peripherals. However, it suffers from partition-time and execution-time over-privilege issues, due to the limited hardware resources (MPU regions) and the improper way to partition a program.

In this paper, we propose operation-based isolation for bare-metal embedded systems. An *operation* is a logically independent task composed of an entry function and all functions reachable from it. To solve the partition-time over-privilege issue, we utilize the global variables shadowing technique to reduce the needed MPU regions to confine the access of the global variables. To mitigate the execution-time over-privilege issue, we split programs into code compartments (called operation) that only contain necessary functions to perform specific tasks, thereby removing the resources needed by unnecessary functions. We implement a prototype called OPEC, which contains an LLVM-based compiler and a reference monitor. The compiler partitions a program and analyzes the resource dependency for each operation. With the hardware-supported privilege levels and MPU, the reference monitor is responsible for enforcing the privilege and resource isolation at runtime. Our evaluation shows that OPEC can achieve the security guarantees for the privilege and resource isolation with negligible runtime overhead (average 0.23%), moderate Flash overhead (average 1.79%), and acceptable SRAM overhead (average 5.35%).

Yajin Zhou is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EuroSys '22, April 5–8, 2022, RENNES, France

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9162-7/22/04...\$15.00

<https://doi.org/10.1145/3492321.3519573>

CCS Concepts: • Security and privacy → Embedded systems security.

Keywords: security isolation, hardware-assisted security, memory protection unit

ACM Reference Format:

Xia Zhou, Jiaqi Li, Wenlong Zhang, Yajin Zhou, Wenbo Shen, Kui Ren. 2022. OPEC: Operation-based Security Isolation for Bare-metal Embedded Systems. In *Seventeenth European Conference on Computer Systems (EuroSys '22)*, April 5–8, 2022, RENNES, France. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3492321.3519573>

1 Introduction

Bare-metal embedded systems have surged rapidly in recent years. They run a monolithic software that provides both system functionality and application logic without the intervention of the operating system. Bare-metal embedded systems are vulnerable due to the absence of security defenses, such as privilege isolation, data execution prevention (DEP), and address space layout randomization (ASLR) [1]. Those defenses are widely deployed in desktop systems, but they cannot be directly adopted due to insufficient hardware resources and limited hardware security extensions. For instance, desktop systems implement resource isolation with the hardware called Memory Management Unit (MMU). However, bare-metal embedded systems have neither MMU nor sufficient resources to support such isolation. Furthermore, there is no isolation between tasks, i.e., different execution stages of a program.

Nevertheless, the attack techniques targeting desktop systems can be easily transferred to bare-metal embedded systems. Attackers can get full control of the whole system by exploiting a vulnerability in any task. Hence, enforcing secure isolation on bare-metal embedded systems is an effective way to secure the devices.

Previous research provides security isolation to bare-metal embedded systems. However, it either misses or has defects in resource isolation. EPOXY [12] provides privilege isolation by identifying and executing only sensitive instructions at the privileged level. Still, EPOXY does not support flexible resource isolation between tasks. MINION [22] implements privilege and thread-level resource isolation for real-time microcontroller systems. MINION predefines the accessible memory range and permissions for each thread during the

program compartmentalization. Nevertheless, the accessible memory is oversized, and its permission is overset due to hardware constraints, i.e., the limited number of MPU regions. Additionally, threads are absent in many bare-metal embedded systems, so MINION cannot be applied to bare-metal embedded systems directly.

To enforce privilege isolation and more fine-grained resource isolation, ACES [11] adopts a code-module based strategy that partitions a program based on the source files or peripherals. However, ACES suffers from two over-privilege issues. First, ACES allows a compartment to access unnecessary global variables. For a specific compartment, different parts of its accessible global variables may be shared with related compartments. Due to the limited number of MPU regions, ACES directly grants all the related compartments to all of these global variables (named *partition-time over-privilege* issue in Section 3.1). Second, ACES partitions a program without considering the control flow. ACES involves unnecessary functions into a compartment to complete a specific task, e.g., unlocking a smart lock. Therefore, a compartment can access unnecessary resources at runtime (named *execution-time over-privilege* issue in Section 3.1). Furthermore, ACES limits the number of accessible peripherals of a compartment due to the limited number of MPU regions, which brings inflexibility when partitioning a program.

In this paper, we propose operation-based security isolation. It provides both privilege isolation and resource isolation to bare-metal embedded systems. An *operation* is a *logically independent task composed of an entry function and all functions reachable by that entry function*. For instance, as shown in Figure 1, the PinLock application can be split to six tasks (operations). As in Listing 1, when PinLock receives the correct pin code, the operation `Unlock_Task` invokes the function `do_unlock` to perform unlocking. When executing inside one operation, the code can only access necessary resources, i.e., the global data and peripherals needed to complete the task. The code cannot access other resources. This method provides the security guarantee that a vulnerable task cannot compromise other ones or the whole system.

For the partition-time over-privilege issue, we solve it through the *global data shadowing* technique. We make a shadow copy of each shared global variable of one operation and put it into a continuous memory space that is only accessible by that operation. By doing so, the continuous memory space only contains global data necessary for a specific operation, which avoids incorporating redundant resources. For the execution-time over-privilege issue, we mitigate it through an approach that one operation only includes the functions needed to perform a specific task. Indeed, a function may contain some basic blocks that are not executed at runtime, which could also cause the execution-time over-privilege issue. Furthermore, our system virtualizes the MPU regions to use the MPU flexibly and facilitate the operation

```

1 void Unlock_Task() {
2   HAL_UART_Receive_IT(&PinRxBuffer); /*buggy*/
3   result = hash(PinRxBuffer);
4   if(compare(result, KEY))
5     do_unlock();
6 }
7
8 void Lock_Task() {
9   HAL_UART_Receive_IT(&PinRxBuffer); /*buggy*/
10  if(PinRxBuffer[0]=='0')
11    do_lock();
12 }
13
14 int main(void) {
15   System_Init(); /*Task1: configure core peripherals*/
16   Uart_Init(); /*Task2: configure UART*/
17   Key_Init(); /*Task3: compute the hash of PIN,
18              and save it to KEY*/
19   Init_Lock(); /*Task4: Init lock*/
20   while (1)
21   {
22     Unlock_Task(); /*Task5: perform the unlock task*/
23     Lock_Task(); /*Task6: perform the lock task*/
24   }
25 }

```

Listing 1. Main logic of the application PinLock.

partitioning. Explicitly, our system supports reconfiguring the MPU on demand at runtime.

We implement a prototype called OPEC. It consists of two components, i.e., OPEC-Compiler and OPEC-Monitor. The former is an LLVM-based compiler that is responsible for partitioning the program into different operations and identifying needed resources for each operation. The latter is responsible for global data shadowing, stack isolation, and operation switching. Specifically, OPEC-Compiler gets the application source code and the operation entry function list specified by developers. Then it performs a static analysis to identify the global data and peripherals needed by each operation to generate the *operation policy*. After that, it uses the *operation policy* to produce the final program image. OPEC-Monitor is linked to the image during the compiling time to enforce the isolation policy. The operation switch is triggered by a software interrupt (the SVC instruction), which will be handled by OPEC-Monitor. OPEC-Monitor first performs data synchronization based on the corresponding policy. Then, it enforces resource access permissions by setting up the MPU.

We evaluate the effectiveness of OPEC with six representative applications and the CoreMark benchmark [13]. The evaluation shows that our system effectively enforces privilege isolation and reduces the resources that can be accessed for an operation at runtime. Compared to the vanilla system, an average of 37.79% of the resources are accessible at runtime. We further evaluate the performance overhead. The average runtime overhead is 0.23%. OPEC consumes low Flash and moderate SRAM resources, i.e., 1.79% and 5.35% for Flash and SRAM, respectively.

In summary, the contributions of our work are as follows:

- We analyze the over-privilege issues of the previous security isolation for bare-embedded systems, which motivates our work (Section 3).

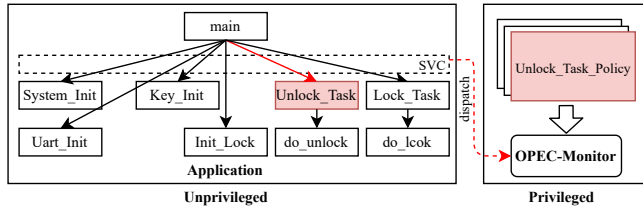


Figure 1. Example of the operation partitioning, switching, and runtime enforcement to PinLock. The PinLock can be divided into six tasks, each task can be regarded as an operation. At runtime, SVC instructions trigger the operation switch, and the OPEC-Monitor of OPEC performs the switch and updates the MPU configurations according to the policy to enforce the security isolation.

- We propose operation-based security isolation and solves or mitigates the over-privilege issues (Section 4 and Section 5).
- We prototype our system OPEC and evaluate it with six representative applications on two different development boards. The experiment shows that OPEC enforces privilege isolation and resource isolation for bare-metal embedded systems with a negligible impact on performance. (Section 6).

2 Background

2.1 Memory Address Space and Privilege Levels

OPEC works towards bare-metal embedded systems [12], whose system libraries and application code are statically linked and executed on low-end microcontrollers, e.g., ARM Cortex-M CPU families. In this paper, we use the popular ARMv7-M architecture [4] as a reference to describe our system. This choice complies with previous work [11, 12].

The address space of the ARMv7-M architecture contains code, data, peripherals, and so forth., as shown in Figure 2. For instance, the code usually resides in the lower 512 MB address space, ranging from 0x00000000 to 0x1FFFFFFF. The program data, including the stack and global variables, are in the SRAM region. For peripherals, they are either at the peripheral region or the external device region. Despite a large address space, only a small portion is used in a real system. A bare-metal embedded system usually has a few MBs of Flash and KBs of SRAM.

The ARMv7-M architecture has two privilege levels for software execution, i.e., privileged and unprivileged levels. The program can execute some security-sensitive instructions at the privileged level, such as configuring the hardware directly. The unprivileged level is reserved for the application code. The unprivileged application can execute a SVC instruction to make a supervisor call to escalate to the privileged level. Although the microcontroller supports privilege isolation, it is not fully leveraged by bare-metal embedded systems. Both system libraries and applications are running

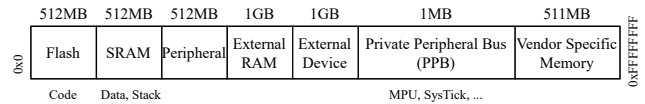


Figure 2. Memory layout of the ARMv7-M architecture.

at the *privileged level*. Our system intends to provide privilege isolation. Note that only privileged software is allowed to access memory range within Private Peripheral Bus (PPB), where core peripherals such as MPU and SysTick timer reside. Unprivileged access will trigger a bus fault.

2.2 Memory Protection Unit (MPU)

Memory Protection Unit (MPU) is used to provide physical memory access control. It defines several regions by specifying the base address, size, and access permissions at the privileged and unprivileged levels. Each region is labeled with a number. If two memory regions have an overlapped memory range, the access to this range is confined by the region labeled with the highest number. Accessing the memory range prohibited by the MPU will trigger a memory management fault. However, there are some restrictions on the region configuration. The region size must be the power of two, and the smallest permitted region size is 32 bytes. The region’s base address must be aligned with its size. Therefore, it is challenging to enforce precise access control for a large number of scattered address spaces by the limited number of MPU regions. One MPU region can be further split into eight equal-sized sub-regions, and each can be enabled or disabled individually. If a sub-region of a region is disabled, a lower-numbered MPU region that overlaps this memory range confines the memory access to it.

3 Overall Design and Threat Model

3.1 Motivation

Our system intends to provide security isolation to bare-metal embedded systems. By doing so, one compromised task cannot access arbitrary resources (data and peripherals). The data may contain security-critical system states. Peripherals, as the interface between the device and the outside world, may directly impact the physical world. For instance, by writing a special peripheral register, the attacker can control a robot arm’s movement speed causing safety issues.

Though the previous system ACES [11] intends to provide security isolation, it suffers from two over-privilege issues. We use two examples derived from the PinLock application to illustrate two different types of over-privilege issues. In this paper, we refer to an isolated partition as a domain, i.e., an operation in our paper and a compartment in ACES.

Partition-time Over-privilege results from the limitation of MPU. Figure 3 shows the partition-time over-privilege issue when using the MPU to isolate global variables. ACES solves the over-privilege caused by the size and alignment restriction of the MPU region by rearranging the addresses of

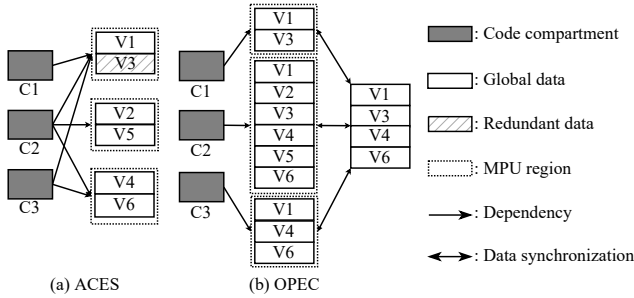


Figure 3. Example of the partition-time over-privilege issue derived from the PinLock.

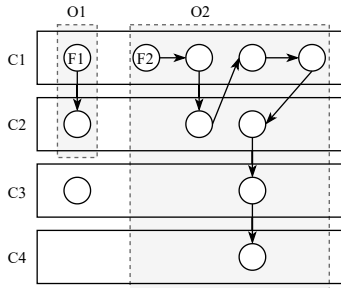


Figure 4. Example of the execution-time over-privilege issue derived from the PinLock.

global variables. However, it fails to solve the over-privilege issue caused by limited MPU regions. Different compartments may share variables, and most of the shared variables would become a separate region. It leads to exceeding the available MPU regions when a compartment has too many shared variables among different compartments. Therefore, ACES merges some regions to solve this issue at the cost of introducing the over-privilege issue. In Figure 3(a), compartment C2 groups V1 and V3 to save the use of regions. It leads to the over-privilege issue for compartment C3 since C3 only needs to access V1.

Our system solves this problem through the global data shadowing technique that each compartment has a shadow copy of the shared global variables (Figure 3(b)). These variables are rearranged inside a continuous memory range that is confined by one MPU region.

Execution-time Over-privilege is caused by including unnecessary code when executing a task. As shown in Figure 4, the compartment C1 contains multiple functions. However, when performing a task (e.g., unlocking a smart lock), only the function F1 in the compartment C1 executes. Resources that are accessed by other functions in compartment C1 become redundant. If the code in compartment C1 is compromised, it can access these redundant resources, which leads to execution-time over-privilege.

Moreover, the code-module based way to divide compartments [11] also induces frequent compartment switches because the execution flow of a specific task may cross multiple compartments. For instance, when performing task two in Figure 4, five compartment switches between compartments

are needed. Our system considers the execution flow when partitioning a program. It mitigates the execution-time over-privilege issue while reducing the overhead of compartment switches simultaneously.

3.2 System Overview

The workflow of our system is shown in Figure 5. It consists of two key stages, *compiler-assisted operation partition* and *hardware-assisted operation isolation*. These two stages are involved by the following two components, i.e., OPEC-Compiler and OPEC-Monitor, respectively. The former is an LLVM-based compiler, which analyzes the program and builds the image with the operation isolation property. The latter enforces privilege and resource isolation between operations at runtime. It is linked to the application code during compiling.

3.3 Assumptions and Security Benefits

Our system assumes a strong threat model that the attacker can gain the arbitrary code execution capability to construct the primitives to read from and write to arbitrary memory locations. These primitives could be achieved through hijacking the control flow of the program, e.g., corrupting function pointers, and then by leveraging Return-Oriented Programming (ROP) to construct powerful attack primitives. However, attackers cannot inject code since the writable memory regions are not executable. We also assume the program is buggy, which may contain vulnerabilities that attackers can exploit to compromise the whole system. Note that intra-domain protections to control-flow integrity and data integrity are orthogonal to our work.

We assume that the hardware has two privilege levels (privileged and unprivileged) and a memory protection unit (MPU). In addition, we assume that the system is a bare-metal one whose program image is a statically linked binary. It does not support dynamically linked libraries. It is consistent with the bare-metal embedded systems in the real world. These assumptions comply with the previous system [11].

OPEC enforces the least-privilege principle to bare-metal embedded systems and provides the following protections. First, it ensures the isolation between operations. The compromise of one operation cannot directly take over the whole system. Second, it enforces the least-privilege principle of operation. The compromise of one operation can only manipulate the resources that are accessible in that operation. Third, the sanitization of changes to global data *alleviates* the cross-domain threat that corrupts safety-critical global variables¹. Our system reduces the security consequences of bare-metal embedded systems after being compromised, providing a practical defense with a powerful threat model.

¹Note that if attackers can directly manipulate safety-critical peripherals that are accessible in a compromised operation, the sanitization cannot prevent this attack.

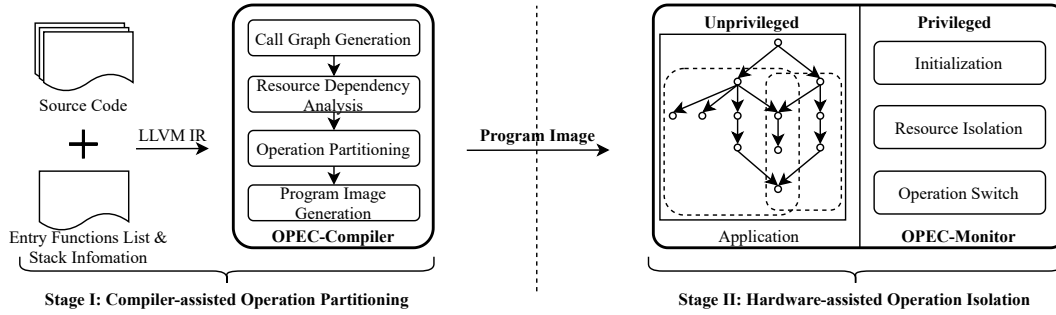


Figure 5. The workflow of OPEC.

4 Compiler-assisted Operation Partitioning

This section presents how OPEC-Compiler partitions operations and generates the program image with our isolation property. Our system takes the program source code and the list of operation entry functions as inputs to produce the final image (Figure 5). In our prototype, the OPEC-Compiler is built on LLVM 9.0 [28], with new passes to perform the program analysis and instrumentation.

4.1 Call Graph Generation

An operation is a specific task of the bare-metal system. From the program’s perspective, *an operation is a subtree in the call graph, with a developer-provided function as the root*. With the call graph, we can easily traverse all the functions from the root node. However, building a precise call graph is challenging due to the indirect function calls (*icalls* for short) [29]. OPEC-Compiler leverages the Andersen point-to analysis implemented by SVF [38] to determine the potential legal targets of function pointers. Nevertheless, there are some icalls that cannot be resolved by SVF. In this case, we use type-based analysis to find the potential targets. We consider two function types identical if the number of arguments, the type of the structure argument, the type of the pointer argument, and the type of the return value are the same. The indirect function call edges are then added to the call graph to ensure its soundness. Note that the results of the point-to analysis are conservative and over-approximated, which contains false positives. Otherwise, an unsound call graph will bring dependency miss to operations.

4.2 Resource Dependency Analysis

After constructing the call graph, the next step is to analyze the needed resources of each function. The resources include global variables and peripherals. OPEC-Compiler leverages state-of-the-art static analysis techniques to perform a conservative analysis.

Global Variables OPEC-Compiler leverages the forward slicing to identify accessed global variables in each function. Global variables can be accessed directly and indirectly. Our system processes them differently.

Direct accesses can be identified by recognizing all the load and store instructions that use the global variables as

operands. OPEC-Compiler uses the built-in def-use analysis of the LLVM to identify. Indirect access happens when a variable is accessed through a pointer. OPEC-Compiler performs an inter-procedural point-to analysis using the state-of-the-art tool SVF [38]. The analysis identifies local and global variables targets of a pointer. We further filter out the local targets pointed to by pointers and leave the global ones. Note that the size of an array pointed to by data pointers should be determined at compile time. We have not encountered arrays with statically unknown sizes.

OPEC-Compiler also records the pointer fields of a global variable by leveraging its type. This information is further used for updating pointer fields of shadow variables when switching an operation (Section 5.3).

Peripherals are accessed through memory-mapped addresses (Figure 2) in bare-metal systems. For a specific System-on-Chip (SoC), memory addresses of peripherals are constant. OPEC-Compiler uses this information to identify peripheral access. We obtain the addresses of peripherals from the SoC datasheet. OPEC-Compiler uses the backward slicing at IR level to examine whether the operand of a load/store instruction contains a constant memory address. It then compares the address with a predefined list of peripheral addresses retrieved from the datasheet to check whether the address is peripheral. If so, the OPEC-Compiler marks it as a peripheral accessible by the function where the store/load the instruction belongs to. Furthermore, we divide the accessible peripherals into two lists for general peripherals and core peripherals, since accessing core peripherals (on PPB) requires the code to run at the privileged level (Section 2).

4.3 Operation Partitioning

This step aims to partition a program into operations and identify all resources needed by each operation. OPEC-Compiler depends on the entry functions list provided by the developers to partition a program (Figure 5). The process of selecting the entry function of an operation is straightforward. Take the application PinLock as an example (Listing 1). This application can be divided into six tasks. We can treat each task as an operation and choose the root node of each task in the call graph as the entry function. Apart from the operations generated through the entry function list, OPEC-Compiler

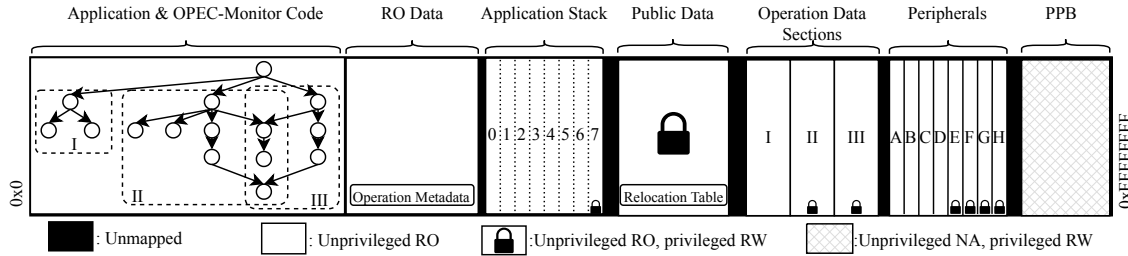


Figure 6. The memory layout of the program armed with OPEC. For operation I, it can modify the stack except for the sub-region 7 of the stack used by the previous operation, its operation data section, and four peripherals. RO: Read-only; RW: Read and write; NA: No access.

also considers the function `main` as a default operation. The default operation can access the resources necessary to the function `main`. Note that the operation entries cannot be functions having variable-length arguments or within an interrupt handling routine.

For each entry function, `OPEC-Compiler` performs the Depth-First-Search (DFS) algorithm to traverse the call graph from the entry function to determine the functions that operation contains. When reaching another operation entry function, the `OPEC-Compiler` performs backtracking. Note that two different operations can share functions. Our system supports recursion since it can be grouped into one operation.

After that, the `OPEC-Compiler` groups all the member functions callable from one operation entry to an operation and merges all needed resources to produce the resource dependency. In our design, each peripheral is protected by an individual MPU region. To save the use of the limited MPU regions, `OPEC-Compiler` will first sort the peripherals needed by one operation at the ascending order of their start addresses. Then the adjacent peripherals will be merged and protected by an individual MPU region. Even so, one operation may need regions exceeding the limitation of the MPU to access the peripherals. In this scenario, `OPEC` virtualizes the usage of MPU regions to solve this issue. The details are in Section 5.2. Alternatively, the developers can choose to split one large operation into smaller ones since smaller operations may access fewer peripherals and need fewer MPU regions. Finally, `OPEC-Compiler` generates a policy file that contains accessible resources of each operation.

4.4 Program Image Generation

After partitioning the program, the `OPEC-Compiler` generates the operation data sections and MPU configurations, then instruments the program to produce the final image.

Operation Data Section `OPEC` uses the global data shadowing technique to isolate global variables. Specifically, each operation has an exclusive *operation data section* that contains all the global variables it needs. An operation can access global variables only within its operation data section. Each operation data section is confined by one MPU region preventing variables from being corrupted by other operations.

Our system divides the global variables into two types, *internal* and *external*. The internal ones are accessed by only one operation and directly placed into the corresponding operation data section. The external ones are accessed by two or more operations. Each external variable has shadow copies in corresponding data sections. Those shadow copies are accessed through a *variables relocation table*. Specifically, the `OPEC-Compiler` creates a data pointer for each external variable at the variables relocation table. The data pointer either points to the original variable or the shadow copy. We will discuss the details of using the relocation table for data accessing in Section 5.2.

In our design, each operation data section is protected by one MPU region. However, the MPU has a rigid address alignment requirement (Section 2). Improper placement of operation data sections will cause external memory fragments. To reduce the external fragments, the `OPEC-Compiler` first sorts the operation data sections according to their sizes in descending order. Then it calculates the start addresses of these sections and places them accordingly.

Operation Metadata The metadata of each operation contains five parts, i.e., MPU configurations, stack information, sanitization value of critical global variables, a peripheral list, and a variable relocation table. The MPU configurations, which are generated by the `OPEC-Compiler` after the building of operation data sections, are used to confine the memory access permissions of each operation. The stack information and sanitization value are provided by the developers. The former is used to annotate the pointers in the entry function arguments, which facilitates the `OPEC-Monitor` performing stack synchronization during the operation switching. The latter is used for global variables sanitization. The peripheral list is used as an allow list when accessing the peripherals.

Code Instrumentation `OPEC-Compiler` inserts the initialization routines before entering the function `main` of the application. We will illustrate the initialization in Section 5.1. Moreover, the `OPEC-Compiler` inserts `SVC` instructions before and after the call site of each operation entry function. The application code will escalate to the privileged level after executing the `SVC` instruction. After that, the control flow will be transferred to the operation switch routine. When the

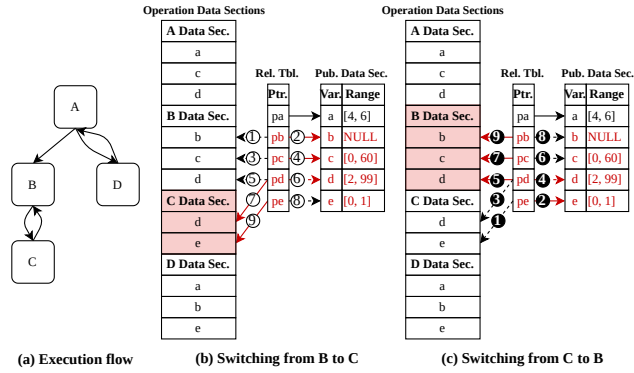


Figure 7. External global variables synchronization.

program enters into or exits from an operation, the operation switch is invoked to set the new execution environment or recover from the previous execution context. After instrumentation, the library that contains the OPEC-Monitor is linked to the application code to produce the final image.

5 Hardware-assisted Operation Isolation

This section illustrates the OPEC-Monitor. It enforces privilege isolation, resource isolation, and operating switching based on the policy generated in the previous stage.

5.1 Initialization

OPEC-Monitor initializes the system before running the application code. First, it initializes the operation data section of each operation by copying the initial value of each global variable to its shadow copy. Second, it enables the exception handling of the supervisor call (SVC), memory management fault, and bus fault. Finally, it drops the privilege and runs the unprivileged application code to ensure privilege isolation.

5.2 Resource Isolation

Figure 6 shows the memory layout and the access permissions of the final image generated by OPEC-Compiler. The data of OPEC-Monitor and the relocation tables of each operation are placed in the memory region that is only writable at the privileged level. It prevents the isolation policy from being tampered with by a compromised unprivileged operation. The operation metadata excluding relocation tables is stored in the read-only memory.

The eight MPU regions are arranged as follows. Region 0 sets all memory ranges as read-only. Region 1 enables the execution of the unprivileged application code. Region 2 and 3 enforce read and write permissions to the stack and the operation data section. Region 4 to 7 are reserved for peripheral access. Note that unprivileged adversaries cannot execute security-sensitive instructions of OPEC-Monitor because of the privilege isolation as described in Section 5.1. In the following, we will present how our system enforces resource isolation at runtime.

Global Variables OPEC-Monitor synchronizes shared variables in operation data sections and modifies the variable relocation table to ensure that the program accesses correct shadow copies. The data synchronization process during the operation switch is illustrated in Figure 7. When entering into or returning from one operation, OPEC-Monitor first identifies whether the global variable needs to be synchronized. As shown in Figure 7(b)(c), global variable a, b, c, e, and f are shared variables. When entering into C from B, OPEC-Monitor first writes back the value of b, c, and d from their shadow copies in B’s data section to the public data section (1⇒2, 3⇒4, and 5⇒6). Then OPEC-Monitor further copies the value of d and e from the public data section to their shadow copies in C’s data section (6⇒7 and 8⇒9). When returning to B from C, OPEC-Monitor first writes back the value of d and e from C’s data section to the public data section (1⇒2 and 3⇒4). Then the value of b, c, d is further copied to their shadow copies in B’s data section (4⇒5, 6⇒7, and 8⇒9). Note that OPEC-Monitor does not synchronize a during the operation switch because neither B nor C accesses a.

Before synchronizing, OPEC-Monitor performs data sanitization to check whether the value is legitimate by comparing the value with the developer-provided valid value range. If the check fails, there may exist data corruption inside the operation. OPEC-Monitor will abort the program and prevent the value of the corrupted shadow copy from propagating to other operations. This sanitization ensures that the new value is in safe ranges specified by developers. For instance, the speed of a moving robotic arm should be limited.

Stack The challenge of stack protection stems from accessing local variables on the previous stack frame of another operation. ACES [11] handles this challenge by using a micro-emulator. It uses one MPU region to enforce the access permission for the stack and disables as many sub-regions of this MPU region as possible to prevent access to previous portions of the stack. When the program accesses the previous portion of the stack, a memory access fault occurs. Then it invokes the micro-emulator to check against an allow list and enables this access. However, this solution requires profiling the program to get the precise stack access information, which may be imprecise.

OPEC implements precise protection for the stack by leveraging the sub-region feature of the MPU and the data relocating technique. In our design, the stack is protected by an individual MPU region. OPEC-Monitor divides the stack into eight equally sized portions. Each portion of the stack corresponds to a sub-region of the MPU region. The OPEC-Compiler analyzes the parameters of the entry function of each operation to get the size of the arguments saved on the stack and the size of data pointed by the pointer-type arguments. If these data need to be accessed by the new operation but saved in the previous operation’s stack, OPEC-Monitor will relocate

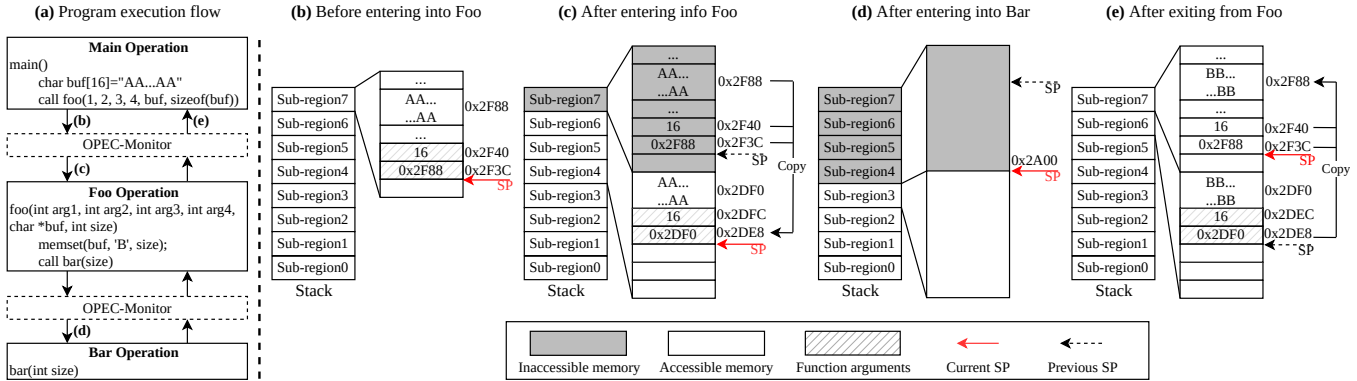


Figure 8. Illustration of stack protection by OPEC.

them when switching to a new operation. Note that the stack information (including the number of arguments and size of the buffer) is provided by the developers (Figure 5).

Specifically, OPEC-Monitor first copies the data pointed by the pointer-type arguments along with the arguments saved on the stack to the first available sub-region of the stack. Next, it redirects the pointer-type arguments to point to their duplicated copies on its stack. OPEC-Monitor disables the previous sub-regions of the stack to protect the stack from being corrupted by other operations. Then it saves the previous stack pointer and updates its value. Finally, OPEC-Monitor disables the sub-regions corresponding to the previous stack frame to prevent access. The current prototype of our system cannot handle nested pointer-type arguments of operation entry functions. In the future, the deep copy can be leveraged to solve this issue. Note that pointer-type arguments can be aliased. However, this can lead to duplicated copying and incur extra runtime overhead.

Figure 8 shows an example of the data relocation process during the operation switching. The default operation has a local variable `buf` which is saved at address `0x2F88`. Before entering into the operation `Foo`, the last two arguments of the entry function are pushed into the stack (others are saved in registers). One is the start address of `buf`, another is its size. Note that an operation can use multiple sub-regions of the stack. This strategy significantly mitigates the issue of insufficient stack space. In Figure 8(c), the operation `Foo` occupied three sub-regions. After exiting from the `Foo`, OPEC-Monitor copies back the data pointed by pointer-type arguments, i.e., copying back the array pointed by `buf`, and restores the stack pointer as shown in Figure 8(e).

Peripherals OPEC reserves four MPU regions for each operation to access four peripherals and protect every peripheral by an individual MPU region. However, one operation may need more than four MPU regions for accessing peripherals because: 1) the operation may need to access multiple peripherals; 2) one peripheral may need two more MPU regions due to the MPU region alignment requirement. To solve this problem, OPEC-Monitor *virtualizes* the MPU regions. OPEC-Monitor uses the memory management

fault handler to update the MPU configurations dynamically. Before updating the MPU, OPEC-Monitor verifies whether it is legitimate access by checking the peripheral address against the peripheral list of the current operation. If so, OPEC-Monitor will select one of the four reserved MPU regions and update its configurations. Note that OPEC-Monitor uses the round-robin algorithm to determine which MPU region should be swapped out.

Generally, memory access to code, the operation data section, and stack occurs more frequently than peripherals. Therefore, virtualizing the MPU regions that confine memory access to those three types of memory needs re-configuring the MPU frequently. It leads to OPEC-Monitor's involvement to switch context from the user space to the kernel space frequently and incurs runtime overhead.

For core peripherals that trigger a bus fault when accessed by unprivileged code, OPEC-Monitor *emulates* the execution of the `store/load` instructions. Similarly, OPEC-Monitor utilizes the bus fault handler to handle bus faults triggered by those instructions. OPEC-Monitor first checks whether the fault is triggered by unprivileged access to core peripherals. Then it examines whether the address (the target or source address of the `store` or `load` instructions) causing this fault is permitted to access by the current operation. After that, OPEC-Monitor emulates the `load` or `store` instructions to read from or write to values to those core peripherals. More specifically, to emulate a `store/load` instruction, OPEC-Monitor parses the target/source address from the instruction, writes the value to/reads the value from the parsed address, and updates the corresponding registers finally.

Heap OPEC-Compiler cannot determine which part of the heap memory is used by a function statically. Therefore, if one function uses a variable that resides in the heap memory, the whole heap memory is allowed to be accessed by this function. Moreover, OPEC-Monitor places the heap memory into a separate section rather than operation data sections, which reduces performance overhead by avoiding copying or writing back its content when switching the operation.

5.3 Operation Switch

Operation switch is triggered by executing `SVC` instructions before and after the call site of each operation entry function. Then `OPEC-Monitor` handles the `SVC` and executes the operation switch routine. `OPEC-Monitor` uses the *operation context* to save the context of previous operations so that the context can be restored. The operation context information includes MPU configurations, the stack pointer value, and the peripheral list. This operation context is only writable at the privileged level, protecting it from unauthorized changes.

When entering into a new operation, `OPEC-Monitor` copies global shadow variables and relocates variables on the stack. In the meanwhile, `OPEC-Monitor` examines the pointers field information recorded by `OPEC-Compiler` (Section 4.2). Specifically, `OPEC-Monitor` checks whether these pointers point to the shadow variables of another operation's data section. If so, `OPEC-Monitor` redirects those pointers to the shadow variables of the new operation. After that, it configures the MPU and returns to the first instruction of the new operation.

When exiting from one operation, `OPEC-Monitor` sanitizes and synchronizes the shadow variables of the current operation. Next, it recovers the stack for the previous operation and clears the value of general-purpose registers. Finally, it gets the saved context information of the target operation from the operation context and restores the MPU configuration. Therefore, task dependencies are ensured by synchronizing shared global variables and relocating stacks.

6 Evaluation

In this section, we evaluate our OPEC prototype to answer the following three research questions:

- R1: What are the security benefits provided by OPEC?
- R2: What is the performance overhead of OPEC?
- R3: What are the advantages of OPEC compared to the state-of-the-art security isolation?

We test OPEC with six representative IoT applications and the CoreMark benchmark [13] on two boards, a STM32F4-Discovery [37] development board and a STM32479I-EVAL evaluation board [35]. Both boards have an ARM Cortex-M4 core. The STM32479I-EVAL board has diverse peripherals, such as a camera and an ethernet interface. Whereas five applications are implemented by STMicroelectronics [36], PinLock is adopted from the tested applications used by ACES [11]. PinLock is a smart lock running on the STM32F4-Discovery board. It first prompts the user to enter the pin code to lock. Then the program receives the input pin from the serial port. After that, PinLock hashes the input pin and compares the result to a correct pin. Finally, it sends a message to the serial port to indicate whether the result is correct or not. Animation reads pictures from an SD card and displays those pictures on an LCD screen to demonstrate a moving butterfly. FatFs-uSD implements a FAT file system on an SD card. Then it writes some fixed content to a newly

created file in the file system. After that, it reads the file and checks whether the content is correct. LCD-uSD presents the pictures pre-stored in an SD card with fade-in and fade-out visual effects. TCP-Echo runs a TCP echo server based on lwIP [15], which is a lightweight implementation of TCP/IP protocol. It receives TCP packets sent from a client running on a desktop and replies to them. Camera uses the camera on the STM32479I-EVAL board to take a photo after the user presses the button. The picture is saved to a USB flash disk. CoreMark measures the performance of microcontrollers. It contains implementations of algorithms, including list processing, matrix manipulation, and state machine.

We generate two individual binaries for each application. The first one built from the vanilla application serves as the baseline, and the second one is built with OPEC. To evaluate our system, we first use the PinLock application as the case study to illustrate how OPEC constrain a compromised operation, whereas ACES cannot. Then, we evaluate the effectiveness of the security isolation of OPEC through several static metrics. Next, we measure the performance overhead induced by OPEC and compare it with ACES to show how lightweight OPEC is. After that, we evaluate the partition-time and execution-time over-privilege issues of OPEC and ACES. Finally, we measure the efficiency of resolving the icalls.

6.1 PinLock Case Study

We use the PinLock application as an example to demonstrate how OPEC solves the partition-time over-privilege issue and protects the system from a compromised compartment. As illustrated in Listing 1, the PinLock has six tasks. Virtually, task `Unlock_Task` and `Lock_Task` are two independent tasks that should be divided two into different compartments (or operations). Both `Lock_Task` and `Unlock_Task` invoke the function `HAL_UART_Receive_IT`, which is defined in the vendor-provided Hardware Abstraction Libraries (HAL), to receive a user input pin from the serial port. We assume there is a vulnerability in the function `HAL_UART_Receive_IT`. An attacker with the arbitrary memory write ability can exploit this vulnerability to overwrite the correct pin. The pin is finally saved to a global buffer `PinRxBuffer`. For the `Unlock_Task` function, it hashes the global variable `PinRxBuffer`. The hash result is then compared to a global variable `KEY`, which is the hash value of the correct pin. If the input pin is correct, the function `do_unlock` will be called. For the `Lock_Task` function, it checks whether the first byte of the `PinRxBuffer` is 0 and performs locking by invoking `do_lock`. Therefore, `PinRxBuffer` is shared by these two compartments.

To save the MPU regions usage, ACES groups the two global variables `PinRxBuffer` and `KEY` into one data region, which leads to the partition-time over-privilege issue (Section 3.1). Once `Lock_Task` is compromised, attackers can overwrite `KEY` and perform unlocking directly to break the isolation. However, OPEC prevents the compromised task

Table 1. The metrics of the security evaluation. The percent in the parentheses is compared to the baseline; a smaller number is better. **#OPs.**: the number of operations. **#Avg. Funcs:** the average number of functions in an operation. **#Pri. Code:** the size of code runs in the privileged level. **#Avg. GVars:** the average size of the accessible global variables.

Application	#OPs	#Avg. Funcs	#Pri. Code(%)	#Avg. GVars(%)
PinLock	6	14.00	8344(8.86)	72.86(44.43)
Animation	8	85.00	8398(5.10)	1422.44(35.61)
FatFs-uSD	10	77.00	8364(7.39)	1134.09(52.07)
LCD-uSD	11	67.82	8362(4.55)	1126.50(34.19)
TCP-Echo	9	57.67	8646(5.76)	14989.2(45.82)
Camera	9	61.11	8428(3.75)	1428.40(28.52)
CoreMark	9	17.11	8414(12.87)	1100.50(48.10)
Average	8.86	54.24	8422.29(6.90)	3039.14(41.25)

from accessing unneeded resources since its operation data section does not contain the shadow copy of KEY.

6.2 Security Evaluation

Privileged Code OPEC enforces privilege isolation by running the application code at the unprivileged level, whereas running OPEC-Monitor at the privileged level. Note that all the code executes at the privileged level in baseline because there is no privilege isolation. OPEC reduces the privileged code running at the privileged level. As the results show in Table 1, the average percentage of the privileged code size is 6.60% compared to the baseline.

As we discussed in Section 5.2, reading or writing the core registers requires the code to run at the privileged level. To solve this issue, ACES lifts the compartment to the privileged level if this compartment needs to access the core peripherals, which executes application code at the privileged level.

Accessible Global Variables We measure the accessible global variable size of each operation to understand the effectiveness of resource isolation. As shown in Table 1, OPEC can effectively confine the average percentage of accessible global variables of each operation to 41.25%. Specifically, FatFs-uSD’s average percentage of accessible global variables is high. FatFs-uSD implements a FatFs file system on an SD card to manage its storage. The program has two large structure variables, MyFile and SDFatFs, which are used as a file object and a file system object, respectively. These two variables are shared among several operations, which leads to the large average percentage of this metric of this application. PinLock, TCP-Echo, and CoreMark have a relatively high value in this metric. For PinLock, all its writable global variables are shared. For TCP-Echo, the large-size variables used for handling incoming packets are shared among four operations because of the false positives induced by point-to-analysis. Besides, the memory pools used for allocating memory are shared among five operations. For Coremark, there

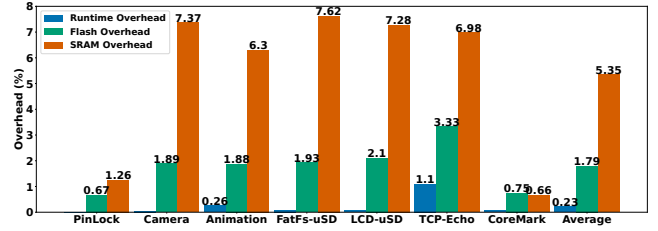


Figure 9. Performance overhead of OPEC.

are two large buffers shared among operations, incurring the average accessible global variables.

6.3 Performance Overhead

Runtime Overhead Since bare-metal embedded systems are resource-constrained devices, a lightweight security isolation scheme can protect the availability of the devices and save energy compared to the previous isolation scheme. To measure the runtime overhead caused by OPEC, we record the timestamp before and after executing the code. The executed time is the difference value between these two timestamps. We take advantage of a peripheral named Data Watchpoint and Trace (DWT) [5] to measure the runtime overhead. In particular, DWT measures the number of used clock cycles at a specific point. We first record the number of clock cycles used by the baseline and the application built with OPEC. After that, the latter is further divided by the former to calculate the final ratio. Note that DWT does not influence the execution time of applications.

PinLock stops profiling after 100 successful unlocks and locks. It receives correct and wrong pin code sent alternately from a desktop. Animation shows 11 pictures from the SD card one by one on the LCD screen and stops profiling. LCD-uSD stops profiling after displaying all 6 pictures from the SD card. It displays each picture in a short time. FatFs-uSD finishes profiling once it reads a previously written message from the created file. TCP-Echo stops running after handling 5 valid TCP packets and 45 invalid packets. Moreover, Camera stops running after the program saves the captured picture into the USB flash disk. CoreMark stops profiling after the benchmark executes the finish function. As illustrated in Figure 9, the average runtime overhead induced by OPEC is 0.23%, and the maximum overhead is 1.1%. The results show that OPEC incurs a low runtime overhead. The first reason is that OPEC adopts an operation-based partition methodology that follows the execution flow of the application, which avoids frequent domain switching. Second, OPEC synchronizes only shared variables during the operation switch. Third, all tested applications except CoreMark wait on I/O, which incurs the execution time of the baseline binaries. Note that the evaluation of TCP-Echo does not comply with ACES, which tests TCP-Echo with 1000 TCP packets, due to the SRAM limit. The overhead incurred by OPEC would be about 75% if TCP-Echo is tested with 1000 TCP packets as well as 9000 invalid packets.

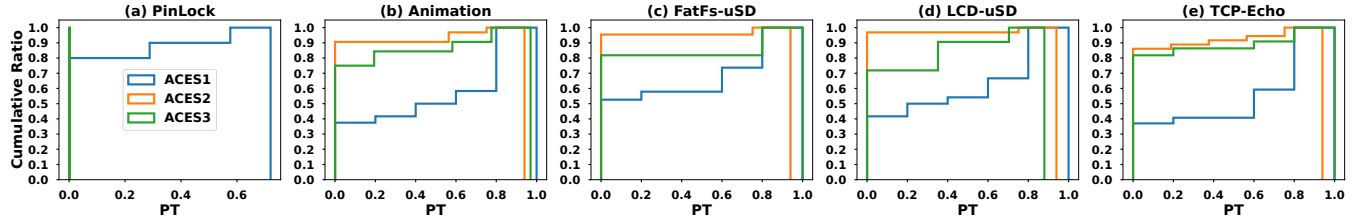


Figure 10. The cumulative ratio of PT for the five applications tested in ACES under three partitioning strategies. The legends in the last four figures are the same as those in the first one. **PT:** Partition-Time over-privilege value.

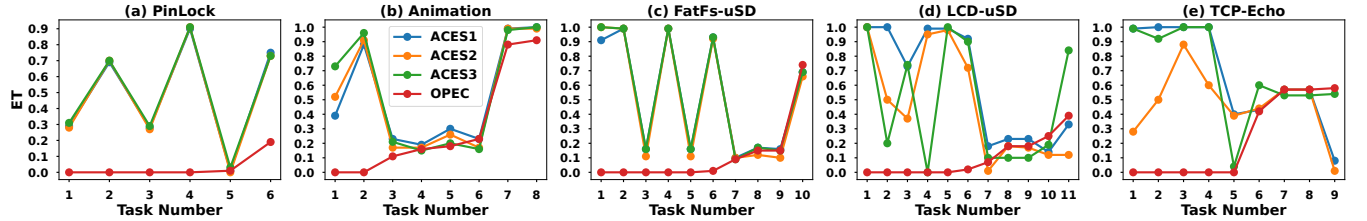


Figure 11. The evaluation results of ET for the five applications tested in ACES under three partitioning strategies. The legend in the first and last three figures are the same as those in the second one. **ET:** Execution-Time over-privilege value

Flash Overhead To measure the overall increased memory, we use a tool readelf [7] to collect each section’s memory information of the final binary. The result is further divided by the Flash size of the device to get the increased memory ratio. The Flash size of the STM32F4-Discovery board and the STM32479I-EVAL board is 1 MB and 2 MB, respectively. As shown in Figure 9, OPEC causes a minimal Flash overhead. The average Flash overhead induced by OPEC is 1.79%, and the maximum is 3.33%. Note that the operation metadata used by the OPEC-Monitor accounts for the most Flash overhead to six applications.

SRAM Overhead We measure the SRAM overhead by calculating the increased memory caused by operation data sections. The final results are divided by the size of the SRAM. The SRAM size of the STM32F4-Discovery board and the STM32479I-EVAL board is 192 KB and 288 KB, respectively. As illustrated in Figure 9, OPEC causes a moderate SRAM overhead. The maximum SRAM overhead is 7.62%. The average SRAM overhead is 5.35%. The operation data sections and their fragments required by the MPU region account for the most SRAM overhead. As discussed in Section 2, the MPU region size must be the power of two. This limitation leads to the internal fragments of operation data sections, i.e., one operation data section probably contains some unused memory. Such SRAM overhead is hard to eliminate.

6.4 Comparison to ACES

We compare OPEC with ACES [11], the state-of-the-art security isolation, for the over-privilege issues and performance overhead. We use five applications evaluated by ACES for comparison and evaluate the three supported partition strategies, i.e., filename (ACES1), filename without optimization (ACES2), and peripheral (ACES3).

Partition-time Over-privilege This issue is caused by shared global variables between compartments and the limitation of the MPU. We use a new metric, partition-time over-privilege value (PT for short), to quantitatively measure the degree of the partition-time over-privilege issue. The PT of a domain (a compartment or an operation) is the percentage of the size of the global variables that are *unnneeded* by it to the size of its accessible global variables. We consider one global variable *unnneeded* by a domain if and only if no function within the domain has a data dependency on this variable. Note that some domains may not need to access any global variable but can access unnneeded ones. Using the ratio rather than the numerator of PT can cover this situation.

The PT value can be calculated by this equation:

$$PT = \frac{\sum_i^n var2size(unnneeded_var_{f_i})}{\sum_i^n var2size(accessible_var_{f_i})} \quad (1)$$

where f_i is a function of a domain, n is the number of functions that a domain contains, and $var2size$ is the function that calculates the size of a set of variables. If a domain does not access any global variable or it does not suffer from the partition-time over-privilege issue, its PT is 0.

We evaluate the PT of each compartment of the five applications used by ACES [11] under the three different partition strategies. Figure 10 presents the cumulative ratio of PT for the compartments of different applications. Take Figure 10(b)-ACES1 as an example. The ratio of compartments whose PT is smaller than 0.2 is nearly 40%. The ratio of compartments having 0.8 PT is nearly 40% (i.e., 1-0.6). The results reveal that all strategies supported by ACES, except for the application PinLock partitioned by the strategy ACES2 and ACES3, suffer from the partition-time over-privilege issue. Note that our system avoid this issue because OPEC adopts the global variables shadowing technique, which avoids incorporating unnneeded global variables into the operation data section.

Table 2. Comparison of Runtime Overhead, Flash Overhead, Sram Overhead, and Privileged Application Code between OPEC and ACES.

Application	Policy	RO(X)	FO(%)	SO(%)	PAC(%)
PinLock	OPEC	1.00	0.74	1.40	0.00
	ACES-1	1.00	1.03	0.28	40.9
	ACES-2	1.01	0.94	0.25	3.42
	ACES-3	1.00	0.99	0.18	4.51
Animation	OPEC	1.00	1.88	6.30	0.00
	ACES-1	1.19	2.51	1.21	12.13
	ACES-2	5.70	3.53	4.30	0.34
	ACES-3	1.13	2.71	0.95	0.23
FatFs-uSD	OPEC	1.00	1.92	7.62	0.00
	ACES-1	1.90	1.57	0.43	13.86
	ACES-2	1.95	2.04	0.43	0.64
	ACES-3	1.25	1.59	0.42	0.44
LCD-uSD	OPEC	1.00	2.10	7.28	0.00
	ACES-1	1.78	2.52	4.44	23.95
	ACES-2	5.69	2.90	0.91	0.56
	ACES-3	1.72	2.10	1.24	0.45
TCP-Echo	OPEC	1.01	3.65	6.56	0.00
	ACES-1	2.17	3.67	0.64	10.75
	ACES-2	3.69	0.94	6.28	0.64
	ACES-3	1.22	0.99	0.49	0.52

Execution-time Over-privilege This issue is caused by including unnecessary code when executing a task (Section 3). Similarly, we use a new metric, execution-time over-privilege value (ET for short), to measure the degree of the execution-time over-privilege of a task. The basic idea is to calculate the used global variables when running a specific task. The ET of a task is the percentage of the size of the global variables that are actually *unused* by it during execution to the size of its *needed* global variables. We do not consider measuring the code that may execute since all unprivileged code is executable at runtime. If one function is executed in that task, the global variables needed by this function are *used*. Otherwise, they are *unused*. Note that some tasks may not use global variables but need them. Using the ratio rather than the numerator of ET can cover this situation.

To evaluate the execution-time over-privilege of a task under different partitioning strategies, we need to measure the size of its used global variables and its needed global variables. A task’s used global variables are all the global data dependency of the functions executed in the task. For OPEC, since each operation is actually a task, the needed global variables of a task are all the global data dependency of the operation. For ACES, the needed global variables are all the global data dependency of the functions within the compartments involved during execution. We can calculate the ET of a task under a specific partitioning methodology by the following equation:

$$ET = 1 - \frac{\sum_i^n \text{var2size}(\text{used_var}_{d_i})}{\sum_i^n \text{var2size}(\text{needed_var}_{d_i})} \quad (2)$$

Table 3. Efficiency of icall analysis. **#Icall**: the number of icalls. **#SVF**: the number of resolved icalls by the point-to analysis. **Time(s)**: the time used to perform the point-to analysis. **#Type**: the number of resolved icalls by the type-based analysis. **#Avg.**: the average number of targets of the icalls. **#Max**: the maximum number of targets of the icalls.

Application	#Icall	#SVF	Time(s)	#Type	#Avg.	#Max
PinLock	1	0	0.06	1	1.00	1
Animation	18	12	0.34	6	1.33	2
FatFs-uSD	18	12	0.25	6	1.72	3
LCD-uSD	19	11	0.33	0	0.58	1
TCP-Echo	29	15	5.26	13	1.28	5
Camera	35	30	6.45	5	1.77	5
CoreMark	1	1	0.12	0	2.00	2

where d_i is the domain that a task involves, n is the number of domains that a task involves.

To acquire the executed functions within a task, we need to trace the execution of the tested applications. Note that we need to trace the execution at the function granularity. We use the GDB [16] to single-stepping the whole program. Specifically, we use a script for GDB to automatically execute the tested application step by step and record the execution information, such as the line number of the source code to the logs. After that, we parse the logs and extract the executed functions.

As the results illustrated in Figure 11, our system effectively mitigates the execution-time over-privilege issue. However, there are some tasks in **LCD-uSD** and **TCP-Echo** partitioned by OPEC that have higher ET value than ACES. The reason is that those tasks contain more unnecessary code, which can be divided into two categories. The first category is the *untaken branch*. This category is prevalent, including error handling code and unmatched peripheral configuring functions. The second category is the *spurious icall target*. Due to the over-approximation of the point-to analysis, there are false positives when resolving the icalls. Thus these functions will not execute at runtime.

Performance Overhead As illustrated in Table 2, OPEC incurs overall lower runtime overhead than ACES. Moreover, the Flash overhead of OPEC is slightly smaller than the ACES. Our system incurs more SRAM overhead than ACES because OPEC uses a global variable shadowing technique. On the contrary, ACES moves global variables and merges them into a separate region, which does not increase global variables.

Privileged Application Code As presented in Table 2, ACES runs some application code at the privileged level. The reason is that this code needs privilege to access core peripherals (Section 2). However, the unprivileged application is vulnerable and may be compromised by attackers to break the security isolation. OPEC avoids it through the store/load instruction emulation technique (Section 5.2).

6.5 Efficiency of the Icall Analysis

This section evaluates the efficiency of our icall analysis because it affects building a precise function call graph which is essentially important in the operation partitioning stage (Section 4). However, it is infeasible to get the ground truth of the targets for each icall by running the tested applications because even legitimate targets may not execute at runtime. Therefore, we use the average and maximum targets for each application to measure the efficiency of icall analysis. Table 3 presents the results. In summary, the time cost for performing the point-to analysis via SVF is less because all the tested applications have a small code size. As previously discussed, for those icalls that SVF cannot resolve, we use the type-based analysis to analyze the potential targets. Among the applications, the maximum number of targets of those icalls is 5 (in TCP-Echo and Camera). For the average targets of the icalls, the maximum average is 1.77 (in Camera). Note that there are still some unresolved icalls. For LCD-uSD, there are eight unresolved icalls in an IRQ handler function running at the privileged level, which does not interfere with the unprivileged operations. For TCP-Echo, there is one unresolved icall in the function `udp_input`. Since TCP-Echo processes only TCP packets, some source files related to handling UDP packets are removed manually. Besides, TCP-Echo will not execute this icall. The results show that the false positives of the icalls' targets have a limited impact.

7 Discussion and Future Work

Confused Deputy Attacks Attackers may compromise and manipulate an ordinary operation to feed malicious inputs to a critical operation, such as the `Unlock_Task` operation of PinLock, to trigger the execution of security-sensitive functions. Security isolation is vulnerable to this kind of attacks [9] [18]. OPEC provides data sanitization to shared variables between operations to mitigate such attacks.

Defects of Static Analysis OPEC uses static analysis to build a sound function call graph and to determine the resource dependency of each function. For icalls and implicit data usage through pointers, OPEC leverages the point-to analysis to find out their potential legitimate targets. In particular, we use the SVF [38] to conduct the point-to analysis. However, the SVF guarantees soundness but sacrifices precision, i.e., the results are over-approximated and contain false positives. Although it may introduce the execution-time over-privilege issue, it avoids program execution errors caused by missing dependency. If one of the operation entry functions is wrong, the resource dependency of that operation will be incomplete. When accessing those missed resources confined by the MPU, a memory management fault will be triggered (Section 2.2). The failure of static analysis leads to either incomplete or redundant dependency. The former case triggers a memory management fault. The latter case

worsens the isolation. The static analysis is not the contribution of our work, and the improvements of the static analysis could be transparently integrated into OPEC.

Heap As discussed in Section 5.2, OPEC supports dynamic memory allocation, i.e., heap. However, if the program uses the heap implemented by the static libraries, e.g., glibc [6], OPEC-Compiler cannot identify such access. If the program uses such kind of heap, it needs to be modified by replacing dynamically allocated memory objects with global variables. In our experiment, four-sevenths of the tested applications use the heap implemented by the static libraries to allocate memory. We can replace heap objects with global variables among the three-fourths of applications because each application requests heap memory only once. The last application TCP-Echo uses memory pools to allocate memory dynamically. OPEC can exclude unneeded memory pools from one operation's resource dependency. Only correlative memory pools are placed in a separate data section. In the future, OPEC can adopt a secure heap allocator to protect the heap.

Concurrency Our system is aimed to provide security isolation to bare-metal embedded systems, which execute single monolithic programs with a single thread. To support multi-threading systems, OPEC needs extra engineering efforts. For the single-core systems, during the context switch, OPEC needs to: (1) write back the shadow copies of the previous thread's operation and synchronize the shadow copies of the new thread's operation; (2) reconfigure the MPU.

For multi-core systems, OPEC also needs to ensure the consistency of global variables shared by operations of two different threads since each thread has independent shadow copies. There are two different solutions: (1) OPEC synchronizes the shared global variables after/before the thread enters/leaves the critical section; (2) the OS scheduler binds the threads sharing global variables to the same CPU core.

Other Hardware Platforms Applying OPEC to other hardware platforms faces three main challenges. First, the target hardware platform is required to have a memory protection unit, which has enough regions enforcing the physical memory permissions similar to the ARM MPU, e.g., RISC-V PMP [33]. Second, the target architecture should be supported by the LLVM. Third, developers need to modify the exception handlers to adjust the load/store instruction emulation and the operation switch procedures.

8 Related Work

Besides being related to the research on enforcing security isolation on bare-metal embedded systems, OPEC is also associated with the work towards providing security defenses and ensuring the correct execution of the embedded systems.

Security Isolation Clements et al. [12] proposes EPOXY, which enforces privilege isolation on bare-metal embedded systems by executing sensitive instructions at the privileged level. However, EPOXY neglects resource isolation. Kim et

al. [22] proposes MINION, which extends the previous work and enforces thread-level memory isolation on real-time microcontroller systems. Furthermore, Clements et al. [11] proposes ACES that enforces privilege isolation and sub-thread level memory isolation. Huo et al. [21] extends ACES by adopting a machine learning-assisted compartmentalization scheme. TrustLite [24] introduces execution-aware MPU to isolate and prevent trusted components from unauthorized modifications. TyTAN [8] provides security isolation by using the hardware extensions proposed by TrustLite. However, they require a customized hardware extension which is difficult to deploy. There are other works that enforce security isolation on commodity software [10, 17, 26, 27, 40, 41]. Alejandro et al. [30] enhances existing isolation schemes by securing DMA operations of embedded applications. OPEC proposes privilege isolation and improved resource isolation to bare-metal embedded systems.

Security Protection Techniques Various systems intend to deploy protection or mitigation techniques widely used on desktop systems to deeply embedded systems. EPOXY [12] uses DEP and a modified SafeStack to prevent control-flow hijacking attacks. It also uses diversification to introduce uncertainty to the produced firmware, which raises the cost of attackers to analyze the firmware. Abbasi et al. [1] quantitatively investigate the mitigation techniques adopted in deeply embedded systems. It further proposes uArmor that implements DEP and stack canaries for those systems. Kwon et al. [25] realizes an efficient eExecute-Only-Memory on Cortex-M microcontrollers by leveraging the unprivileged memory instructions of the ARMv7-M architecture. Zhou et al. [46] also takes advantage of this feature to design a shadow stack on low-end embedded systems. Almakhdhub et al. [3] proposes uRAI to protect the return address integrity of microcontroller-based embedded systems. Moreover, it enforces software fault isolation (SFI) on the privileged exception handlers. ARM TrustZone-M provides a secure world for isolating security-sensitive memory and peripherals for ARMv8-M architecture microcontrollers. However, it is insufficient to isolate multiple tasks only with this hardware feature. Rust [34] ensures memory safety. However, it cannot solve partition-time over-privilege and execution-time over-privilege issues (Section 3.1). These advanced security protection techniques are orthogonal to the security isolation provided by OPEC. They can be integrated into our system to protect the code and data inside one operation.

Correctness Assurance Some research that focuses on diagnosing embedded systems. Kim et al. [23] proposes MAYDAY to facilitate the post-accident analysis of drones. Fang et al. [14] uses the record and replay technique to troubleshoot the errors of embedded systems. Niesler et al. [31] propose HERA that exploits a hardware feature named Flash Patch and Breakpoint unit to hotpatch the real-time embedded systems. Yi et al. [19] proposes RapidPatch that ports the eBPF

virtual machine to embedded systems, which allows heterogeneous devices to apply patches generated from the same patch source file. Similar works [42–44] also introduce hot patching to embedded systems. Xu et al. [45] and Huber et al. [20] aim to recover the embedded systems even when they are compromised. Another category of research enables remote attestation to ensure the code integrity, control-flow integrity, and data integrity of the embedded systems [2, 32, 39]. OPEC enforces the operation isolation on bare-metal embedded systems and constrains the compromised component.

9 Conclusion

In this paper, we propose lightweight and fine-grained security isolation for bare-metal embedded systems. We prototype our system, OPEC, which enforces privilege isolation and fine-grained resource isolation. At compile time, OPEC-Compiler partitions the program into operations and generates the policy file, which describes the resource dependency of each operation. At runtime, OPEC-Monitor enforces the hardware-assisted security isolation by applying the policy through the MPU.

OPEC solves the partition-time over-privilege via global variable shadowing. OPEC-Compiler creates shadow copies for shared global variables in the corresponding operation data section at compile time, and OPEC-Monitor synchronizes the value during operation switch. To confine the access to peripherals scalably, OPEC virtualizes the MPU regions reserved for general peripherals and emulates store/load instructions accessing core peripherals. The instruction emulation avoids executing unprivileged application code at the privileged level. OPEC mitigates the execution-time over-privilege through operation-based program partitioning.

To evaluate our prototype, we test OPEC with six IoT applications and the CoreMark benchmark on two different boards. We propose two different metrics for evaluating the partition-time over-privilege issue and execution-time over-privilege issue, respectively. The results demonstrate the security benefits of our system. Moreover, OPEC induces low runtime overhead and moderate memory overhead.

Acknowledgments

The authors would like to thank the anonymous reviewers and our shepherd Kiwan Maeng for their insightful comments and feedback. This work was supported by the National Natural Science Foundation of China under Grant U21A20464, the Fundamental Research Funds for the Central Universities (Zhejiang University NGICS Platform), Leading Innovative and Entrepreneur Team Introduction Program of Zhejiang (2018R01005). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of funding agencies.

A Artifact Appendix

A.1 Abstract

Our artifact includes the source code of OPEC, a docker image containing the compiled OPEC, and the experimental scripts to reproduce the tables (Table 1, 2, and 3) and the figures (Figure 9, 10, and 11) in Section 6.

A.2 Description & Requirements

A.2.1 How to access.

- The repository that contains the source code of OPEC and its experimental scripts is publicly available at GitHub: <https://github.com/XiaZhouZero/OPEC>.
- Archived (DOI): <https://zenodo.org/badge/latestdoi/452701528>.
- The docker image that contains the compiled OPEC is publicly available at DockerHub: <https://hub.docker.com/r/zhouxzju/opec>.

A.2.2 Hardware dependencies.

- One desktop with 6 CPU cores, 16GB memory, and 60GB free hardware disk.
- One STM32F407 board and one STM32479I-Eval board.

A.2.3 Software dependencies.

- Ubuntu 18.04 and its packages: `zlib1g-dev`, `git`, `build-essential`, `gcc-arm-none-eabi`, `make`, `texinfo`, `bison`, `flex`, `cmake`, `ninja-build`, `libusb-dev`, `python3-pip`, `textlive-full`, `clang`, `libusb-dev`, `openocd`
- Python 3.6.9 and its packages: `networkx==2.5`, `pydot`, `matplotlib`, `serial`, `pydotplus`, `prettytable`

A.2.4 Benchmarks.

We evaluate the OPEC prototype with the CoreMark benchmark: <https://www.eembc.org/coremark/>.

A.3 Set-up

Please follow the instructions, which are available at <https://github.com/XiaZhouZero/OPEC/blob/main/README.md>, to set up the experimental environment and build OPEC.

A.4 Evaluation workflow

A.4.1 Major Claims.

- (C1): OPEC enforces privilege isolation for bare-metal embedded systems. This is proven by the security evaluation in Section 6.2 whose results are presented in Table 1.
- (C2): OPEC enforces resource isolation by confining the access to peripherals and global variables. The confinement to accessing peripherals are described in Section 5.2. The average reduced accessible global variables of each tested application is evaluated in Section 6.2 whose results are presented in Table 1.
- (C3): OPEC is a lightweight security isolation scheme for bare-metal embedded systems. This is proven by

the performance evaluation in Section 6.3 (whose results are illustrated in Figure 9) and the comparison to ACES in Section 6.4 (whose results are illustrated in Table 2).

- (C4): OPEC solves the partition-time over-privilege issue (whose results are depicted in Figure 10) and reduces the execution-time over-privilege issue (whose results are depicted in Figure 11).
- (C5): OPEC performs point-to analysis to resolve the indirect function calls (*icall* for short) in the programs. The evaluation of the *icall* analysis efficiency is in Section 6.5 whose results are in Table 3.

A.4.2 Experiments.

Experiment (Security Evaluation): [20 human-minutes + 20 computer-minutes]: This experiment measures some metrics of each tested application: the number of operations, the average number of functions of operations, the size of privileged code, and the average size of accessible global variables of operations.

[How to]

[Preparation] Build all the tested applications.

[Execution] Enter into the `experiments/table1` directory and run `make` command on the shell.

[Results] Table 1 will be printed on the console.

Experiment (Performance Evaluation): [1 human-hour + 1 computer-hour]: This experiment measures the runtime overhead, flash overhead, and SRAM overhead incurred by OPEC.

[How to]

[Preparation] Build all the tested applications. Run each tested application on the STM32F407 board or the STM32479I-Eval board to collect the runtime results. Note that only the runtime overhead results are pre-stored in the directory. Please contact the authors if one has any trouble in setting up the testing environment and evaluating the runtime overhead by himself.

[Execution] Enter into the `experiments/figure9` directory and run `make` command on the shell.

[Results] Figure 9 will be generated in the `experiments/-figure9` directory.

Experiment (Partition-time Over-privilege Evaluation): [20 human-minutes + 20 computer-minutes]: This experiment measures the partition-time over-privilege issue of OPEC and ACES.

[How to]

[Preparation] Build all the tested applications.

[Execution] Enter into the experiments/figure10 directory and run make command on the shell.

[Results] Figure 10 will be generated in the experiments/figure10 directory.

Experiment (Execution-time Over-privilege Evaluation): [20 human-minutes + 20 computer-minutes]: This experiment measures the execution-time over-privilege issue of OPEC and ACES.

[How to]

[Preparation] Build all the tested applications.

[Execution] Enter into the experiments/figure11 directory and run make command on the shell.

[Results] Figure 11 will be generated in the experiments/figure11 directory.

Experiment (Comparison to ACES): [20 human-minutes + 20 machine-minutes]: This experiment compares the performance overhead between OPEC and ACES.

[How to]

[Preparation] Build the five tested applications evaluated by ACES. Note that the runtime overhead results of ACES are acquired from ACES's paper <https://www.usenix.org/conference/usenixsecurity18/presentation/clements>.

[Execution] Enter into the experiments/table2 directory and run make command on the shell.

[Results] Table 2 will be printed on the console.

Experiment (Icall Evaluation): [20 human-minutes + 20 machine-minutes]: This experiment evaluates the efficiency of the icall analysis.

[How to]

[Preparation] Build all the tested applications.

[Execution] Enter into the experiments/table3 directory and

run make command on the shell.

[Results] Table 3 will be printed on the console.

A.5 Notes on Reusability

Function `do_region_switch` performs the virtualization of MPU regions reserved for peripherals. Its code is in the `operation-rt.c` file at the `compiler/operation-rt` directory.

Function `do_peripheral_rw` emulates the execution of load/store thumb2 instructions that access core peripherals on the private peripheral bus. Its code is in the `operation-rt.c` file at the `compiler/operation-rt` directory.

The source files of the LLVM pass used for analyzing resource dependency are at the `compiler/analysis/def-use` directory.

The source files of the LLVM pass used for instrumenting the program are at the `compiler/llvm` directory.

References

- [1] Ali Abbasi, Jos Wetzels, Thorsten Holz, and Sandro Etalle. 2019. Challenges in designing exploit mitigations for deeply embedded systems. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. 31–46.
- [2] Tigist Abera, N Asokan, Lucas Davi, Jan-Erik Ekberg, Thomas Nyman, Andrew Paverd, Ahmad-Reza Sadeghi, and Gene Tsudik. 2016. C-FLAT: control-flow attestation for embedded systems software. In *Proceedings of the 23rd ACM SIGSAC Conference on Computer and Communications Security*. 743–754.
- [3] Naif Saleh Almkhndhub, Abraham A Clements, Saurabh Bagchi, and Mathias Payer. 2020. uRAI: Securing Embedded Systems with Return Address Integrity. In *Network and Distributed Systems Security (NDSS) Symposium*.
- [4] Arm. 2022. ARMv7-M Architecture Reference Manual. Retrieved March 16, 2022 from <https://developer.arm.com/documentation/ddi0403/latest>
- [5] Arm. 2022. Data Watchpoint and Trace Unit. Retrieved March 16, 2022 from <https://developer.arm.com/documentation/ddi0439/b/Data-Watchpoint-and-Trace-Unit>
- [6] Arm. 2022. GNU Arm Embedded Toolchain. Retrieved March 16, 2022 from <https://developer.arm.com/tools-and-software/open-source-software/developer-tools/gnu-toolchain/gnu-rm>
- [7] binutils. 2022. . Retrieved March 16, 2022 from <https://man7.org/linux/man-pages/man1/readelf.1.html>
- [8] Ferdinand Brasser, Brahim El Mahjoub, Ahmad-Reza Sadeghi, Christian Wachsmann, and Patrick Koeberl. 2015. TyTAN: Tiny trust anchor for tiny devices. In *Proceedings of the 52nd annual design automation conference*. 1–6.
- [9] Sven Bugiel, Lucas Davi, Alexandra Dmitrienko, Thomas Fischer, Ahmad-Reza Sadeghi, and Bhargava Shastry. 2012. Towards Taming Privilege-Escalation Attacks on Android. In *Network and Distributed Systems Security (NDSS) Symposium*.
- [10] Yaohui Chen, Sebassujeen Reymondjohnson, Zhichuang Sun, and Long Lu. 2016. Shreds: Fine-Grained Execution Units with Private Memory. In *Proceedings of the 2016 IEEE Symposium on Security and Privacy (SP 16)*. 56–71.
- [11] Abraham A Clements, Naif Saleh Almkhndhub, Saurabh Bagchi, and Mathias Payer. 2018. ACES: Automatic Compartments for Embedded Systems. In *27th USENIX Security Symposium (USENIX Security 18)*. 65–82.

- [12] Abraham A Clements, Naif Saleh Almkhduh, Khaled S Saab, Prashast Srivastava, Jinkyu Koo, Saurabh Bagchi, and Mathias Payer. 2017. Protecting bare-metal embedded systems with privilege overlays. In *Proceedings of the 2017 IEEE Symposium on Security and Privacy (SP 17)*. 289–303.
- [13] EEMBC. 2022. EEMBC, “Coremark - industry-standard benchmarks for embedded. Retrieved March 16, 2022 from <https://www.eembc.org/coremark/>
- [14] Kaiming Fang and Guanhua Yan. 2020. IoTReplay: Troubleshooting COTS IoT Devices with Record and Replay. In *2020 IEEE/ACM Symposium on Edge Computing (SEC)*. 193–205.
- [15] Free Software Foundation. 2022. lwIP - A Lightweight TCP/IP stack. Retrieved March 16, 2022 from <https://savannah.nongnu.org/projects/lwip/>
- [16] GNU. 2022. GDB: The GNU Project Debugger. Retrieved March 16, 2022 from <https://www.gnu.org/software/gdb/>
- [17] Khilan Gudka, Robert NM Watson, Jonathan Anderson, David Chisnall, Brooks Davis, Ben Laurie, Ilias Marinou, Peter G Neumann, and Alex Richardson. 2015. Clean application compartmentalization with SOAAP. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. 1016–1031.
- [18] Norm Hardy. 1988. The Confused Deputy: (or why capabilities might have been invented). *ACM SIGOPS Operating Systems Review* 22, 4 (1988), 36–38.
- [19] Yi He, Zhenhua Zou, Kun Sun, Zhuotao Liu, Ke Xu, Qian Wang, Chao Shen, Zhi Wang, and Qi Li. 2022. RapidPatch: Firmware Hotpatching for Real-Time Embedded Devices. In *31th USENIX Security Symposium (USENIX Security 22)*.
- [20] Manuel Huber, Stefan Hristozov, Simon Ott, Vasil Sarafov, and Marcus Peinado. 2020. The Lazarus Effect: Healing Compromised Devices in the Internet of Small Things. In *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*. 6–19.
- [21] Dongdong Huo, Chao Liu, Xiao Wang, Mingxuan Li, Yu Wang, Yazhe Wang, Peng Liu, and Zhen Xu. 2020. A Machine Learning-Assisted Compartmentalization Scheme for Bare-Metal Systems. In *International Conference on Information and Communications Security*. 20–35.
- [22] Chung Hwan Kim, Taegy Kim, Hongjun Choi, Zhongshu Gu, Byoungyoung Lee, Xiangyu Zhang, and Dongyan Xu. 2018. Securing Real-Time Microcontroller Systems through Customized Memory View Switching. In *Network and Distributed Systems Security (NDSS) Symposium*.
- [23] Taegy Kim, Chung Hwan Kim, Altay Ozen, Fan Fei, Zhan Tu, Xiangyu Zhang, Xinyan Deng, Dave Jing Tian, and Dongyan Xu. 2020. From Control Model to Program: Investigating Robotic Aerial Vehicle Accidents with MAYDAY. In *29th USENIX Security Symposium (USENIX Security 20)*. 913–930.
- [24] Patrick Koerber, Steffen Schulz, Ahmad-Reza Sadeghi, and Vijay Varadharajan. 2014. TrustLite: A security architecture for tiny embedded devices. In *Proceedings of the Ninth European Conference on Computer Systems*. 1–14.
- [25] Donghyun Kwon, Jangseop Shin, Giyeol Kim, Byoungyoung Lee, Yeongpil Cho, and Yunheung Paek. 2019. uXOM: Efficient eExecute-Only Memory on ARM Cortex-M. In *28th USENIX Security Symposium (USENIX Security 19)*. 231–247.
- [26] Shen Liu, Dongrui Zeng, Yongzhe Huang, Frank Capobianco, Stephen McCamant, Trent Jaeger, and Gang Tan. 2019. Program-mandering: Quantitative privilege separation. In *Proceedings of the 26th ACM SIGSAC Conference on Computer and Communications Security*. 1023–1040.
- [27] Yutao Liu, Tianyu Zhou, Kexin Chen, Haibo Chen, and Yubin Xia. 2015. Thwarting memory disclosure with efficient hypervisor-enforced intradomain isolation. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. 1607–1619.
- [28] LLVM. 2022. The LLVM Compiler Infrastructure. Retrieved March 16, 2022 from <https://llvm.org/>
- [29] Kangjie Lu and Hong Hu. 2019. Where does it go? refining indirect-call targets with multi-layer type analysis. In *Proceedings of the 26th ACM SIGSAC Conference on Computer and Communications Security*. 1867–1881.
- [30] Alejandro Mera, Yi Hui Chen, Ruimin Sun, Engin Kirda, and Long Lu. 2022. D-Box: DMA-enabled Compartmentalization for Embedded Applications. In *Network and Distributed Systems Security (NDSS) Symposium*.
- [31] Christian Niesler, Sebastian Surminski, and Lucas Davi. 2021. HERA: Hotpatching of Embedded Real-time Applications. In *Proceedings of THE 28th Network and Distributed System Security Symposium*.
- [32] Ivan De Oliveira Nunes, Karim Eldefrawy, Norrathep Rattanavipanon, and Gene Tsudik. 2020. APEX: A Verified Architecture for Proofs of Execution on Remote Devices under Full Software Compromise. In *29th USENIX Security Symposium (USENIX Security 20)*. 771–788.
- [33] RISC-V. 2022. RISC-V Physical Memory Protection. Retrieved March 16, 2022 from <https://riscv.org/technical/specifications/>
- [34] Rust. 2022. Rust Programming Language. Retrieved March 16, 2022 from <https://www.rust-lang.org/>
- [35] STMicroelectronics. 2022. STM32479I-EVAL Evaluation Board. Retrieved March 16, 2022 from <https://www.st.com/en/evaluation-tools/stm32479i-eval.html>
- [36] STMicroelectronics. 2022. STM32Cube MCU Full Package. Retrieved March 16, 2022 from https://github.com/STMicroelectronics/STM32CubeF4/tree/master/Projects/STM32469I_EVAL
- [37] STMicroelectronics. 2022. STM32F4-Discovery Board. Retrieved March 16, 2022 from <https://www.st.com/en/evaluation-tools/stm32f4discovery.html>
- [38] Yulei Sui and Jingling Xue. 2016. SVF: interprocedural static value-flow analysis in LLVM. In *Proceedings of the 25th international conference on compiler construction*. 265–266.
- [39] Zhichuang Sun, Bo Feng, Long Lu, and Somesh Jha. 2020. OAT: Attesting operation integrity of embedded devices. In *Proceedings of the 2020 IEEE Symposium on Security and Privacy (SP 20)*. 1433–1449.
- [40] Nikos Vasilakis, Ben Karel, Nick Roessler, Nathan Dautenhahn, André DeHon, and Jonathan M Smith. 2017. Towards fine-grained, automated application compartmentalization. In *Proceedings of the 9th Workshop on Programming Languages and Operating Systems*. 43–50.
- [41] Nikos Vasilakis, Ben Karel, Nick Roessler, Nathan Dautenhahn, André DeHon, and Jonathan M Smith. 2018. BreakApp: Automated, Flexible Application Compartmentalization. In *Network and Distributed Systems Security (NDSS) Symposium*.
- [42] Michael Wahler and Manuel Oriol. 2014. Disruption-free software updates in automation systems. In *Proceedings of the 2014 IEEE Emerging Technology and Factory Automation (ETFA)*. 1–8.
- [43] Michael Wahler, Stefan Richter, Sumit Kumar, and Manuel Oriol. 2011. Non-disruptive large-scale component updates for real-time controllers. In *2011 IEEE 27th International Conference on Data Engineering Workshops*. 174–178.
- [44] Michael Wahler, Stefan Richter, and Manuel Oriol. 2009. Dynamic software updates for real-time systems. In *Proceedings of the 2nd International Workshop on Hot Topics in Software Upgrades*. 1–6.
- [45] Meng Xu, Manuel Huber, Zhichuang Sun, Paul England, Marcus Peinado, Sangho Lee, Andrey Marochko, Dennis Mattoon, Rob Spiger, and Stefan Thom. 2019. Dominance as a new trusted computing primitive for the internet of things. In *Proceedings of the 2019 IEEE Symposium on Security and Privacy (SP 19)*. 1415–1430.
- [46] Jie Zhou, Yufei Du, Zhuojia Shen, Lele Ma, John Criswell, and Robert J Walls. 2020. Silhouette: Efficient protected shadow stacks for embedded systems. In *29th USENIX Security Symposium (USENIX Security 20)*. 1219–1236.