Sphinx: A Secure Architecture Based on Binary Code Diversification and Execution Obfuscation

Michel A. Kinsy, Donato Kava, Alan Ehret, Miguel Mark

Adaptive and Secure Computing Systems Laboratory

Boston University, Boston, USA

Abstract— Sphinx, a hardware-software co-design architecture for binary code and runtime obfuscation. The Sphinx architecture uses binary code diversification and self-reconfigurable processing elements to maintain application functionality while obfuscating the binary code and architecture states to attackers. This approach dramatically reduces an attacker's ability to exploit information gained from one deployment to attack another deployment. Our results show that the Sphinx is able to decouple the program's execution time, power and memory and I/O activities from its functionality. It is also practical in the sense that the system (both software and hardware) overheads are minimal.

I. Introduction

Over the last several decades, computers have become more and more important to society and our daily lives. Almost all modern activity, from banking to research to small talk, has been connected to computers and the Internet. Considering how much computers have permeated our lives, it is necessary that computer systems are secured against malicious activity. This need for secure computing systems has produced extensive research efforts in many areas of computer science and engineering. Despite these efforts, the frequent reports of security breaches and cyber-attacks demonstrate that the work is on-going. One threat that has plagued security researchers and software developers for years is the possibility of reverse engineering programs and binary code to learn system vulnerabilities and modify the code to circumvent other security measures.

Few hardware systems have been developed to protect software from reverse engineering. Those methods that have been proposed are often centered around encrypting the program and then decrypting before execution. Depending on when the program is decrypted, an attacker may be able to use probes to obtain the plaintext program, such as from the bus between the CPU and the memory [1]. One suggestion to avoid this attack is to decrypt the instructions immediately prior to execution as is done in [2]. While this approach is promising in terms of blocking reverse engineering, it does not take into consideration side-channel attacks which could pose a threat.

Side Channel Attacks (SCAs) are a category of hardware attacks in which the attacker uses unintentional outputs (called side channels) to discover hidden information. Examples of side channel attacks include analyzing power, timing, and electromagnetic radiation to correlate the values with the secret information [3]. There are many proven examples of these attacks in [4], [5], [6]. Side channel attacks provide a major attack surface because they observe unintentional output sources which are usually ignored by system and software designers. In order to fully hide a program from attackers, all side channels would have to be completely obscured.

Many counter measures have been suggested to help defend against side channel attacks. Some works present changes to algorithms that attempt to tax the processor similarly in all cases, obscuring the power, electromagnetic (EM), and time channels. Other suggestions include oblivious RAM (ORAM), which is a hardware structure that randomizes memory access to prevent eavesdroppers from discovering patterns in the memory access [7]. Secure processors, like ASCEND [8] and AEGIS [9], attempt to implement impenetrable security measures in hardware. ASCEND tries to block side-channel attacks by activating every hardware module on each clock cycle and accessing memory at regular intervals even if no access is pending.

Such security measures do not come cheap; all of these systems suffer high cost and significant slowdown to execution.

Numerous security mechanisms have been proposed to prevent reverse engineering and to secure hardware against side channel attacks; however, each of the existing systems is found to be lacking in some way. The Sphinx system has a secure processor that can provide security against reverse engineering and side channel attacks without incurring excessive time and power overheads.

The general idea behind the Sphinx architecture is the minimization of attack-surface via software-hardware obfuscation. Every time a program (e.g., written in C or C++) is compiled, a different binary is generated. Every time a binary is executed, the power, execution time and memory activity profiles are all different. In essence, for the same application, the binary code and architecture states are obfuscated and the compute system operates differently to attackers. Figure 1 illustrates this behavior changing nature of the architecture.
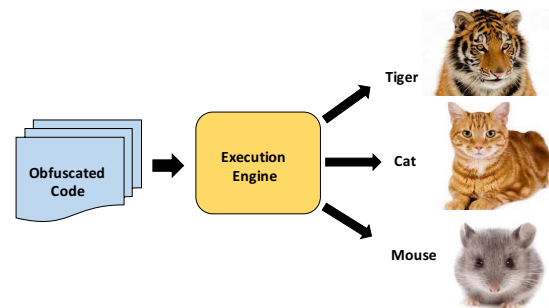

Fig. 1. Illustration of the Sphinx Architecture: The observed execution behavior depends on the context.

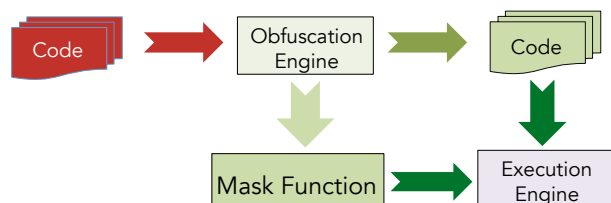## II. Obfuscation-based Secure Architecture Design

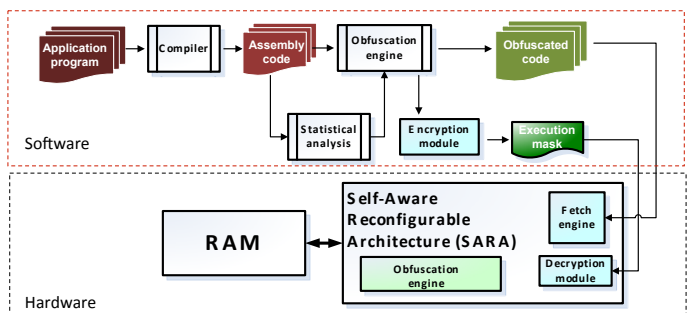Fig. 2. Sphinx system functionality overview.

Fig. 3. Sphinx obfuscation-based secure system

Sphinx is a software-hardware co-design framework to efficiently execute a program while providing moving target defense against code reverse engineering and side-channel based attacks. Figure 2 depicts the system functionality. The Sphinx secure system has two parts: a software obfuscator and a hardware/processing unit. Figure 3 shows the Sphinx system with its software and hardware components. The obfuscator inserts random machine instructions in the program at compile-time and creates a binary mask that indicates which instructions are real and which ones are falsified. The goal of using obfuscation is twofold: first, to render assembly or binary code analysis attacks harder - not impossible, since Barak et al., [10] have already shown that the impossibility of indistinguishable obfuscation; second, to give runtime adaptation range to the hardware. On the software side, the code is compiled and analyzed. The result of the analysis is used to generate a similar profile obfuscation assembly code to be added to the original. Figure 4 shows the obfuscation process. The *obf.asm* file goes through the normal compilation process. The binary mask is encrypted and loaded into the processor's memory with the obfuscated program. To run the program the processor transfers the encrypted mask to the execution unit and decrypts it using a securely stored key. In this work, we assume the secrecy of the key is provided by a Physical(ly) Unclonable Functions (PUF) technique. Then, the execution unit runs the program throwing out fake instructions as indicated by the mask.
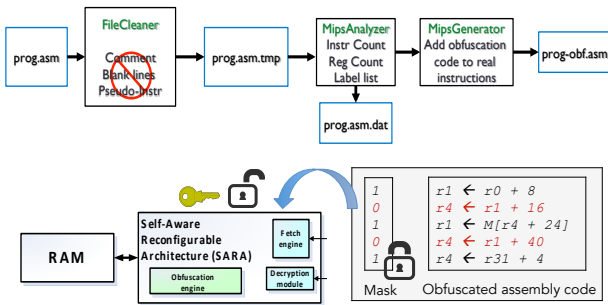

Fig. 4. The Sphinx software obfuscation process

The Sphinx core execution unit is a Self-Aware Reconfigurable Architecture (SARA). The key idea for the SARA design approach is to allow the hardware to have multiple ways of executing the same instruction with different time, power and memory/IO profiles, illustrated in Figure 5. The software level obfuscation gives reconfiguration range to the hardware and aids the architecture in isolating the program's execution time, power and memory and I/O activities from its functionality. Figure 5 shows the effects of user-defined entropy levels on the distribution of real and falsified instructions. Effectively, the Sphinx architecture provides the following capabilities (a) performance/timing-awareness for timing obfuscation, (b) power-awareness for power obfuscation and (c) self-organized data storage for memory and I/O obfuscation.
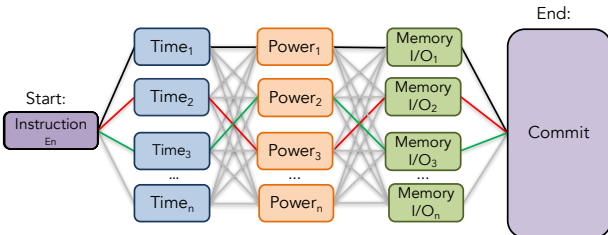

Fig. 5. Instruction execution profile candidates

III. EVALUATIONS

To test the Sphinx system, we implement a Verilog version of the hardware. We use a number of common benchmark suites (SPLASH-2, PARSEC, SPEC CINT2006). The selected benchmark applications
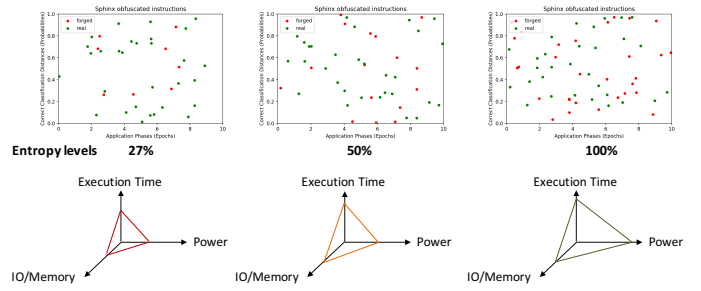

Fig. 6. Effects of user-defined entropy levels on the distribution of real and falsified instructions

are compiled, obfuscated and executed on the emulated hardware. The benchmarks are also run on an unsecured but otherwise similar processor to perform a comparative study. Figure 7 highlights some of the performance results of the Sphinx system.
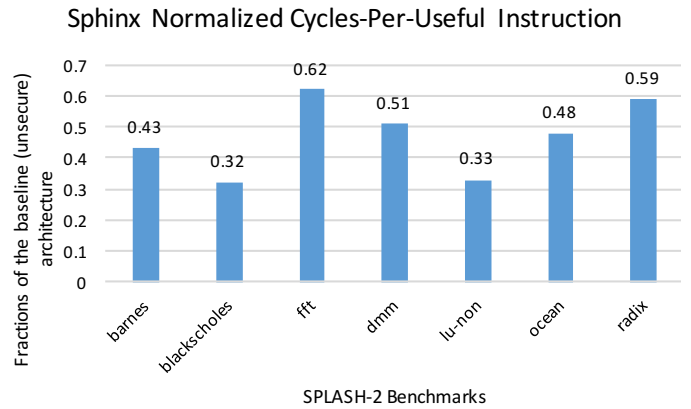

Fig. 7. Splash-2 benchmarks on the Sphinx system.

REFERENCES

[1] M. G. Kuhn, "Cipher instruction search attack on the bus-encryption security microcontroller ds5002fp," *IEEE Transactions on Computers*, no. 10, pp. 1153–1157, 1998.

[2] J. Cappaert, B. Preneel, B. Anckaert, M. Madou, and K. De Bosschere, "Towards tamper resistant code encryption: Practice and experience," *lecture notes in Computer Science*, vol. 4991, pp. 86–100, 2008.

[3] Y. Zhou and D. Feng, "Side-channel attacks: Ten years after its publication and the impacts on cryptographic module security testing," 2005.

[4] H. Bar-El, H. Choukri, D. Naccache, M. Tunstall, and C. Whelan, "The sorcerer's apprentice guide to fault attacks," *Proceedings of the IEEE*, vol. 94, no. 2, pp. 370–382, Feb 2006.

[5] E. Hess, N. Janssen, B. Meyer, and T. Schütze, "Information leakage attacks against smart card implementations of cryptographic algorithms and countermeasures–a survey," in *EUROSMART Security Conference*, vol. 130. Citeseer, 2000.

[6] J. Fan, X. Guo, E. De Mulder, P. Schaumont, B. Preneel, and I. Verbauwhede, "State-of-the-art of secure ecc implementations: A survey on known side-channel attacks and countermeasures." *HOST*, vol. 2010, pp. 76–87, 2010.

[7] E. Stefanov, E. Shi, and D. Song, "Towards practical oblivious ram," *arXiv preprint arXiv:1106.3652*, 2011.

[8] C. W. Fletcher, M. v. Dijk, and S. Devadas, "A secure processor architecture for encrypted computation on untrusted programs," in *Proceedings of the seventh ACM workshop on Scalable trusted computing*. ACM, 2012, pp. 3–8.

[9] G. E. Suh, C. Fletcher, D. Clarke, B. Gassend, M. van Dijk, and S. Devadas, "Author retrospective aegis: architecture for tamper-evident and tamper-resistant processing," in *25th Anniversary International Conference on Supercomputing Anniversary Volume*. ACM, 2014, pp. 68–70.

[10] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang, "On the (im) possibility of obfuscating programs," 2010.