Making Secure Processors OS- and Performance-Friendly

SIDDHARTHA CHHABRA, BRIAN ROGERS, and YAN SOLIHIN

North Carolina State University

and

MILOS PRVULOVIC

Georgia Institute of Technology

In today's digital world, computer security issues have become increasingly important. In particular, researchers have proposed designs for secure processors that utilize hardware-based memory encryption and integrity verification to protect the privacy and integrity of computation even from sophisticated physical attacks. However, currently proposed schemes remain hampered by problems that make them impractical for use in today's computer systems: lack of virtual memory and Inter-Process Communication support as well as excessive storage and performance overheads. In this article, we propose (1) address independent seed encryption (AISE), a counter-mode-based memory encryption scheme using a novel seed composition, and (2) bonsai Merkle trees (BMT), a novel Merkle tree-based memory integrity verification technique, to eliminate these system and performance issues associated with prior counter-mode memory encryption and Merkle tree integrity verification schemes. We present both a qualitative discussion and a quantitative analysis to illustrate the advantages of our techniques over previously proposed approaches in terms of complexity, feasibility, performance, and storage. Our results show that AISE+BMT reduces the overhead of prior memory encryption and integrity verification schemes from 12% to 2% on average for single-threaded benchmarks on uniprocessor systems, and from 15% to 4% for coscheduled benchmarks on multicore systems while eliminating critical system-level problems.

Categories and Subject Descriptors: C.1.0 [**Processor Architectures**]: General

General Terms: Security, Performance, Design

Additional Key Words and Phrases: Secure processor architectures, memory encryption, memory integrity verification, virtualization

This work is supported in part by the National Science Foundation through grants CCF-0347425, CCF-0447783, and CCF-0541080.

An early version of this article titled "Using Address Independent Seed Encryption and Bonsai Merkle Trees to Make Secure Processors OS- and Performance-Friendly" appeared in Proceedings of the 40th Annual IEEE/ACM International Symposium on Microarchitecture (Micro-40) [Rogers et al. 2007]

Authors' addresses: Siddhartha Chhabra, Brian Rogers and Yan Solihin, North Carolina State University, Raleigh, NC 27695-7256; email: {schhabr, bmrogers, solihin}@ncsu.edu; Milos Prvulovic, Georgia Institute of Technology, Atlanta, GA 30332-0280; email: milos@cc.gatech.edu

Permission to make digital or hard copies part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from the Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org. © 2009 ACM 1544-3566/2009/03-ART16 \$5.00 DOI 10.1145/1498690.1498691 <http://doi.acm.org/10.1145/1498690.1498691>

ACM Transactions on Architecture and Code Optimization, Vol. 5, No. 4, Article 16, Publication date: March 2009.

ACM Reference Format:

Chhabra, S., Rogers, B., Solihin, Y., and Prvulovic, M. 2009. Making secure processors OS- and performance-friendly. *ACM Trans. Architect. Code Optim.* 5, 4, Article 16 (March 2009), 35 pages. DOI = 10.1145/1498690.1498691 <http://doi.acm.org/10.1145/1498690.1498691>

1. INTRODUCTION

With the tremendous amount of digital information stored on today's computer systems, and with the increasing motivation and ability of malicious attackers to target this wealth of information, computer security has become an increasingly important topic. An important research effort towards such computer security issues focuses on protecting the privacy and integrity of computation to prevent attackers from stealing or modifying critical information. This type of protection is important for enabling many important features of secure computing such as enforcement of digital rights management, reverse engineering and software piracy prevention, and trusted distributed computing.

One important emerging security threat exploits the fact that most current computer systems communicate data in its plaintext form along wires between the processor chip and other chips, such as the main memory. Also, the data is stored in its plaintext form in the main memory. This presents a situation where, by dumping the memory content and scanning it, attackers may gain a lot of valuable sensitive information, such as passwords [Kumar 2004]. Another serious and feasible threat is physical or hardware attacks, which involve placing a bus analyzer that snoops data communicated between the processor chip and other chips [Huang 2003, 2002]. Although physical attacks may be more difficult to perform than software-based attacks, they are very powerful, as they can bypass any software security protection employed in the system. The proliferation of mod-chips that bypass Digital Rights Management protection in game systems has demonstrated that given sufficient financial payoffs, physical attacks are very realistic threats.

Recognizing these threats, computer architecture researchers have recently proposed various types of secure processor architectures [Gassend et al. 2003; Gilmont et al. 1999; Lie et al. 2003; Lie et al. 2000; Rogers et al. 2006; Shi and Lee 2006; Shi et al. 2004; Shi et al. 2005; Shi et al. 2004; Suh et al. 2003a, 2003b; Yan et al. 2006; Yang et al. 2003; Zhang et al. 2005]. Secure processors assume that off-chip communication is vulnerable to attack and that the chip boundary provides a natural security boundary. Under these assumptions, secure processors seek to provide private and tamper-resistant execution environments [Suh et al. 2003b] through memory encryption [Gilmont et al. 1999; Lie et al. 2003; Lie et al. 2000; Rogers et al. 2006; Shi et al. 2004; Shi et al. 2005; Shi et al. 2004; Suh et al. 2003a, 2003b; Yan et al. 2006; Yang et al. 2003; Zhang et al. 2005] and memory integrity verification [Gassend et al. 2003; Lie et al. 2000; Rogers et al. 2006; Shi and Lee 2006; Shi et al. 2004; Shi et al. 2004, Suh et al. 2003a, 2003b, Yan et al. 2006; Zhang et al. 2005]. The chip industry also recognizes the need for secure processors, as evident, for example, in the recent effort by IBM in the SecureBlue project [IBM 2006] and Dallas Semiconductor [Semiconductor tion]. Memory encryption protects computation privacy from passive

attacks, where an adversary attempts to silently observe critical information, by encrypting and decrypting code and data as it moves on and off the processor chip. Memory integrity verification protects computation integrity from active attacks, where an adversary attempts to modify values in off-chip storage or communication channels, by computing and verifying message authentication codes (MACs) as code and data moves on and off the processor chip.

Unfortunately, current memory encryption and integrity verification designs are not yet suitable for use in general-purpose computing systems. In particular, we show in this article that current secure processor designs are incompatible with important features such as virtual memory, interprocess communication (IPC), in addition to having large performance and storage overheads. The challenges are detailed as follows:

Memory Encryption. Recently proposed memory encryption schemes for secure processors have utilized counter-mode encryption due to its ability to hide cryptographic delays on the critical path of memory fetches. This is achieved by applying a block cipher to a seed to generate a cryptographic pad, which is then bit-wise XORed with the memory block to encrypt or decrypt it. A seed is selected to be independent from the data block value so that pad generation can be started while the data block is being fetched.

In counter-mode encryption, the choice of seed value is critical for both security and performance. The security of counter-mode *requires* the uniqueness of each pad value, which implies that each seed must be unique. In prior studies [Rogers et al. 2006; Shi et al. 2004; Shi et al. 2005; Shi et al. 2004; Suh et al. 2003b; Yan et al. 2006; Yang et al. 2003; Zhang et al. 2005], to ensure that pads are unique across different blocks in memory (spatial uniqueness), the block address is used as one of the seed's components. To ensure that pads are unique across different values of a particular block over time (temporal uniqueness), a counter value, which is incremented on each write back is also used as a seed component. From the performance point of view, if most cache misses find the counters of the missed blocks available on-chip, either because they are cached or predicted, then seeds can be composed at the cache miss time, and pad generation can occur in parallel while fetching the blocks from memory.

However, using the address (virtual or physical) as a seed component causes a significant system-level dilemma in general-purpose computing systems that must support virtual memory and IPC. A virtual memory mechanism typically involves managing pages to provide process isolation and sharing between processes. It often manages the main memory by extending the physical memory to swap memory located on the disk.

Using the physical address as a seed component creates re-encryption work on page swapping. When a page is swapped out to a disk and then back into memory, it will likely reside at a new physical address. This requires the blocks of the page to be decrypted using their previous physical addresses and re-encrypted with their new physical addresses. In addition, encrypted pages in memory cannot be simply swapped out to a disk, as this creates potential pad reuse between the swapped out page and the new page at that physical address in memory. This leaves an open problem as to how to protect pages on a disk. We

could entrust the OS to encrypt and decrypt swapped pages in software if the OS is assumed to be authentic, trusted, and executing on the secure processor. However, this is likely not the most desirable solution because it makes the secure processor's hardware-based security mechanisms contingent on a secure and uncompromised OS. Alternatively, we could rely on hardware to re-encrypt swapped pages; however, this solution has its own set of problems. First, this requires supporting two encryption methods in hardware. Second, there is the issue of who can request the page re-encryptions, and how these requests are made, which requires an extra authentication mechanism.

Using virtual address as a seed component can lead to vulnerable pad reuse because different processes use the same virtual addresses. While we can prevent this by adding process ID to the seed [Yan et al. 2006], this solution creates a new set of serious system-level problems. First, this renders process IDs nonreusable, and current OSes have a limited range of process IDs. Second, shared memory-based IPC mechanisms are infeasible to use (e.g., mmap). The reason is that different processes access a shared page in memory using different combinations of virtual address and process ID. This results in different encryptions and decryptions of the shared data. Third, other OS features that also utilize page sharing cannot be supported. For example, process forking cannot utilize the copy-on-write optimization because the page in the parent and child are encrypted differently. This also holds true for shared libraries. This lack of IPC support is especially problematic in the era of chip multiprocessor systems (CMPs). Finally, storage is required for virtual addresses at the lowest level on-chip cache, which is typically physically indexed and tagged.

The root cause of problems when using address in seed composition is that address is used as a fundamental component of memory management. Using address also as a basis for security intermingles security and memory management in undesirable ways.

Memory Integrity Verification. Recently proposed memory integrity verification schemes for secure processors have leveraged a variety of techniques [Gassend et al. 2003; IBM 2006; Lie et al. 2000; Shi and Lee 2006; Shi et al. 2004; Suh et al. 2003a, 2003b; Yan et al. 2006]. However, the security of Merkle tree-based schemes [Gassend et al. 2003] has been shown to be stronger than other schemes because every block read from memory is verified individually (as opposed to Suh et al. [2003b]), and data replay attacks can be detected in addition to spoofing and splicing attacks, which are detectable by simply associating a single MAC per data block [Lie et al. 2000]. In Merkle tree memory integrity verification, a tree of MAC values is built over the memory. The root of this tree never goes off-chip, as a special on-chip register is used to hold its current value. When a memory block is fetched, its integrity can be checked by verifying its chain of MAC values up to the root MAC. Since the on-chip root MAC contains information about every block in the physical memory, an attacker cannot modify or replay any value in memory.

Despite its strong security, Merkle tree integrity verification suffers from two significant issues. First, since a Merkle tree built over the main memory computes MACs on memory events (cache misses and writebacks) generated by the processor, it covers the physical memory but not the swap memory that resides

on disk. Hence, although Merkle tree schemes can prevent attacks against values read from memory, there is no protection for data brought into memory from the disk. This is a significant security vulnerability, since by tampering with swap memory on disk, attackers can indirectly tamper with main memory. One option would be to entrust the OS to protect pages swapped to and from the disk; however, as with memory encryption, it requires the assumption of a trusted OS. Another option, as discussed in Suh et al. [2003a], is to associate one Merkle tree and on-chip secure root per process. However, managing multiple Merkle trees results in extra on-chip storage and complexity.

Another significant problem is the storage overhead of internal Merkle tree nodes in both the on-chip cache and main memory. To avoid repeated computation of internal Merkle tree nodes as blocks are read from memory, a popular optimization lets recently accessed internal Merkle tree nodes be cached on-chip. Using this optimization, the verification of a memory block only needs to proceed up the tree until the first cached node is found. Thus, it is not necessary to fetch and verify all Merkle tree nodes up to the root on each memory access, significantly improving memory bandwidth consumption and verification performance. However, our results show that Merkle tree nodes can occupy as much as 50% of the total L2 cache space, which causes the application to suffer from a large number of cache capacity misses.

The Case for Secure CMPs. Increasing chip densities and transistor counts have provided microarchitects rich opportunities to improve single-core performance through various microarchitectural innovations like exploiting instruction level parallelism via dynamic scheduling and issuing multiple instructions. However, the performance benefits achievable using a single processor will hit a ceiling due to fundamental circuit limitations and limited amounts of instruction level parallelism [Olukotun et al. 1996]. This motivates the need to better utilize the increasing transistor counts. CMPs are the current design of choice to exploit the increasing transistor counts by placing multiple simple cores on a single die [Olukotun et al. 1996]. CMPs are particularly attractive for high-performance servers where commercial workloads having large amounts of thread level parallelism like web and database applications have become most popular [Barroso et al. 2000].

Existing CMP designs [Barroso et al. 2000; Laudon and Spracklen 2007; Sinharoy et al. 2005] are typically organized with private L1 caches per core and some combination of shared and private lower-level caches, such as L2 and possibly L3 caches. All cores on the chip typically share a single, common memory bus and off-chip main memory. The memory encryption and integrity verification mechanisms proposed in this article as well as those proposed in prior secure processor studies can be applied to such CMP architectures in the same manner as in uniprocessor systems. Since these mechanisms exist at the edge of the processor chip boundary, at the interface to the memory bus, they are independent of the on-chip cache hierarchy. However, CMPs present secure processor architects with new challenges primarily due the fact that the memory hierarchy design becomes even more critical. The limited off-chip bus bandwidth can prove to be an even more precious resource as multiple threads running simultaneously on multiple cores issue requests to the memory

over the same bus. Despite the growing importance and application of CMPs, prior works on secure processors have not evaluated their mechanisms on CMP architectures.

Contributions. In this article, we investigate system-level issues in secure processors, and propose mechanisms to address these issues that are simple yet effective. Our first contribution is address independent seed encryption (AISE), which decouples security and memory management by composing seeds using logical identifiers instead of virtual or physical addresses. The logical identifier of a block is the concatenation of a logical page identifier with the page offset of the block. Each page has a logical page identifier which is distinct across the entire memory and over the lifetime of the system. It is assigned to the page the first time the page is allocated or when it is loaded from disk. AISE provides better security, since it provides complete seed/pad uniqueness for every block in the system (both in the physical and swap memory). At the same time, it also easily supports virtual memory and shared-memory, based IPC mechanisms, and simplifies page swap mechanisms by not requiring decryption and re-encryption on a page swap. AISE also lends itself to easy support for virtualization. We show that using our mechanisms, virtualization can be easily supported without requiring modifications to the guest OSes running in the virtualized environment.

The second contribution of this article is a novel and efficient extension to Merkle tree based memory integrity verification that allows extending the Merkle tree to protect off-chip data (i.e., both physical and swap memory) with a single Merkle tree and secure root MAC over the physical memory. Essentially, our approach allows pages in the swap memory to be incorporated into the Merkle tree so that they can be verified when they are reloaded into memory.

Next, we propose bonsai Merkle trees (BMTs), a novel organization of the Merkle tree that naturally leverages counter-mode encryption to reduce its memory storage and performance overheads. We observe that if each data block has a MAC value computed over the data and its counter, a replay attack must attempt to replay an old data, MAC, and counter value together. A Merkle tree built over the memory is able to detect any changes to the data MAC, which prevents any undetected changes to counter values or data. Our key insight is that: (1) there are many more MACs of data than MACs of counters, since counters are much smaller than data blocks, (2) a Merkle tree that protects counters prevents any undetected counter modification, (3) if counter modification is thus prevented, the Merkle tree does not need to be built over data MACs, and (4) the Merkle tree over counters is much smaller and significantly shallower than the one over data. As a result, we can build such a BMT over the counters, which prevents data replay attacks using a much smaller tree for less memory storage overhead, fewer MACs to cache, and a better worst-case scenario if we miss on all levels of the tree up to the root. As our results show, BMT memory integrity verification reduces the performance overhead significantly, from 12.1% to 1.8% across all SPEC 2000 benchmarks [Standard Performance Evaluation Corporation 2004], along with reducing the storage overhead in memory from 33.5% to 21.5%.

Finally, we provide a complete evaluation of the proposed schemes for a CMP architecture. As our results show, the proposed mechanisms maintain a distinct advantage for secure CMPs, reducing the performance overhead from 15.1% to 4.1%.

In the remainder of this article, we discuss related work in Section 2. Section 3 describes our assumed attack model. Section 4 describes our proposed encryption technique, while Section 5 describes our proposed integrity verification techniques in detail. Section 6 shows our experimental setup, and Section 7 discusses our results and findings. Finally, Section 8 summarizes our main contributions and results.

2. RELATED WORK

Research on secure processor architectures [Gassend et al. 2003; Gilmont et al. 1999; IBM 2006; Lie et al. 2003; Lie et al. 2000; Rogers et al. 2006; Shi and Lee 2006; Shi et al. 2004; Shi et al. 2005; Shi et al. 2004; Suh et al. 2003a, 2003b; Yan et al. 2006; Yang et al. 2003; Zhang et al. 2005] consists of memory encryption for ensuring data privacy and memory integrity verification for ensuring data integrity. Early memory encryption schemes utilized direct encryption modes [Gilmont et al. 1999; Lie et al. 2003; Lie et al. 2000; Suh et al. 2003a], in which a block cipher such as AES [FIPS Publication 197 2001] is applied directly on a memory block to generate the plaintext or ciphertext when the block is read from or written to memory. Since, on a cache miss for a block, the block must first be fetched on chip before it can be decrypted, the long latency of decryption is added directly to the memory fetch latency, resulting in execution time overheads of up to 35% (almost 17%, on average) [Yang et al. 2003]. In addition, there is a security concern for using direct encryption because different blocks having the same data value would result in the same encrypted value (ciphertext). This property implies that the statistical distribution of plaintext values matches the statistical distribution of ciphertext values, and may be exploited by attackers.

As a result of these concerns, recent studies have leveraged counter-mode encryption techniques [Rogers et al. 2006; Shi et al. 2004; Shi et al. 2005; Shi et al. 2004; Suh et al. 2003b; Yan et al. 2006; Yang et al. 2003; Zhang et al. 2005]. Counter-mode encryption overlaps decryption and memory fetch by decoupling them. This decoupling is achieved by applying a block cipher to a seed value to generate a cryptographic pad. The actual encryption or decryption is performed through an XOR of the plaintext or ciphertext with this pad. The security of counter-mode depends on the guarantee that each pad value (and thus each seed) is only used once. Consequently, a block's seed is typically constructed by concatenating the address of the block with a per-block counter value, which is incremented each time the block is encrypted [Shi et al. 2005; Suh et al. 2003b; Yan et al. 2006; Yang et al. 2003]. If the seed components are available on chip at cache miss time, decryption can be started while the block is fetched from memory. Per-block counters can be cached on chip [Suh et al. 2003b; Yan et al. 2006; Yang et al. 2003] or predicted [Shi et al. 2005].

Several different approaches have previously been studied for memory integrity verification in secure processors. These approaches include a MAC-based scheme where a MAC is computed and stored with each memory block when the processor writes to memory, and the MAC is verified when the processor reads from memory [Lie et al. 2000]. In Suh et al. [2003b], a Log Hash scheme was proposed where the overhead of memory integrity verification is reduced by checking the integrity of a series of values read from memory at periodic intervals during a program's execution using incremental, multiset hash functions. Merkle tree-based schemes have also been proposed where a tree of MAC values is stored over the physical memory [Gassend et al. 2003]. The root of the tree, which stores information about every block in memory, is kept in a secure register on-chip. Merkle tree integrity verification is often preferable over other schemes because of its security strength. In addition to spoofing and splicing attacks, replay attacks can also be prevented. We note that the Log Hash scheme can also prevent replay attacks, but as shown in Shi et al. [2004], the long time intervals between integrity checks can leave the system open to attack.

The proposed scheme in this study differs from prior studies in the following ways. Our memory encryption avoids intermingling security with memory management by using logical identifiers (rather than address) as seed components. Our memory integrity verification scheme extends Merkle tree protection to the disk in a novel way, and our BMT scheme significantly reduces the Merkle tree size. The implications of this design will be discussed in detail in the following sections.

3. ATTACK MODEL AND ASSUMPTIONS

As in prior studies on hardware-based memory encryption and integrity verification, our attack model identifies two regions of a system. The secure region consists of the processor chip itself. Any code or data on-chip (e.g., in registers or caches) is considered safe and cannot be observed or manipulated by attackers. The nonsecure region includes all off-chip resources, primarily including the memory bus, physical memory, and the swap memory in the disk. We do not constrain attackers' ability to attack code or data in these resources, so they can observe any values in the physical and swap memory and on all off-chip interconnects. Attackers can also act as a man-in-the-middle to modify values in the physical and swap memory and on all off-chip interconnects.

Note that memory encryption and integrity verification cover code and data stored in the main memory and communicated over the data bus. Information leakage through the address bus is not protected, but separate protection for the address bus, such as proposed in Gao et al. [2006], Zhuang et al. [2004], and Zhuang et al. [2004], can be employed in conjunction with our scheme.

We assume that a proper infrastructure is in place for secure applications to be distributed to end users for use on secure processors. Finally, we also assume that the secure processor is executing applications in the steady state. More specifically, we assume that the secure processor already contains the cryptographic keys and code necessary to load a secure application, verify its digital signature, and compute the Merkle tree over the application in memory.

Fig. 1. Counter-mode based memory encryption.

4. MEMORY ENCRYPTION

4.1 Overview of Counter-Mode Encryption

The goal of memory encryption is to ensure that all data and code stored outside the secure processor boundary is in an unintelligible form, not revealing anything about the actual values stored. Figure 1 illustrates how this is achieved in counter-mode encryption. When a block is being written back to memory, a seed is encrypted using a block cipher (e.g., AES) and a secret key, known only to the processor. The encrypted seed is called a cryptographic pad, and this pad is combined with the plaintext block via a bitwise XOR operation to generate the ciphertext of the block before the block can be written to memory. Likewise, when a ciphertext block is fetched from memory, the same seed is encrypted to generate the same pad that was used to encrypt the block. When the block arrives on-chip, another bitwise XOR with the pad restores the block to its original plaintext form. Mathematically, if P is the plaintext, C is the ciphertext, E is the block cipher function, and K is the secret key, the encryption performs $C = P \oplus E_K(Seed)$. By XORing both sides with $E_K(Seed)$, the decryption yields the plaintext $P = C \oplus E_K(Seed)$.

The security of counter-mode encryption relies on ensuring that the cryptographic pad (and hence the seed) is unique each time a block is encrypted. The reason for this is that suppose two blocks having plaintexts P_1 and P_2 , and ciphertext C_1 and C_2 , have the same seeds, that is $Seed_1 = Seed_2$. Since the block cipher function has a one-to-one mapping, then their pads are also the same, that is $E_K(Seed_1) = E_K(Seed_2)$. By XORing both sides of $C_1 = P_1 \oplus E_K(Seed_1)$ and $C_2 = P_2 \oplus E_K(Seed_2)$, we obtain the relationship of $C_1 \oplus C_2 = P_1 \oplus P_2$, which means that if any three variables are known, the other can be known, too. Since ciphertexts are known by the attacker, if one plaintext is known or can be guessed, then the other plaintext can be obtained. Therefore, the security requirement for seeds is that they must be globally unique, both spatially (across blocks) and temporally (versions of the same block over time).

The performance of counter-mode encryption depends on whether the seed of a code/data block that misses in the cache is available at the time the cache miss

is determined. If the seed is known by the processor at the time of a cache miss, the pad for the code/data block can be generated in parallel with the off-chip data fetch, hiding the overhead of memory encryption.

Two methods to achieve the global uniqueness of seeds have been studied. The first is to use a global counter as the seed for all blocks in the physical memory. This global counter is incremented each time a block is written back to memory. The global counter approach avoids the use of address as a seed component. However, when the counter reaches its maximum value for its size, it will wrap around and start to reuse its old values. To provide seed uniqueness over time, counter values cannot be reused. Hence, when the counter reaches its maximum, the secret key must be changed, and the entire physical memory along with the swap memory must be decrypted with the old key and re-encrypted with the new secret key. This re-encryption is very costly and frequent for the global counter approach [Yan et al. 2006], and can only be avoided by using a large global counter, such as 64 bits. Unfortunately, large counters require a large on-chip counter cache storage in order to achieve a good hit rate and overlap decryption with code/data fetch. If the counter for a missed code/data cache block is not found in the counter cache, it must first be fetched from memory along with fetching the code/data cache block. Decryption cannot begin until the counter fetch is complete, which exposes decryption latency and results in poor performance.

To avoid the fast growth of global counters, which leads to frequent memory re-encryption, prior studies use per-block counters [Shi et al. 2005; Suh et al. 2003b; Yan et al. 2006; Yang et al. 2003], which are incremented each time the corresponding block is written back to memory. Since each block has its own counter, the counter increases at an orders-of-magnitude slower rate compared to the global counter approach. To provide seed uniqueness across different blocks, the seed is composed by concatenating the per-block counter, the block address, and chunk id.¹ This seed choice also meets the performance criterion, since block addresses can be known at cache miss time, and studies have shown that frequently needed block counters can be effectively cached on-chip [Suh et al. 2003b; Yan et al. 2006; Yang et al. 2003] or predicted [Shi et al. 2005] at cache miss time.

However, this choice for seed composition has several significant disadvantages due to the fact that block address, which was designed as an underlying component of memory management, is now being used as a component of security. Because of this conflict between the intended use of addresses and their function in a memory encryption scheme, many problems arise for a secure processor when block address (virtual or physical) is used as a seed component. We discuss these problems in the next section.

¹A *chunk* refers to the unit of encryption/decryption in a block cipher, such as 128 bits (16 bytes). A cache or memory block of 64 bytes contains four chunks. Seed uniqueness must hold across chunks, hence the *chunk id*, referring to which chunk being encrypted in a block, is included as a component of the seed.

Fig. 2. Virtual Memory management allows virtual pages of different processes to map to a common physical page for sharing purpose (1), the same virtual pages in different processes to map to different physical pages (2), and some virtual pages to reside in the swap memory in the disk (3).

4.2 Problems with Current Counter-Mode Memory Encryption

Most general purpose computer systems today employ virtual memory, illustrated in Figure 2. In a system with virtual memory, the system gives an abstraction that each process can potentially use all addresses in its virtual address space. A paging mechanism is used to translate virtual page addresses (that a process sees) to physical page addresses (that actually reside in the physical and swap memory). The paging mechanism provides process isolation by mapping the same page address of different processes to different physical pages (circle (2)), and *sharing* by mapping virtual pages of different processes to the same physical page (circle (1)). The paging mechanism often extends the physical memory to the swap memory area in disks in order to manage more pages. The swap memory holds pages that are not expected to be used soon (circle (3)). When a page in the swap memory is needed, it is brought in to the physical memory, while an existing page in the physical memory is selected to be replaced into the swap memory.

The use of physical address in the seed causes the following complexity and possible security problems. The mapping of a virtual page of a process to a physical frame may change dynamically during execution, due to page swaps. Since the physical address changes, the entire page must be first decrypted using the old physical addresses and then re-encrypted using the new physical addresses on a page swap. In addition, pages encrypted based on physical address cannot be simply swapped to disk or pad reuse may occur between blocks in the swapped out page and blocks located in the page's old location in physical memory. This leaves an open problem as to how to protect pages on disk.

The use of virtual address has its own set of critical problems. Seeds based on virtual address are vulnerable to pad reuse, since different processes use the same virtual addresses and could easily use the same counter values. Adding process ID to the seed solves this problem but creates a new set of system-level issues. First, process IDs can now no longer be reused by the OS, and current OSes have a limit on the range of possible process IDs. Second, shared-memory IPC mechanisms cannot be used. Consider that a single physical page may be mapped into multiple virtual pages in either a single process or in multiple processes. Since each virtual page will see its own process ID and

virtual address combination, the seeds will be different and will produce different encryption and decryption results. Consequently, `mmap/munmap` (based on shared-memory) cannot be supported, and these are used extensively in `glibc` for file I/O and memory management, especially for implementing threads. This is a critical limitation for secure processors, especially in the age of CMPs. Third, other OS features that also utilize page sharing cannot be supported. For example, process forking cannot utilize the copy-on-write optimization because the page in the parent and child are encrypted differently. This also holds true for shared libraries. Finally, since virtual addresses are often not available beyond the L1 cache, extra storage may be required for virtual addresses at the lowest level on-chip cache.

One may attempt to augment counter-mode encryption with special mechanisms to deal with paging or IPC. Unfortunately, they would likely result in great complexity. For example, when physical address is used, to avoid seed/pad reuse in the swap memory, an authentic, secure OS running on the secure processor could encrypt and decrypt swapped pages in software. However, this solution is likely not desirable, since it makes the secure processor's hardware-based security mechanisms contingent on a secure and uncompromised OS. OS vulnerabilities may be exploited in software by attackers to subvert the secure processor. Alternatively, we could rely on hardware to re-encrypt swapped pages, however, this solution has its own set of problems. First, this requires supporting two encryption methods in hardware. A page that is swapped out must first be decrypted (using counter mode) and then encrypted (using direct mode) before it is placed in the swap memory, while the reverse must occur when a page is brought from the disk to the physical memory. Second, there is the issue of who can request the page re-encryptions, and how these requests are made, which requires an extra authentication mechanism. Another example, when virtual address is used, is that shared memory IPC and copy-on-write may be enabled by encrypting all shared pages with direct encryption, while encrypting everything else with counter-mode encryption. However, this also complicates OS handling of IPC and copy-on-write and, at the same time, complicates the hardware, since it must now support two modes of encryption. Therefore, it is arguably better to identify and deal with the root cause of the problem: address is used as a fundamental component of memory management, and using the address also as a basis for *security* intermingles security and memory management in undesirable ways.

4.3 Address Independent Seed Encryption

In light of the problems caused by using address as a seed component, we propose a new seed composition mechanism that we call address-independent seed encryption (AISE), that is free from the problems of address-based seeds. The key insight is that rather than using addresses as a seed component alongside a counter, we use logical identifiers instead. These logical identifiers are truly unique across the entire physical and swap memory and over time.

Conceptually, each block in memory must be assigned its own logical identifier. However, managing and storing logical identifiers for the entire memory

would be quite complex and costly (similar to global counters). Fortunately, virtual memory management works on the granularity of pages (usually 4 Kbytes) rather than words or blocks. Any block in memory has two components: page address, which is the unit of virtual memory management, and page offset. Hence, it is sufficient to assign logical identifiers to pages, rather than to blocks. Thus, for each chunk in the memory, its seed is the concatenation of a *logical page identifier* (LPID), the page offset of the chunk's block, the block's counter value, and the chunk id.

To ensure complete uniqueness of seeds across the physical and swap memory and over time, the LPID is chosen to be a unique value assigned to a page when it is first allocated by the system. The LPID is unique for that page across the system lifetime, and never changes over time. The unique value is obtained from an on-chip counter called the global page counter (GPC). Once a value of the GPC is assigned to a new page, it is incremented. To provide true uniqueness over time, the GPC is stored in a nonvolatile register on chip. Thus, even across system reboots, hibernation, or power optimizations that cut power off to the processor, the GPC retains its value. Rebooting the system does not cause the counter to reset and start reusing seeds that have been used in the past boot. The GPC is also chosen to be large (64 bits) so that it does not overflow for millenia, easily exceeding the lifetime of the system. We present further details on the assignment of LPID to pages in Section 4.6.

One may have concerns for how the LPID scheme can be used in systems that support multiple page sizes, such as when super pages (e.g., 16 MB) are used. However, the number of page offset bits for a large page always exceeds the number of page offset bits for a smaller page. Hence, if we choose the LPID portion of the seed to have as many bits as needed for the smallest page size supported in the system, the LPID still covers the unit of virtual memory management (although sometimes unnecessarily covering some page offset bits) and provides seed uniqueness for the system.

The next issue we address is how to organize the storage of LPIDs of pages in the system. One alternative is to add a field for the LPID in page table entries and TLB entries. However, this approach significantly increases the page table and TLB size, which is detrimental to system performance. Additionally, the LPID is only needed for accesses to off-chip memory, while TLBs are accessed on each memory reference. Another alternative would be to store the LPIDs in a dedicated portion of the physical memory. However, this solution also impacts performance, since a memory access now must fetch the block's counter and LPID in addition to the data, thus increasing bandwidth usage. Consequently, we choose to costore LPIDs and counters by taking an idea from the split counter organization [Yan et al. 2006]. We associate each counter block with a page in the system, and each counter block contains one LPID and all block counters for a page.

Figure 3 illustrates the organization, assuming 32-bit virtual addresses, a 4-Kbyte page size, 64-byte blocks, 64-bit LPID, and a 7-bit counter per block. A virtual address is split into the high 20-bit virtual page address and 12-bit page offset. The virtual page address is translated into the physical address, which is used to index a counter cache. Each counter block stores the 64-bit LPID of

Fig. 3. Organization of logical page identifiers.

a page, and 64 7-bit counters where each counter corresponds to one of the 64 blocks in a page. If the counter block is found, the LPID and one counter are used for constructing the seed for the address, together with the 8 high order bits of the page offset (6-bit block offset and 2-bit chunk id). Padding is added to make the seed 128 bits, which corresponds to the chunk size in the block cipher. Note that the LPID and counter block can be found using simple indexing for a given physical address.

In contrast to using two levels of counters in Yan et al. [2006], we only use small per-block (minor) counters. We eliminate the major counter and use the LPID instead. If one of the minor counter overflows, we need to avoid seed reuse. To achieve that, we assign a new LPID for that page by looking up the GPC, and re-encrypt only that page. Hence, the LPID of a page is no longer static. Rather, a new unique value is assigned to a page when a page is first allocated and when a page is re-encrypted.

4.4 Dealing with Swap Memory and Page Swapping

In our scheme, no two pages share the same LPID and hence seed uniqueness is guaranteed across the physical and swap memory. In addition, once a unique LPID value is assigned to a page, it does not change until the page needs to be re-encrypted. Hence, when a page is swapped out to the disk, it retains a unique LPID and does not need to be re-encrypted or specially handled. The virtual memory manager can just move a page from the physical memory to the swap memory. The page's LPID and block of counters can be moved to the swap memory as well to free up as much memory as possible, or moved to a region in the kernel memory if one wants to reduce I/O activities.

When an application suffers a page fault, the virtual memory manager locates the page and its block of counters in the disk, then brings it into the physical memory. The block of counters (including LPID) are placed at the appropriate physical address in order for the block to be directly indexable and storable by the counter cache. Therefore, the only special mechanism that needs

to be added to the page swapping mechanism is proper handling of the page's counter blocks. Since no re-encryption is needed, moving the page in and out of the disk can be accomplished with or without the involvement of the processor (e.g., we could use DMA).

Finally, while in AISE, we rely on the OS to swap pages between the main memory and the disk, we do so in an indirect way. Specifically, we only rely on the OS to move already protected data, which has been previously encrypted and authenticated in hardware, around in memory and between memory and disk. Thus, we avoid the more dangerous form of OS-dependence where the system relies on the OS to directly perform cryptographic functions in certain cases (i.e., directly encrypting, decrypting, or authenticating data values). For example, we do not want to rely on the OS to encrypt and decrypt data moved between memory and the disk because the OS has direct access to plaintext values in this case, and a compromised OS could access and/or tamper with plaintext values. In our scheme, the OS has no direct access to plaintext, and hence, the security of the data afforded by our scheme is not contingent on an uncompromised OS.

4.5 Dealing with Page Sharing

Page sharing is problematic to support if virtual address is used as a seed component, since different processes may try to encrypt or decrypt the same page with different virtual addresses. With our LPID scheme, the LPID is unique for each page and can be directly looked up using the physical address. Therefore, all page sharing uses can naturally be facilitated without any special mechanisms.

4.6 Virtualization support for Secure Processors

Virtualization was first introduced in the 1960s to allow partitioning of expensive mainframe resources, particularly the main memory. Modern computers with increased processing power have led to a resurgence of interest in Virtualization Technology as a way to multiplex the real machine into multiple virtual machines (VMs) with each VM running a separate Operating System instance. The virtualization layer (hypervisor or virtual machine monitor) is responsible for controlling and allocating hardware resources to the VMs.

Virtualization increasingly finds applications in supporting server and workload consolidation [Marty and Hill 2007], distributed Web servers [Whitaker et al. 2002] and secure computing platforms [Garfinkel et al. 2003] among others. The growing importance and applications of virtualization make it imperative that secure processor designs continue to support virtualization in a manner similar to current systems. We show that AISE allows virtualization to be supported naturally by requiring minimal changes to the hypervisor or the virtual machine monitor (VMM) for virtualizing the LPID. AISE allows support for all variants of virtualization, namely, full virtualization, paravirtualization, and hardware assisted virtualization.

Virtualization Model. In a virtualized environment, the guest operating systems run on top of the VMM. The VMM itself can either run directly on top of

Fig. 4. VMM models.

Fig. 5. Portion of Page Fault Handling Routine of an AISE-compliant OS.

the hardware (Type 1 or native VM) or on top of a host operating system (Type 2 or hosted VM). The two models are shown in Figure 4.

Type 2 VMM uses the existing host operating system’s abstractions to implement its services. The two levels of indirection in Type 2 VMM result in poor performance when compared to Type 1 VMM, which runs directly on the hardware. The rest of our discussion assumes a Type 1 VMM similar to the one used by VMWare’s ESX server [Waldspurger 2002] and Xen [2003].

OS Changes to Support AISE. Before we present the proposed virtualization of the LPID, we discuss the details of how LPIDs are assigned to pages. As discussed in Section 4.3, a page is assigned an LPID when it is first loaded to the main memory. This assignment requires the OS to access the nonvolatile, on-chip GPC register when a page is loaded from disk to main memory and to assign this counter value as the LPID for the page.

We propose a new privileged instruction for this purpose, which we call RDAGPC (Read and Assign GPC). This instruction reads (Load) the GPC register, associates (Assign) this GPC value as the LPID for the page being loaded from disk, and subsequently increments the GPC register value to prevent the reuse of LPIDs. RDAGPC is a privileged instruction so only the Operating System kernel running at the highest priority level (CPL 0) has the permission to execute it. More specifically, the page fault handler of the OS will execute this instruction after determining that the required page was not in main memory and needs to be loaded from the disk (*major page fault*). Figure 5 shows a part of the page fault handling routine, which is responsible for handling major page faults that needs to be modified in order to exploit AISE capabilities.

Based on this model, we now discuss how virtualization, in its three main varieties, can be supported on an AISE-enabled secure processor.

4.6.1 Virtualizing the LPID. Full Virtualization. Full virtualization using binary translation is the most widely used variant of virtualization today [Adams and Agesen 2006]. In Full virtualization, the guest OSes run unmodified except that they run at a lower privilege level (CPL 1), and the VMM runs at the highest priority level (CPL 0) and is responsible for managing the hardware resources. When a guest OS attempts to execute a privileged instruction, it causes a trap into the VMM, and the VMM is responsible for emulating the instruction. The instruction we introduced for assigning LPIDs to pages loaded to memory from the disk, RDAGPC is a privileged instruction and will therefore result in a trap to the VMM when executed by a guest OS.

The guest OS is not allowed to access the physical memory directly. The guest OS assigns virtual frames for the virtual pages and the VMM is responsible for mapping these virtual frames to real frames. The guest OS maintains its own copy of virtual page tables, which contain the mappings from virtual page numbers to virtual frame numbers. The VMM shadows the page tables of the guest OSes. The shadowed page structure maintains a mapping from the virtual page numbers to the real page numbers. Coherency of the shadow structures is maintained via tracing, wherein the guest OS can read its page table but any write to the page table results in a trap to the VMM.

On a page fault in the guest OS, the guest OS executes its page fault handling routine and constructs a mapping from the virtual page number to a virtual frame number. As part of the page fault routine, AISE compliant OSes execute the RDAGPC instruction to assign the LPID. When the guest OS executes this instruction, it causes a trap to the VMM. The VMM, knowing that it is responsible for assigning the LPIDs ignores the RDAGPC instruction and returns control to the guest OS. Next, the guest OS attempts to write the newly constructed mapping to its page table. This again results in a trap to the VMM. The VMM now constructs an appropriate shadow page table entry by mapping the virtual page to a real frame. As part of this process, the VMM may be required to load a page from the disk, if the page is not already present in the main memory. If the page is loaded from disk, the VMM executes the RDAGPC instruction to associate an LPID for this page. Once the shadow page table entry is created, the VMM resumes the guest OS execution. The guest OS can now write the mapping it constructed to its page table. The same guest OS, when run in a nonvirtualized environment will continue to assign the LPIDs to pages loaded from the disk.

Hence, by virtualizing the on-chip GPC register and making the VMM responsible for managing and assigning LPIDs on behalf of the guest OSes, full virtualization is easily supported with AISE, without requiring any changes to an AISE compliant OS and requiring minimal changes to the VMM.

Paravirtualization. In paravirtualization, guest OSes are modified to know that they are running inside a VM. The VMM provides a hypercall interface to provide services to the VMs. The guest OSes access all hardware state through hypercalls by voluntarily trapping to the VMM. Each guest OS knows that it does not have the entire physical memory. Instead, each guest OS requests physical memory from the VMM. Each guest OS is statically assigned its share

of memory during initialization and is responsible for managing its own memory in a way similar to self-paging [Hand 1999].

Unlike full virtualization, there is no concept of virtual frame numbers as now the guest OS performs its own paging. When the guest OS requires a new page table, possibly due to a new process creation, it allocates the page table from its own memory store. Any writes to the page table must be validated by the hypervisor via hypercalls.

On a page fault, the guest OS executes its page fault handler. If the requested physical frame is loaded from disk, the RDAGPC instruction is executed. RDAGPC, being a privileged instruction, traps to the VMM via a hypercall. The VMM then emulates this instruction by actually executing it on behalf of the guest OS. Once the RDAGPC instruction is completed, the physical frame loaded from the disk will have an associated LPID and the guest OS is restarted. Hence, similar to full virtualization, by virtualizing the on-chip GPC register, paravirtualization can be easily supported without requiring any changes to an AISE compliant OS and requiring minimal changes to the VMM.

Hardware Assisted Virtualization. Intel [2005] and AMD [2005] have recently introduced hardware virtualization where privileged instructions automatically trap to the VMM. The first generation of hardware-assisted virtualization techniques do not provide support for memory virtualization. On present generation hardware, the VMM is responsible for virtualizing memory. Hence, AISE can support virtualization as described above. We believe that when hardware supports memory virtualization, it should be straightforward to extend the support with AISE. The current hardware allows the VMM to specify unconditional traps. Assuming that the next generation hardware continues to allow the VMM to specify unconditional traps, virtualization can be supported with AISE in the environment by specifying accesses to GPC register to trap to the VMM.

4.7 Advantages of AISE

Our AISE scheme satisfies the security and performance criteria for counter-mode encryption seeds, while naturally supporting virtual memory management features and IPC without much complexity. The LPID portion of the seed ensures that the blocks in every page, both in the physical memory and on disk, are encrypted with different pads. The page offset portion of the seed ensures that each block within a page is encrypted with a different pad. The block counter portion of the seed ensures that the pad is unique each time a single block is encrypted. Finally, since the global page counter is stored in nonvolatile storage on-chip, the pad uniqueness extends across system boots.

From a performance perspective, AISE does not impose any additional storage or runtime overheads over prior counter-mode encryption schemes. AISE allows seeds to be composed at cache miss time, since both the LPID and counter of a block are co-stored in memory and cached together on-chip. Storage overhead is equivalent to the already-efficient split counter organization, since LPID replaces the major counter of the split counter organization and does not add extra storage. On average, a 4-Kbyte page only requires 64 bytes of storage

for the LPID and counters, representing a 1.6% overhead. Similar to the split counter organization, AISE does not incur entire-memory re-encryption when a block counter overflows. Rather, it only incurs re-encryption of a page when overflow occurs.

From a complexity perspective, AISE allows pages to be swapped in and out of the physical memory without involving page re-encryption (unlike using physical address), while allowing all types of IPC and page sharing (unlike using virtual address). AISE can be naturally extended to provide support for all variants of virtualization without requiring any modifications to an AISE compliant OS and with minimal changes to the existing VMMs.

To summarize, memory encryption using our AISE technique retains all of the latency-hiding ability as proposed in prior schemes, while eliminating the significant problems that arise from including address as a component of the cryptographic seed.

5. MEMORY INTEGRITY VERIFICATION

The goal of a memory integrity verification scheme is to ensure that a value loaded from some location by a processor is equal to the most recent value that the processor last wrote to that location. There are three types of attacks that may be attempted by an attacker on a value at a particular location. Attackers can replace the value directly (spoofing), exchange the value with another value from a different location (splicing), and replay an old value from the same location (replay). As discussed in XOM [Gilmont et al. 1999], if for each memory block a MAC is computed using the value and address as its input, spoofing and splicing attacks would be detectable. However, replay attacks can be successfully performed by rolling back both the value and its MAC to their older versions. To detect replay attacks, Merkle tree verification has been proposed [Gassend et al. 2003]. A Merkle tree keeps hierarchical MACs organized as a tree, in which a parent MAC protects multiple child MACs. The root of the tree is stored on-chip at all times so that it cannot be tampered by attackers. When a memory block is fetched, its integrity can be verified by checking its chain of MAC values up to the root MAC. When a cache block is written back to memory, the corresponding MAC values of the tree are updated. Since the on-chip MAC root contains information about every block in the physical memory, an attacker cannot modify or replay any value in the physical memory.

5.1 Extended Merkle Tree Protection

Previously proposed Merkle tree schemes, which only cover the physical memory, as shown in Figure 6(a), compute MACs on memory events (cache misses and write backs) generated by the processor. However, I/O transfer between the physical memory and swap memory is performed by an I/O device or DMA and is not visible to the processor. Consequently, the standard Merkle tree protection only covers the physical memory but not the swap memory. This is a significant security vulnerability, since by tampering with the swap memory in the disk, attackers can indirectly tamper with the main memory. We note that it would

Fig. 6. Our novel Merkle tree organization for extending protection to the swap memory in disk.

be possible to entrust a secure OS, with the job of protecting pages swapped to and from the disk in software. However, this solution requires the assumption of a secure and untampered OS, which may not be desirable. Also, as discussed in Suh et al. [2003a], it would be possible to compute the Merkle tree over the virtual address space of each process to protect the process in both the memory and the disk. However, this solution would require one Merkle tree and on-chip secure root MAC per process, which results in extra on-chip storage for the root MACs and complexity in managing multiple Merkle trees.

This security issue clearly motivates the need to extend the Merkle tree protection to all off-chip data, both in the physical and swap memory, as illustrated in Figure 6(b). To help explain our solution, we define two terms: page Merkle subtree and page root. A page Merkle subtree is simply the subset of all the MACs of the Merkle tree, which directly cover a particular page in memory. A page root is the top-most MAC of the page Merkle subtree. Note that the page Merkle subtree and page root are simply MAC values that make up a portion of the larger Merkle tree over the entire physical memory.

To extend Merkle tree protection to the swap memory, we make two important observations. First, for each memory page, its page root is sufficient to verify the integrity of all values on the page. The internal nodes of the page Merkle subtree can be recomputed and verified as valid by comparing the computed page root with the stored, valid page root. Secondly, the physical memory is covered entirely by the Merkle tree, and hence, it provides secure storage. From these two observations, we can conclude that as long as the page roots of all swap memory pages are stored in the physical memory, then the entire swap memory integrity can be guaranteed. To achieve this protection, we dedicate a small portion of the physical memory to store page root MACs for pages currently on disk, which we refer to as the page root directory. Note that while our scheme requires a small amount of extra storage in main memory for the page root directory, the on-chip Merkle tree operations remain the same and a single on-chip MAC root is still all we require to maintain the integrity of the

entire tree. Furthermore, as shown in Figure 6(b), the page root directory itself is protected by the Merkle tree. The implication of this design is that every page in memory is associated with a page root. If a page needs to be swapped out to the disk, then we can maintain the integrity of its page root by retaining it in memory (in the page root directory), which is secure, since it is protected by the Merkle tree over memory. The page roots are themselves stored on a memory page that also has its own page root. Thus, we can continue this process of swapping pages out to disk, and retaining their page roots in the Merkle tree protected memory to allow the later verification of any pages reloaded into memory from the swap space on disk.

To illustrate how our solution operates, consider the following example. Suppose that the system wants to load a page B from swap memory into physical memory currently occupied by a page A. The integrity verification proceeds as follows. First, the page root of B is looked up from the page root directory and brought on chip. Since this lookup is performed using a regular processor read, the integrity of the page root of B is automatically verified by the Merkle tree. Second, page A is swapped out to the disk and its page root is installed at the page root directory. This installation updates the part of the Merkle tree that covers the directory, protecting the page root of A from tampering. Third, the page Merkle subtree of A is invalidated from on-chip caches in order to force future integrity verification for the physical frame where A resided. Next, the page root of B is installed in the proper location as part of the Merkle tree, and the Merkle tree is updated accordingly. Finally, the data of page B can be loaded into the physical frame. When any value in B is loaded by the processor, the integrity checking will take place automatically by verifying data against the Merkle tree nodes, at least up to the already-verified page root of B.

5.2 Bonsai Merkle Trees

We introduce BMT, a novel Merkle tree organization designed to significantly reduce their performance overhead for memory integrity verification. To motivate the need for our BMT approach, we note a common optimization that has been studied for Merkle tree verification is to cache recently accessed and verified MAC values on-chip [Gassend et al. 2003]. This allows the integrity verification of a data block to complete as soon as a needed MAC value is found cached on-chip. The reason being, since this MAC value has previously been verified and is safe on-chip, it can be trusted as if it were the root of the tree. The resulting reduction in memory bandwidth consumption significantly improves performance compared to fetching MAC values up to the tree root on every data access. However, the sharing of on-chip cache between data blocks and MAC values can significantly reduce the amount of available cache space for data blocks. In fact, our experiments show that for memory-intensive applications, up to 50% of a 1-MB L2 cache can be consumed by MAC values during application execution, severely degrading performance. It is likely that MACs occupy such a large percentage of cache space because MACs in upper levels of a Merkle tree have high temporal locality when the verification is repeated due to accesses to the data blocks that the MAC covers.

Before we describe our BMT approach, we motivate it from a security perspective. BMTs exploit certain security properties that arise when Merkle tree integrity verification is used in conjunction with counter-mode memory encryption. We make two observations. First, the Merkle tree is designed to prevent data replay attacks. Other types of attacks, such as data spoofing and splicing, can be detected simply by associating a single MAC value with each data block. Second, in most proposed memory encryption techniques using counter-mode, each memory block is associated with its own counter value in memory [Shi et al. 2004; Shi et al. 2005; Suh et al. 2003b; Yan et al. 2006; Yang et al. 2003]. Since a block's counter value is incremented each time a block is written to memory, the counter can be thought of as a version number for the block. Based on these observations, we make the following claim:

In a system with counter-mode encryption and Merkle tree memory integrity verification, data values do not need to be protected by the Merkle tree as long as (1) each block is protected by its own MAC, computed using a keyed hashing function (e.g., HMAC based on SHA-1), (2) the block's MAC includes the counter value and address of the block, and (3) the integrity of all counter values is guaranteed.

To support this claim, we provide the following argument. Let us denote the plaintext and ciphertext of a block of data as P and C , its counter value as ctr , the MAC for the block as M , and the secret key for the hash function as K . The MAC of a block is computed using a keyed cryptographic hash function H with the ciphertext and counter as its input, that is, $M = H_K(C, ctr)$. Integrity verification computes the MAC and compares it against the MAC that was computed in the past and stored in the memory. If they do not match, integrity verification fails. Since the integrity of the counter value is guaranteed (a requirement in the claim), attackers cannot tamper with ctr without being detected. They can only tamper with C to produce C' , and/or the stored MAC to produce M' . However, since the attacker does not know the secret key of the hash function, they cannot produce a M' to match a chosen C' . In addition, due to the noninvertibility property of a cryptographic hash function, they cannot produce a C' to match a chosen M' . Hence, $M' \neq H_K(C', ctr)$. Since, during integrity verification, the computed MAC is $H_K(C', ctr)$, while the stored one is M' , integrity verification will fail and the attack detected. In addition, attackers cannot replay both C and M to their older version because the old version satisfies $M^{old} = H_K(C^{old}, ctr^{old})$, while the integrity verification will compute the MAC using the fresh counter value whose integrity is assumed to be guaranteed ($H_K(C^{old}, ctr)$), which is not equal to $H_K(C^{old}, ctr^{old})$. Hence, replay attacks would also be detected.

The claim is significant because it implies that we only need the Merkle tree to cover counter blocks, but not code or data blocks. Since counters are a lot smaller than data (a ratio of 1:64 for 8-bit counters and 64-byte blocks), the Merkle tree to cover the block counters is substantially smaller than the Merkle tree for data. Figure 7(a) shows the traditional Merkle tree, which covers all

Fig. 7. Reduction in size of bonsai Merkle trees compared to standard Merkle trees.

data blocks, while Figure 7(b) shows our BMT that only covers counters, while data blocks are now only covered by their MACs.

Since the size of the Merkle tree is significantly reduced, and since each node of the Merkle tree covers more data blocks, the amount of on-chip cache space required to store frequently accessed BMT nodes is significantly reduced. To further reduce the cache footprint, we do not cache data block MACs. Since each data block MAC only covers four data blocks, it has a low degree of temporal reuse compared to upper level MACs in a standard Merkle tree. Hence, it makes sense to only cache BMT nodes but not data block MACs, as we will show in Section 7.

Overall, BMTs achieve the same security protection as in previous schemes where a Merkle tree is used to cover the data in memory (i.e., data spoofing, splicing, and replay protection), but with much less overhead.

6. EXPERIMENTAL SETUP

6.1 Machine Models

We use SESC [Renau et al. 2004], an open source execution driven simulator, to evaluate the performance of our proposed memory encryption and integrity verification approaches. For uniprocessor evaluations, we model a 2-GHz, 3-issue, out-of-order processor with split L1 data and instruction caches. Both caches have a 32-KB size, 2-way set associativity, and 2-cycle round-trip hit latency. The L2 cache is unified and has a 1-MB size, 8-way set associativity, and 10-cycle round-trip hit latency. For counter mode encryption, the processor includes a 32-KB, 16-way set-associative counter cache at the L2 cache level. All caches have a 64-B block size and use LRU replacement. We assume a 1-GB main memory with an access latency of 200 processor cycles. We model a memory bus with a bandwidth of 10-GBytes/s. The encryption/decryption engine simulated is a 128-bit AES engine with a 16-stage pipeline and a total latency of 80 cycles, while the MAC computation models HMAC [Krawczyk et al. 1997] based on SHA-1 [FIPS Publication 180-1 1995] with 80-cycle latency [Kgil et al. 2004]. Counters are composed of a 64-bit LPID concatenated with a 7-bit block counter. So a 64-B counter cache block contains one LPID value along with

64 block counters (enough for a 4-KB memory page). The default authentication code size used is 128 bits.

For the CMP evaluation, we model a two-core CMP system where each core has private L1 data and instruction caches. The L2 cache and all lower levels of the memory hierarchy are shared by both cores. To better match current CMP configurations, we have changed two of the main system parameters from the uniprocessor model. The L2 cache size is increased to 2 MB and the memory bus bandwidth is increased to 20 GBytes/s. All other system parameters are the same as the uniprocessor case.

6.2 Benchmarks

We use 21 C/C++ SPEC2K benchmarks [Standard Performance Evaluation Corporation 2004] for our uniprocessor evaluations. We only omit Fortran 90 benchmarks, which are not supported on our simulator infrastructure. In each figure, we show the individual result for benchmarks that have an L2 cache miss rate higher than 20%, but the average is calculated across all 21 benchmarks that we simulate.

For our CMP evaluations, we have created 26 pairs of benchmarks using the SPEC2K benchmarks. Each pair consists of two SPEC2K benchmarks, which are spawned as two separate threads on each of the two cores of the modeled CMP system. To capture different memory behaviors, we classify the benchmarks into two categories: those that when run alone have L2 cache miss rates of less than 20% and those that have L2 cache miss rates of 20% or higher. We select a few benchmarks from each group and match them so that all combinations are represented. In the first type of benchmark pairs, the benchmarks in a pair are both taken from the low miss rate group: *perlbnk_twolf* and *twolf_vpr*. In the second type of benchmark pairs, one benchmark in a pair is taken from the low miss rate group while another is taken from the high miss rate group: *apsi_bzip2*, *gzip_applu*, *gzip_apsi*, *gzip_art*, *perlbnk_art*, *perlbnk_swim*, *swim_gzip*, *swim_twolf*, *twolf_swim*, *vpr_applu*, *vpr_art*, *applu_gzip*, and *swim_perlbnk*. The last type of benchmark pairs are ones in which both benchmarks in a pair are taken from high miss rate group: *apsi_art*, *apsi_earthquake*, *apsi_mcf*, *art_mcf*, *art_swim*, *mcf_art*, *mcf_swim*, *swim_art*, *swim_mcf*, *earthquake_apsi*, and *mcf_apsi*.

For each simulation, we use the reference input set and simulate for 1 billion instructions after fast forwarding for 5 billion. For CMP simulations, instructions are skipped only for the first benchmark in the benchmark pair and the simulation ends when the combined number of instructions simulated for the benchmark pair reaches 1 billion. In our experiments, we ignore the effect of page swaps as the overhead due to page swaps with our techniques is negligible. Finally, for evaluation purposes, we use timely but nonprecise integrity verification, that is, each cache block is immediately verified as soon as it is brought on chip, but we do not delay the retirement of the instruction that brings the block on chip if verification is not completed yet. Note that all of our schemes (AISE and BMT) are compatible with both nonprecise and precise integrity verification.

Fig. 8. Performance overhead comparison of AISE with BMT versus global counter scheme with traditional Merkle tree.

Fig. 9. Performance overhead comparison of AISE versus the global counter scheme.

7. EVALUATION

To evaluate our approach, we first present simulations results for the performance of our AISE and BMT schemes on a uniprocessor system. Next, we show the results of our schemes on CMP systems, and finally we present several sensitivity studies.

7.1 Uniprocessor Evaluation Results

In our first experiment, we compare AISE+BMT to another memory encryption and integrity verification scheme, which can provide the same type of system-level support as our approach (e.g., shared memory IPC, virtual memory support, etc.). Figure 8 shows these results of AISE+BMT compared to the 64-bit global counter scheme plus standard Merkle tree protection (global64+MT), where the execution time overhead is shown normalized to a system with no protection. While the two schemes offer similar system level benefits, the performance benefit of our AISE+BMT scheme is tremendous. The average execution time overhead of global64+MT is 25.9% with a maximum of 151%, while the average for AISE+BMT is a mere 1.8% with a maximum of only 13%. This figure shows that our AISE+BMT approach overwhelmingly provides the best of both worlds in terms of support of system-level issues and performance overhead reduction, making it more suitable for use in real systems.

To better understand the results from the previous figure, we next present figures which break the overhead into encryption vs. integrity verification components. Figure 9 shows the normalized execution time overhead of AISE

Fig. 10. Performance overhead comparison of AISE with our BMT vs. AISE with the Standard Merkle tree.

compared to the global counter scheme with 32-bit and 64-bit counters (note that only encryption is being performed for this figure). As the figure shows, AISE by itself is significantly better from a performance perspective than the global counter scheme (1.6% average overhead vs. around 4% and 6% for 32- and 64-bit global counters). Recall also that 64-bit counters, which should be used to prevent frequent entire-memory re-encryptions [Yan et al. 2006], require a 12.5% memory storage overhead. Note that we do not show results for counter-mode encryption using address plus block counter seeds, since the performance will be essentially equal to AISE if same-sized block counters are used. Since AISE supports important system level mechanisms not supported by address-based counter-mode schemes, and since the performance and storage overheads of AISE are superior to the global counter scheme, our AISE approach is an attractive memory encryption option for secure processors.

To see the overhead due to integrity verification, Figure 10 shows the overhead of AISE only (the same as the AISE bar on the previous figure), AISE plus a standard Merkle tree (AISE+MT), and AISE plus our BMT scheme (AISE+BMT). Note that we use AISE as the encryption scheme for all cases so that the extra overhead due to the different integrity verification schemes is evident. Our first observation is that integrity verification due to maintaining and verifying Merkle tree nodes is the dominant source of performance overhead, which agrees with other studies [Rogers et al. 2006; Yan et al. 2006]. From this figure, it is also clear that our BMT approach outperforms the standard Merkle tree scheme, reducing the overhead from 12.1% in AISE+MT to only 1.8% in AISE+BMT. Even for memory intensive applications, such as art, mcf, and swim, the overhead using our BMT approach is less than 15%, while it can be above 60% with the standard Merkle tree scheme. Also, for every application except for swim, the extra overhead of AISE+BMT compared to AISE is negligible, indicating that our BMT approach removes almost all of the performance overhead of Merkle tree-based memory integrity verification. We note that Yan et al. [2006] also obtained low average overheads with their memory encryption and integrity verification approach; however, for more memory-intensive workloads, such as art, mcf, and swim, their performance overheads still approached 20% and they assumed a smaller, 64-bit MAC size. Since our BMT

Fig. 11. L2 cache pollution.

scheme retains the security strength of standard Merkle tree schemes, the improved performance of BMTs is a significant advantage.

To understand why our BMT scheme can outperform the standard Merkle tree scheme by such a significant amount, we next present some important supporting statistics. Figure 11 measures the amount of “cache pollution” in the L2 cache due to storing frequently accessed Merkle tree nodes along with data. The bars in this figure show the *average* portion of L2 cache space that is occupied by data blocks during execution. For the standard Merkle tree, we found that on average data occupies only 68% of the L2 cache, while the remaining 32% is occupied by Merkle tree nodes. In extreme cases (e.g., art and swim), almost 50% of the cache space is occupied by Merkle tree nodes. Note that for 128-bit MACs, the main memory storage overhead incurred by Merkle tree nodes stands at 25%, so if the degree of temporal locality of Merkle tree nodes is equal to data, then only 25% of the L2 cache should be occupied by Merkle Tree nodes. Thus, it appears that Merkle tree nodes have a higher degree of temporal locality than data. Intuitively, this observation makes sense because for each data block that is brought into the L2 cache, one or more Merkle tree nodes will be touched for the purpose of verifying the integrity of the block. With our BMT approach, on the other hand, data occupies 98% of the L2 cache, which means that the remaining 2% of the L2 cache is occupied by BMT nodes. This explains the small performance overheads of our AISE+BMT scheme. Since the ratio of the size of a counter to a data block is 1:64, the footprint of the BMT is very small, so as expected it occupies an almost negligible space in the L2 cache. Furthermore, since data block MACs are not cached, they do not take up L2 cache space.

Next, we look at the (local) L2 cache miss rate and bus utilization of the base unprotected system, the standard Merkle tree, and our BMT scheme, shown in Figure 12. The figure shows that while the L2 cache miss rates and bus utilization increase significantly when the standard Merkle tree scheme is used (average L2 miss rate from 37.8% to 47.5%, bus utilization from 14% to 24%), our BMT scheme only increases L2 miss rates and bus utilization slightly (average L2 miss rate from 37.8% to 38.5% and bus utilization from 14% to 16%). These results show that the impact of reduced cache pollution from Merkle tree nodes results in a sizable reduction in L2 cache miss rates and bus

(a) L2 cache miss rate

(b) Bus utilization

Fig. 12. L2 cache miss rate and bus utilization of an unprotected system, standard Merkle tree, and our BMT scheme.

utilization, and thus, the significant reduction of performance overheads seen in Figure 10.

7.2 CMP Evaluation Results

In our first experiment on a CMP, we compare the overheads of our scheme (AISE+BMT) to the 64-bit global counter scheme with standard Merkle tree protection (global64+MT). Figure 13 shows the percentage degradation in the combined IPC of both benchmarks that run on different cores, for global64+MT and AISE+BMT, relative to a base system with no protection. While the two schemes offer similar system level benefits, the performance overheads of our scheme are significantly lower. The average execution time overhead of global64+MT is 24.3%, while the average for AISE+BMT is a 4.1%. Hence, as in the uniprocessor case, our scheme provides the best of both worlds in terms of supporting critical system features and low performance overheads.

Next, we compare the overheads of our AISE scheme in conjunction with a standard Merkle tree (AISE+MT) to AISE with our BMT scheme (AISE+BMT). Figure 14 shows the percentage degradation in combined IPC for AISE+MT and AISE+BMT relative to a base system with no protection. The combined IPC for the benchmark pair is calculated as the harmonic mean of the IPCs of the individual benchmarks when run together as separate threads on each core of our modeled CMP system. As can be seen, AISE+BMT maintains a large advantage over AISE+MT for CMPs, reducing the average IPC degradation from 15.1% to 4.1%. In addition, several benchmark pairs, such as *mcf.art*,

Fig. 13. IPC degradation comparison of AISE with our BMT versus 64-bit global counter scheme with the standard Merkle tree.

Fig. 14. IPC degradation comparison of AISE with our BMT versus AISE with the standard Merkle tree.

swim_art, and *mcf_apsi*, suffer from large IPC degradations of 35% or more. While no benchmark pair suffers an IPC degradation of more than 20% with our AISE+BMT scheme, and only three benchmark pairs suffer more than 10% degradation. This is important since CMP systems [Barroso et al. 2000; Laudon and Spracklen 2007; Sinharoy et al. 2005] are likely to find widespread applications in server settings [Marty and Hill 2007]. In such environments, excessive overheads due to security mechanisms are not likely to be tolerated. In addition to the lower average overhead, AISE+BMT achieves a much lower standard deviation of overheads of 5%, compared to that of AISE+MT, which has a standard deviation of overheads of 14.5%. The lower average and standard deviation of overheads give server vendors more confidence about the stability of performance of their systems when they use AISE+BMT.

We now present supporting statistics, similar to those shown in our evaluation for uniprocessor systems, to understand why BMT outperforms the standard Merkle tree scheme in CMP systems. Figure 15, shows the average portion of L2 cache space that is occupied by data blocks during execution. For the standard Merkle tree scheme, on an average, data occupies 63% of the L2 cache space and the remaining 37% is occupied by Merkle tree nodes. On the

Fig. 15. L2 cache pollution.

(a) L2 cache miss rate

(b) Bus utilization

Fig. 16. L2 cache miss rate and bus utilization of an unprotected system, standard Merkle tree, and our BMT scheme.

other hand, for our BMT scheme data occupies 98.2% of the L2 cache space. This explains the small overheads of AISE+BMT even in a CMP system.

Finally, Figure 16 shows the L2 cache miss rate (local) and the off-chip bus utilization of AISE+MT and AISE+BMT compared to a base system with no protection. The figure shows that L2 miss rate and bus utilization increases

Fig. 17. Performance overhead comparison across MAC size.

significantly for AISE+MT (average L2 miss rate from 26% to 36% and bus utilization from 20% to 41%). On the other hand, for AISE+BMT, the L2 cache miss rate and bus utilization increase only slightly (average L2 miss rate virtually unchanged at 26.3% and bus utilization from 20% to 21.4%). These results are attributed primarily to the reduced L2 cache pollution from the Merkle tree nodes and explain the significant reduction in IPC degradation.

7.3 Sensitivity Studies

In this section, we present two sensitivity studies for the uniprocessor environment. In the first case study, we examine the sensitivity of the standard Merkle tree (MT) and our BMT schemes to MAC size variations. In the second case study, we examine the sensitivity of the MT and BMT schemes to cache size variations.

7.3.1 Sensitivity to MAC Size. The level of security of memory integrity verification increases as the MAC size increases, since collision rates decrease exponentially with every 1-bit increase in the MAC size. Security consortiums, such as NIST, NESSIE, and CRYPTREC, have started to recommend the use of longer MACs, such as SHA-256 (256-bit) and SHA-384/512 (512 bits). However, it is possible that some uses of secure processors may not require a very high cryptographic strength, relieving some of the performance burden. Hence, Figure 17 shows both the average execution time overhead and fraction of L2 cache space occupied by data across MAC sizes, ranging from 32 bits to 256 bits. The figure shows that as the MAC size increases, the execution time overhead for MT increases almost exponentially from 3.9% (32-bit) to 53.2% (256-bit). In contrast, for BMT, the overhead remains low, ranging from 1.4% (32-bit) to 2.4% (256-bit). The overheads are related to the amount of L2 cache available to data, which is reduced from 89.4% (32-bit) to 36.3% (256-bit) for MT, but is only reduced from 99.5% (32-bit) to 94.9% (256-bit) for our BMT. Overall, it is clear that while large MACs cause serious performance degradation in standard Merkle trees, they do not cause significant performance degradation for our enhanced BMT scheme.

7.3.2 Sensitivity to Cache Size. Figure 18 shows the average performance overheads for L2 cache sizes of 512 KB, 1 MB, and 2 MB. The figure shows that as the cache size increases, the average overheads for both the schemes decreases,

Fig. 18. Sensitivity to cache size.

with standard Merkle tree benefitting more than BMT. This is expected as in the standard Merkle tree scheme, the Merkle tree nodes cause thrashing of data blocks and an increased cache size helps reduce this thrashing, whereas our BMT scheme which has very little thrashing to begin with. We make two other important observations from this figure. One, BMT overheads are stable across cache sizes (from almost negligible for a cache size of 2 MB to 2.4% for a cache size of 512 KB). On the other hand, standard Merkle tree overheads vary significantly (from 2.3% for a cache size of 2 MB to 17.1% for a cache size of 512 KB). Secondly, the overheads of the standard Merkle tree with a cache size of 2 MB (2.3%) are the same as BMT overheads with a much smaller cache size of 512 KB (2.4%). To summarize, BMT offers performance stability across cache sizes, even in memory constrained environments. With the growth of number of cores on a chip in future CMPs, it is important for a security scheme to achieve small overheads even when the cache size per core is relatively small.

7.4 Storage Overheads in Main Memory

An important metric to consider for practical implementation is the required total storage overhead in memory for implementing a memory encryption and integrity verification scheme. For our approach, this includes the storage for counters, the page root directory, and MAC values (Merkle tree nodes and per-block MACs). The percentage of total memory required to store each of these security components for the two schemes: global64+MT and AISE+BMT across MAC sizes varying from 32-bits to 256-bits is shown in Table I.

Since each data block (64B) requires effectively 8-bits of counter storage (one 7-bit block counter plus 1-bit of the LPID), the ratio of counter to data storage is only 1:64 (1.6%) versus 1:8 (12.5%) if 64-bit global counters are used. This counter storage would occupy 1.23% of the main memory of the secure processor with 128-bit MACs. The page root directory is also small, occupying 0.31% of main memory with 128-bit MACs. The most significant storage overhead comes from Merkle tree nodes, which grow as the MAC size increases. The traditional Merkle tree suffers the most, with overhead as high as 25% of the main memory with 128-bit MACs and 50% for 256-bit MACs. The overhead for our BMT is both smaller and increases at a much slower rate as the MAC size increases (i.e., 20% overhead for 128-bit MACs and 33% for 256-bit MACs). The reason our BMT still has significant storage overheads is because of the per-block MACs (BMT nodes themselves require a very small storage). These overheads are still significant; however, our scheme is compatible with several techniques proposed in Gassend et al. [2003] that can reduce this overhead, such as using a single

Table I. MAC and Counter Memory Overheads

		MT	Page Root	Counters	Total
256b	global64+MT	49.83%	0.35%	5.54%	55.71%
MAC	AISE+BMT	33.50%	0.51%	1.02%	35.03%
128b	global64+MT	24.94%	0.26%	8.31%	33.51%
MAC	AISE+BMT	20.02%	0.31%	1.23%	21.55%
64b	global64+MT	12.48%	0.15%	9.71%	22.34%
MAC	AISE+BMT	11.11%	0.17%	1.36%	12.65%
32b	global64+MT	6.24%	0.08%	10.41%	16.73%
MAC	AISE+BMT	5.88%	0.09%	1.45%	7.42%

MAC to cover not one block but several blocks. However, the key point here is that AISE+BMT is more storage-efficient than global64+MT irrespective of the MAC size used. AISE+BMT uses $1.6\times$ less memory compared to global64+MT with 256-bit MACs with the gap widening to $2.3\times$ with 32-bit MACs. Hence, our scheme maintains a distinct storage advantage over global64+MT across varying levels of security.

8. CONCLUSIONS

We have proposed and presented a new counter-mode encryption scheme, which uses AISE, and a new Bonsai Merkle Tree integrity verification scheme (BMT). AISE is compatible with general computing systems that use virtual memory and IPC, and it is free from other issues that hamper schemes associated with counter-based seeds. AISE can easily be extended to support the different variants of virtualization without requiring any changes to an AISE compliant OS and requiring minimal changes to existing VMMs. Despite the improved system-level support, with careful organization, AISE performs as efficiently as prior counter-mode encryption.

We also found that the Merkle tree does not need to cover the entire physical memory, but only the part of the memory that holds counter values. This discovery allows us to construct BMTs, which take less space in the main memory, but more importantly much less space in the L2 cache, resulting in a significant reduction in the overheads from 12.1% to 1.8% for single threaded SPEC 2000 benchmarks and from 15% to 4% for multi-threaded benchmarks, along with a reduction in storage overhead in memory from 33.5% to 21.5%.

REFERENCES

- ADAMS, K. AND AGESEN, O. 2006. A comparison of software and hardware techniques for x86 virtualization. In *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, New York, 2–13.
- AMD. 2005. AMD64 Virtualization Codenamed Pacifica Technology: Secure Virtual Machine Architecture Reference Manual.
- BARHAM, P., DRAGOVIC, B., FRASER, K., HAND, S., HARRIS, T., HO, A., NEUGEBAER, R., PRATT, I., AND WARFIELD, A. 2003. Xen and the art of virtualization. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles*. ACM, New York, 164–177.
- BARROSO, L. A., GHARACHORLOO, K., McNAMARA, R., NOWATZYK, A., QADEER, S., SANO, B., SMITH, S., STETS, R., AND VERGHESE, B. 2000. Piranha: A scalable architecture based on single-chip

- multiprocessing. In *Proceedings of the 27th International Symposium on Computer Architecture*. ACM, New York, 282–293.
- FIPS PUBLICATION 180-1. 1995. Secure Hash Standard. National Institute of Standards and Technology, Federal Information Processing Standards.
- FIPS PUBLICATION 197. 2001. Specification for the Advanced Encryption Standard (AES). National Institute of Standards and Technology, Federal Information Processing Standards.
- GAO, L., YANG, J., CHROBAK, M., ZHANG, Y., NGUYEN, S., AND LEE, H.-H. S. 2006. A low-cost memory remapping scheme for address bus protection. In *Proceedings of the 15th International Conference on Parallel Architectures and Compilation Techniques*. ACM, New York, 74–83.
- GARFINKEL, T., ROSENBLUM, M., AND BONEH, D. 2003. Flexible OS support and applications for trusted computing.
- GASSEND, B., SUH, G. E., CLARKE, D., VAN DIJK, M., AND DEVADAS, S. 2003. Caches and hash trees for efficient memory integrity verification. In *Proceedings of the 9th International Symposium on High Performance Computer Architecture*. IEEE Computer Society, Los Alamitos, CA, 295.
- GILMONT, T., LEGAT, J.-D., AND QUISQUATER, J.-J. 1999. Enhancing security in the memory management unit. In *Proceedings of the 25th EuroMicro Conference*. IEEE Computer Society, Los Alamitos, CA, 449–456.
- HAND, S. M. 1999. Self-paging in the Nemesis operating system. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation*. USENIX Association, Berkeley, CA, 73–86.
- HUANG, A. 2002. The Trusted PC: Skin-Deep Security. *IEEE Computer* 35, 10, 103–105.
- HUANG, A. 2003. *Hacking the Xbox: An Introduction to Reverse Engineering*. No Starch Press, San Francisco, CA.
- IBM. 2006. IBM Extends Enhanced Data Security to Consumer Electronics Products.
- INTEL CORPORATION. 2005. Intel Virtualization Technology Specifications for the IA-32 Intel Architecture.
- KGIL, T., FALK, L., AND MUDGE, T. 2004. Chiplock: Support for secure microarchitectures. *SIGARCH Comput. Architect. News* 33, 1, 134–143.
- KRAWCZYK, H., BELLARE, M., AND CANETI, R. 1997. HMAC: Keyed-hashing for message authentication. <http://www.ietf.org/rfc/rfc2104.txt>.
- KUMAR, A. 2004. Discovering passwords in memory. White Paper, Paladion Networks, Mumbai, India.
- LAUDON, J. AND SPRACKLEN, L. 2007. The coming wave of multithreaded chip multiprocessors. *Int. J. Parall. Program.* 35, 3, 299–330.
- LIE, D., MITCHELL, J., THEKKATH, C. A., AND HOROWITZ, M. 2003. Specifying and verifying hardware for tamper-resistant software. In *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE Computer Society, Washington, DC, 166.
- LIE, D., THEKKATH, C. A., MITCHELL, M., LINCOLN, P., BONEH, D., MITCHELL, J., AND HOROWITZ, M. 2000. Architectural support for copy and tamper resistant software. In *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, New York, 168–177.
- MARTY, M. R. AND HILL, M. D. 2007. Virtual hierarchies to support server consolidation. In *Proceedings of the 34th International Symposium on Computer Architecture*. ACM, New York, 46–56.
- OLUKOTUN, K., NAYFEH, B. A., HAMMOND, L., WILSON, K., AND CHANG, K.-Y. 1996. The case for a single-chip multiprocessor. *SIGOPS Operat. Syst. Rev.* 30, 5, 2–11.
- RENAU, J., TUCK, J., CEZE, L., STRAUSS, K., AND TORELLAS, J. 2004. SESC. <http://sesc.sourceforge.net>.
- ROGERS, B., CHHABRA S., SOLIHIN, Y. AND PRVULOVIC, M. 2007. Using address independent seed encryption and bonsai merkle trees to make secure processors os- and performance- friendly In *Proceedings of the 40th International Symposium on Microarchitecture*. IEEE Computer Society, Los Alamitos, CA, 183–196.
- ROGERS, B., SOLIHIN, Y., AND PRVULOVIC, M. 2006. Efficient data protection for distributed shared memory multiprocessors. In *Proceedings of the 15th International Conference on Parallel Architectures and Compilation Techniques*. ACM, New York, 84–94.
- SEMICONDUCTOR, M. 2007 (last modification). DS5002FP Secure Microprocessor Chip.

- SHI, W. AND LEE, H.-H. S. 2006. Authentication control point and its implications for secure processor design. In *Proceedings of the 39th Annual International Symposium on Microarchitecture*. IEEE Computer Society, Los Alamitos, CA, 103–112.
- SHI, W., LEE, H.-H. S., GHOSH, M., AND LU, C. 2004. Architectural support for high speed protection of memory integrity and confidentiality in multiprocessor systems. In *Proceedings of the 13th International Conference on Parallel Architectures and Compilation Techniques*. IEEE Computer Society, Los Alamitos, CA, DC, 123–134.
- SHI, W., LEE, H.-H. S., GHOSH, M., LU, C., AND BOLDYREVA, A. 2005. High efficiency counter mode security architecture via prediction and precomputation. *SIGARCH Comput. Architect. News* 33, 2, 14–24.
- SHI, W., LEE, H.-H. S., LU, C., AND GHOSH, M. 2004. Towards the issues in architectural support for protection of software execution. *SIGARCH Comput. Architect. News* 33, 1, 6–15.
- SINHARAY, B., KALLA, R. N., TENDLER, J. M., EICKEMEYER, R. J., AND JOYNER, J. B. 2005. Power5 system microarchitecture. *IBM J. Res. Dev.* 49, 4/5, 505–521.
- STANDARD PERFORMANCE EVALUATION CORPORATION. 2004. <http://www.spec.org>.
- SUH, G. E., CLARKE, D., GASSEND, B., VAN DIJK, M., AND DEVADAS, S. 2003a. Aegis: Architecture for tamper-evident and tamper-resistant processing. In *Proceedings of the 17th International Conference on Supercomputing*. ACM, New York, 160–171.
- SUH, G. E., CLARKE, D., GASSEND, B., VAN DIJK, M., AND DEVADAS, S. 2003b. Efficient memory integrity verification and encryption for secure processors. In *Proceedings of the 36th International Symposium on Microarchitecture*. IEEE Computer Society, Los Alamitos, CA, 339.
- WALDSPURGER, C. A. 2002. Memory resource management in vmware esx server. *SIGOPS Operat. Syst. Rev.* 36, SI, 181–194.
- WHITAKER, A., SHAW, M., AND GRIBBLE, S. D. 2002. Denali: Lightweight virtual machines for distributed and networked applications. In *Proceedings of the UNENIX Annual Technical Conference*.
- YAN, C., ROGERS, B., ENGLENDER, D., SOLIHIN, Y., AND PRVULOVIC, M. 2006. Improving cost, performance, and security of memory encryption and authentication. *SIGARCH Comput. Architect. News* 34, 2, 179–190.
- YANG, J., ZHANG, Y., AND GAO, L. 2003. Fast secure processor for inhibiting software piracy and tampering. In *Proceedings of the 36th Annual International Symposium on Microarchitecture*. IEEE Computer Society, Washington, DC, 351.
- ZHANG, Y., GAO, L., YANG, J., ZHANG, X., AND GUPTA, R. 2005. Senss: Security enhancement to symmetric shared memory multiprocessors. In *Proceedings of the 11th International Symposium on High-Performance Computer Architecture*. IEEE Computer Society, Los Alamitos, CA, 352–362.
- ZHUANG, X., ZHANG, T., LEE, H.-H. S., AND PANDE, S. 2004. Hardware assisted control flow obfuscation for embedded processors. In *Proceedings of the International Conference on Compilers, Architecture and Synthesis for Embedded Systems*. ACM, New York, 292–302.
- ZHUANG, X., ZHANG, T., AND PANDE, S. 2004. Hide: An infrastructure for efficiently protecting information leakage on the address bus. *SIGOPS Operat. Syst. Rev.* 38, 5, 72–84.

Received March 2008; revised September 2008; accepted September 2008