Energy-Security Tradeoff in a Secure Cache Architecture Against Buffer Overflow Attacks

Inoue, Koji

PRESTO, Japan Science and Technology Agency | Department of Informatics, Kyushu University

<http://hdl.handle.net/2324/6169>

出版情報 : Proc. of the The International Workshop on Architectural Support for Security and Anti-Virus (WASSA04), pp.77-85, 2004-10. The International Workshop on Architectural Support for Security and Anti-Virus

バージョン :

権利関係 :

Energy-Security Tradeoff in a Secure Cache Architecture Against Buffer Overflow Attacks

Koji Inoue^{† §}

[†]Department of Informatics, Kyushu University.
6-1 Kasuga-Koen, Kasuga, Fukuoka 816-8580 Japan

[§]PRESTO, Japan Science and Technology Agency,
4-1-8 Honcho Kawaguchi, Saitama 332-0012 Japan

inoue@c.csce.kyushu-u.ac.jp

ABSTRACT

In this paper, we propose a cache architecture, called *SCache*, to detect buffer-overflow attacks at run time. Furthermore, the energy-security efficiency of *SCache* is discussed. *SCache* generates replica cache lines on each return-address store, and compares the original value loaded from the memory stack to the replica one on the corresponding return-address load. The number and the placement policy of the replica line strongly affect both energy and vulnerability. In our evaluation, it is observed that *SCache* can protect more than 99.3% of return-address loads from buffer-overflow attacks, while it increases total cache energy consumption by about 23%, compared to a well-known low-power cache.

1. INTRODUCTION

As the popularity of mobile computing devices and the advance in internet information services, considering energy-security efficiency of computer systems becomes more important. Although the internet is a much useful instrument, it also gives an opportunity for attacking remote connected devices to malicious persons. On the other hand, reducing energy consumption is an inevitable design constraint for mobile devices such as laptop computers and cellular phones, because it affects directly the battery lifetime. Fundamentally, in order to defend own computer system, some extra tasks which do not contribute program-execution results are required, resulting in wasted energy consumption. However, only few attempts have so far been made at the tradeoff between energy and security.

In this paper, we propose a cache architecture, called *Secure Cache (SCache)*, to prevent buffer-overflow attacks. Moreover, we evaluate the energy overhead caused by *SCache* based on a 0.18 μ m SRAM design, and discuss the tradeoff between energy and security. The attackers attempt to alter the procedure return address by causing a buffer overflow which breaks the structure of memory stack. *SCache* detects the

return-address corruption without any software supports at run time. When a return address is pushed onto the memory stack, *SCache* generates one or more replicas of the return-address value, and saves them into the other cache space. Since on-chip caches give a large impact on both performance and energy consumption, researches have proposed a number of approaches to lowering cache energy [6][7][8][10]. In contrast to prior work, this paper focuses on the tradeoff between energy and security.

This paper is organized as follows: Section 2 explains briefly the mechanism of buffer overflow, and introduces related work to solve the buffer-overflow problem. Section 3 proposes *SCache* architecture, and the organization and operation are explained in detail. In Section 4, we evaluate vulnerability and energy consumption of several *SCache* models, and discuss the energy-security tradeoff. Finally, in Section 5, we conclude this paper.

2. STACK SMASHING ATTACK

In this section, we explain the mechanism of the stack smashing, and show related work to prevent the buffer-overflow attack.

2.1. Buffer-Overflow Vulnerability

To attack vulnerable computer systems, at least two processes have to be done: injecting an attack code and hijacking the program-execution control. The buffer overflow makes it possible to achieve both the issues simultaneously. This is one of the main reasons why the buffer-overflow vulnerability is commonly exploited for attacking target computers. For example, the malicious programs such as Code Red worm in 2001 and Braster in 2003, which raged in the world, utilize this defect. Figure 1 illustrates the percentage of CERT advisories relating to the buffer-overflow

Figure 1: buffer overflow advisories (from [9])

vulnerability for each year. We see from Figure 1 that 50% of advisories in 2001 relate to this weak point.

The buffer overflow is caused by writing an inordinately large amount of data into a buffer. Unfortunately, the C programming language does not perform automatically array-bound checks, and this defect mainly exists in the standard C library such as *strcpy()*. Therefore, many programs have the possibility to suffer from the buffer-overflow vulnerability.

The overflow breaks memory stack structure as depicted in Figure 2, and this operation is called *stack smashing*. In this figure, we assume that the function *f()* calls the function *g()*, which includes a vulnerable operation *strcpy()*, as shown in the rightmost figure. The state of the memory stack immediately after the function call *g()* is depicted in the leftmost figure. The stack consists of the function parameters, the return address to the caller, the previous frame pointer, and the local array variable *buf*. In *g()*, if the size of the string pointed by *s* is larger than the memory size allocated for the local variable *buf*, a buffer overflow takes place when the *strcpy()* function is executed. As a result, the contiguous stack contents are overwritten. If the pointer *s* points a malicious string which is meticulously constructed by an attacker, the attack code is injected into the memory stack and the return address is altered to the top of the injected code, as shown in the middle one in Figure 2. The corrupted return address is set to the program counter (PC) when the execution of *g()* completes. As a result, the program-execution control is finally hijacked by the injected attack code.

Figure 2: Stack Smashing

2.2. Related Work

The most straightforward way to solve the buffer-overflow attack is to prohibit the execution of codes stored in data segments. For example, AMD Athlon64 employs this protection. However, some programs attempt to generate an executable code at run time, e.g. just-in-time compiler. To support such operations, the microprocessors should be able to execute the instructions stored in a data segment.

So far, many techniques to address the buffer-overflow attacks have been proposed. They can be classified into two types: static and dynamic. The static approach generates a secure object code based on source code analysis. For example, the paper [11] formulates detection of buffer overflows as an integer range analysis problem in order to find the potential of stack smashing. *SASI* introduced in [4] inserts reference-monitor codes into application programs to observe program-execution behavior. For instance, a bound-checking code as a reference monitor may detect buffer-overflow attacks. *StackGuard*, which is a patch to gcc, is another static approach to defending the stack smashing[3]. Each return address is pushed onto the stack with a “canary word” which is a randomly generated value. The canary is allocated to the next stack entry of the return address, and its copy is also stored into a general-purpose register. Therefore, we can detect stack smashing by comparing the canary value read from the stack memory with that saved in the register. This approach stands on the assumption that the canary word is altered whenever a return address corruption takes place.

One of the main drawbacks of the static approach is code compatibility, because it requires a code translation or re-compilation. On the other hand, the

dynamic approach does not have this negative effect. A dynamically loadable library (DLL) called *libsafe* is used to check vulnerabilities before un-trusted standard C libraries are executed[1]. Similarly, *libverify* is another DLL which injects a verification code at the start of the process execution via a binary re-writing[1]. *StackGhost* exploits cleverly the register window of the SPARC architecture[5]. In such architecture, a return address is saved in the register window. Only when a register-window overflow occurs, the return address is stored into a memory stack by OS support. Although the memory stack access is potentially vulnerable, OS can insert some operation to protect the return address, e.g. implementing a canary word as well as StackGuard. Another approach called *SRAS* (*Secure Return Address Stack*) is an LIFO fashion small memory embedded in the microprocessor core[9]. SRAS is a straightforward but very efficient architectural support to prevent the buffer-overflow attacks. A return address is pushed onto not only the memory stack but also SRAS, and they are compared to detect stack smashing when the corresponding return instruction is executed. Our research has been started from this paper.

3. THE SCACHE ARCHITECTURE

3.1 Overview

order to achieve such a hardware protection, SCache attempts to make one or more copies of the return address when it is stored into the cache. We call a cache line including the copy of the return address a *replica line*. Actually, the replica line is not a complete copy of whole cache line. Only the return address is copied. Since the replica line can be placed in the cache set indexed by the reference address of the current return-address store, we can generate at most $Asso-1$ replica lines where $Asso$ is the cache associativity. Any cache accesses, except return-address stores, can not overwrite the replica lines. Therefore, the replica lines are treated as read only. When a return address is popped off the memory stack, SCache selects one of the replica data and compares it with the popped original return address. If they are exactly the same, we can ensure that the popped return address is safe. Otherwise, it means that a return-address corruption takes place, thus a signal to report security status is sent to the processor in order to terminate the current program execution.

Figure 3 shows the structure of a four-way set-associative SCache. Here, we assume that the number of replica lines to be generated for each return-address store (noted as N_{rep}) is two. Unlike conventional caches, one-bit replica flag (R_flag) is added to each tag entry. The R-flag is set to one if the corresponding cache line is a replica, otherwise it is reset to zero. Moreover, a multiplexer to select a replica data and a 32-bit comparator for examining return-address corruption are required. Figure 4 illustrates its operation on return-address load/store hits. When

a return address is pushed, the cache works as follows.

- W1.** With the same manner as conventional caches, the set indexed by the reference address is accessed for tag checking. The corresponding R-flags are also read in parallel.
- W2.** Then the return address is stored into the cache-hit line. We call the target line *the master line*. Namely, the master line includes the original value of the return address.
- W3.** If some replica lines already exist in the indexed set, the return address is stored into them in order to avoid a coherence problem. The replica lines can be detected by examining the tag and the R-flag at the step W1, because the corresponding replica lines have the same tag information and a valid R-flag.
- W4.** New replica lines are generated by writing the return-address value into other non-replica line(s) and setting the corresponding R-flag(s) to one. The tag and offset information of the original write operation performed at the step W2 is used to make the replica line(s).

Note that the example assumes a cache hit. However, even on a miss, the cache works in the same manner after the line replacement is completed. As explained in Section 3.1, SCache **does not make a complete copy of whole master line**, but just writes one word return-address value to the several lines in parallel. In other words, the contents of the replica lines are not exactly the same as those of the master line. In addition, a replica line can include several return addresses if the offset addresses are different. On the other hand, at a return-address load, the cache examines whether or not the popped return address is safe as follows.

- R1.** The cache activates all the ways to read cache lines, tags, and R-flags. Then, the master line, which has the matching tag and the invalid R-flag, is selected. The return-address data read from the master line is sent to the processor.
- R2.** The replica lines having the matching tag are searched in the referenced set. If there are several replica lines corresponding to the current return-address load, one of them is randomly selected. Otherwise, the cache reports to the processor that the loaded return address may be unsafe, and completes the current cache access.

Figure 4: Operation on return-address accesses

- R3.** The return address read from the master line is compared with that obtained from the selected replica line. For this examination, only one replica data is selected and compared with the original one, because all of the replica lines include the same copy of the original return-address value. If the 32-bit comparison result is not a match, the cache alarms to terminate the program execution due to the insecure return-address load.

In conventional write accesses, the tag checks and the data write (W1 and W2) are performed sequentially. Although SCache requires two more steps, W3 and W4, they can be executed in parallel with the step W2. For read accesses, SCache completes to read the target data at the step R1 as well as conventional caches. In addition, the hardware components for SCache, a replica-line MUX and a 32-bit word comparator, do not affect the cache critical paths. Accordingly, SCache does not worsen the cache-access time.

For normal accesses (i.e. non-return-address loads or stores), SCache operates as the same as conventional caches except that the **invalid** R-flag is included in the cache-hit condition. Therefore, the normal accesses do not modify the replica lines. For instance, let us consider a write operation to the next stack entry of the return address. Here, we assume that the write target entry is allocated to the same cache line as the return-address value. In this scenario, the write operation is performed only to the master line, which has the invalid R-flag. The master line does not have any protection mechanism, thus the write operation can be completed.

The microprocessor needs to output a signal to indicate whether the current access targets a return address. This can be easily achieved by checking the source (or destination) operand of the current memory reference [9]. For many microprocessors, the return address is located to a special register, e.g. R31.

3.3 Impact on Energy and Security

SCache attempts to protect return addresses by generating replica lines. However, for each return-address store, the cache needs to write the return-address value into several locations, i.e. one to the master line and one to each replica line, thereby increasing write-access energy. Furthermore, on return-address loads, the cache examines all the ways to find replica lines. Namely, SCache can not reduce read-access energy as well as already proposed low-power caches which attempt to avoid unnecessary way activation [6][10]. Another drawback of SCache is the energy overhead caused by degrading cache-hit rates. Since generating replica lines pollutes effective cache area, the energy consumed for cache-line replacements is increased.

On the other hand, from the security point of view, SCache can detect return-address corruption whenever at least one replica line exists. However, replica lines may be evicted from the cache, because they are also candidates for the cache-line replacement on misses. The most straightforward way to solve this issue is to prohibit evicting the replica lines from the cache. In this approach, we need to carefully release the replica lines with an appropriate timing. This is because a too-early-release degrades the efficiency of security, while a too-late-release pollutes unnecessarily the effective cache area. A simple strategy is to release the locked replica lines when the corresponding return-address load is issued to the cache. However, in this approach, still we have a possibility to take the too-early-releases. A number of microprocessors employ various predictions to gain performance, e.g. branch prediction, value prediction, etc. In such high-performance microprocessors, some instructions may be squished when a miss-prediction takes place. Therefore, the return-address load which is squished in later releases rashly the corresponding replica lines in the cache. Another approach to achieve secure operations is that we allow to evict the replica lines but try to make the replica lines reside in the cache as long as possible. We can consider at least two approaches: increasing the number of replica lines to be generated ($Nrep$) and employing the MRU

algorithm for replica-line placements. However, they affect negatively cache energy consumption and miss rates.

4. EVALUATION

In this section, we define six SCache models, and discuss the ability of proposed approach. The SCache models differ in the number and the placement policy of replica lines. First, we explain how the energy and vulnerability are evaluated. Next, we measure them for the SCache models, and discuss the energy-security tradeoff. The performance overhead caused by SCache is also evaluated.

4.1 Experimental Setup

We extended SimpleScalar tool set (ver.3.0d) [12] to support the SCache approach, and executed seven integer programs and four floating-point programs from the SPEC2000 benchmark suite [13]. The small input data set was used to complete the whole program execution. In this evaluation, we assume that the L1 data cache size is 16 KB, the line size is 32 B, and the associativity is 4. Furthermore, a 4-way superscalar out-of-order execution is assumed. For other parameters, we used the simplescalar default value defined in [2]. In this section, we refer to the issued return-address load (or store) to the cache as the *IRA load* (or *store*). We use the following equation to evaluate the efficiency for security,

$$Vulnerability = (Nv-rald / Nrald) * 100,$$

where $Nrald$ and $Nv-rald$ are the total number of IRA loads in the program execution and that of insecure IRA loads (i.e. return-address loads without any replica line), respectively. On the other hand, to evaluate the energy overhead, we use the energy model,

$$E_{total} = E_{rd} + E_{wt} + E_{wb} + E_{mp},$$

where E_{rd} and E_{wt} are the total energy consumed for cache read and write accesses, respectively. E_{wb} is another energy overhead caused by evicting dirty lines from the cache due to replica placement. E_{mp} is the energy dissipated for cache-line replacements. We assume that the energy consumed for a next-level memory access is ten times larger than that for an L1 cache read access. Based on a 0.18 μ m CMOS technology, we designed a 4KB SRAM array and estimated energy consumption. The circuits have been optimized to meet 3.0 ns access time. After the layout, we have measured energy consumption by performing

Table 1: Cache-Miss Rates

Model Bench	#IRA Load(<i>Nrald</i>)	CONV	LRU1-NR	LRU1	LRU2	MRU1	MRU2	ALL
164.zip	4,930,467	5.22%	5.23%	5.22%	5.22%	5.22%	5.23%	5.25%
175.vpr	5,627,709	3.53%	3.59%	3.56%	3.63%	3.59%	3.66%	3.74%
176.gcc	37,519,156	4.26%	6.06%	4.29%	4.37%	4.33%	4.43%	4.64%
181.mcf	992,419	20.02%	20.05%	20.02%	20.03%	20.05%	20.06%	20.10%
197.parser	45,466,527	4.13%	4.25%	4.18%	4.44%	4.23%	4.55%	5.07%
255.vortex	22,101,265	1.75%	1.83%	1.79%	1.91%	1.82%	1.94%	2.32%
256.bzip	18,147,017	2.31%	2.31%	2.31%	2.32%	2.31%	2.32%	2.45%
177.mesa	4,727,396	0.14%	0.15%	0.15%	0.16%	0.15%	0.16%	1.08%
179.art	32,466	42.93%	42.93%	42.93%	42.93%	42.93%	42.93%	42.93%
183.equake	3,580,827	2.44%	2.45%	2.44%	2.46%	2.45%	2.47%	2.52%
188.ammp	6,307,839	36.27%	36.29%	36.28%	36.31%	36.28%	36.30%	36.38%

IRA: Issued Return Address, CONV: Conventional

Hspice circuit simulations with extracted load capacitances. First, we obtained the energy for accessing 1-bit memory cell that includes sensing and pre-charging. Next, we calculated the average energy for each operation based on the total number of bits to be accessed. Then we multiplied the energy by the number of events occurred during the program execution.

Since the cache associativity is assumed as four, we can generate at most three replica lines for each return-address store, i.e. $N_{rep} = 3$. Furthermore, there are two options for the replica line placement, LRU and MRU. Here, we define six SCache models: *LRU1-NR*, *LRU1*, *LRU2*, *MRU1*, *MRU2*, and *ALL*. LRU1 and LRU2 generate one and two replica line(s) with the LRU placement algorithm, respectively. MRU1 and MRU2 place replica line(s) on the MRU location except the master line. The ALL model makes the maximum number of replica lines. In these models, the replica lines are treated as the same as normal lines, i.e. they can be evicted from the cache. On the other hand, LRU1-NR prohibits evicting the replica lines, and they are released when the corresponding return-address load is issued to the cache. We compare the SCache models with a conventional low-power way-predicting cache, noted as *CONV*. This model attempts to activate only the hit way which includes the reference data by employing an MRU-base way prediction [6], thereby saving

Figure 5: Vulnerability

cache-access energy. Here, we do not take the energy overhead caused by the way prediction, e.g. accessing an MRU table, into account. Note that the SCache models also perform the MRU-base way prediction. However, on each return-address load, they need to activate all the ways in spite of the correct way-prediction, as explained in Section 3.3.

4.2 Vulnerability

Figure 5 shows the vulnerability of the SCache models. We should notice that conventional caches without any consideration for stack smashing have 100% of vulnerability. The number of IRA loads, *Nrald*, and cache-miss rates are also presented in Table 1.

Figure 6: Energy Consumption

First, we discuss the effects of the number of replica lines, N_{rep} . In this simulation, we assumed that the LRU policy is employed for the cache-line replacement on misses. Therefore, replica lines generated in LRU1 are easily evicted from the cache due to conflicts. As we expected, LRU2 produces better results than LRU1 for all benchmarks due to the increased number of replica lines. We see the same situation for the MRU based SCache, MRU1 and MRU2. If they can perform a strict MRU placement, the number of replica lines does not affect the vulnerability. However, when a return address is stored into the cache, all existing replica lines with the same tag information are updated in order to avoid coherence problem, as explained in Section 3.2. If a corresponding replica line already exists at the LRU location, MRU1 works as the same as LRU1. Therefore, MRU1 is more vulnerable than MRU2. For all but *181.mcf* and *179.art*, the most secure model ALL can protect more than 99.3% of IRA loads from stack smashing.

Next, we discuss the impact of the replica-line placement algorithm. The MRU strategy constantly achieves higher security than the LRU models if the number of replica lines to be generated at each return-address store is the same. This is because the MRU placement makes the replica lifetime longer as well as increasing the number of replica lines. However,

against our expectation, LRU1-NR does not work well for some benchmarks. One of the reasons of this result is a hasty release of replica data caused by squished return-address loads.

4.3 Energy Consumption

Figure 6 shows energy consumption and its breakdown for the SCache models. All results are normalized to *CONV*. From the figure, we see that increasing the number of replica lines worsens energy efficiency. The ALL model increases energy consumption by about 23% in the worst case, *197.parser*. On the other hand, replica-line placement algorithm does not give a large impact on energy.

The SCache models have the same energy overhead for read accesses. Since all of the ways in SCache are activated on each IRA load, the read-access energy depends on not the number of replica lines generated but the total number of IRA loads. In contrast, the energy dissipated for write accesses *Ewt* increases with the increase in the number of replica lines generated. This situation can be seen for *Emp* due to the increased cache-miss rates. For instance, in case of *177.mesa*, the ALL model worsens the cache-miss rate from 0.14% to 1.08%, thereby increasing the energy for cache-line replacements *Emp*. For the programs with small energy overhead, *181.mcf* and *179.art*, it is observed that the increase in *Erd*, *Ewt*,

Figure 7: Energy-Security Tradeoff

and *Emp* are trivial. For these benchmarks, the total energy is originally dominated by *Emp*, 51% for *181.mcf* and 62% for *179.art*, due to the higher cache-miss rates as reported in Table 1. Therefore, the energy overhead caused by the replica lines is relatively hidden.

4.4 Energy-Security Tradeoff

In order to evaluate both energy and security at the same time, we introduce the energy-vulnerability product (EVP). We multiply the energy overhead by the number of vulnerable IRA loads. Furthermore, we measure the $E^2 \cdot V$ (or $E \cdot V^2$) product in order to consider more energy-oriented (or security-oriented) applications. Figure 7 shows the evaluation results for the five SCache models (LRU1-NR is not included). All the results are normalized to LRU1. From the figure, it is observed that each SCache model shows different characteristics. For instance, MRU1 produces the best performance if we see the energy-oriented metric $E^2 \cdot V$, while MRU2 or ALL give better results for security-oriented applications $E \cdot V^2$. This means that there is a tradeoff between energy and security, thus it is very important to explore the design space for coping both high security and low energy consumption. For the SCache approach, we conclude that MRU1 which increases the energy consumption only by 10% in the worst case but achieves relatively higher security should be selected if energy consumption is the primary design constraint. In contrast, MRU2 or ALL can protect more than 99% of IRA loads for many benchmarks, thus they are suitable for security-oriented applications.

Figure 8: Performance

4.5 Performance Overhead

Finally, we evaluate the impact of the SCache approach on processor performance. As shown in Table 1, increasing the number of replica lines, N_{rep} , worsens cache-hit rates, thus the processor performance will be degraded. This negative effect appears clearly on the MRU-based models. Figure 8 reports the performance overhead caused by the SCache scheme. For the ALL model, the increase in execution time is at most 1.1%, *177.mesa*. Furthermore, we see that in many cases the performance overhead is less than 0.4%. Therefore, we believe that the performance degradation caused by SCache is negligible.

5. CONCLUSIONS

In this paper, we have proposed a secure cache architecture, called SCache. The cache makes it possible to detect stack smashing at run time. The stack smashing alters a function return address for transferring the program-execution control to an injected malicious code. By making one or more replica lines in the large cache area, we can protect the return address. We have also considered an energy-security tradeoff by evaluating energy and vulnerability of six SCache models which differ in the number and the placement algorithm of replica lines. As a result, it has been observed that the MRU1 model is good for energy-oriented applications, while the MRU2 and ALL models are suitable for security-oriented applications.

In this evaluation, we have estimated energy consumption based on an SRAM layout design which does not include peripheral circuits for SCache, e.g. a selector for replica lines, control logic, and so on. Our ongoing work is to design a complete SCache core. Another future work is to explore the SCache design space with various cache configurations, and to establish an optimization technique to find the best point to maximize the energy-security efficiency.

ACKNOWLEDGMENTS

I would like to thank Prof. Shingi Tomita, Prof. Hiroto Yasuura, and all other members of PREST “information infrastructure and applications” research group for discussing at technical meetings. The VLSI chip in this study has been fabricated in the chip fabrication program of VLSI Design and Education Center (VDEC), the University of Tokyo in collaboration with Hitachi Ltd. and Dai Nippon Printing Corporation. This research was supported in part by the Grant-in-Aid for Creative Basic Research, 14GS0218, and for Encouragement of Young Scientists (A), 14702064.

REFERENCES

- [1] A.Baratloo, N.Singh, and T.Tsai, “Transparent Run-Time Defense Against Stack Smashing Attacks,” Proc. of 2000 USENIX Annual Technical Conference, June 2000.
- [2] D.Burger and T.M.Austin, “The SimpleScalar Tool Set, Version 2.0,” Univ. of Wisconsin-Madison Computer Sciences Department Technical Report #1342, June, 1997.
- [3] C.Cowan, C.Pu, D.Maier, H.Hinton, J.Walpole, P.Bakke, S.Beattie, A.Grier, P.Wagle, and Q.Zhang, “StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks,” Proc. of 7th USENIX Security Symposium, Jan, 1998.
- [4] U.Erlingsson and F.B.Schneider, “SASI Enforcement of Security Policies: A Retrospective,” Proc. of the workshop on New security paradigm, 1999.
- [5] M.Frantzen and M.Shuey, “StackGhost: Hardware Facilitated Stack Protection,” Proc. of the 10th USENIX Security Symposium, Aug. 2001.
- [6] K.Inoue, T.Ishihara, and K.Murakami, “Way-Predicting Set-Associative Cache for High Performance and Low Energy Consumption,” Proc. of the Int. Symp.on Low Power Electronics and Design, pp. 273--275, Aug. 1999.
- [7] M.B.Kamble and K.Ghose, “Analytical Energy Dissipation Models For Low Power Caches,” Proc. of the Int. Symp. on Low Power Electronics and Design, pp.143--148, Aug. 1997.
- [8] J.Kin, M.Gupta, and W.H.Mngione-Smith, “The Filter Cache: An Energy Efficient Memory Structure,” Proc. of the 30th Int. Symp. on Microarchitecture, pp.184--193, Dec. 1997.
- [9] R.B.Lee, D.K.Karig, J.P.McGregor, and Z.Shi, “Enlisting Hardware Architecture to Thwart Malicious Code Injection,” Proc. of the Int. Conf. on Security in Pervasive Computing, Mar. 2003.
- [10] M.D.Powell, A.Agarwal, T.N.Vijaykumar, B.Falsafi, and K.Roy, “Reducing Set-Associative Cache Energy via Way-Prediction and Selective Direct-Mapping,” Proc. of the 34th Int. Symp. on Microarchitecture, pp.54--65, Dec. 2001.
- [11] D.Wagner, J.S.Foster, E.A.Brewer, and A.Aiken, “A First Step Towards Automated Detection of Buffer Overrun Vulnerabilities,” Proc. of the Network and Distributed System Security Symposium, Feb. 2000.
- [12] SimpleScalar Tool Sets, <http://www.simplescalar.com/>.
- [13] SPEC(Standard Performance Evaluation Corporation), <http://www.specbench.org/>