SCALPEL: Exploring the Limits of Tag-enforced Compartmentalization

NICK ROESSLER and ANDRÉ DEHON, University of Pennsylvania, USA

2

We present Secure Compartments Automatically Learned and Protected by Execution using Lightweight metadata (SCALPEL), a tool for automatically deriving compartmentalization policies and lowering them to a tagged architecture for hardware-accelerated enforcement. SCALPEL allows a designer to explore high-quality points in the privilege-reduction vs. performance overhead tradeoff space using analysis tools and a detailed knowledge of the target architecture to make best use of the available hardware. SCALPEL automatically implements hundreds of compartmentalization strategies across the privilege-performance tradeoff space, all without manual tagging or code restructuring. SCALPEL uses two novel optimizations for achieving highly performant policies: the first is an algorithm for packing policies into working sets of rules for favorable rule cache characteristics, and the second is a rule prefetching system that allows it to exploit the highly predictable nature of compartmentalization rules. To create policies, SCALPEL introduces a quantitative privilege metric (the Overprivilege Ratio) that is used to drive its algorithmic compartment generation. We implement SCALPEL on a FreeRTOS stack and target a tag-extended RISC-V core. Our results show that SCALPEL-created policies can reduce overprivilege by orders of magnitude with hundreds of logical compartments while imposing low overheads (<5%).

CCS Concepts: • Security and privacy → Embedded systems security; Mobile platform security; Software security engineering; Access control;

Additional Key Words and Phrases: Security, least privilege, tagged architectures, privilege quantification, privilege-performance tradeoffs, automated compartmentalization

ACM Reference format:

Nick Roessler and André DeHon. 2021. SCALPEL: Exploring the Limits of Tag-enforced Compartmentalization. *J. Emerg. Technol. Comput. Syst.* 18, 1, Article 2 (September 2021), 28 pages.
<https://doi.org/10.1145/3461673>

1 INTRODUCTION

Modern software stacks are notoriously vulnerable. Operating systems, device drivers, and countless applications, including most embedded applications, are written in unsafe languages and run

This research was funded by DARPA contract HR0011-18-C-0011 under the System Security Integrated Through Hardware and Firmware (SSITH) program. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not reflect the official policy or position of the U.S. Government.

Authors' address: N. Roessler and A. DeHon, University of Pennsylvania, 200 S. 33rd St., Philadelphia, PA, 19104; emails: nick.roessler@protonmail.com, andre@acm.org.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

© 2021 Copyright held by the owner/author(s).
1550-4832/2021/09-ART2 \$15.00
<https://doi.org/10.1145/3461673>

in large, monolithic protection domains where any single vulnerability may be sufficient to compromise an entire machine. Privilege separation is a defensive approach in which a system is separated into components, and each is limited to (ideally) just the privileges it requires to operate. In such a separated system, a vulnerability in one component (e.g., the networking stack) is isolated from other system components (e.g., sensitive process credentials), making the system substantially more robust to attackers, or at least increasing the effort of exploitation in cases where it is still possible.

While the principle of least privilege [44] is a powerful guiding force in secure system design, in practice it is often at odds with system performance. Given the limited hardware resources that have been allocated for security, privilege separation has typically relied on coarse-grained, process-level separation in which the virtual memory system is used to provide isolation. For example, some security-critical and network-facing software such as OpenSSH [9] and Dovecot [19] have been manually decomposed into multiple intercommunicating-but-isolated processes to contain the effects of vulnerabilities, should they be found.

While these select cases are a win for privilege reduction, the prevailing wisdom has been that only coarse-grained privilege separation is feasible in practice, given the high cost of virtual memory context switching. Indeed, all modern OSs run on insecure but performant monolithic kernels, with more functionality frequently moving into the highly privileged kernel to reduce such costs; privilege separated microkernels, in contrast, remain plagued with the perception of high overheads and have seen little adoption. IoT and embedded systems—which we now find ourselves surrounded by in our everyday lives—have fallen even farther behind in security than their general-purpose counterparts. They are also written in memory unsafe languages, typically C, often lack basic modern exploit mitigations [1], and many run directly on bare metal with no separation between any parts of the system at all.

However, there has recently been a surge of interest—both academic and in industry—in architectural and hardware support for new security primitives. For example, ARM recently announced that it will integrate hardware capability support (CHERI [55]) into its chip designs, Oracle has released SPARC processors with coarse-grained memory tagging support (ADI [37]), and NXP has announced it will use Dover’s CoreGuard [46], among many others [4]. One interesting and practical use case for these primitives is privilege separation enforcement. In this article, we build privilege separation policies for a fine-grained, hardware-accelerated security monitor design (the PIPE architecture [17, 51], Section 2.1).

A flexible, tag-based hardware security monitor, like the PIPE, provides an exciting opportunity to enforce fine-grained, hardware-accelerated privilege separation. At a bird’s-eye view, one can imagine using metadata tags on code and data to encode logical protection domains, with rules dictating which memory operations and control-flow transitions are permitted. The PIPE leaves tag semantics to software interpretation, meaning one can express policies ranging from coarse-grained decompositions, such as a simple separation between “trusted” and “untrusted” components, to hundreds or thousands of isolated compartments depending on the privilege reduction and performance characteristics that are desired.

To explore this space, we present **SCALPEL (Secure Compartments Automatically Learned and Protected by Execution using Lightweight metadata)**, a tool that enables the rapid self-learning of low-level privileges and the automatic creation and implementation of compartmentalization security policies for a tagged architecture. At its back-end, SCALPEL contains a policy compiler that decouples logical compartmentalization policies from their underlying concrete enforcement with the PIPE architecture. The back-end takes as input a particular compartmentalization strategy, formulated in terms of C-level constructs and their allowed privileges, and then automatically tags a program image to instantiate the desired policy. To ease policy creation and

exploration, the SCALPEL front-end provides a tracing mode, compartment-generation algorithms, and analysis tools, to help an engineer quickly create, compare, and then instantiate strategies using the back-end. These tools build on a range of similar recent efforts that treat privilege assessment quantitatively and compartment generation algorithmically [11, 34, 38], allowing SCALPEL's automation to greatly assist in the construction of good policies, a task that would otherwise be costly in engineering time. In cases where human expertise is available for additional fine-tuning, SCALPEL easily integrates human supplied knowledge in its policy exploration; for example, a human can add additional constraints to the algorithms, such as predefining a set of boundaries, adjusting the weight on an object, or specifying that a particular object is security-critical and should not be exposed to additional, unnecessary code.

Additionally, SCALPEL presents two novel techniques for optimizing security policies to a tagged architecture. The first is a policy-construction algorithm that directly targets the rule cache characteristics of an application: The technique is based on packing rules needed for different program phases into sets that can be cached favorably. While we apply this technique on SCALPEL's compartmentalization policies, the core idea could be used to improve the performance of other policies on tagged architectures. Additionally, we show that this same technique can be used to pack an *entire policy* into a fixed-size set such that no rule cache misses will be taken besides compulsory misses—this makes it possible to achieve real-time guarantees while using a hardware security monitor like the PIPE, which may be of particular value to embedded, real-time devices and applications. Second, we design and evaluate a rule prefetching system that exploits the highly predictable nature of compartmentalization rules; by intelligently prefetching rules before they are needed, we show that the majority of stalled cycles spent waiting for policy evaluation can be avoided.

We evaluate SCALPEL and its optimizations on a typical embedded IoT environment consisting of a FreeRTOS stack targeting a PIPE-extended RISC-V core. We implement our policies on several applications, including an HTTP web server, the bzip (de)compressor, an H264 video encoder, the GNU Go engine, and the libXML parsing library. Using SCALPEL, we show how to automatically derive compartmentalization strategies for off-the-shelf software that balance privilege reduction with performance and that hundreds of isolated compartments can be simultaneously enforced with acceptable overheads on a tagged architecture.

To summarize, SCALPEL combines (1) hardware support for fine-grained metadata tagging and policy enforcement with (2) compartmentalization and privilege analysis tools, which together allow a thorough exploration of the level of privilege separation that can be achieved with hardware tagging support. Our primary contributions are:

- A tool that automatically creates and instantiates tag-based compartmentalization policies on real software without manual refactorings.
- Compartment-generation algorithms and analysis tools that quantify the privilege exposure and performance of a wide range of automatically generated compartmentalization alternatives, providing the security engineer with a variety of privilege-performance design points to explore and evaluate.
- New techniques for using a tagged architecture for the rapid self-learning of privileges on unmodified software.
- A new technique for optimizing a security policy to a tagged architecture by directly targeting its rule cache characteristics.
- A new rule prefetching technique for tagged architectures in which additional rules are identified and prefetched on a cache miss.

2 BACKGROUND AND RELATED WORK

2.1 The PIPE Architecture

Tag-based hardware security monitors can be used to improve software security by detecting and preventing violations of security policies at runtime [10, 14, 17, 24, 37]. The **PIPE (Processor Interlocks for Policy Enforcement)** [51] is a software/hardware co-designed processor extension for hardware-accelerated security monitor enforcement. The core idea is that word-sized metadata tags are associated with data words in the system, including on register values, words stored in memory, and also on the program counter. As each instruction executes on the primary processor core (referred to as the application or *AP core*; see Figure 1), the tags relevant to the instruction are used to validate the operation against a software security monitor, typically in parallel with instruction execution.

This policy evaluation is performed on a dedicated coprocessor, the **policy execution (PEX)** core. The semantics of tags are entirely determined by how the policy software interprets the tag bits, allowing the expression of a rich range of security policies. The software monitor determines if a particular set of tags represents a valid operation, and if so, it also produces new tags for the result words of that operation. Prior work has shown this model can express a range of useful security policies, such as heap safety, stack safety, dynamic tainting, control-flow integrity, and information-flow control [5, 6, 17, 43].

To accelerate the behavior of the software security monitor, an implementation of the PIPE architecture will include a hardware tag rule cache of metadata rules. When a rule misses in the cache, it is evaluated on the PEX core and then inserted into the rule cache. In the future, when a rule hits in the cache, it can be validated without re-executing the monitor software or interpreting the tag bits. Details for rule cache design are provided in Reference [17]. The AP core can run ahead of rule resolution; a Write Queue holds unvalidated writes until the write instruction and its predecessors can be validated by the rule cache. With the rule cache, when the cache hit rate is high, the processor can run with little performance impact resulting from policy enforcement. To keep the hit rate high, policies should be designed with temporal locality in mind. For privilege separation compartmentalization policies, this property will be driven by the number of identifiers that are used for encoding protection domains on code and data objects, as well as their temporal locality characteristics. This interplay of policy design and architecture is explored in Section 7.

Last, we note that the privilege separation policies we derive could likely be ported to other tagged architectures such as Oracle ADI [37], LBA [10], Harmoni [14], or FADE [24] using approaches like those outlined in Reference [43]; however, SCALPEL uses the PIPE and its architectural characteristics for concrete evaluation.

2.2 The Protection-performance Tradeoff

While the PIPE can express memory safety policies [17, 43], fine-grained enforcement of all memory accesses can become expensive for some workloads. Compartmentalization policies represent an alternative design point that can very flexibly tune performance-protection tradeoffs through changing compartment sizes and intelligently drawing boundaries for high performance. With a small number of tags, one can separate out trusted from untrusted components such as ARM TrustZone [4] or OS from the application as in References [49, 50], but ultimately we are interested in exploring finer-grained separations. For example, one question we may pose is how tightly can we compartmentalize a software system with tag support while maintaining a certain rule cache hit rate, say, 99.9%.

Walking the line between protection and overhead costs is a well-known problem space. Dong et al. [18] observed that different decomposition strategies for web browser components produced

Fig. 1. PIPE architecture.

wildly different overhead costs, which they manually balanced against domain code size or prior bug rates. Mutable Protection Domains [38] proposes dynamically adjusting separation boundaries in response to overhead with a custom operating system and manually engineered boundaries. Several recent works have proposed more quantitative approaches to privilege separation. Program-Mandering [34] uses optimization techniques to find good separation strategies that balance reducing sensitive information flows with overhead costs, but requires manual identification of sensitive data, and ACES [11] similarly measures the average reduction in write exposure to global variables as a property of compartmentalizations. While these systems begin to automate portions of the compartment formation problem that SCALPEL builds upon, they all still rely on manual input. SCALPEL takes a policy derivation approach with a much stronger emphasis on automation: It uses analysis tools and performance experiments to explore the space of compartmentalizations, then automatically optimizes and lowers them to its hardware backend, a tag-extended RISC-V architecture, for enforcement.

2.3 Automatic Privilege Separation

The vast majority of compartmentalization work to date has been manual, demanding a security expert manually identify and refactor the code into separate compartments. This includes the aforementioned projects such as OpenSSH [9] and Dovecot [19], and even MicroKernel design [20, 31, 41] using standard OS process isolation, and runtime protection for embedded systems with metadata tags [49, 50]. Academic compartmentalization work has also relied on manual or semi-manual techniques for labeling and partitioning [9, 25, 28, 54].

In contrast, one goal for SCALPEL is *automation*; that is, to apply tag-based privilege-separation defenses to applications without expensive refactorings or manual-tagging; automated efforts relieve the labor-intensive costs of prior manual compartmentalization frameworks. Additionally, automation is important to ease the task of protecting existing software with the PIPE—SCALPEL decouples policy creation from enforcement by automatically lowering an engineer’s C-level compartmentalization strategies to the instruction-level enforcement provided by the PIPE.

ACES [11] is an automated compartmentalization tool for embedded systems and shares similarities with SCALPEL. It begins with a **program dependence graph (PDG)** representation of an application and a security policy (such as *Filename*, or one of several other choices), which indicates a potential set of security boundaries. It then lowers the enforcement of the policy to a target microcontroller device to meet performance and hardware constraints. The microcontroller it targets supports configurable permissions for eight regions of physical memory using a light-weight **Memory Protection Unit (MPU)**; protection domains in the desired policy are merged together until they can be enforced with these eight regions. Unlike ACES, SCALPEL targets a tagged architecture to explore many possible policies, some of which involve hundreds of protection domains, for fine-grained separation, far beyond what can be achieved with the handful of segments supported by conventional MPUs [33].

Towards Automatic Compartmentalization of C Programs on Capability Machines [52] is also similar to SCALPEL. In this work, the compiler translates each compilation unit of an application into a protection domain for enforcement with the CHERI capability machine [55]. This allows finer-grained separation than can be afforded with a handful of memory segments, but provides no flexibility in policy exploration to tune performance and security characteristics like SCALPEL does. To summarize, *SCALPEL is a complete tool for automatically compartmentalizing unmodified software for hardware acceleration, including automatically self-learning the required privileges, systematically exposing the range privilege-performance design-points through algorithmic exploration, and optimizing policies for good rule cache performance.* It complements and extends prior work along four axes: (1) quantitatively scoring the overprivilege in compartmentalization strategies, (2) providing complete automatic generation of compartments without manual input, (3) offering decomposition into much larger numbers of compartments (hundreds), and (4) automatically identifying the privilege-performance tradeoff curves for a wide-range of compartmentalization options.

3 THREAT MODEL AND ASSUMPTIONS

We assume a standard [11, 39] but powerful threat model for conventional C-based systems, in which an attacker may exploit bugs in either FreeRTOS or the application to gain read/write primitives on the system, which they may use to hijack control-flow, corrupt data, or leak sensitive information. Attackers supply inputs to the system, which, depending on the application, may include through a network connection or through files to be parsed or encoded. We assume both FreeRTOS and the application are compiled statically into a single program image with no separation before our compartmentalization; as such, a vulnerability in any component of the system may lead to full compromise. We assume that the test cases for the tracing mode are trusted, i.e., an attacker may not supply malicious tests.

The protection supplied by SCALPEL isolates memory read and write instructions to the limited subset of objects dictated by the policy and also limits program control-flow operations to valid entry points within domains as dictated by the policy. Additionally, SCALPEL is composed with a write-xor-execute [3] tag memory permissions policy that guarantees that instructions (executable code) are not writeable, and any writeable value is not executable. This means attackers cannot inject new executable code into the system. These constraints prevent bugs from reaching the full system state and limit the impacts of attacks to their contained compartments.

4 COMPARTMENTALIZATION TAG POLICY FORMULATION

In this section, we sketch our general policy model for compartmentalizing software using a tagged architecture. The goal of the compartmentalization policies is to decompose a system into separate logical protection domains, with runtime enforcement applied to each to ensure that memory accesses and control-flow transfers are permitted according to the valid operations granted to that domain. How do we enforce policies like these with a tagged architecture?

The PIPE provides hardware support for imposing a security monitor on the execution of each instruction. Whether or not each instruction is permitted to execute can depend on the tags on the relevant pieces of architectural state (Section 2.1). For example, we may have a stored private key that should only be accessible to select cryptographic modules. We can mark the private key object with a `private_key` tag and the instructions in the signing function with a `crypto_sign` tag. Then, when the signing function runs and the PIPE sees a load operation with instruction tag `crypto_sign` and data tag `private_key`, it can allow the operation. However, if a video processing function whose instructions are labeled `video_encode` tries to access the private key, the PIPE will

see a load operation with instruction tag `video_encode` and data tag `private_key` and disallow the invalid access.

In general, to enable compartmentalization policies, we place a *Domain-ID* label (an integer identifier) on each instruction in executable memory indicating the logical protection domain to which the instruction belongs; this enables rules to conditionally permit operations upon their tagged domain grouping, which serves as the foundation for dividing an application's code into isolated domains. Similarly, we tag each object with an *Object-ID* (also an integer identifier) to demarcate that object as a unique entity onto which privileges can be granted or revoked. For static objects, such as global variables and memory mapped devices, these object identifiers are simply placed onto the tags of the appropriate memory words at load time. Objects that are allocated dynamically (such as from the heap) require us to decide how we want to partition out and grant privileges to those objects. We choose to identify all dynamic objects that are allocated from a particular program point (e.g., a call to `malloc`) as a single object class, which we will refer to simply as an object. For example, all strings allocated from a particular `char * name = malloc(16)` call are the same object from SCALPEL's perspective; this formulation is particularly well-suited to the PIPE, because it enables rules in the rule cache to apply to all such dynamic instances. It also means that all dynamic objects allocated from the same allocation site must be treated the same way in terms of their privilege—dynamic objects could be differentiated further (such as by the calling context of the program point) to provide finer separation, but we leave such exploration to future work. As a result of these subject and object identification choices, the number of subjects and objects in a system is fixed at compile time.

Between pairs of subjects and objects (or in the case of a call or return, between two subjects), we would like to grant or deny *operations*. Accordingly, the tag on each instruction in executable memory also includes an *opgroup* field that indicates the operation type of that instruction. We define four *opgroups* and each instruction is tagged with exactly one *opgroup*: *read*, *write*, *call*, and *return*. For example, in the RISC-V ISA, the `sw`, `sh`, `sb`, and so on, instructions would compose the *write* *opgroup*.

When an instruction is executed, the security monitor determines if the operation is legal based upon (1) the *Domain-ID* of the executing instruction, (2) the type of operation $op \in \{read, write, call, return\}$ being executed, and (3) the *Object-ID* of the accessed word of memory (for loads and stores), or the *Domain-ID* of the target instruction (for calls and returns). As a result, the set of permitted operations can be expressed as a set of triples (*subject*, *operation*, *object*) with all other privileges revoked (default deny). In this way, the security monitor check can be viewed as a simple lookup into a privilege table or access-control matrix whose dimensions are set by the number of *Domain-IDs*, *Object-IDs*, and the four operation types. Such a check can be efficiently implemented in the security monitor software as single hash table lookup; once validated in software, a privilege of this form is represented as a single rule that is cached in the PIPE rule cache for hardware-accelerated, single-cycle privilege validation. Additionally, we define a fifth *unprivileged* *opgroup*, which is placed on instructions that do not represent privileges in our model (e.g., `add`); these instructions are always permitted to execute.

We define a *compartmentalization* as an assignment of each instruction to a *Domain-ID*, an assignment of each object to an *Object-ID*, and a corresponding set of permitted operation triples (*Domain-ID*, *op*, *Object-ID*). The SCALPEL backend takes a compartmentalization as an input and then automatically lowers it to a tag policy kernel suitable for enforcement with the PIPE. In this way, SCALPEL decouples policy construction from the underlying tag enforcement. The *opgroup* mapping is the same across all compartmentalizations.

In addition to these privilege checks, the SCALPEL backend also applies three additional defenses to support the enforcement of the compartmentalization. The first is a write-xor-execute

policy that prevents the execution of writeable data and the overwrite of executable code so an attacker cannot inject new executable code into the system. The second is that the words of memory inside individual heap objects that store allocator metadata (e.g., the size of the block) are tagged with a special *ALLOCATOR* tag. The allocator itself is placed in a special *ALLOCATOR* compartment that is granted the sole permission to access such words; as a result, sharing heap objects between domains permits only access to the data fields of those objects and not the inlined allocator metadata. Last, SCALPEL is built with the Clang compiler and uses LLVM's [23] static analysis to compute the set of instructions that are valid call and return entry points. These are tagged with special *CALL-TARGET* and *RETURN-TARGET* tags, and we apply additional clauses to the rules to validate that each taken control-flow transfer is both to a permitted domain and to a legal target instruction; this means that when a call or return privilege is granted, it is only granted for valid entry and return points.

An advantage of this policy design is that privilege enforcement is conducted entirely in the tag plane and software does not require refactoring to be protected with SCALPEL. Last, we note that there are multiple ways to encode compartmentalization policies on a tagged architecture. For example, the current compartment context could be stored on the program counter tag and updated during domain transitions, rather than from being inferred from the currently executing code. Some of these alternate formulations may work better with different concrete tagging architectures. However, for the PIPE, these formulations are largely equivalent to the above static formulation combined with localizing code into compartments (and making some decisions about object ownership), and we choose the static variant for a slight reduction in policy complexity; the choice is not particularly significant, and SCALPEL could produce policies for many such formulations.

5 THE TRACING POLICY

While a motivated developer or security engineer could manually construct a *compartmentalization* for a particular software artifact and provide it to the SCALPEL back-end, SCALPEL seeks to assist in policy derivation by providing a tracing mode (similar to e.g., AppArmor [35]) as well as a set of analysis tools for understanding the tradeoffs in different decomposition strategies. To this end, we build a compartmentalization *tracing policy*, which collects a lower-bound on the privileges exercised by a program as well as rule cache statistics we use later for policy optimization. While the PIPE architecture was designed for *enforcing* security policies, in this case, we repurpose the same mechanism for fine-grained, programmable dynamic analysis. SCALPEL's tracing policy provides several significant practical advantages over other approaches by (1) greatly simplifying tracing by running as a policy replacement on the same hardware and software, (2) directly using the PIPE for hardware-accelerated self-learning of low-level privileges (3), and making it possible to run in real environments and on unmodified software.

For the tracing policy, code and objects should be labeled at the finest granularity at which a security engineer may later want to decompose them into separate domains. On the code side, we find that function-level tracing provides a good balance of performance and precision, and so in this work SCALPEL tags each function with a unique *Domain-ID* during tracing. As a result, our SCALPEL implementation considers functions to be the smallest unit of code that can be assigned to a protection domain. Note that this is a design choice, and the PIPE could collect finer-grained (instruction-level) privileges at a higher cost to the tracing overhead.

On the object side, the tracing policy also assigns an *Object-ID* to each primitive object in the system. For software written in C, this includes a unique *Object-ID* for each global variable, a unique *Object-ID* for each memory-mapped device/peripheral in the memory map (e.g., Ethernet, UART), and a unique *Object-ID* associated with each allocation call site to an allocator, as discussed

vIPNetworkUpCalls:	prvIPTimerReload:	prvIPTimerStart:
Call prvIPTimerReload	Call prvIPTimerStart	Call vTaskSetTimeOutState
Call vDNSInitialise	Return vIPNetworkUpCalls	Return prvIPTimerReload
Return vDHCPPProcess	Return prvIPTask	Return prvIPTimerCheck
Write global_xNetworkUp	Return vIPReloadDHCPTimer	Return prvCheckNetworkTimers
	Write global_xARPTimer	Read global_xARPTimer
	Write global_xDHCPTimer	Read global_xDHCPTimer
	Write global_xTCPTimer	Read global_xTCPTimer
		Write global_xARPTimer
		Write global_xDHCPTimer
		Write global_xTCPTimer

Fig. 2. The set of privileged operations recorded by the tracing policy for three functions in the FreeRTOS TCP/IP networking stack. SCALPEL uses a tag-based hardware security monitor to automatically self-learn and then enforce fine-grained privileges like these at runtime.

in Section 4. All data memory words in a running system receive an *Object-ID* from one of these static or dynamic cases.

With these identifiers in place, the tracing policy is then used to record the observed dynamic behavior of the program. The PIPE invokes a software miss handler when it encounters a rule that is not in the rule cache. When configured as the tracing policy, the miss handler simply records the new privileges it encounters—expressed as interactions of *Domain-IDs*, operation types, and *Object-IDs*—as valid privileges that the program should be granted to perform; it then installs a rule so the program can use that privilege repeatedly without invoking the miss handler again. Unlike other policies, the tracing policy never returns a policy violation. We show an example output from the tracing policy in Figure 2 with identifiers printed as strings from their source program objects. From this tracing information, we form a fine-grained function and object graph, where each function and each object is a node and each privilege is an edge between nodes.

In addition to collecting privileges, the tracing policy also periodically records the rules that were encountered every N_{epoch} instructions, which we set to one million. As we will see in later sections (Section 7), this provides the SCALPEL analysis tools with valuable information about rule co-locality that it uses to construct low-overhead policies.

Dynamic Analysis Limitations: SCALPEL uses dynamic analysis to capture the observed low-level operations performed by a program. Observing dynamic behavior is important for SCALPEL to capture performance statistics to build performant policies (Section 7). However, this also means that our captured traces represent a lower bound of the true privileges that might be exercised by a program. There are a number of ways to handle this issue, and SCALPEL is agnostic to that choice. In cases where extensive test suites are available or can be constructed, one might use precise SCALPEL; that is, the traced program behavior serves as a ground truth for well-behaved programs and any violations produce fail-stop behavior; some simpler embedded systems applications may fit into this category. For higher usability on more complex software, one could either run SCALPEL in *alert mode* (described below), or SCALPEL could be combined with static analysis techniques for a hybrid policy design. In that case, the policy construction proceeds exactly as described in this article for capturing important performance effects, but the allowed interactions between *Domain-IDs* and *Object-IDs* would be relaxed to the allowed sets as found by static analysis. The best choice among these options will depend on security requirements, the quality and availability of test suites, and the tolerable failure rate of the protected application. We consider these issues orthogonal to SCALPEL’s primary contributions.

Runtime Modes: SCALPEL has two primary runtime modes: *alert mode* and *enforcement mode*. In *alert mode*, SCALPEL does not terminate a program if a policy violation is encountered;

instead, it produces a detailed log of the privilege violations that have been observed; this mode could provide near real-time data for intrusion detection and forensics in the spirit of Transparent Computing [13]. Alternatively, in `strict` mode, any policy violation produces fail-stop behavior.

6 PRIVILEGE QUANTIFICATION MODEL

In practice, one likely wants to deploy compartmentalizations that are coarser than the tracing policy granularity (i.e., individual functions and C-level objects) to reduce the number of tags, rules, and thus runtime costs associated with policy enforcement. Importantly, the tracing policy leads to a natural privilege quantification model we can use to compare these relaxed decompositions against the finest-grained function/object granularity. We can think of each rule in the tracing policy (*Domain-ID*, *op*, *Object-ID*) as a *privilege*, to which we can assign a *weight*. The *least privilege* of an application is the sum of the lower-bound privileges that it requires to run; without any of these, the program could not perform its task. For any coarser-grained compartmentalization, we can compute its privilege by counting up the number of fine-grained privileges it permits, which will include additional privileges beyond those in the least-privilege set. This enables us to compute the *overprivilege ratio (OR)*, which we define as the ratio of the privileges allowed by a particular compartmentalization compared to the least-privilege minimum; i.e., an OR of 2.0 means that twice as many privileges are permitted as strictly required. While crude, the OR provides a useful measure of how much privilege decomposition has been achieved, both to help understand where various compartmentalization strategies lie in the privilege-performance space and as an objective function for SCALPEL automatic policy derivation. For our weighting function, we choose to weight each object and function by its size in bytes; this helps account for composite data structures such as a struct that may have multiple fields and should count for additional privilege. Optionally, a developer can manually adjust the weights of functions or objects relative to other components in the system and interactively rerun the algorithms to easily tune the produced compartmentalizations.

7 POLICY EXPLORATION

To assist in creating and exploring compartment policies, SCALPEL provides three compartment generation algorithms. The first and simplest such approach, presented in Section 7.1, generates compartment boundaries based upon the static source code structure, such as taking each compilation unit or source code file as a compartment. The second algorithm, presented in Section 7.2, instead takes an algorithmic optimization approach that uses the tracing data to group together collections of mutually interacting functions. This algorithm is parameterized by a target domain size, allowing it to expose many design points, ranging from hundreds of small compartments to several large compartments. This is an architecture-independent approach that broadly has the property that larger compartments need fewer rules that will compete for space in the rule cache. Last, in Section 7.3, we present a second algorithmic approach that specifically targets producing efficient policies for the PIPE architecture; it targets packing policies into working sets of rules for improved cache characteristics. This algorithm uses both the tracing data and the cache co-locality data (Section 5) to produce optimized compartmentalization definitions and is the capstone algorithm proposed in SCALPEL.

7.1 Syntactic Compartments

A simple and commonly used approach [11, 52] for defining compartment boundaries is to mirror the static source code structure into corresponding security domains—we call these the *syntactic domains*. We define the OS syntactic domain by placing all of the FreeRTOS code into one compartment (*Domain-ID* 1) and all of the application code into a second compartment

(*Domain-ID* 2). This decomposition effectively implements a kernel/userspace separation for an embedded application that does not otherwise have one. Similarly, the *directory* syntactic domains are constructed by placing the code that originates from each directory of source code into a separate domain, i.e., *Domain-ID* i is assigned to the code generated from the i th directory of code. Programmers typically decompose large applications into separate, logically isolated-but-interacting modules, and the directory domains implement these boundaries for such systems. Last, the *file* and *function* syntactic domains are constructed by assigning a protection domain to each individual source code file or function that composes the program. Note that each syntactic domain is a strict decomposition from the one before it; for example, compilation units are a sub-decomposition of the OS/application boundary.

For the syntactic compartments, objects are labeled at the fine, individual object granularity (a fresh *Object-ID* for each global variable and heap allocation site); afterwards, all objects with identical permission classes based upon the tracing data are merged together. For example, if two global variables are accessed only by *Domain-ID* 1, then they can be joined into a single *Object-ID* with no loss in privilege; however, if one is shared and one is not, then they must be assigned separate *Object-IDs* to encode their differing security policies.

A second use we find for the syntactic code domains is applying syntactic constraints to other algorithms: for example, we can generate compartments algorithmically but disallow merging code across compilation units to maintain interfaces and preserve the semantic module separation introduced by the programmer. These results are presented in Section 9.4.

7.2 Domain-size Compartments

While the syntactic domains allow us to quickly translate source code structure into security domains, we are ultimately interested in exploring privilege-performance tradeoffs in a more systematic and data-driven manner than can be provided by the source code itself. We observe that the output of the tracing policy (Section 5) is a rich trove of information—a complete record of the code and object interactions including their dynamic runtime counts—on top of which, we can run optimization algorithms to produce compartments.

Because optimal clustering is known to be NP-Hard [2], we start with the fine-grained function and object graph and employ a straightforward greedy clustering algorithm that groups together sets of mutually interacting functions into domains while reducing unnecessary overprivilege. The algorithm is parameterized by C_{max} , the maximum number of bytes of code that are permitted per cluster. The algorithm works as follows: Upon initialization, each function is placed into a separate compartment C_i with size $C_{i\text{size}}$ taken to be the size (in bytes) of that function. At each step of the algorithm, two compartments C_A and C_B are chosen to merge together; the size of the resulting compartment is simply the sum of the sizes of the compartments being merged: $C_{AB\text{size}} = C_{A\text{size}} + C_{B\text{size}}$. To determine which two compartments to merge at each merge step, we compute the ratio of a utility function to that of a cost function for all pairs and select the pair with the highest ratio. The utility function we use is the number of cross-compartment calls and returns found by the tracing policy between those two compartments (i.e., their call affinity). The cost function we use is the increase in privilege (as given by Section 6) that would result from the merge: That is, we would like to identify mutually interacting functions with high affinity and group them together, while reducing unnecessary overprivilege. The algorithm terminates when no legal merges remain; that is, no candidates A and B maintain $C_{A\text{size}} + C_{B\text{size}} \leq C_{max}$.

After completion, each cluster C_i is translated into security domain *Domain-ID* i and objects are processed in the same manner as described in Section 7.1. We show results of the Domain-size algorithm on our HTTP web server application in Figure 3. First, Figure 3(a) shows how the number of compartments trends with the C_{max} parameter on the HTTP web server application. As can be

Fig. 3. Impact of clustering algorithms for the HTTP web server (Section 9.1) on 1,024-entry rule cache. Top Row: The Domain-size algorithm. Figure 3(a) shows how the number of compartments in final web server design trends with the C_{max} parameter. Figure 3(b) shows how many rules are required to represent the compartmentalization at a given C_{max} value. Figure 3(c) shows the runtime rule cache miss rate of the compartmentalization produced from that C_{max} value. Last, Figure 3(d) shows the overprivilege ratio of the design. Bottom Row: The Working-set algorithm. Similar to the top row, Figures 3(e), (f), (g), (h) show how the number of compartments, rules, the rule cache miss rate, and the overprivilege ratio all trend with WS_{max} parameter in the Working-set algorithm. Note how the Working-set algorithm outperforms the Domain-size algorithm in terms of the number of compartments, rules, and overprivilege ratio for the same rule cache miss rate.

seen, the Domain-size algorithm coupled with a target size parameter exposes a wide range of compartmentalization strategies, from just a few compartments up to hundreds. Figure 3(b) shows the total number of rules that are required to represent the compartmentalization policy at that granularity. With fewer compartments there are fewer unique subject domains, object domains, and interaction edges, which means fewer total rules are required for the design. Figure 3(c) shows the dynamic runtime rule cache miss rate of the compartmentalization produced from that C_{max} value. Last, Figure 3(d) shows how the Overprivilege Ratio trends with target domain size.

7.3 Working-set Compartments

The Domain-size compartment algorithm allows us to explore a wide range of compartmentalization strategies independent of the security architecture, but it is not particularly well-suited to the PIPE. The utility function that drives cluster merge operation is the number of dynamic calls and returns between those clusters. For enforcement mechanisms that impose a cost per domain transition (such as changing capabilities [53] or changing page tables between processes when using virtual memory process isolation), such a utility function would be a reasonable choice, as it does lead to minimizing the number of cross-compartment interactions. Grouping together code and data in this way does reduce the number of tags, rules, and thus cache characteristics of enforcing the compartmentalization on the PIPE, but there is only a broad correlation (Figure 3(c)).

Fig. 4. An example of how the rule savings is calculated for merging the S_1 and S_2 domains together. In this example, there are five rules (privilege triples) in Working Set 1 before the merge and three rules afterwards, for a total of two rules saved. However, S_2 did not have write access to O_1 before the merge, so overprivilege is also introduced by the merge. Assuming all components of the system have a uniform weight of one, then the utility for this merge would be two (two rules saved) and the cost would be one (one additional privilege exposed), for a ratio of $2/1 = 2$. Note that following the depicted subject merge, merging objects O_1 and O_2 would be chosen next by the algorithm, as it would save an additional rule at no further increase in privilege; in this way, the Working-set algorithm simultaneously constructs both subject and object domains.

For the PIPE, there is no cost to change domains, provided the required rules are already cached; instead, what matters is *rule locality*. As a result, to produce performant policies for the PIPE, we instead would like to optimize the runtime rule cache characteristics rather than minimizing the number of domain transitions. To this end, we construct an algorithm based on reducing the set of rules required by each of a program's phases so each set will fit comfortably into the rule cache for favorable cache characteristics.

How do we identify program phases such that we can consider their cache characteristics? Recall that the tracing policy records the rules that it encounters during each epoch of 1M instructions (Section 5). We consider the set of rules encountered during each epoch to compose a *working set*. As an intuitive, first-order analysis, if we can keep the rules in each working set below the cache size and the product of those rules and the miss handling time small compared to the epoch length, then the overhead for misses in the epoch will be small. As we will see, since not all rules are used with high frequency in an epoch, it is not strictly necessary to reduce the rules in the epoch below the cache size. While there is prior work on program phase detection [45, 48], SCALPEL takes a simple epoch-based approach that we see is adequate to generate useful working sets; integrating more sophisticated phase detection into SCALPEL would be interesting future work and would only improve the benefits of the PIPE protection policy.

The Working-set algorithm targets a maximum number of rules allowed per working set, WS_{max} . We construct the Working-set algorithm in a similar fashion to the Domain-size algorithm (Section 7.2) starting with the fine-grained function and object graph, except that we consider clustering of both subjects and objects simultaneously under a unified cost function. The algorithm works as follows: Upon initialization, each function is placed into a subject domain S_i and each primitive object is placed into a separate object domain O_i . We then initialize the rules in each working set to those found by the tracing policy during that epoch. At each step of the algorithm, either a pair of subjects or a pair of objects are chosen for merging together. The pair that is chosen is the pair with the highest ratio of a utility function to that of a cost function across all pairs. In contrast to the Domain-size algorithm, the utility function we use is the sum of the rules that would be saved across all working sets that are currently over the target rule limit WS_{max} . We show an example of how the rule delta calculation is performed in Figure 4. After performing a merge, the new, smaller set of rules that would be required for each affected working set is

Table 1. Architectural Modeling Parameters

Model	Cost (cycles)
Cyc_{L1} (64 KB, 4-way)	1
Cyc_{L2} (512 KB, 8-way)	3
Cyc_{DRAM} (2 GB)	100
Cyc_{policy_eval}	300
Cyc_{PIPE} (DMHC [16], 1,024)	1

calculated, and then the process repeats. The Working-set algorithm uses the same cost function as the Domain-size algorithm, i.e., the increase in privilege that would result from combining the two subjects or objects into a single domain. As a result, the Working-set algorithm attempts to reduce the number of rules required by the program during each of its program phases down to a cache-friendly number while minimizing the overprivilege. The algorithm is run until the number of rules in all the working sets is driven below the target value of WS_{max} .

Like the Domain-size algorithm, we can vary the value of WS_{max} to produce a range of compartmentalizations at various privilege-performance tradeoffs. If we set our WS_{max} target to match the actual rule cache size, then we will pack the policy down to fit comfortably in the cache and produce a highly performant policy; however, we find that this tight restriction is not strictly necessary—Figure 3(g) shows how the rule cache miss rate trends with the target WS_{max} value, achieving an almost linear reduction in miss rate with WS_{max} .

The core advantage of the Working-set algorithm is that it is able to coarsen a compartmentalization in only the key places where it actually improves the runtime cache characteristics of the application, while maintaining the majority of fine-grained rules that do not actually contribute meaningfully to the rule cache pressure. In Figures 3(e), (f), (g), (h), we show how the number of compartments, the number of rules, the rule cache miss rate, and the overprivilege ratio trend with the WS_{max} parameter. In contrast to the Domain-size algorithm, we can see that many more compartments and rules are maintained as the algorithms are driven to smaller and smaller rule cache miss rates, demonstrating the advantages of the Working-set algorithm in intelligently producing compartmentalization policies at much lower levels of overprivilege. For example, at the WS_{max} of 1,024 the Working-set algorithm achieves the same rule cache miss rate as the Domain-size algorithm does at a C_{max} of 16,364, but it has more than twice as many total rules and an Overprivilege Ratio that is twice as small, a much more privilege-restricted design for the same overhead. We illustrate the differences between the algorithms more directly in Section 9 (Evaluation).

8 PERFORMANCE MODEL

Our SCALPEL evaluation targets a single-core, in-order RISC-V CPU that is extended with the PIPE tag-based hardware security monitor. To match a typical, lightweight embedded processor [32], we assume 64 KB L1 data and instruction caches and a unified 512 KB L2 cache. To this, we add a 1,024 entry DMHC [16] PIPE rule cache. The application is a single, statically linked program image that includes both the FreeRTOS operating system as well as the program code. The image is run on a modified QEMU that simulates data and rule caches inline with the program execution to collect event statistics. SCALPEL is built on top of an open-source PIPE framework that includes tooling for creating and running tag policies [30]. The architectural modeling parameters we use are given in Table 1. We use the following model for baseline execution time:

$$T_{baseline} = N_{inst} + N_{L1miss} \times Cyc_{L2} + N_{L1Dmiss} \times Cyc_{L2} + N_{L2miss} \times Cyc_{DRAM}.$$

Beyond the baseline, SCALPEL policies add overhead time to process misses:

$$T_{SCALPEL} = T_{baseline} + N_{PIPE_{miss}} \times C_{yc_{policy_eval}}.$$

We take $C_{yc_{policy_eval}}$ to be 300 cycles based on calibration measurements from our hash lookup implementation.

Last, we calculate overhead as:

$$Overhead = \frac{T_{SCALPEL} - T_{baseline}}{T_{baseline}} \times 100\%. \quad (1)$$

9 EVALUATION AND RESULTS

In this section, we present the results of our SCALPEL evaluation. Section 9.1 details the applications we use to conduct our experiments. Section 9.2 shows statistics about the applications and the results of the tracing policy. Section 9.3 shows the privilege-performance results of SCALPEL's Domain-size and Working-set algorithms. Section 9.4 shows the Syntactic Domains and the results of applying the syntactic constraints to the Working-set algorithm. Last, Section 9.5 shows how SCALPEL's Working-set rule clustering technique can be used to pack entire policies for real-time systems.

9.1 Applications

Web Server: One application we use to demonstrate SCALPEL is an HTTP web server built around the FreeRTOS + FAT + TCP demo application [47]. Web servers are common portals for interacting with embedded/IoT devices, such as viewing baby monitors or configuring routers. Our final system includes a TCP networking stack, a FAT file system, an HTTP web server implementation, and a set of CGI programs that compose a simple hospital management system. The management system allows users to login as patients or doctors, view their dashboard, update their user information, and perform various operations such as searches, prescribing medications, and checking prescription statuses. All parts of the system are written in C and are compiled together into a single program image. To drive the web server in our experiments, we use curl [12] to generate web requests. The driver program logs in as a privileged or unprivileged user, performs random actions available from the system as described above, and then logs out. For the tracing policy, we run the web server for 500 web requests with a 0.25 s delay between requests, which we observe is sufficient to cover the web server's behavior. For performance evaluation, we run five trials of 100 requests each and take the average.

libXML Parsing Library: Additionally, we port the libXML2 [40] parsing library to our FreeRTOS stack. To drive the library, we construct a simple wrapper around the xmlTextReader SAX interface, which parses plain XML files into internal XML data structures. For our evaluation experiments, we run it on the MONDIAL XML database [22], which contains aggregated demographic and geographic information. It is 1 MB in size and contains 22K elements and 47K attributes. Parsing structured data is common in many applications and is also known to be error-prone and a common source of security vulnerabilities: libXML2 has had 65 CVEs including 36 memory errors between 2003 and 2018 [15]. Our libXML2 is based on version 2.6.30. Timing-dependent context switches causes nondeterministic behavior; we run the workload five times and take the average.

H264 Video Encoder, bzip2, GNU Go: Additionally, we port three applications from the SPEC benchmarks that have minimal POSIX dependencies (e.g., processes or filesystems) to our FreeRTOS stack. Porting the benchmarks involved translating the main function to a FreeRTOS task, removing their reliance on configuration files by hardcoding their execution settings, and integrating them with the FreeRTOS dynamic memory allocator. The H264 Video Encoder is

Table 2. Application Statistics and Results of the Tracing Policy

Application	Lines of Code	Live Functions	Live Objects	Total Rules	Monolithic OR
bzip2	8K	128	109	2,880	39
Web Server	49K	1,231	218	12,025	96
H264	53K	363	692	19,641	244
GNU Go	198K	3,288	10,532	30,077	187
libXML	290K	260	384	10,221	538

Fig. 5. The impact of the WS_{max} parameter on the rule cache miss rate for a 1,024-entry rule cache. The max value corresponds to the tracing-level granularity (Table 2), and the solid lines show how the rule cache miss rate trends with WS_{max} . As can be seen, the SCALPEL algorithms allow a designer to generate compartmentalization designs that target any desired rule cache miss rate. The dashed lines show the even lower rule cache miss rate that is achieved by prefetching rules, which we describe in Section 10.

based on 464.h264ref, the bzip2 compression workload is based on 401.bzip2, and the GNU Go implementation is based on 445.gobmk. Video encoders are typical for any systems with cameras (baby monitors, smart doorbells), compression and decompression are common for data transmission, and search implementations may be found in simple games or navigation devices. We run the H264 encoder on the reference SSS.yuv, a video with 171 frames with a resolution of 512×320 pixels. We run bzip2 on the reference HTML input and the reference blocking factors. We run GNU Go in benchmarking mode, where it plays both black and white, on an 11×11 board with four random initial stones. Timing-dependent context switches causes nondeterministic behavior; we run each workload five times and take the average.

9.2 Application Statistics, the Tracing Policy, and Rule Cache Miss Rates

In Table 2, we show the application statistics of our workloads and the results of the tracing policy. First, to give a broad sense for the application sizes, we show the total lines of code; this column includes only the application, on top of which there is an additional 12K lines of core FreeRTOS code. Next, we show the total number of live functions and objects logged by the tracing policy during the program's execution. These subjects and objects compose the fine-grained privileges that SCALPEL enforces. In the Total Rules column, we show the total number of unique rules generated during the entire execution of the program under the tracing policy granularity (Section 5). While this number indicates the complexity of the program's data and control graph, it is not necessarily predictive of the cache miss rate, which depends on the dynamic rule locality. We show the rule cache miss rate in Figure 5: The rightmost point (max) corresponds to the miss rate at the tracing policy granularity. As can be seen, the web server has fewer rules than libXML, but also has a lower cache hit rate due to the larger amount of logic that runs at its steady-state web serving phases (such as receiving network requests, parsing them, running CGI programs, and sending output). In

Fig. 6. The range of compartmentalizations produced from SCALPEL’s algorithms on a 1,024-entry rule cache. Each point corresponds to a single specific concrete compartmentalization that is run for performance evaluation and characterized by its runtime overhead (Y axis) and aggregate Overprivilege Ratio (X axis). The Domain-size line shows compartments that are generated from the various values of C_{max} to the Domain-size algorithm, and the Working-set line shows compartments that are generated from the various values of WS_{max} to the Working-set algorithm. As can be seen, SCALPEL allows a security engineer to rapidly create and evaluate many compartmentalization strategies to explore design tradeoffs without the excessive labor required for manual separations. Note that the Working-set algorithm dominates the Domain-size algorithm, with a particularly strong advantage at the low-end of the overhead spectrum.

contrast, H264 has more total rules, but exhibits more locality as it spends long program phases on a small subset of code and data (e.g., doing motion estimation), a much more rule-cache-friendly workload. Very simple workloads, such as bzip2, require only a couple thousand rules and have effectively no rule cache misses even at the tracing-level granularity.

9.3 Privilege-performance Tradeoffs

A key question we would like to answer is how we can trade off privilege for performance on a per-application basis using the range of SCALPEL compartment generation algorithms (Section 7). First, Figure 5 shows how the rule cache miss rate trends with the WS_{max} parameter to the Working-set algorithm: As can be seen, it allows a designer to target any desired rule cache miss rate for an application.

To explore these compartmentalization options, in Figure 6, we show the privilege-performance curves (where each compartmentalization design is scored by its overhead and Overprivilege Ratio) generated from both the Domain-size algorithm and the Working-set algorithms. The top-left point in the Domain-size algorithm corresponds to the tracing-level granularity; this point enforces the full, fine-grained access control matrix, but also imposes large overheads; for example, on the web server application, the cost of enforcement is $>100\%$. The other points in this line correspond to larger values of the C_{max} parameter, which produces fewer, larger compartments for more favorable runtime overheads; however, as can be seen, these coarser compartmentalizations also introduce additional overprivilege. The Working-set lines in the plots correspond to the range of compartmentalizations produced from the Working-set algorithm and its WS_{max} parameter. The top-left point corresponds to the maximum value of WS_{max} where no clustering is performed, and the bottom-right point corresponds to packing the rules in each working set to the rule cache size (1,024), producing designs that have very favorable performance characteristics but more overprivilege.

Note that in both cases the curves have a very steep downward slope, meaning large improvements in runtime performance can be attained with little increases in privilege; the curves eventually flatten out, at which point additional decreases in overhead come at the expense of larger

(a) The privilege-performance points generated by the Syntactic Domains (Sec. 7.1). Unlike the Domain-Size and Working-Set algorithms, the Syntactic Domains only provide a handful of choices, none of which are competitive with the alternate algorithms.

(b) Impact of syntactic constraints. The algorithmic compartment algorithms (Sec. 7.2, 7.3) can optionally take an additional constraint, such as a syntactic boundary, that must be respected while creating the compartmentalization. In this case, we show the application of three syntactic domains to the Working-Set algorithm.

Fig. 7. Privilege-performance impact of syntactic domains and constraints on the HTTP web server running on a 1,024-entry rule cache.

amounts of overprivilege. Note that the Working-set compartments strictly dominate the Domain-size compartments, producing more privilege reduction at lower costs than the Domain-size counterparts. As can be seen, SCALPEL allows designers to easily explore the tradeoffs in compartmentalization tag policies. These runs represent the default, fully automatic toolflow. A designer can then easily inspect the produced compartmentalization files, tune the privilege weights, and rerun the tools interactively as time and expertise allow.

9.4 Syntactic Compartments and Syntactic Constraints

In Figure 7(a), we show the syntactic compartments (Section 7.1) on the HTTP web server application. Unlike the Domain-size and Working-set algorithms, which are parameterized to produce a wide range of compartmentalization options, the syntactic compartments only provide a handful of decomposition choices. And, as can be seen, none of the options are competitive compared to the Domain-size or Working-set decompositions, which suggests that it is indeed useful to approach the compartmentalization problem with more sophisticated techniques.

However, it is also true that software engineers often decompose their own projects into modules, and those modules' boundaries bear semantic information about code interfaces and relationships. For example, the web server application has the core FreeRTOS code in one directory, the TCP/IP networking stack in another directory, the web server application (CGI app) in another directory, and the FAT filesystem implementation in another separate directory. When the algorithmic compartment generation algorithms (Sections 7.2, 7.3) optimize for privilege-performance design points, they have the full freedom to reconstruct boundaries in whatever way they find produces better privilege-performance tradeoffs. However, if we would like to preserve the original syntactic boundaries during the algorithmic optimization process, we can add additional constraints, such as a *syntactic constraint*, which limits the set of legal merges allowed by the algorithms. For example, under the *file* syntactic constraint, two global variables can only be merged if they originate from the same source file. This allows SCALPEL to optimize privilege separation while respecting the interfaces to that module. We note that a compartmentalization that is a strict sub-decomposition of another compartmentalization is never less secure.

In Figure 7(b), we show the application of the syntactic domains as constraints to the Working-set algorithm. The OS restriction adds little additional overhead to the produced design points but

Table 3. The Overprivilege Ratio (OR) of the Applications When They Are Packed for Real-time Performance to the Given Total Rule Count

Application	Real-time Rule Target				
	512	1,024	2,048	4,096	8,192
bzip2	2.82	1.34	1.01	1.00	1.000
Web Server	8.92	5.76	2.71	1.35	1.005
H264	11.9	2.81	1.46	1.05	1.002
libXML	12.3	7.46	2.61	1.18	1.000
Gnu Go	29.4	12.7	3.47	1.43	1.033

When packed in this way, they can be (1) loaded in constant time and (2) experience no additional runtime rule resolutions, making them suitable for real-time systems.

guarantees a cleaner separation of the OS and application than may be found by the algorithms naturally. However, the *file* constraint is very restrictive, reducing the number of moves available to the algorithms to such a large extent that many of the WS_{max} targets fail to generate.¹ These examples illustrate the benefits of the rapid exploration enabled by SCALPEL, and we note that a manually constructed constraint can be a very convenient method for interacting with SCALPEL's automation.

9.5 Packing Policies for Real-time Systems

The ideas presented in the Working-set algorithm (Section 7.3) can be used to pack an entire security policy (i.e., the complete set of rules that compose the policy) into a single, fixed-size set of rules. For this construction, we simply take the union of all rules required to represent the policy and present it to the Working-set algorithm as a single working set—the entire policy will then be packed down to a number of rules equal to WS_{max} . Importantly, this means that the policy can be loaded in a constant amount of time, and assuming the WS_{max} matches the rule cache size, then no additional runtime rule resolutions will occur, giving the system predictable runtime characteristics suitable for real-time systems. We show the results of this technique in Table 3 when applied to a range of rule targets. The overprivilege points generated from this technique could be used to decide on a particular rule cache size for a specific embedded application to achieve target protection and performance characteristics. Note that the working-set cached case achieves lower OR at the same 1,024-entry rule capacity, since it only needs to load one working set at a time. It will take a larger rule cache to achieve comparably low OR. However, it is worth noting that, if the rule memory does not need to act as a cache, then it can be constructed more cheaply than a dynamically managed cache, meaning the actual area cost is lower than the ratio of rules, and might even favor the fixed-size rule memory. Furthermore, if one is targeting a particular application, then the tag bits can also be reduced to match the final compartment and object count (e.g., can be 8b instead of a full word width), which will further decrease the per rule area cost.

10 PREFETCHING

Finally, we consider one last performance optimization to reduce the overhead costs of SCALPEL's policies: rule prefetching. During the normal operation of the PIPE, rules are evaluated and

¹Note that the Working-set algorithm can only perform merge operations between code or object domains that are found to co-reside in the same epoch, which means that they may not reach the same minimum overhead as the corresponding syntactic domain.

```

short GetCRC16(char *pbyData, int stLength){
    short bTableValue;
    short wCRC = 0;

    while(stLength-- != 0){
        bTableValue = ((wCRC & 0x00FF) ^ *pbyData++);
        wCRC = (((crc16_table_high[bTableValue]) << 8) +
                (crc16_table_low[bTableValue] ^ ((wCRC >> 8)
                & 0x00FF)));
    }
    return wCRC;
}

```

(a) The GetCRC16 function in the FreeRTOS FAT Filesystem.

(b) The Rule-Successor Graph for the GetCRC16 function. Each rule (privilege triple) is shown as a rectangle with three fields corresponding to the subject, operation and object tags.

Fig. 8. The Rule-successor Graph, a data structure used by SCALPEL's prefetching system.

installed into the rule cache only when an application misses on that rule. When such a miss occurs, the PEX core awakens, evaluates the rule, and finally installs it into the rule cache. Much like prefetching instructions or data blocks for processor caches [36], there is an opportunity for the PEX core to preemptively predict and install rules into the cache. Such a technique can greatly reduce the number of runtime misses that occur, provided that the required rules can reliably be predicted and prefetched before they are needed. In this section, we explore the design and results of a rule prefetching system.

10.1 The Rule-successor Graph

The core data structure of our prefetching design is the *Rule-successor Graph*. The Rule-successor Graph is a directed, weighted graph that records the immediate temporal relationships of rule evaluations. A rule is represented as a node in the graph, and a weighted edge between two nodes indicates the relative frequency of the miss handler evaluating the source rule followed immediately by evaluating the destination rule.

Figure 8(a) shows an example function from the FreeRTOS FAT filesystem, and Figure 8(b) shows the Rule-successor Graph for its function-entry rule. When this function is called, it issues loads and stores to the task's stack, and then it issues loads to the `crc16_table_high` and `crc16_table_low` global variables in exactly that order; this deterministic sequence is learned and encoded in the Rule-successor Graph. Many kinds of rule relationships are highly predictable, such as rules that are required for sequential instructions in the same basic block.

However, data or control-flow dependent program behavior can produce less predictable rule sequences—for example, a return instruction can have many low-weighted rule successors if that function is called from many locations within a program. In this example, `GetCRC16` has two callers and may return to either, although one is much more common than the other; similarly, `GetCRC16` also accepts a data pointer `pbyData` that could produce data-dependent rule sequences depending on the object it points to, although in this program it always points to the task's stack, which does not require another rule. Last, if `stLength` were 0, then the program would take an alternate control-flow path and several of the rules would be skipped. Like other architectural optimizations such as caches and branch predictors [21], optimistic prefetching accelerates

common-case behavior, but may have a negative impact on performance when the prediction is wrong.

A program's Rule-successor Graph can be generated from the miss handler software with no other changes to the system. To do so, the miss handler software simply maintains an account of the last rule that it handled. When a new miss occurs, the miss handler software updates the Rule-successor Graph by updating the weight from the last rule to the current rule (and adding any missing nodes). Finally, the record of which rule was the last rule is updated to the current rule, and the process continues.

10.2 Generating Prefetching Policies

A *prefetching policy* is a mapping from each individual rule (called the *source rule*) to a list of rules (the *prefetch rules*) that are to be prefetched by the miss handler when a miss occurs on that source rule. Prefetching policies are generated offline using a program's Rule-successor Graph; the goal is to determine which rules (if any) should be prefetched from each source rule on future runs of that program.

To find good candidate prefetch rules for each source rule, we deploy a Breadth-first Search algorithm on the Rule-successor Graph to discover high-likelihood, subsequent rules. Each such search begins on a source rule with an initial probability $p = 1.0$. When a new node (rule) is explored by the search algorithm, its relative probability is calculated by multiplying the current probability by the weight of the edge taken. When a new, unexplored rule is discovered, it is added to a table of explored nodes, and its depth and probability are recorded with it. If a rule is already in the table when it is explored from a different path, then the running probability is added to the value in the table to reflect the sum of the probabilities of the various paths on which the rule may be found.

The algorithm terminates searching on any path in which the probability falls below a minimum threshold value. We set this value to 0.1%, which we observe sufficiently captures the important rules across our benchmarks. After search is complete, the table of explored nodes is populated and ready to be used for deriving prefetching policies. To test the impact of various degrees of prefetching, we add a pruning pass in which any rules below a target probability p_{min} are discarded from the table. For example, if p_{min} is set to the maximum of 1.0, then rules are only included in the prefetching set if they are always observed to occur in the Rule-successor Graph following the source rule. Lower p_{min} values cause more speculative rules to be prefetched; in the best case they may avert more misses, but in the worst case they may pollute the rule cache by evicting a more important rule. In Figure 8(b), the bottom left rule has a probability of 0.78. If no rules remain after pruning, then no prefetch rules are found. Otherwise, the remaining rules are sorted to compose the final list of prefetch rules. They are sorted by depth (smallest first), then within the same depth by probability (highest first) to order the rules in a sequence most likely to be useful. We vary the values of p_{min} from 1.0 to 0.25 to explore the impact of various levels of prefetching on final performance.

10.3 Prefetching Cost Model

When the PIPE misses on a rule, it traps and alerts the PEX core for rule evaluation. In SCALPEL, a rule evaluation is a hash table lookup that checks the current operation against a privilege table (Section 4). When prefetching is enabled, we choose to store the prefetch rules in the privilege hash table along with the source rule to which they belong. When a miss occurs, the miss handler performs the initial privilege lookup on the source rule and installs it into the PIPE cache, allowing the AP core to resume processing. Afterwards, the PEX core continues to asynchronously load and install the rest of the prefetch rules in that hash table entry. Assuming a cache line size of 64 B

Fig. 9. The results of the SCALPEL's rule prefetching system.

and a rule size of 28 B (five 4 B input tags and two 4 B output tags), then two rules fit in a single cache line. As such, the first prefetch rule can be prepared for insertion immediately following the resolution of the source rule. We assume a 10 cycle install time into the PIPE cache for each rule installation. For each subsequent cache line (which can hold up to two rules), we add an additional cost of 20 cycles for a DRAM CAS operation, in addition to the 10 cycle insertion time for each rule. We set the maximum number of prefetch rules to seven so all eight rules (including the source rule) may fit onto a single same DRAM page, assuming a 2,048 B page size. Our data shows that this number is sufficient to capture a majority of the benefits of prefetching on our workloads.

We begin by looking at the structure of the Rule-successor Graph Figure 9(a) from our various applications to get a sense for the number and likelihood of prefetch rules per source rule that might be prefetched; the more high-likelihood rules there are, the more benefits we expect to see from prefetching. In this graph, the X axis shows the p_{min} cutoff probability in the range of [0.25,1], and the Y axis shows the average number of rules per source rule that have at least the given cutoff probability. The data shows that there are around five rules per source rule that can be prefetched even at the maximum p_{min} value of 1.0 (i.e., they always follow the source rule during tracing) in our benchmarks; H264 and bzip2 are more predictable than the other three benchmarks with values closer to 10. At lower values of p_{min} , more rules make the cutoff, although the slope is low (less than one rule on average per 10% decrease in likelihood) meaning there are significantly diminishing returns on prefetching larger numbers of rules.

Next, to see the results of prefetching on rule cache miss rate, we show the prefetching cases as dashed lines in Figure 5. To see the final impact on program overhead, Figure 9(b) shows the resulting privilege-performance curves generated from the various prefetching policies. The red line shows the baseline (no prefetching) case, and the other lines show the prefetching policies generated from the various values of p_{min} . All of the prefetching cases strictly dominate the baseline case on privilege and performance; the yellow line ($p_{min} = 1.0$) captures a majority of the benefits, but each additional relaxation to p_{min} continues to lower enforcement costs for the same privilege reduction level in diminishing amounts. On the high end of the overhead spectrum (i.e., the tracing-level granularity), the prefetching system reduces the overhead from an average of 105% to only 27%, a 3.9 $\times$ reduction in overhead costs. This shows the predictable nature of the PIPE rules in a compartmentalization policy and the large benefit of prefetching rules to reduce costs. On the lower end of the overhead spectrum, the benefits of prefetching are less pronounced but

```

static char first_name[USER_NAME_LENGTH] = { 0 };
static char last_name[USER_NAME_LENGTH] = { 0 };
static char type[2] = { 0 };
static char address[USER_ADDRESS_LENGTH] = { 0 };
static char condition[MEDICAL_NAME_LENGTH] = { 0 };

void ParseSearchCgiString(char *CgiArgs){
    CgiArgValue(first_name, USER_NAME_LENGTH, "Firstname", CgiArgs);
    CgiArgValue(last_name, USER_NAME_LENGTH, "Lastname", CgiArgs);
    CgiArgValue(address, USER_ADDRESS_LENGTH, "address", CgiArgs);
    CgiArgValue(type, USER_TYPE_LENGTH, "type", CgiArgs);
    // Incorrect length, overflow by 64 bytes:
    CgiArgValue(condition, USER_ADDRESS_LENGTH, "condition", CgiArgs);
}

```

(a) A buffer overflow vulnerability in the Web Server's search functionality. The vulnerable CGI program can be reached and triggered by any user who browses to the search.html page. This vulnerability allows an attacker to corrupt the user_auth and session_table variables.

Symbol	Address	Size	Source
first_name	0x80500cec	16	src/search-results.c:53
last_name	0x80500d0c	16	src/search-results.c:54
address	0x80500d2c	128	src/search-results.c:56
condition	0x80500dac	64	src/search-results.c:57
user_auth	0x80500dec	4	src/auth.c:39
session_table	0x80500df0	32	src/auth.c:35

(b) The symbol table showing the addresses and sizes of several of the symbols in the Web Server's data section.

Fig. 10. An exploitable vulnerability in the Web Server application.

do enable the system to achieve even finer privilege separation at the same costs: At an overhead of 10%, the prefetching cases allow for 20% more rules and an OR that is 12% smaller.

11 SECURITY, OVERPRIVILEGE, AND WORK-FACTOR

Vulnerabilities such as memory safety errors permit a program to perform behaviors that violate its original language-level abstractions, e.g., they allow a program to perform an access to a memory location that is either outside the object from which the pointer is derived or has since been freed and is therefore temporally expired. An exploit developer has the task of using such a vulnerability to corrupt the state of the machine and to redirect the operations of the new, emergent program such as to reach new states that violate underlying security boundaries or assumptions, such as changing an authorization level, leaking private data, or performing system calls with attacker-controlled inputs. In practice, bugs come in a wide range of expressive power, and even memory corruption vulnerabilities are often constrained in one or more dimensions, e.g., a typical contiguous overflow error may only write past the end of an existing buffer [42], or an off-by-one-error allows an attacker to write a pointer value past the end of an array but gives the attacker no control of the written data [8]. Modern exploits are typically built from *exploit chains*, in which a series of bugs are assembled together to achieve arbitrary code execution [8], and complex exploits take many man-months of effort to engineer [7] even in the monolithic environments in which they run.

The privilege separation defenses imposed by SCALPEL limit the reachability of memory accesses and control-flow instructions to a small subset of the full machine's state. These restrictions affect the attacker's calculus in two ways: First, they may lower the impact of bugs sufficiently to disarm them entirely, i.e., rendering them unable to impart meaningful divergence from the original program. Second, they may vastly increase the number of exploitation steps and bugs required to reach a particular target from a given vulnerability: An attacker must now perform repeated confused deputy attacks [26] at each stage to incrementally reach a target; when privilege reduction is high, these available operations become substantially limited, thus driving up attacker costs and defeating attack paths for which no exploit can be generated due to the imposed privilege separation limitations.

We illustrate these ideas with a vulnerability example from the Web Server application in Figure 10. A buffer overflow in the search CGI code is reachable from the server's web interface and can cause the program to write beyond the end of the condition buffer onto objects located at higher addresses, which are the user_auth and session_table variables. Corrupting user_auth

Table 4. The Relationship between WS_{max} , the Overprivilege Ratio, and the Exploitability of a Vulnerability in the Web Server

WS_{max}	max	3,200	2,800	2,400	2,000	1,800	1,600	1,400	1,200	1,024	512
Overprivilege Ratio (OR)	1.00	1.02	1.04	1.08	1.17	1.26	1.36	1.53	1.92	2.43	8.42
Protect user_auth	✓	✓	✓	✓	✓	✓	✓	✓	X	X	X
Protect session_table	✓	✓	✓	✓	✓	✓	✓	✓	✓	X	X
HashTableEqual call targets	1	1	1	1	7	7	11	12	17	67	241

We validate the ✓ and X vulnerability assignments by running the Web Server with the policy shown and triggering the buffer overflow vulnerability. Last, we show the number of call targets reachable from the HashTableEqual function under the given compartmentalization if its code pointer is corrupted.

can allow an unprivileged user to escalate their privileges. However, the fault is entirely contained if user_auth is tagged with an *Object-ID* for which CgiArgValue does not have write permission, because any out-of-bounds write will incur a policy violation. In Table 4, we show the range of compartmentalizations generated from the Working-set algorithm. Row 1 shows the compartmentalization's Overprivilege Ratio, and row 2 shows whether the user_auth overwrite is prevented (which we verify against our policy implementation by triggering the buffer overflow to classify as ✓ or X in the table). If that write is not prevented, then an attacker can (1) escalate their privileges, and (2) there is also a possibility to corrupt the subsequent session_table as well if *that object* is also writable from CgiArgValue. The session_table is a structure that contains a hash table root node, which includes a code pointer session_table->compare. Like the user_auth object, this object is protected if the CgiArgValue code does not have permission to write to it. We show this relationship in row 3. If it can be corrupted, then it could provide additional footing to compromise the contained compartment, such as through hijacking the program's control flow by overwriting the session_table->compare field.

While we have illustrated that these specific vulnerabilities are eliminated at specific higher compartmentalization levels and lower ORs, we expect this trend to hold for other vulnerabilities—as OR lowers, at some point each specific vulnerability based on a privilege violation is eliminated. Each vulnerability may, in general, be eliminated at a different OR. Consequently, we expect lower OR to generally correlate well with lower vulnerability. Last, in row 4, we show the total number of legal call targets that are permitted by the domain containing HashTableEqual (the only function in the program that performs indirect calls using session_table->compare) to show the reachability of such a control-flow hijack. What this shows is that even if the code pointer is corrupted, the attacker is limited to only a handful of options to continue their attack, which for many of our domains is around 10 or less; furthermore, even *those targets* are all functions related to the hash table operations, which would require further steps still to reach other parts of the system. In other words, both examples show there is a relationship between the overprivilege permitted to each component of a system and the effort expended by exploit developers to weaponize their bugs to reach their targets. Deeper study on these relationships would be fruitful ground for future work.

12 COMPARISONS WITH RELATED EMBEDDED SYSTEM SECURITY WORK

Hex-Five's MultiZone Security [27] is a state-of-the-art compartmentalization framework for RISC-V. However, it requires a developer to manually decompose the application into separated binaries called “zones,” each of which are very coarse-grained—the recommended decomposition is one zone for FreeRTOS, one for the networking stack, and one or several for the application. MultiZone Security requires hundreds of cycles to switch contexts, which is negligible when only employed at millisecond intervals, but the overprivilege is very high, as large software components still have

no separation; as a result, MultiZone Security achieves a privilege reduction that falls in between the OS and dir syntactic points shown in Figure 7(a). SCALPEL imposes significantly finer-grained separation and provides substantially easier policy development and exploration. MINION [29] is another compartmentalization tool for embedded systems. However, it also enforces only very coarse-grained separation by switching between a small number of memory views and provides no method for exploring policies to tune protection and performance characteristics.

ACES [11] is closer to SCALPEL in terms of providing automatic separation for applications, however it targets enforcement using the handful of segments provided by the MPU. ACES has negligible overhead for some applications, but 20%–30% overhead is more typical, with some applications requiring over 100% overhead. As a close comparison point, we ran the Domain-size algorithm with a few modifications to target four code and four object domains; the resulting design for the HTTP web server application has an OR of 28.7 compared to SCALPEL’s OR of 1.26 (at a WS_{max} of 1,800 for a comparable overhead), which is more than 20× more separation at that level; including prefetching, that same level of protection can be enforced at an overhead that is 3× lower. As a result, SCALPEL shows that a hardware tag-based security monitor can be used to provide unprecedented levels of privilege separation for embedded systems.

13 CONCLUSION

We presented SCALPEL, a compartmentalization framework for tagged architectures. The SCALPEL back-end is a policy compiler that automatically lowers compartmentalization policies to the PIPE architecture for hardware-accelerated enforcement. The SCALPEL front-end (1) provides a tracing mode for automatically learning a program’s privileges with hardware-accelerated dynamic analysis and (2) contains compartment generation algorithms based on quantitative privilege reduction for automatically generating compartmentalization strategies. These allow a designer to easily generate and evaluate a wide range of compartmentalization designs in the privilege-performance space, interactively tune them, and select among the privilege-performance tradeoffs that meet their requirements. The capstone algorithm proposed in SCALPEL directly targets the rule cache miss rate to generate highly performant policies at any level of allowed over-privilege; SCALPEL also introduces a rule prefetching mechanism to further reduce costs by up to 4×. These allow SCALPEL to use the PIPE architecture for fine-grained privilege separation, which we demonstrate allows the enforcement of hundreds of compartments at low overhead.

REFERENCES

- [1] Ali Abbasi, Jos Wetzels, Thorsten Holz, and Sandro Etalle. 2019. Challenges in designing exploit mitigations for deeply embedded systems. In *IEEE European Symposium on Security and Privacy (EuroS&P’19)*. IEEE, 31–46.
- [2] Daniel Aloise, Amit Deshpande, Pierre Hansen, and Preyas Popat. 2009. NP-hardness of Euclidean sum-of-squares clustering. *Mach. Learn.* 75, 2 (01 May 2009), 245–248. DOI: <https://doi.org/10.1007/s10994-009-5103-0>
- [3] Anmibe. 2010. CPU Features: Non-Executable Memory. Retrieved from <https://wiki.ubuntu.com/Security/CPUFeatures>.
- [4] ARM. 2016. TrustZone technology for ARM v8-M Architecture. Retrieved from <https://developer.arm.com/documentation/100690/latest/>.
- [5] Arthur Azevedo de Amorim. 2017. *A Methodology for Micro-policies*. Ph.D. Dissertation. University of Pennsylvania. Retrieved from <http://www.seas.upenn.edu/~aarthur/thesis.pdf>.
- [6] Arthur Azevedo de Amorim, Maxime Dénès, Nick Giannarakis, Cătălin Hrițcu, Benjamin C. Pierce, Antal Spector-Zabusky, and Andrew Tolmach. 2015. Micro-policies: Formally verified, tag-based security monitors. In *36th IEEE Symposium on Security and Privacy (Oakland S&P’15)*. IEEE Computer Society, 813–830. DOI: <https://doi.org/10.1109/SP.2015.55>
- [7] Ian Beer. [n.d.]. An iOS zero-click radio proximity exploit odyssey. Retrieved from <https://googleprojectzero.blogspot.com/2020/12/an-ios-zero-click-radio-proximity.html>.

- [8] Ian Beer. 2019. In-the-wild iOS Exploit Chain 1. Retrieved from <https://googleprojectzero.blogspot.com/2019/08/in-wild-ios-exploit-chain-1.html>.
- [9] Andrea Bittau, Petr Marchenko, Mark Handley, and Brad Karp. 2008. Wedge: Splitting applications into reduced-privilege compartments. In *5th USENIX Symposium on Networked Systems Design and Implementation (NSDI'08)*. USENIX Association, Berkeley, CA, 309–322.
- [10] Shimin Chen, Michael Kozuch, Theodoros Strigkos, Babak Falsafi, Phillip B. Gibbons, Todd C. Mowry, Vijaya Ramachandran, Olatunji Ruwase, Michael P. Ryan, and Evangelos Vlachos. 2008. Flexible hardware acceleration for instruction-grain program monitoring. In *35th International Symposium on Computer Architecture (ISCA'08)*. IEEE, 377–388. Retrieved from <http://www.cs.cmu.edu/~lba/papers/LBA-isca08.pdf>.
- [11] Abraham A. Clements, Naif Saleh Almahdhub, Saurabh Bagchi, and Mathias Payer. 2018. ACES: Automatic compartments for embedded systems. In *27th USENIX Security Symposium (USENIX Security'18)*. USENIX Association, 65–82. Retrieved from <https://www.usenix.org/conference/usenixsecurity18/presentation/clements>.
- [12] The curl project. [n.d.]. curl: command line tool and library. Retrieved 15 Oct, 2020 from <https://curl.se/>.
- [13] DARPA. [n.d.]. Transparent Computing. Retrieved 1 Oct, 2020 from <https://www.darpa.mil/program/transparent-computing>.
- [14] Daniel Y. Deng and G. Edward Suh. 2012. High-performance parallel accelerator for flexible and efficient run-time monitoring. In *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'12)*. IEEE Computer Society, 1–12. Retrieved from <http://tsg.ece.cornell.edu/lib/exe/fetch.php?media=pubs:flex-dsn2012.pdf>.
- [15] CVE Details. [n.d.]. CVE Details: Libxml2 Vulnerability Statistics. Retrieved 5 Oct, 2020 from https://www.cvedetails.com/product/3311/Xmlsoft-Libxml2.html?vendor_id=1962.
- [16] Udit Dhawan and André DeHon. 2013. Area-efficient near-associative memories on FPGAs. In *International Symposium on Field-programmable Gate Arrays*. 191–200. Retrieved from http://ic.ece.upenn.edu/abstracts/dmhc_fpga2013.html.
- [17] Udit Dhawan, Cătălin Hrițcu, Rafi Rubin, Nikos Vasilakis, Silviu Chiricescu, Jonathan M. Smith, Thomas F. Knight, Jr., Benjamin C. Pierce, and André DeHon. 2015. Architectural support for software-defined metadata processing. In *International Conference on Architectural Support for Programming Languages and Operating Systems*. 487–502. Retrieved from http://ic.ece.upenn.edu/abstracts/sdmp_asplos2015.html.
- [18] Xinshu Dong, Hong Hu, Prateek Saxena, and Zhenkai Liang. 2013. A quantitative evaluation of privilege separation in web browser designs. In *European Symposium on Research in Computer Security*. Springer, 75–93.
- [19] Dovecot. [n.d.]. Dovecot Mail Server. Retrieved 12 Oct, 2020 from <https://github.com/dovecot/core>.
- [20] Kevin Elphinstone and Gernot Heiser. 2013. From L3 to seL4 what have we learnt in 20 years of L4 microkernels? In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP'13)*. ACM, New York, NY, 133–150. DOI : <https://doi.org/10.1145/2517349.2522720>
- [21] Joseph A. Fisher and Stefan M. Freudenberger. 1992. Predicting conditional branch directions from previous runs of a program. *ACM SIGPLAN Not.* 27, 9 (1992), 85–95.
- [22] Institute for Informatics Georg-August-Universität Göttingen. 1999. The MONDIAL Database. Retrieved 25 Oct, 2020 from <https://www.dbis.informatik.uni-goettingen.de/Mondial>.
- [23] LLVM Foundation. [n.d.]. The LLVM Compiler Infrastructure. Retrieved from 4 Oct, 2020 <https://llvm.org/>.
- [24] Sotiria Fytraki, Evangelos Vlachos, Yusuf Onur Koçberber, Babak Falsafi, and Boris Grot. 2014. FADE: A programmable filtering accelerator for instruction-grain monitoring. In *20th IEEE International Symposium on High Performance Computer Architecture*. 108–119. DOI : <https://doi.org/10.1109/HPCA.2014.6835922>
- [25] Khilan Gudka, Robert N. M. Watson, Jonathan Anderson, David Chisnall, Brooks Davis, Ben Laurie, Ilias Marinos, Peter G. Neumann, and Alex Richardson. 2015. Clean application compartmentalization with SOAAP. In *22nd ACM SIGSAC Conference on Computer and Communications Security (CCS'15)*. ACM, New York, NY, 1016–1031. DOI : <https://doi.org/10.1145/2810103.2813611>
- [26] Norm Hardy. 1988. The confused deputy (or why capabilities might have been invented). *SIGOPS Oper. Syst. Rev.* 22, 4 (Oct. 1988), 36–38.
- [27] HEX-Five. 2020. *MultiZone Security Reference Manual*. HEX-Five. Retrieved on 14 Oct, 2020 from <https://github.com/hex-five/multizone-sdk/blob/master/manual.pdf>.
- [28] Terry Ching-Hsiang Hsu, Kevin Hoffman, Patrick Eugster, and Mathias Payer. 2016. Enforcing least privilege memory views for multithreaded applications. In *ACM Conference on Computer and Communication Security*. DOI : <https://doi.org/10.1145/2976749.2978327>
- [29] Chung Hwan Kim, Taegyu Kim, Hongjun Choi, Zhongshu Gu, Byoungyoung Lee, Xiangyu Zhang, and Dongyan Xu. 2018. Securing real-time microcontroller systems through customized memory view switching. In *Proceedings of the 25th Network and Distributed System Security Symposium (NDSS '18)*.
- [30] Draper Laboratory. [n.d.]. Hope-tools Github Repository. Retrieved 1 June, 2019 from <https://github.com/draperlaboratory/hope-src>.

- [31] Jochen Liedtke. 1995. On micro-kernel construction. In *15th ACM Symposium on Operating Systems Principles*. 237–250.
- [32] Arm Limited. 2013. Arm Cortex-A53 Specification. Retrieved 5 Oct, 2020 from <https://developer.arm.com/ip-products/processors/cortex-a/cortex-a53>.
- [33] Arm Limited. 2016. ARMv8-M Architecture Reference Manual. Retrieved 6 Oct, 2020 from <https://developer.arm.com/documentation/ddi0553/ab/>.
- [34] Shen Liu, Dongrui Zeng, Yongzhe Huang, Frank Capobianco, Stephen McCamant, Trent Jaeger, and Gang Tan. 2019. Program-mandering: Quantitative privilege separation. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'19)*. ACM, New York, NY.
- [35] Canonical Ltd. [n.d.]. AppArmor. Retrieved 11 Sept, 2020 from <https://wiki.ubuntu.com/AppArmor>.
- [36] Sparsh Mittal. 2016. A survey of recent prefetching techniques for processor caches. *ACM Comput. Surv.* 49, 2 (2016), 1–35.
- [37] Oracle. [n.d.]. Introduction to SPARC M7 and Application Data Integrity (ADI). Retrieved 3 Dec, 2019 from https://swisdev.oracle.com/_files/What-Is-ADI.html.
- [38] Gabriel Parmer and Richard West. 2011. Mutable protection domains: Adapting system fault isolation for reliability and efficiency. *IEEE Trans. Softw. Eng.* 38, 4 (2011), 875–888.
- [39] Marios Pomonis, Theofilos Petsios, Angelos D. Keromytis, Michalis Polychronakis, and Vasileios P. Kemerlis. 2017. kR'X: Comprehensive kernel protection against just-in-time code reuse. In *Proceedings of the Twelfth European Conference on Computer Systems*. 420–436.
- [40] The GNOME Project. [n.d.]. The XML C Parser and toolkit of Gnome. Retrieved 4 Oct, 2020 from <http://www.xmlsoft.org/>.
- [41] Richard F. Rashid and George G. Robertson. 1981. Accent: A communication oriented network operating system kernel. In *8th ACM Symposium on Operating Systems Principles (SOSP'81)*. ACM, New York, NY, 64–75. DOI: <https://doi.org/10.1145/800216.806593>
- [42] Nick Roessler. 2018. Exploiting LaTeX with CVE-2018-17407. Retrieved from <https://nickroessler.com/latex-cve-2018-17407/>.
- [43] Nick Roessler and André DeHon. 2018. Protecting the stack with metadata policies and tagged hardware. In *IEEE Symposium on Security and Privacy (Oakland S&P'18)*. IEEE Computer Society.
- [44] Jerry H. Saltzer and Mike D. Schroeder. 1975. The protection of information in computer systems. *Proc. IEEE* 63, 9 (Sept. 1975), 1278–1308.
- [45] Andreas Sembrant. 2012. *Efficient Techniques for Detecting and Exploiting Runtime Phases*. Ph.D. Dissertation. Uppsala University.
- [46] NXP Semiconductors. 2018. NXP Selects Dover Microsystems' State-of-the-Art CoreGuard Cybersecurity Technology for Future Embedded Platforms. Retrieved from <https://media.nxp.com/news-releases/news-release-details/nxp-selects-dover-microsystems-state-art-coreguard-cybersecurity>.
- [47] Amazon Web Services. [n.d.]. HTTP Web Server Example. Retrieved 30 Sept, 2020 from https://freertos.org/FreeRTOS-Plus/FreeRTOS_Plus_TCP/HTTP_web_Server.html.
- [48] Timothy Sherwood, Erez Perelman, Greg Hamerly, Suleyman Sair, and Brad Calder. 2003. Discovering and exploiting program phases. *IEEE Micro* 23, 6 (2003), 84–93.
- [49] Jia Song. 2014. *Security Tagging for a Real-time Zero-kernel Operating System: Implementation and Verification*. Ph.D. Dissertation. University of Idaho.
- [50] Jia Song and Jim Alves-Foss. 2013. Security tagging for a zero-kernel operating system. In *46th Hawaii International Conference on System Sciences (HICSS'13)*. IEEE, 5049–5058. Retrieved from <http://www.computer.org/csdl/proceedings/hicss/2013/4892/00/4892f049.pdf>.
- [51] Gregory T. Sullivan, André DeHon, Steven Milburn, Eli Boling, Marco Ciaffi, Jothy Rosenberg, and Andrew Sutherland. 2017. The Dover inherently secure processor. In *IEEE International Symposium on Technologies for Homeland Security (HST'17)*. IEEE, 1–5.
- [52] Stylianos Tsampas, Akram El-Korashy, Marco Patrignani, Dominique Devriese, Deepak Garg, and Frank Piessens. 2017. Towards automatic compartmentalization of C programs on capability machines. In *Workshop on Foundations of Computer Security*. 1–14.
- [53] R. N. M. Watson, R. M. Norton, J. Woodruff, S. W. Moore, P. G. Neumann, J. Anderson, D. Chisnall, B. Davis, B. Laurie, M. Roe, N. H. Dave, K. Gudka, A. Joannou, A. T. Markettos, E. Maste, S. J. Murdoch, C. Rothwell, S. D. Son, and M. Vadera. 2016. Fast protection-domain crossing in the CHERI capability-system architecture. *IEEE Micro* 36, 5 (Sept. 2016), 38–49. DOI: <https://doi.org/10.1109/MM.2016.84>
- [54] Emmett Witchel, Junghwan Rhee, and Krste Asanović. 2005. Mondrix: Memory isolation for Linux using Mondriaan memory protection. In *20th ACM Symposium on Operating Systems Principles (SOSP'05)*. ACM, New York, NY, 31–44. DOI: <https://doi.org/10.1145/1095810.1095814>

- [55] Jonathan Woodruff, Robert N. M. Watson, David Chisnall, Simon W. Moore, Jonathan Anderson, Brooks Davis, Ben Laurie, Peter G. Neumann, Robert Norton, and Michael Roe. 2014. The CHERI capability model: Revisiting RISC in an age of risk. In *International Symposium on Computer Architecture (ISCA'14)*. 457–468. DOI: <https://doi.org/10.1109/ISCA.2014.6853201>

Received October 2020; revised March 2021; accepted April 2021