Speculative Virtual Verification: Policy-Constrained Speculative Execution

Michael E. Locasto Stelios Sidiroglou Angelos D. Keromytis
Network Security Lab, Department of Computer Science, Columbia University
{*locasto, stelios, angelos*}*@cs.columbia.edu*

Abstract

A key problem facing current computing systems is the inability to autonomously manage security vulnerabilities as well as more mundane errors. Since the design of computer architectures is usually performance-driven, hardware often lacks primitives for tasks in which raw speed is not the primary goal. There is little architectural support for monitoring execution at the instruction level, and no mechanisms for assisting an automated response.

We propose *speculative virtual verification* (SVV), a model for the speculative execution of code based on complex security and safety constraints. We introduce architectural enhancements to support this framework, including the ability to supply an automated response by rewriting the instruction stream. Finally, since SVV is a novel approach to executing software, we briefly consider some important challenges for SVV-based systems.

1 Introduction

Software faults and vulnerabilities continue to pose significant problems to achieving reliable and secure software. The lack of comprehensive and low-cost protection mechanisms presents a critical problem for computing systems. It is unlikely that static analysis techniques or improved programming practices will provide a complete solution to the types of errors that threaten system stability or create exploitable vulnerabilities. Even systems that dynamically monitor process execution often impose a noticeable performance cost. Furthermore, these systems may reinvent the same primitives because the hardware does not supply them. However, even if such capabilities existed, system security is often a matter of *policy*; these utilities would need some level of flexibility to be applicable and remain useful in a wide variety of diverse and evolving environments. Finally, systems lack the capability to respond intelligently to both attacks and non-malicious faults.

The ability for computing systems to detect and correct faults and vulnerabilities would greatly improve their stability and security. The main contribution of this paper is the proposal of a set of architectural components that provide a basis for such systems by speculatively executing the entire instruction stream. In much the same way that a superscaler processor speculatively executes past a branch instruction and discards the mis-predicted code path, we propose that processors operate on the instruction stream in two phases. The first phase executes instructions, optimistically "speculating" that the results of these computations are benign. The second phase makes the effects of the speculated instruction stream visible to the OS and application software layers and potentially rewrites the instruction stream if it has been deemed harmful.

1.1 Speculative Execution

Speculative execution is a technique used in microprocessors to execute the instructions in a code branch before the evaluation of the branch conditional is finished. The need to perform speculative execution arises in pipelined processors because the conditional instruction that the branch depends on has proceeded deeply into the pipeline but has not

been evaluated by the time the processor is ready to fetch additional instructions. While a complete discussion of the strategies for dealing with branch predication is beyond the scope of this paper, a basic overview of the subject and pointers to other material are available in [10, 6]. Our proposal differs from these techniques by introducing another layer of speculative execution in which the acceptance of a particular execution path is not based on the evaluation of a branch conditional, but rather a higher-order constraint on a set of instructions.

1.2 Motivation and Feasibility

We are motivated by work on constructing an emulator [23] to supervise program execution in response to exploits and errors. Unfortunately, using an emulator incurs a sizable performance hit because the emulator executes every program instruction in software. The first way to ease this burden, which was adopted in [23], is to only emulate the portions of the program demonstrated to be vulnerable, thereby reducing the time that is spent in the emulator. The second approach is to eliminate the emulation penalty altogether by executing the process directly on the CPU. Unfortunately, adopting this approach currently means relinquishing the monitoring capabilities that the emulator provides. Therefore, we advocate adding monitoring mechanisms to processors so that a certain level of safety is relatively inexpensive. In order to address more complex attacks, we also propose that execution can be delegated to the software emulator if needed.

Arguing for the widespread adoption of fundamental changes to hardware is a controversial proposition. We believe the hardware necessary to support our system is easily implementable. Indeed, large parts of the system are already present in modern processors to support thread level speculation (TLS). The design parameters of general-purpose microprocessors have traditionally been driven by raw performance. Instead, we advocate design parameters aimed at more high-level feature support.

2 System Design

The core features of SVV form a two level monitoring environment. The first level includes hardware mechanisms for monitoring instruction execution (bounds checking, taint-tracking [26], SRAS [13], transfer control validation [12], *etc.*). The second level of monitoring is provided by the Policy Constraint Unit (PCU) and the Virtual Emulator Registration Unit (VERU). Instructions are filtered by the PCU according to some policy constructed by the programmer, compiler, or runtime profiling. The policy could range from filtering on a particular class of instructions (integer vs. load/store) to more complex constraints that require keeping state. The design of a constraint language to express these policies is future work. The VERU stores an address for code that should be executed if the PCU identifies a sequence of instructions that require more resources than the hardware can easily provide. Finally, the Verification Buffer (VB) and the Instruction Rewrite Unit (IRWU) provide some basic support for an automatic response capability.

2.1 SVV Execution Model

The execution model for SVV (see Figure 1) is similar to normal program execution. Instructions are fetched, decoded, issued to functional units (possibly out of order), executed, and gathered in a reorder buffer (ROB) to be committed in program order. However, at each stage, instructions are filtered by the PCU and monitored by hardware-level security mechanisms. Additionally, the VB accumulates completed instructions as they leave the ROB and commits them only if they pass the monitoring tests.

Instructions proceed through the pipeline in three major scenarios. First, the instruction may be harmless. In this case, it proceeds normally to the ROB, graduates when appropriate, moves to the VB, and is committed. Second, an instruction may be harmful as determined by the monitoring mechanisms (*e.g.*, it is actually tainted input data, or will write input data to the code area of the process address space) or the PCU. In this case, the IRWU flushes the scope of the harmful instruction and constructs a 'safe' version of the flushed code. The proces-

sor then executes this alternate instruction stream, including a return to the normal path of execution. The third scenario enables an emulator to be loaded on the CPU and supervise code execution. If the PCU decides that a particular sequence of instructions requires more complex supervision, it can invoke execution of this emulator. Note that there is no requirement for the software invoked by the VERU to be an emulator. The VERU simply holds an address and transfers control to the code at this address. Such an approach enables a more general response mechanism than software emulation. For example, the VERU may transfer control to an OS routine that kills the process, or suspends the process and transfers it to an isolated host for analysis, auditing, intrusion detection, or debugging.

2.2 Scope of SVV

The single largest obstacle for SVV is a three part problem and involves determining the scope of supervision. First, even though SVV is meant to run continuously, some applications (especially those working in a power-constrained environment) may wish to avoid the overhead associated with constant monitoring. In addition, hardware is fundamentally limited in the number of virtual execution contexts it can support concurrently. Finally, it is likely that the basic monitoring mechanisms, while capable of stopping large classes of attacks, may be unable to cope with more sophisticated attacks (some forms of DoS, multi-step attacks, information leaks, improperly set permissions, phishing attacks, *etc.*) or analysis tasks that require a large amount of state (anomaly or intrusion detection via data mining).

To address the latter two problems, we use the VERU to register a software emulator that can perform high-level monitoring of an instruction stream. An emulator has the flexibility to be more intrusive and is easily customizable. This hybrid approach to monitoring is more promising than an approach based solely on hardware or software. To address the first problem, SVV can be selectively invoked. Control over this invocation can be handled by the OS (a new system call to invoke or halt the SVV hardware) or the compiler (new assembly instructions can delimit an SVV monitored code region).

2.3 Automated Response

Automating a response strategy is difficult, as it is often unclear what a program should do in response to an error or attack. A response system is forced to anticipate the intent of the programmer, even if that intent was not well expressed or even well-formed. Ideal computing systems would recover from attacks and errors without human intervention. However, the state of the art is far from mature, and most existing response mechanisms are external to the system they protect. Some simply crash the process that was attacked (and do nothing to fix the fault, thereby ensuring that the system is still vulnerable when it is rebooted). Other systems may restrict network connectivity or resource consumption. SVV includes the ability to rewrite a vulnerable sequence of instructions without recompilation. This mechanism is general enough that a wide variety of response techniques can be implemented, such as: data structure repair [5], failure oblivious computing [21], and error virtualization [23]. The rewritten instruction stream can be propagated to the code section of the process address space to protect future execution. The new instruction sequence could be applied (with OS support) to the on-disk binary as a rudimentary patch.

Furthermore, compilers can be augmented to provide "alternative execution paths" to some code sections. These alternatives can be driven by explicit program code, programmer annotation, purely compiler-generated, taken from profiling information for the application, or gathered by the processor itself from previous runs of the same code block as a form of machine learning.

3 Related Work

SVV draws on ideas from computer architecture, fault-tolerant computing, and computer security. We examined some hardware support [24] for an *x86* emulator (STEM) that supervise the execution of vulnerable code slices [23]. The approach of SVV is akin to systems [22, 20, 19] that utilize a secondary host machine as a sandbox or instrumented honeypot: work is offloaded to this host and affects the primary host as little as possible. The

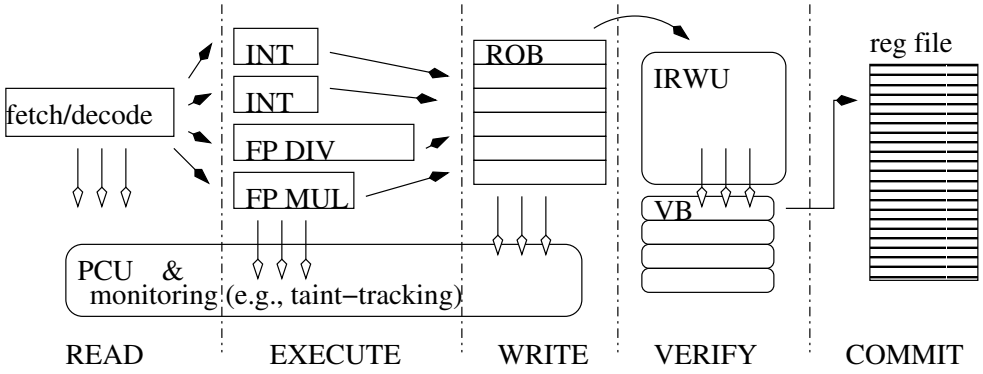

Figure 1: *High-level execution model for SVV*. Here, a simplified pipeline for a superscaler processor is modified to add an extra verification stage as well as policy-driven hardware-based monitoring mechanisms. The IRWU can optionally rewrite the instruction stream and cause the new version (stored in the VB) to be executed. Traditional hardware components are shown as full rectangles, new components are rounded. Not shown is the VERU, which holds the address for an emulator capable of higher-level supervision.

work most closely related to ours is Oplinger and Lam's proposal [18] for using TLS to improve software reliability. Their key idea is to execute an application's monitoring code in parallel with the primary computation and roll back the computation "transaction" depending on the results of the monitoring code.

Evers *et al.* [6] investigate the predictability of branches and provide an overview of various branch prediction schemes that have been proposed to ameliorate the cost of incorrect predictions. Wang *et al.* [27] explore an interesting result: about 50% of mispredicted branches do no affect correct program behavior.

3.1 Secure Hardware

Incorporating security mechanisms in hardware has traditionally been limited to providing implementations of cryptographic algorithms. McGregor and Lee [16] also investigate protecting cryptographic secrets. Recent efforts such as Cerium [3] and XOM [14, 14, 15] focused on providing a trusted computing base (TCB) and tamper-resistant architecture that can attest to the validity of a particular computation [9]. Of a more focused scope is Lee *et al.*'s proposal [13] of a hardware–based return stack (SRAS) to frustrate buffer overflow attacks. Suh *et al.* [26] propose hardware extensions to thwart control-transfer attacks by tracking "tainted" input data (as identified by the OS). If the processor detects the use of this tainted data as a jump address or an executed instruction, it raises an exception.

3.2 Execution Supervision Environments

Virtual machine emulation of operating systems or processor architectures to provide a sandboxed environment is an active area of research. Virtual machine monitors (VMMs) are employed in a number of security–related contexts, from autonomic patching of vulnerabilities [22] to intrusion detection [8]. MiSFIT [25] is a tool that constructs a sandbox by instrumenting applications at the assembly language level. Program shepherding [12] works on uninstrumented IA-32 binaries and validates branch instructions to prevent transfer of control to injected code.

Other protection mechanisms include compiler techniques like Stackguard [4] and safer libraries, such as *libsafe* and *libverify* [1]. Tools exist to verify and supervise code during development or debugging; of these tools, Purify[1] and Valgrind [17] are popular choices.

In work inspired by the ideas fundamental to artifi-

[1] `http://www.rational.com/products/purify_unix/index.jtmpl`

cial system diversity [7], Holland, Lim, and Seltzer [11] introduce the idea of automatically generating randomized architectures to support system security. Since synthesizing the hardware for such every such generated architecture is an untenable approach, they recommend using VMMs to provide the necessary execution environments.

3.3 Recovery and Repair

A key feature of SVV is the use of instruction stream re-writing as a basic building block for an adaptive response mechanism. Other recent work that examines repair mechanisms includes failure-oblivious computing [21] and data structure repair [5]. Candea and Fox propose a different approach: design software systems such that they employ crashing as the normal halting mode and use recursive microreboots to safely restart [2].

4 Conclusions

We have described the architectural components needed to support a new execution model for secure and reliable computing: *speculative virtual verification* (SVV). This model complements previous work on trustworthy and tamper-resistant computing architectures but is not meant as a replacement for the capabilities such systems provide. There are a wide variety of challenging problems to be addressed in the construction, testing, and deployment of SVV. We intend to study these issues and implement SVV in a variety of execution environments, including *x86* emulators, the Java Virtual Machine, and simulators for the MIPS and ARM architectures.

There is no silver bullet for system security, and SVV is not meant to address all possible attacks. However, we believe that the current arms race between attackers and system designers is unwinnable. As a fresh start, we advocate modifying general-purpose processors to (a) provide implicit supervision functionality, (b) export a policy-driven monitoring mechanism, and (c) provide the foundation for an automatic response capability via instruction stream rewriting.

References

[1] A. Baratloo, N. Singh, and T. Tsai. Transparent Run-Time Defense Against Stack Smashing Attacks. In *Proceedings of the USENIX Annual Technical Conference*, June 2000.

[2] G. Candea and A. Fox. Crash-Only Software. In *Proceedings of the 9th Workshop on Hot Topics in Operating Systems (HOTOS-IX)*, May 2003.

[3] B. Chen and R. Morris. Certifying Program Execution with Secure Processors. In *Proceedings of the 9th Workshop on Hot Topics in Operating Systems*, pages 133–138, May 2003.

[4] 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. 1998.

[5] B. Demsky and M. C. Rinard. Automatic Data Structure Repair for Self-Healing Systems. In *Proceedings of the 1st Workshop on Algorithms and Architectures for Self-Managing Systems*, June 2003.

[6] M. Evers, S. J. Patel, and Y. N. Patt. An Analysis of Correlation and Predictability: What Makes Two-Level Branch Predictors Work. In *Proceedings of the 25th International Symposium on Computer Architecture*, June 1998.

[7] S. Forrest, A. Somayaji, and D. Ackley. Building Diverse Computer Systems. In *Proceedings of the 6th Workshop on Hot Topics in Operating Systems*, pages 67–72, 1997.

[8] T. Garfinkel and M. Rosenblum. A Virtual Machine Introspection Based Architecture for Intrusion Detection. In 10th *ISOC Symposium on Network and Distributed Systems Security (SNDSS)*, February 2003.

[9] T. Garfinkel, M. Rosenblum, and D. Boneh. Flexible OS Support and Applications for Trusted Computing. In *Proceedings of the 9th Workshop on Hot Topics in Operating Systems*, pages 145–150, May 2003.

[10] J. L. Hennessy and D. A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann Publishers, 3rd edition, 2003.

[11] D. A. Holland, A. T. Lim, and M. I. Seltzer. An Architecture a Day Keeps The Hacker Away. In *Proceedings of the Workshop on Architectural Support for Security and Anti-Virus (WASSA)*, October 2004.

[12] V. Kiriansky, D. Bruening, and S. Amarasinghe. Secure Execution Via Program Shepherding. In *Proceedings of the* 11th *USENIX Security Symposium*, August 2002.

[13] R. B. Lee, D. K. Karig, J. P. McGregor, and Z. Shi. Enlisting Hardware Architecture to Thwart Malicious Code Injection. In *Proceedings of the International Conference on Security in Pervasive Computing (SPC-2003), Lecture Notes in Computer Science, Springer Verlag*, March 2003.

[14] D. Lie, C. Thekkath, and M. Horowitz. Implementing an Untrusted Operating System on Trusted Hardware. In *Proceedings of the* 19th *ACM Symposium on Operating Systems Principles (SOSP)*, October 2003.

[15] D. Lie, C. Thekkath, M. Mitchell, and P. Lincoln. Architectural Support for Copy and Tamper Resistant Software. In *Proceedings of the* 9th *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS IX)*, 2000.

[16] J. P. McGregor and R. B. Lee. Protecting Cryptographic Keys and Computations via Virtual Secure Coprocessing. In *Proceedings of the Workshop on Architectural Support for Security and Anti-Virus (WASSA)*, October 2004.

[17] N. Nethercote and J. Seward. Valgrind: A Program Supervision Framework. In *Electronic Notes in Theoretical Computer Science*, volume 89, 2003.

[18] J. Oplinger and M. S. Lam. Enhancing Software Reliability with Speculative Threads. In *Proceedings of the* 10th *International Conference on Architectural Support for Programming Languages and Operating Systems(ASPLOS X)*, October 2002.

[19] H. Patil and C. N. Fischer. Efficient Turn-time Monitoring Using Shadow Processing. In *Proceedings of the* 2nd *International Workshop on Automated and Algorithmic Debugging*, 1995.

[20] J. C. Reynolds, J. Just, L. Clough, and R. Maglich. On-Line Intrusion Detection and Attack Prevention Using Diversity, Genrate-and-Test, and Generalization. In *Proceedings of the* 36th *Hawaii International Conference on System Sciences (HICSS)*, 2003.

[21] M. Rinard, C. Cadar, D. Dumitran, D. Roy, T. Leu, and J. W Beebee. Enhancing server availability and security through failure-oblivious computing. In *Proceedings of the* 6th *Symposium on Operating Systems Design and Implementation (OSDI)*, December 2004.

[22] S. Sidiroglou and A. D. Keromytis. A Network Worm Vaccine Architecture. In *Proceedings of the IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE), Workshop on Enterprise Security*, pages 220–225, June 2003.

[23] S. Sidiroglou, M. E. Locasto, S. W. Boyd, and A. D. Keromytis. Building a Reactive Immune System for Software Services. In *Proceedings of the USENIX Annual Technical Conference (To appear)*, April 2005.

[24] S. Sidiroglou, M. E. Locasto, and A. D. Keromytis. Hardware Support For Self-Healing Software Services. In *Proceedings of the Workshop on Architectural Support for Security and Anti-Virus (WASSA)*, pages 37–43, October 2004.

[25] C. Small and M. Seltzer. MiSFIT: A Tool for Constructing Safe Extensible C++ Systems. *IEEE Concurrency*, 6(3):33–41, 1998.

[26] G. E. Suh, J. W. Lee, D. Zhang, and S. Devadas. Secure Program Execution via Dynamic Information Flow Tracking. In *Proceedings of the* 11th *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XI)*, October 2004.

[27] N. Wang, M. Fertig, and S. J. Patel. Y-Branches: When You Come to a Fork in the Road, Take It. In *Proceedings of the* 12th *International Conference on Parallel Architectures and Compilation Techniques*, September 2003.