HOP: Hardware makes Obfuscation Practical

Kartik Nayak*, Christopher W. Fletcher†, Ling Ren‡, Nishanth Chandran§, Satya Lokam§, Elaine Shi‖ and Vipul Goyal§

*UMD – kartik@cs.umd.edu
†UIUC – cwfletch@illinois.edu
‡MIT – renling@mit.edu
§Microsoft Research – {nichandr, satya, vipul}@microsoft.com
‖Cornell University – rs2359@cornell.edu

Abstract— Program obfuscation is a central primitive in cryptography, and has important real-world applications in protecting software from IP theft. However, well known results from the cryptographic literature have shown that software only virtual black box (VBB) obfuscation of general programs is impossible. In this paper we propose HOP, a system (with matching theoretic analysis) that achieves simulation-secure obfuscation for RAM programs, using secure hardware to circumvent previous impossibility results. To the best of our knowledge, HOP is the *first* implementation of a provably secure VBB obfuscation scheme in any model under any assumptions.

HOP trusts only a hardware single-chip processor. We present a theoretical model for our complete hardware design and prove its security in the UC framework. Our goal is both provable security and practicality. To this end, our theoretic analysis accounts for all optimizations used in our practical design, including the use of a hardware Oblivious RAM (ORAM), hardware scratchpad memories, instruction scheduling techniques and context switching. We then detail a prototype hardware implementation of HOP. The complete design requires 72% of the area of a V7485t Field Programmable Gate Array (FPGA) chip. Evaluated on a variety of benchmarks, HOP achieves an overhead of $8\times \sim 76\times$ relative to an insecure system. Compared to all prior (not implemented) work that strives to achieve obfuscation, HOP improves performance by more than three orders of magnitude. We view this as an important step towards deploying obfuscation technology in practice.

I. Introduction

Program obfuscation [29], [4] is a powerful cryptographic primitive, enabling numerous applications that rely on intellectually-protected programs and the safe distribution of such programs. For example, program obfuscation enables a software company to release software patches without disclosing the vulnerability to an attacker. It could also enable a pharmaceutical company to outsource its proprietary genomic testing algorithms, to an untrusted cloud provider, without compromising its intellectual properties. Here, the pharmaceutical company is referred to as the "sender" whereas the cloud provider is referred to as the "receiver" of the program.

Recently, the cryptography community has had new breakthrough results in understanding and constructing program obfuscation [21]. However, cryptographic approaches towards program obfuscation have limitations. First, it is well-understood that strong (simulation secure) notions of program obfuscation cannot be realized in general [4] — although they are desired or necessary in many applications such as the aforementioned ones. Second, existing cryptographic constructions of obfuscation (that achieve weaker notions of security, such as indistinguishability obfuscation [22]) incur prohibitive practical overheads, and are infeasible for most interesting application scenarios. For example, it takes ~ 3.3 hours to obfuscate even a very simple program such as an 80-bit point function (a function that is 0 everywhere except at one point) and ~ 3 minutes to evaluate it [37]. Moreover, these cryptographic constructions of program obfuscation rely on new cryptographic assumptions whose security is still being investigated by the community through a build-and-break iterative cycle [14]. Thus, to realize a practical scheme capable of running general programs, it seems necessary to introduce additional assumptions.

In this direction, there has been work by both the cryptography and architecture communities in assuming trusted hardware storing a secret key. However, proposals from the cryptography community to realize obfuscation (and a closely related primitive called functional encryption) have been largely theoretical, focusing on what minimal trusted hardware allows one to circumvent theoretical impossibility and realize simulation-secure obfuscation [27], [15], [17]. Consequently these works have not focused on practical efficiency, and they often require running the program as circuits (instead of as RAM programs) and also utilize expensive cryptographic primitives such as fully homomorphic encryption (FHE) and non-interactive zero knowledge proofs (NIZKs). On the other hand, proposals from the architecture community such as Intel SGX [42], AEGIS [53], XOM [38], Bastion [13], Ascend [18] and GhostRider [40] are more practical, but their designs do not achieve cryptographic definition of obfuscation. In this paper, we close this gap by designing and implementing a practical construction of program obfuscation for RAM programs using trusted hardware.

Problem statement. The problem of obfuscation can be described as follows. A sender, who owns a program, uses an obfuscate procedure to create an obfuscated program. It then sends this obfuscated program to a receiver who can execute the program on inputs of her choice. The obfuscated program
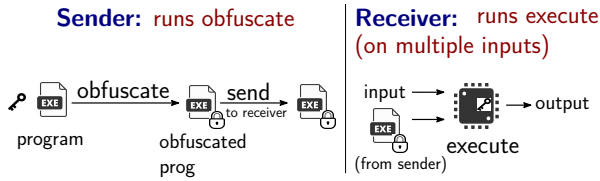
Fig. 1: **Obfuscation Scenario.** The sender obfuscates programs using the obfuscate procedure. It sends (possibly multiple) obfuscated program(s) to the receiver. The receiver can execute any obfuscated program with any input of its choice.

should be functionally identical to the original program. For any given input, the obfuscated program runs for time T (fixed for the program) and returns an output.[1] The receiver only has a black box-like access to the program, i.e., it learns only the program's input/output behavior and the bound on the runtime T. In obfuscation, the inputs/outputs are public (not encrypted).

To make use of a trusted secure processor (which we call a *HOP processor*), our obfuscation model is modified as follows (cf. Figure 1). HOP processors are manufactured with a hardwired secret key. The HOP processor (which is trusted) is given to the receiver, and the secret key is given to the sender. Using the secret key, the sender can create multiple obfuscated programs using the obfuscate procedure and send them to the receiver. The receiver then runs the execute procedure (possibly multiple times) to execute the program with (cleartext) inputs of her choice. As mentioned, the receiver (adversary) learns only the final outputs and nothing else. In other words, we offer virtual blackbox simulation security, where the receiver learns only as much information as if she were interacting with an oracle that computes the obfuscated program. In particular, the receiver should not learn anything from the HOP processor's intermediate behavior such as timing or memory access patterns, or the program's total runtime (since each program always runs for a fixed amount of time set by the sender).

Key distribution with public/private keys. We assume symmetric keys for simplicity. HOP may also use a private/public key distribution scheme common in today's trusted execution technology. The obfuscate and execute operations can be decoupled from the exact setup and key distribution system used to get public/private keys into the HOP processor. A standard setup for key distribution [28], [42] is as follows: First, a trusted manufacturer (e.g., Intel) creates a HOP processor with a unique secret key. Its public key is endorsed/signed by the manufacturer. Second, the HOP processors are distributed to receivers and the certified public keys are distributed to senders (software developers). The modification to our scheme in the public key setting is described in Appendix B. Note that the key goal of obfuscation is to secure the sender's program and this relies on the secrecy of the private key stored in the processor. Thus, it is imperative that the sender and the manufacturer are either the same entity or the sender trusts the manufacturer to not reveal the secret key to another party.

Non-goals. We do not defend against analog side channels

---

[1] T is analogous to a bound on circuit size in the cryptography literature.

such as measuring power analysis or heat dissipation, we also do not defend against hardware fault injection [8], [3], [34]. We assume that the program to be obfuscated is trustworthy and will not leak sensitive information on its own, including through possible software vulnerabilities such as buffer overflows [7]. There exist techniques to mitigate these attacks, and we consider them to be complementary to our work.

Challenges. It may seem that relying on secure hardware as described above easily 'solves' the program obfuscation problem. This is not the case: even with secure hardware, it is still not easy to develop a secure *and* practical obfuscation scheme. The crux of the problem is that many performance optimizations in real systems (and related work in secure processors [18], [40], [45]) hinge on exploiting *program-dependent behavior*. Yet, obfuscation calls for completely hiding all program-dependent behavior. Indeed, we started this project with a strawman processor that gives off the impression of executing any (or every) instruction during each time step – so as to hide the actual instructions being executed. Not surprisingly, this incurs huge ($\sim 10,000\times$; c.f. Section III-B) overheads over an insecure scheme, even after employing a state-of-the-art Oblivious RAM [26], [19] to improve the efficiency of accessing main memory. Moreover, in an obfuscation setting, the receiver can run the same program multiple times for different inputs and outputs. Introducing practical features such as context switching — where the receiver can obtain intermediate program state — enables this level of flexibility but also enables new attacks such as *rewinding* and *mix-and-match* execution. Oblivious RAMs, in particular, are not secure against rewinding and mix-and-match attacks and an important challenge in this work is to protect them against said attacks in the context of the HOP system.

A. Our Contributions

Given the above challenges, a primary goal of this paper is to develop and implement an optimized architecture that is still provably secure by the VBB obfuscation definition. We stress that *all* the performance optimizations made in the paper are included and proven secure in our theoretic analysis: we want our practical design to match the theory to the extent possible. We view this as an important step towards deploying obfuscation technology in practice.

In more detail, we make the following contributions:

1. Theoretical contributions: We provide the first theoretic framework to *efficiently* obfuscate RAM programs directly on secure hardware. One goal here is to avoid implicitly transforming the obfuscated program to its circuit representation (e.g., [17]), as the RAM to circuit transformation can incur a polynomial blowup in runtime [23]. We also wish for our analysis to capture important performance optimizations that matter in an implementation; such as the use of a cryptographic primitive called Oblivious RAM [25], [26], on-chip memory, instruction scheduling, and context switching. As a byproduct, part of our analysis achieves a new theoretical result (extending [27]): namely, how to provide program obfuscation for RAM

programs *directly* assuming only 'stateless' secure hardware.[2] We also show interesting technical subtleties that arise in constructing efficient RAM-model program obfuscation from stateless hardware. In particular, we highlight the different techniques used to overcome all possible forms of *rewinding* and *mix-and-match* attacks (which may be of independent interest). Putting it all together, we provide a formal proof of security for the entire system under the universally composable (UC) simulation framework [10].

2. Implementation with trusted hardware: We design and implement a hardware prototype system (called HOP) that attains the definition of program obfuscation and corresponds to our theoretic analysis. To the best of our knowledge, this effort represents the *first* implementation of a provably secure VBB obfuscation scheme in *any* model under *any* assumptions. For performance, our HOP prototype uses a hardware-optimized Oblivious RAM, on-chip memory and instruction scheduling (our current implementation does not yet support context switching). As mentioned earlier, our key differentiator from prior secure processor work is that our performance optimizations maintain *program privacy* and exhibit no program-dependent behavior. With these optimizations, HOP performs $5\times \sim 238\times$ better than the baseline HOP design across simple to sophisticated programs while the overhead over an insecure system is $8\times \sim 76\times$. The program code size overhead for HOP is only an *additive* constant. Our final design requires 72% area when synthesized on a commodity FPGA device. Of independent interest, we prove that our optimized scheme always achieves to within $2\times$ the performance of a scheme that does not protect the main memory timing channel (Section III-C).

II. RELATED WORK

Obfuscation. The formal study of virtual black-box (VBB) obfuscation was initiated by Hada [29] and Barak *et al.* [4]. Unfortunately, Barak *et al.* showed that it is impossible to achieve program obfuscation for general programs. Barak *et al.* also defined a weaker notion of indistinguishability obfuscation $(i\mathcal{O})$, which avoids their impossibility results. Garg *et al.* [22] proposed a construction of $i\mathcal{O}$ for all circuits based on assumptions related to multilinear maps. However, these constructions are not efficient from a practical standpoint. There are constructions for $i\mathcal{O}$ for RAM programs proposed where the size of the obfuscated program is independent of the running time [6], [11], [36]. However, by definition, these constructions do not achieve VBB obfuscation.

In order to circumvent the impossibility of VBB obfuscation, Goyal *et al.* [27] considered virtual black-box obfuscators on minimal secure hardware tokens. Goyal *et al.* show how to achieve VBB obfuscation for all polynomial time computable functions using *stateless* secure hardware tokens that only perform authenticated encryption/decryption and a single NAND operation. In a related line of work, Döttling *et al.* [17] show a construction for program obfuscation using a single stateless hardware token in universally input-oblivious models

of computation. Bitansky *et al.* [5] show a construction for program obfuscation from "leaky" hardware. Similarly, Chung *et al.* [15] considered basing the closely related primitive of functional encryption on hardware tokens. Unfortunately, all the above works require the obfuscated program run using a universal circuit (or similar model) to achieve function privacy. They do not support running RAM programs directly. This severely limits the practicality of the above schemes, as we demonstrate in Section VI-E.

Oblivious RAMs. To enable running RAM programs directly on secure hardware, we use a hardware implementation of Oblivious RAM (ORAM) to hide access patterns to external memory. ORAM was introduced by Goldreich and Ostrovsky where they explored the use of tamper-proof hardware for software protection [26]. Recently, there has been a lot of work in making ORAMs practical. In this paper, we use an efficient hardware implementation of Path ORAM [52] called Tiny ORAM [20], [19].

Secure processors. Secure processors such as AEGIS [53], XOM [38], Bastion [13] and Intel SGX [42] encrypt and verify the integrity of main memory. Applications such as VC3 [48] that are built atop Intel SGX can run MapReduce computations [16] in a distributed cloud setting while keeping code and data encrypted. However, these secure processors do not hide memory access patterns. An adversary observing communication patterns between a processor and its memory can still infer significant information about the data [43], [58].

There have been some recent secure processor proposals that do hide memory access patterns [18], [41], [40], [45]. Ascend [18] is a secure processor architecture that protects privacy of data against physical attacks when running arbitrary programs. Phantom [41] similarly achieves memory oblivious-ness, and has been integrated with GhostRider [40] to perform program analysis and decide whether to use an encrypted RAM or Oblivious RAM for different memory regions. They also employ a scratchpad wherever applicable. Raccoon [45] hides data access patterns on commodity processors by evaluating all program paths and using an Oblivious RAM in software.

The primary difference between the above schemes and HOP is the following. All of the above schemes focused on protecting input data, while the program is assumed to be public and known to the adversary. GhostRider [40] even utilizes public knowledge of program behavior to improve performance through static analysis. Conversely, obfuscation and HOP protect the program and the input data is controlled by the adversary. We remark, however, that HOP can be extended to *additionally* achieve data privacy simply by adding routines to decrypt the (now private) inputs and encrypt the final outputs before they are sent to the client (now different from the HOP processor owner). Naturally, the enhanced security comes with additional cost. We evaluate this overhead of additionally providing program-privacy by comparing to GhostRider in Section VI-E.

Secure computation. There is a line of work addressing how to build a general purpose MIPS processor for garbled circuits [51], [56]. When one party provides the program, the system is capable of performing private function secure function evaluation (PF-SFE). Similarly, universal circuits [55],

---

[2]Roughly speaking, a HOP processor which allows the host to arbitrary context switch programs on/off the hardware is equivalent to 'stateless' hardware in the language of prior work [27], [15]. This is explained further in Section III.

[35], [33] in combination with garbled circuits (which can be evaluated efficiently with techniques in [30]) or other multiparty computation protocols can be used to hide program functionality from one of the parties. The work of Katz [32] relies on trusted hardware tokens to circumvent the theoretical impossibility of UC-secure multi-party computation under dishonest majority. However, all the above results are in the context of secure computation, which is inherently interactive and only allows one-time use, i.e., for every input, both parties are involved in the computation. On the contrary, obfuscation requires that a party non-interactively execute the obfuscated program several times on multiple inputs.

Heuristic approaches to obfuscation. There are heuristic approaches to code obfuscation for resistance to reverse engineering [58], [31], [47]. These works provide low overheads, but do not offer any cryptographic security.

Terminology: Hardware Tokens. Trusted hardware is widely referred to as hardware *tokens* in the theoretical literature [32], [27], [17], [15]. Secure tokens are typically assumed to be minimal trusted hardware that support limited operations (e.g., a NAND gate in [27]). However, running programs in practice requires full-fledged processors. In this paper, we refer to HOP as "secure hardware" or a "secure processor". As a processor, HOP will store a lot more internal state (e.g., a register file, etc.). We note that from a theoretic perspective, both HOP and 'simple' hardware tokens require a number of gates which is polylogarithmic in memory size.

Terminology: Stateful vs. Stateless tokens. The literature further classifies secure tokens as either stateful tokens or stateless. A stateful token maintains state across invocations. On the other hand, a stateless token, except for a secret key, does not maintain any state across invocations. While HOP maintains state across most invocations for better performance, we will augment HOP to support on-demand context switching — giving the receiver the ability to swap out an obfuscated program for another at any time (Section III-E), which is common in today's systems. In an extreme scenario, the adversary can context switch after every processor cycle. In this case, HOP becomes equivalent to a "stateless" token from a theoretical perspective [27], [15], and our security proof will assume stateless tokens.

III. OBFUSCATION FROM TRUSTED HARDWARE

In this section, we describe the HOP architecture. We will start with an overview of a simple (not practical) HOP processor to introduce some key points. Each subsection after that introduces additional optimizations (some expose security issues, which we address) to make the scheme more practical. We give security intuition where applicable, and formally prove security for the fully optimized scheme in Section IV.

A. Execution On-Chip

Let us start with the simplest case where the whole obfuscated program and its data (working set) fit in a processor's on-chip storage. Then, we may architect a HOP processor to be able to run programs whose working sets don't exceed a given size. In the setup phase, first, the sender correctly determines a value T – the amount of time (in processor cycles) that the program, given any input, runs on HOP. Then, the sender encrypts (obfuscates) the program together using an authenticated encryption scheme. T is authenticated along and included with the program but is public. The obfuscated program is sent to the receiver. The receiver then sends the obfuscated program and her own input to the HOP processor. The HOP processor decrypts and runs the program, and returns a result after T processor cycles. The HOP processor makes no external memory requests during its execution since the program and data fit on chip. Security follows trivially.

B. Adding External Memory

Unfortunately, since on-chip storage is scarce (commercial processors have a few MegaBytes of on-chip storage), the above solution can only run programs with small working sets. To handle this, like any other modern processor, the HOP processor needs to access an external memory, which is possibly controlled by the malicious receiver.

When the HOP processor needs to make an access to this receiver memory, it needs to hide its access patterns. For the purposes of this discussion, the access pattern indicates the processor's memory operations (reads vs. writes), the memory addresses for each access and the data read/written in each access. We hide access pattern by using an Oblivious RAM (ORAM), which makes a polylogarithmic number of physical memory accesses to serve each logical memory request from the processor [52]. The ORAM appears to HOP as an on-chip memory controller that intercepts memory requests from the HOP processor to the external memory. That is, the ORAM is a hardware block on the processor and is trusted. (More formal definitions for ORAM are given in Section IV-A.)

Each ORAM access can take thousands of processor cycles [19]. Executing instructions – once data is present on-chip – is still as fast as an insecure machine (e.g., several cycles). To hide when ORAM accesses are actually needed, HOP must make accesses at a static program-independent frequency (more detail below). As before, HOP runs for T time on all inputs and hence achieves the same privacy as the scheme in Section III-A.

Generating T and security requirements. When accessing receiver-controlled memory, we must change T to represent some amount of work that is independent of the external memory's latency. That is, if T is given in processor cycles, the adversary can learn the true program termination time by running the program multiple times and varying the ORAM access latency each time (causing a different number of logical instructions to complete each time). To prevent this, we change T to mean 'the number of external memory read/writes made with the receiver.'

Integrity. To ensure authenticity of the encrypted program instructions and data during the execution, HOP uses a standard Merkle tree (or one that is integrated with the ORAM [46]) and stores the root of a Merkle tree internally. The receiver cannot tamper with or rewind the memory without breaking the Merkle tree authentication scheme.

Efficiency. While the above scheme can handle programs with large working sets, it is very inefficient. The problem is that

each instruction may trigger multiple ORAM accesses. To give off the impression of running any program, we must provision for this worst case: running each instruction must incur the cost of the worst-case number of ORAM accesses. This can result in $\sim 10,000\times$ slowdown over an insecure processor.[3] The next two subsections discuss two techniques to securely reduce this overhead by over two orders of magnitude. These ideas are based on well-known observations that many programs have more arithmetic instructions than memory instructions, and exhibit locality in memory accesses.

C. Adding Instruction Scheduling

The key intuition behind our first technique is that many programs execute multiple arithmetic instructions for every memory access. For example, an instruction trace may be the following: 'A A A A M A A M', where A, M refer to arithmetic and memory instructions respectively.

Our optimization is to let the HOP processor follow a fixed and pre-defined schedule: N arithmetic instructions followed by one memory access. In the above example, given a schedule of A^4M, the processor would insert two *dummy* arithmetic instructions to adhere to the schedule. A dummy arithmetic instruction can be implemented by executing a `nop` instruction. The access trace observable to the adversary would then be:
$$A\ A\ A\ A\ M\ A\ A\ \mathbf{A}\ \mathbf{A}\ M$$
The bold face **A** letters refer to dummy arithmetic instructions introduced by the processor.

Likewise, if another part of the program trace contains a long sequence of arithmetic instructions, the processor will insert dummy ORAM accesses to adhere to the schedule.

Gains. For most programs in practice, there exists a schedule with $N > 1$ that would perform better than our baseline scheme from Section III-B. For $(N+1)$ instructions, the baseline scheme performs $(N+1)$ arithmetic and memory accesses. With an A^NM schedule, our optimized scheme performs only one memory access which translates to a speedup of $N\times$ in the best case, when the cost of the memory access is much higher than an arithmetic instruction. To translate this into performance on HOP - given that HOP must run for T time - consider the following: If $N > 1$ does improve performance for the given program on all inputs, it means the sender can specify a smaller T for that program, while still having the guarantee that the program will complete given any input. A smaller T means better performance.

Setting N and security intuition. We design all HOP processors to use the same value of N for all programs and all inputs (i.e., N is set at HOP manufacturing time like the private key). More concretely, we set
$$N = \frac{\text{ORAM latency}}{\text{Arithmetic latency}}$$
In other words, the number of processor cycles spent on arithmetic instructions and memory instructions are the same. For typical parameter settings, $N > 1000$ is expected. While

this may sound like it will severely hurt performance given pathological programs, we show that this simple strategy does "well" on arbitrary programs and data, formalized below.

Claim: For *any* program and input, the above N results in $\leq 50\%$ of processor cycles performing dummy work.

We refer the reader to Appendix A for a proof of this claim. We consider this proof to be of independent interest. The claim implies that in comparison to a solution that *does not* protect the main memory timing channel, our fixed schedule introduces a maximum overhead of $2\times$ given any program – whether they are memory or computation intensive. Said another way, even when more sophisticated heuristics than a fixed schedule are used for different applications, the performance gain from those techniques is a factor of 2 at most.

Security. We note that our instruction scheduling scheme does not impact security because we use a fixed, public N for all programs.

D. Adding on-chip Scratchpad Memory

Our second optimization adds a scratchpad: a *small* unit of trusted memory (RAM) inside the processor, accesses to which are not observable by the adversary.[4] It is used to temporarily store a portion of the working set for programs that exhibit locality in their access patterns.

Running programs with a scratchpad. We briefly cover how to run programs using a scratchpad here. More (implementation-specific) detail is given in Section V-A. At a high level, data is loaded into the scratchpad from ORAM/unloaded to ORAM using special (new) CPU instructions that are added to the obfuscated program. These instructions statically determine when to load which data to specified offsets in the scratchpad. Now, the scratchpad load/unload instructions are the only instructions that access ORAM (i.e., are the only 'M' instructions). Memory instructions in the original program (e.g., normal loads and stores) merely lookup the scratchpad inside the processor (these are now considered 'A' instructions). We will assume the program is correctly compiled so that whenever a program memory instruction looks up the scratchpad, the data in question has been put there sometime prior by a scratchpad load/unload instruction.

Security intuition. When the program accesses the scratchpad, it is hidden from the adversary since this is done on-chip. As before, the only adversary-visible behavior is when ORAM is accessed and this will be governed by the program-independent schedule from Section III-C.

Program independence. We note that HOP with a scratchpad is still program independent. Multiple programs can be written (and obfuscated) for the same HOP processor. One minor limitation, however, is that once an obfuscated program is compiled, it must be compiled with 'minimum scratchpad size' specified as a new parameter and cannot be run on HOP

---

[3]Our ORAM latency from Section VI is 3000 cycles. The RISC-V ISA [12] we adopt can trigger 3 ORAM accesses, one to fetch the instruction, 1 or 2 more to fetch the operand, depending on whether the operand straddles an ORAM block boundary.

[4]We remark that we use a software-managed scratchpad (as opposed to a conventional processor cache) as it is easier to determine T when using a scratchpad.

processors that have a smaller scratchpad. This is necessary because having a smaller scratchpad will increase T by some unknown amount. If the program is run on a HOP processor with a larger scratchpad, it will still function but some scratchpad space won't be used.

Gains. In the absence of a scratchpad, the ratio of arithmetic to memory instructions is on average 5:1 for our workloads. When using a scratchpad, a larger amount of data is stored by the processor, thus decreasing memory accesses. This effectively decreases the execution time T of the program and substantially improves performance for programs with high locality (evaluated in Section VI-C).

E. Adding context switching and stateless tokens

A problem with the proposals discussed so far is that once a program is started, it cannot be stopped until it returns a response. But a user may wish to concurrently run multiple obfuscated programs for a practical deployment model. Therefore, we design the HOP processor to support on-demand context switch, i.e., the receiver can invoke a context switch at any point during execution. This, however, introduces security problems that we need to address.

A context switch means that the current program state should be swapped out from the HOP processor and replaced with another program's state. Since such a context switch can potentially happen at every invocation, the HOP processor no longer stores state and is a stateless token. In such a scenario, we design it to encrypt all its internal state, and send this encrypted/authenticated state (denoted $\overline{\text{state}}$) to the receiver (i.e., the adversary) on a context switch. Whenever the receiver passes control back to the token, it will pass back the encrypted state as well, such that the token can "recover" its state upon every invocation.

Challenges. Although on the surface, this idea sounds easy to implement, in reality it introduces avenues for new attacks that we now need to defend against. For the rest of the paper, and in-line with real processors, we assume the only data that remains in HOP is the per-chip secret key (Section I). A notable attack is the *rewinding* attack. In this attack, instead of passing to the token the correct and fresh encrypted $\overline{\text{state}}$ as well as fresh values of memory reads, a malicious receiver can pass old values.[5] The receiver can also *mix-and-match* values from entirely different executions of the same program or different programs. The rest of the section outlines how to prevent the above attacks. We remark that while the below have simple fixes, the problems themselves are easy to overlook and underscore the need for a careful formal analysis. Indeed, we discovered several of these issues while working through the security proof itself.

Preventing mix-and-match. To prevent this attack, we enforce that the receiver must submit an encrypted state $\overline{\text{state}}$, corresponding to an execution at some point t, along with a matching read from time t for the same execution. To achieve this, observe that $\overline{\text{state}}$ is encrypted with a IND-CPA + INT-CTXT-secure authenticated encryption scheme, and that the state carries all necessary information to authenticate the next memory read. The state contains information unique to the specific program, the specific program execution, and to the specific instruction that the token expects.

Preventing rewinding during program execution. An adversary may try to gain more information by rewinding an execution to a previous time step, and replaying it from that point on. To prevent an adversary from learning more information in this way, we make sure that the token simply replays an old answer should rewinding happen — this way, the adversary gains no more information by rewinding. To achieve this, we make sure that any execution for a (program, inp) pair is entirely deterministic no matter how many times you replay it. All randomness required by the token (e.g., those required by the ORAM or memory checker) are generated pseudorandomly based on the tuple (K, H_S, H_R) where K is a secret key hardwired in the token, H_R is a commitment to the receiver's input and $H_S := \text{digest}(\overline{\text{mem}_0})$ is a Merkle root of the program.

Preventing rewinding during input insertion. In our setting, the obfuscated program's inputs inp are chosen by the receiver. Since inputs can be long, it may not be possible to submit the entire input in one shot. As a result, the receiver has to submit the input word by word. Therefore the malicious receiver may rewind to a point in the middle of the input submission, and change parts of the input in the second execution. Such a rewinding causes two inputs to use the same randomness for some part of the execution.

To prevent such an input rewinding attack, we require that the adversary submit a Merkle tree commitment $H_R := \text{digest}(\text{inp})$ of its input inp upfront, before submitting a long input word by word. H_R uniquely determines the rest of the execution, such that any rewinding will effectively cause the token to play old answers (as mentioned above), and the adversary learns nothing new through rewinding.

IV. Formal Scheme

We now give a formal model for the fully optimized HOP processor (i.e., including all subsections in Section III) and prove its security in UC framework. Section IV-A describes the preliminaries. Section IV-B describes the ideal functionality for obfuscation of RAM programs. Sections IV-C and IV-D describe our formal scheme and proof in the UC framework.

A. Preliminaries

The notations used in this section are summarized in Table I. We denote the assignment-operator with :=, while we use = to denote equality. Encryption of data is denoted by an overline, e.g., $\overline{\text{state}} = \text{Enc}_K(\text{state})$, where Enc denotes a IND-CPA + INT-CTXT-secure authenticated encryption scheme and K is the key used for encryption.

---

[5]Here is a possible attack by which the adversary can distinguish between two access patterns. Consider the access pattern $\{a, a\}$ i.e., accessing the same block consecutively. If a tree-based ORAM [49] is used, after the first access, the block is remapped to a new path l' and the new path l' would be subsequently accessed. If the adversary rewinds and executes again, the block may be mapped to a different path l''. Thus, for two different executions, two different paths (l' and l'') are accessed for the second access. Note that for another access pattern $\{a, b\}$ for $a \neq b$, the same paths would be accessed even after rewinding, thus enabling the adversary to distinguish between access patterns.

K	Hardwired secret key stored by the token
mem_0	A program as a list of instructions
inp	Input to the program
mem	Memory required for program execution
outp	Program output
ℓ_{in}, ℓ_{out}, w	Bit-lengths of input, output, and memory word
N	Number of words in memory
T	Time for program execution
RAM.params	$\{T, N, \ell_{in}, \ell_{out}, w\}$
oramstate	State stored by ORAM
sstorestate	State stored by sstore
H_R	Digest of receiver's input, i.e., $\mathsf{digest(inp)}$
H_S	Digest of sender's program, i.e., $\mathsf{digest(\overline{mem_0})}$
H'	Merkle root of the main memory

Universal Composability framework. The Universal Composability framework [10] considers two worlds – 1. *real world* where the parties execute a protocol π. An adversary $\mathcal{A}$ controls the corrupted parties. 2. *ideal world* where we assume the presence of a trusted third party. The parties interact with a trusted third party (also called ideal functionality $\mathcal{F}$) with a protocol ϕ. A simulator $\mathcal{S}$ tries to mimic the actions of $\mathcal{A}$. Intuitively, the amount of information revealed by π in the real world should not be more than what is revealed by interacting with the trusted third party in the ideal world. In other words, we have the following: an environment $\mathcal{E}$ observes one of the two worlds and guesses the world. Protocol π UC-realizes ideal functionality $\mathcal{F}$ if for any adversary $\mathcal{A}$ there exists a simulator $\mathcal{S}$, such that an environment $\mathcal{E}$ cannot distinguish (except with negligible probability) whether it is interacting with $\mathcal{S}$ and ϕ or with $\mathcal{A}$ and π.

Random Access Machines. We now give definitions for Random Access Machine (RAM) programs, a basic processor model for RAM programs. Let $\mathcal{RAM}[\Pi, T, N, \ell_{in}, \ell_{out}, w]$ denote a family of RAM programs with the following public parameters: Π denotes the next instruction circuit; T denotes the number of steps the program will be executed; w denote the bit-width of a memory word; and N, ℓ_{in} and ℓ_{out} denote the memory, input and output lengths respectively (in terms of number of words).

We consider programs $\mathsf{RAM} := \langle \mathsf{cpustate}, \mathsf{mem} \rangle \in \mathcal{RAM}[\Pi, T, N, \ell_{in}, \ell_{out}, w]$ to be a tuple, where cpustate denotes the CPU's initial internal state, and mem denotes an initial memory array. In these programs, for each step of the execution, the next instruction function is executed over the old cpustate and the most recently fetched w bit memory word denoted rdata:

$$(\mathsf{cpustate}, \mathsf{op}) := \Pi(\mathsf{cpustate}, \mathsf{rdata})$$

As a result, cpustate is updated, and a next read/write instruction op is fetched. Initially, rdata is set to 0.

On input inp, the execution of $\mathsf{RAM}[T, N, \ell_{in}, \ell_{out}, w] := \langle \Pi, \mathsf{cpustate}, \mathsf{mem} \rangle$ is defined as the following:

```
rdata := 0
mem[1..ℓin] := inp
for t ∈ [1, 2, ..., T]:
    (cpustate, op) := Π(cpustate, rdata)
    if op = (write, addr, wdata)
        mem[addr] := wdata
    else if op = (read, addr, ⊥)
        rdata := mem[addr]
Output rdata    // rdata stores the output
```

For notational simplicity, we assume that output length ℓ_{out} is small and can be stored in rdata. However, our results can be extended easily to larger values of ℓ_{out}. For succinctness, we denote $(T, N, \ell_{in}, \ell_{out}, w)$ by RAM.params. Wherever its clear from context, we abuse notation to denote $\mathcal{RAM}[\Pi, T, N, \ell_{in}, \ell_{out}, w]$ as $\mathcal{RAM}$.

Oblivious RAM. Let mem denote a memory array that supports two types of operations: a) On $(\mathsf{read}, addr)$, it outputs $\mathsf{mem}[addr]$; b) On $(\mathsf{write}, addr, \mathsf{wdata})$, it sets $\mathsf{mem}[addr] := \mathsf{wdata}$, and outputs $\bot$. In this paper, we define an Oblivious RAM as a *stateful*, probabilistic algorithm that interacts with a memory array mem. It is denoted as $\mathsf{ORAM}^{N,w}$ where N and w are public parameters denoting the memory capacity in terms of number of words, and the bit-width of a word. mem denotes the initial state of the memory, where all but the first N locations are set to 0. An ORAM converts memory contents mem to mem$'$. An ORAM takes two types of inputs: $\mathsf{op} := (\mathsf{read}, addr)$, and $\mathsf{op} := (\mathsf{write}, addr, \mathsf{wdata})$. After receiving input op_i, ORAM interacts with mem$'$, and produces read/write operations into mem$'$ as output, denoted by $\vec{\mathsf{op}}_i$. These operations $\vec{\mathsf{op}}_i$ implicitly define memory contents of mem.

We say that an ORAM algorithm is correct, if for any n, for any input sequence $(\mathsf{op}_1, \ldots, \mathsf{op}_n)$, ORAM outputs correctly. In other words, the memory contents of mem implicitly defined by mem$'$ after execution of $(\vec{\mathsf{op}}_1, \ldots, \vec{\mathsf{op}}_n)$ is identical to the memory contents of mem defined by executing $(\mathsf{op}_1, \ldots, \mathsf{op}_n)$ on mem. We say that an ORAM scheme ORAM is *oblivious* if there exists a polynomial-time simulator Sim such that no polynomial time adversary $\mathcal{A}$ can distinguish between the transcript of the real ORAM execution and a simulated transcript that Sim outputs. Sim is given only N and w, even when the simulated memory access are provided one-by-one to $\mathcal{A}$.

Remark: ORAM initialization. In this paper, we assume an ORAM starts out with a memory array where the first N words are non-zero (reflecting the initial unshuffled memory), followed by all zeros. Most ORAM schemes require an initialization procedure to shuffle the initial memory contents. In this paper, we assume that the ORAM algorithm performs a linear scan of first N memory locations and inserts them into ORAM. This is used by the simulator in our proof to extract the input used for execution of the program. We use the convention that such initialization is performed by the ORAM algorithm upon the first read or write operation — therefore our notation does not make such initialization explicit. This also means that the first ORAM operation will incur a higher overhead than others.

$$\mathcal{F}^{\mathcal{RAM}}_{obf}[\text{sender, receiver}]$$

On receive ("create", RAM) from sender for the first time:

   Create a unique nonce denoted pid

   Store (pid, RAM), send ("create", pid) to receiver

On receive ("execute", pid, inp) from receiver:

   assert (pid, RAM) is stored for some RAM

   outp := RAM(inp), send outp to receiver

Fig. 2: **Ideal Functionality** $\mathcal{F}^{\mathcal{RAM}}_{obf}$. Although there can be multiple instances of this ideal functionality, we omit writing the session identifier explicity without risk of ambiguity. In this paper, we adopt the same UC notational conventions as Pass, Shi, and Tramer [44]. In particular, we parametrize each functionality and protocol by its session identifier, and the identifiers of the parties involved — although in this paper, we omit writing the session identifier explicitly without risk of ambiguity.

B. $\mathcal{F}^{\mathcal{RAM}}_{obf}$: Modeling Obfuscation in UC

The ideal functionality for obfuscation $\mathcal{F}^{\mathcal{RAM}}_{obf}$ is described in Figure 2. The sender sends the description of a RAM program, RAM $\in \mathcal{RAM}$ and a program ID pid, using the "create" query. The functionality stores this program, pid, the sender and receiver. When the receiver invokes "execute" query on an input inp, it evaluates the program on inp, and returns output outp.

C. Scheme Description

We now provide the complete description of our scheme. We model the secure hardware token through the $\mathcal{F}_{token}$ functionality (Figure 3). Our construction realizes $\mathcal{F}^{\mathcal{RAM}}_{obf}$ in the $\mathcal{F}_{token}$-hybrid model [27] and is described in Figure 4.

In order to account for all possible token queries that may be required for an ORAM scheme, $\mathcal{F}_{token}$ relies on an internal, transient instance of $\mathcal{F}_{internal}$ to execute each step of the program evaluation. Each time $\mathcal{F}_{token}$ yields control to the receiver, the entire state of $\mathcal{F}_{internal}$ is destroyed. Whenever the receiver calls back $\mathcal{F}_{token}$ with $\overline{\text{state}}$, $\mathcal{F}_{token}$ once again creates a new, transient instance of $\mathcal{F}_{internal}$, sets its state to the decrypted state, and invokes $\mathcal{F}_{internal}$ to execute next step.

The sender. Let the program to be obfuscated be RAM := ⟨cpustate$_0$, mem$_0$⟩ where mem$_0$ is a list of program instructions. The sender first creates the token containing a hardwired secret key K where $K := (K_1, K_2, K_3)$. K_1 is used as the encryption key for encrypting state, K_2 is used as the key to a pseudorandom function used by the ORAM and K_3 is used as the key for a pseudorandom function used by sstore (described later). This is modeled by our functionality using the "store key" query (Figure 4 line 1). The sender then encrypts mem$_0$ (one instruction at a time) to obtain $\overline{\text{mem}_0}$. It creates a Merkle root $H_S := \text{digest}(\overline{\text{mem}_0})$, which is used by $\mathcal{F}_{token}$ during execution to verify integrity of the program. The sender creates an encrypted header $\overline{\text{header}} := \text{Enc}_{K_1}(\text{cpustate}_0, H_S, \text{RAM.params})$ where RAM.params $= \{T, N, \ell_{\text{in}}, \ell_{\text{out}}, w\}$. The sender sends $\overline{\text{header}}$, $\overline{\text{mem}_0}$, and RAM.params as the obfuscated program to the receiver. As the

obfuscated program consists of only the encrypted program and metadata, for a program of size P bits, the obfuscated program has size $P + O(1)$ bits. In the real world, the sender sends the hardware token implementing functionality $\mathcal{F}_{token}$ to the receiver. The receiver can use the same stateless token to execute multiple obfuscated programs sent by the sender.

$$\mathcal{F}_{token} [\text{sender, receiver}]$$

// Store the secret key K in the token

On receive ("store key", K) from sender:

   Store the secret key K, ignore future "store key" inputs

   Send "done" to sender

// This step commits the <u>receiver</u> to his input through H_R

On receive ("initialize", $\overline{\text{header}}$, H_R) from receiver:

   Parse $K := (K_1, K_2, K_3)$

   (cpustate$_0$, H_S, RAM.params) := $\text{Dec}_{K_1}(\overline{\text{header}})$; abort if fail

   state := {ssid := (H_S, H_R), $time := 0$,

            rdata := 0, cpustate := cpustate$_0$,

            sstorestate := ("init", $H_S, H_R, H' := 0$),

            oramstate := "init", params := RAM.params}

   send $\overline{\text{state}} := \text{Enc}_{K_1}(\text{state})$ to receiver

On receive (_) from $\mathcal{F}_{internal}$: // ORAM queries

   $\overline{\text{state}} := \text{Enc}_{K_1}(\mathcal{F}_{internal}.\text{state})$

   send (_, $\overline{\text{state}}$) to receiver

On receive (_, $\overline{\text{state}}$) from receiver: // ORAM queries

   state := $\text{Dec}_{K_1}(\overline{\text{state}})$, abort if fail

   Instantiate a new instance $\mathcal{F}_{internal}$,

   set $\mathcal{F}_{internal}.\text{state} := \text{state}$, and $\mathcal{F}_{internal}.K := K$

   Send _ to $\mathcal{F}_{internal}$

$$\mathcal{F}_{internal}$$

Define $\mathcal{F}_{internal}.\text{state} \overset{\text{alias}}{:=}$ (ssid, $time$, rdata, cpustate, sstorestate, oramstate, params)

// execute program

On receive ("execute one step") from $\mathcal{F}_{token}$:

  1: assert $time \leq \text{params}.T$

  2: (cpustate, op) $\leftarrow \Pi'(\text{cpustate}, \text{rdata})$

  3: Send op to ORAM[$\text{PRF}_{K_2}(\text{ssid})$, oramstate] $\Leftrightarrow$ sstore[$\text{PRF}_{K_3}(\text{ssid})$, sstorestate] $\Leftrightarrow \mathcal{F}_{token}$, wait for output from ORAM, abort if sstore aborts; /* instantiate ORAM with state oramstate, instantiate sstore with state sstorestate, connect ORAM's communication tape to sstore's input tape, connect sstore's communication tape to caller $\mathcal{F}_{token}$. This represents a multi-round protocol. */

  4: If op = (read, . . .), let rdata := output

  5: $time := time + 1$

  6: If $time = \text{params}.T$: send ("okay", rdata) to $\mathcal{F}_{token}$

      Else send ("okay", $\bot$) to $\mathcal{F}_{token}$

Fig. 3: **Functionality** $\mathcal{F}_{token}$. For succinctness, encryption of some data is represented using an overline on it, e.g., $\overline{\text{state}} = \text{Enc}_{K_1}(\text{state})$, where Enc denotes a IND-CPA + INT-CTXT-secure authenticated encryption scheme. "_" denotes a wildcard field that matches any string.

Prot$_{obf}$[sender, receiver]

Sender:

On receive ("create", RAM = ⟨cpustate$_0$, mem$_0$⟩) from env:

1: If not initialized: $K := (K_1, K_2, K_3) \xleftarrow{\$} \{0,1\}^{3\lambda}$, send ("store key", K) to $\mathcal{F}_{token}$, await "done"

2: $\overline{\text{mem}}_0 := \{\text{Enc}_{K_1}(\text{mem}_0[i])\}_{i \in |\text{mem}_0|}$

3: $H_S := \text{digest}(\overline{\text{mem}}_0)$ // H_S: program Merkle root

4: $\overline{\text{header}} := \text{Enc}_{K_1}(\text{cpustate}_0 || H_S || \text{RAM.params})$

5: Send ($\overline{\text{header}}$, $\overline{\text{mem}}_0$, RAM.params) to receiver

Receiver:

On receive ("execute", pid, inp) from env:

1: Await ($\overline{\text{header}}$, $\overline{\text{mem}}_0$, RAM.params) from sender s.t. RAM.params.H_S = pid if not received already

2: Initialize mem := $\overline{\text{mem}}_0 || \text{inp} || \vec{0}$

3: Send ("initialize", $\overline{\text{header}}$, $H_R := \text{digest}(\text{inp})$) to $\mathcal{F}_{token}$, await $\overline{\text{state}}$ from $\mathcal{F}_{token}$

4: **for** t in $\{1, \ldots, T\}$:

5:    Send ("execute one step", $\overline{\text{state}}$) to $\mathcal{F}_{token}$

6:    Await (oper, $\overline{\text{state}}$) from $\mathcal{F}_{token}$;   // $\overline{\text{state}}$ overwritten with the received value

7:    Until oper = ("okay", _), repeat: //multiple memory requests for the RAM step due to ORAM

8:       perform the operation oper on mem and let the response be res

9:       forward (res, $\overline{\text{state}}$) to $\mathcal{F}_{token}$, and await (oper, $\overline{\text{state}}$) from $\mathcal{F}_{token}$;

10: Parse oper := ("okay", outp), output outp

Fig. 4: **Protocol Prot$_{obf}$.** Realizes $\mathcal{F}^{\mathcal{RAM}}_{obf}$ in the $\mathcal{F}_{token}$-hybrid model.

The receiver. On the receiver's side, the token functionality makes use of an ORAM and a secure store sstore. The token functionality (trusted hardware functionality) is modeled by an augmented RAM machine.

1) **ORAM.** ORAM takes in $[\kappa := \text{PRF}_{K_2}(\text{ssid}), \text{oramstate}]$ (where ssid := (H_S, H_R)) as internal secret state of the algorithm. κ is a session-specific seed used to generate all pseudorandom numbers needed by the ORAM algorithm — recall that all randomness needed by ORAM is replaced by pseudorandomness to avoid rewinding attacks. As mentioned in Section IV-A, we assume that the ORAM initialization is performed during the first read/write operation. At this point, the ORAM reads the first N memory locations to read the program and the input, and inserts them into the ORAM data structure within mem.

2) **Secure store module** sstore. sstore is a stateful deterministic secure storage module that sits in between the ORAM module and the untrusted memory implemented by the receiver. Its job is to provide appropriate memory encryption and authentication. sstore's internal state includes $\kappa := \text{PRF}_{K_3}(\text{ssid})$ and sstorestate. sstorestate contains a succinct digest of program, input and memory to perform memory authentication. κ is a session-specific seed used to generate all pseudorandom numbers for memory encryption.

At the beginning of an execution, sstorestate is initialized

Fig. 5: **Augmented Random Access Machine.** In this figure, cpustate$_i$ is denoted by cs_i and rdata$_i$ is denoted by rd_i.

to sstorestate := $(H_S, H_R, H' := 0)$, where H_S denotes the Merkle root of the encrypted program provided by the sender, H_R denotes the Merkle root of the (cleartext) input and H' denotes the Merkle root of the memory mem. By convention, we assume that if a Merkle tree or any subtree's hash is 0, then the entire subtree must be 0. The operational semantics of sstore is as follows: upon every data access request $(\texttt{read}, addr)$ or $(\texttt{write}, addr, \texttt{wdata})$:

- If $addr$ is in the mem$_0$ part of the memory (the sender-provided encrypted program), interact with mem and use H_S to verify responses. Update H_S appropriately if the request type is $\texttt{write}$.
- If $addr$ is in the inp part of the memory (the receiver-provided input), interact with mem and use H_R to verify responses.
- Otherwise, interact with mem and use H' to verify responses. Update H' appropriately.

Upon successful completion, sstore outputs the data fetched for read requests, and outputs 0 or 1 for write requests. Note that the sstore algorithm simply aborts if any of the responses fail verification.

3) **Augmented Random Access Machines.** We now extend the RAM model to support instruction scheduling and a scratchpad (Sections III-C and III-D). $\mathcal{RAM}$ can be augmented to use a next instruction circuit $\Pi' := \Pi^N$ for a fixed N, with the following modifications:

a) Π' is a combinational circuit, which consists of N next-instruction circuits Π_i cascaded as shown in Figure 5.

b) The Π_i's use an additional shared memory, referred to as scratchpad. Each Π_i (except Π_1) operates on the output of Π_{i-1} and an operand rdata$_{i-1}$ read from scratchpad. The next instruction circuit Π' outputs op$_N$ to retrieve rdata from mem, which is subsequently used by Π_1.

On input inp, the execution of RAM$[T, N, \ell_{\text{in}}, \ell_{\text{out}}, w] := \langle \Pi', \text{cpustate}, \text{mem} \rangle$ is similar to what was defined in Section IV-A but uses Π' as the next instruction circuit. The augmented random access machine $\mathcal{RAM}'$ models a RAM that performs N instructions followed by an ORAM access. If some op$_i$ cannot be served by the scratchpad, subsequent Π_j for $i+1 \le j \le N$ do not update cpustate$_j$ and output op$_N$ = op$_i$ to load the required data in scratchpad.

Remark. For augmented random access machines that uses a scratchpad, rdata would typically be larger than a memory word (e.g. 512 bits).

We now explain how the receiver executes the program using the token described in Figure 3 and protocol in Figure 4.

Program execution. For ease of explanation, let us first assume that the ORAM is initialized and contains the program and input. The execution for any input proceeds in T time steps

(Figure 4 line 4). At each time step, the receiver interacts with the token with two types of queries. For each type of query, $\mathcal{F}_{token}$ decrypts $\overline{\text{state}}$ (aborts if decryption fails), instantiates $\mathcal{F}_{internal}$ with state and forwards the request to $\mathcal{F}_{internal}$. At the end of query, the $\overline{\text{state}}$ is sent to the receiver along with the query response.

- **Execute one step**: This is shown in Figure 3 and Figure 4 line 5. When this query is invoked, $\mathcal{F}_{internal}$ executes the next instruction circuit Π' of the RAM machine to obtain an updated cpustate and an op $\in \{\texttt{read}, \texttt{write}\}$. Once operation op is performed by the ORAM algorithm, $\mathcal{F}_{internal}$ updates state.$time$ to reflect the execution of the instruction (Figure 3 line 5). The message "okay" is then sent to the receiver. At $time = T$, $\mathcal{F}_{internal}$ returns the program output to the receiver (Figure 3 line 6).
- **ORAM queries**: ORAMs can use a multi-round protocol (with possibly different types of queries) to read/write (Figure 3 line 3). It interacts with mem stored at the receiver through $\mathcal{F}_{token}$ (Figure 4 lines 7-9). To account for instantiation of any ORAM, $\mathcal{F}_{token}$ is shown to receive any query from receiver (indicated by wildcard (_) in Figures 3 and 4). These queries are sent to $\mathcal{F}_{internal}$ and vice-versa.

For each interaction with mem, sstore encrypts (resp. decrypts) data sent to (resp. from) the receiver. Moreover, sstore authenticates the data sent by the receiver. This completes the description of execution of the program.

Initialization. To initialize the execution, the receiver first starts by storing the program and input inp in its memory mem := $\overline{\text{mem}_0}||\text{inp}||\vec{0}$. It commits to its input by invoking "initialize" (Figure 4 line 3) and sending a Merkle root of its input ($H_R = \text{digest}(\text{inp})$) along with header := $\text{Enc}_{K_1}(\text{cpustate}_0||H_S||\text{RAM.params})$. $\mathcal{F}_{token}$ initializes the parameters, creates $\overline{\text{state}}$ and sends it to the receiver.

The ORAM and sstore are initialized during the first invocation to "execute one step", i.e., $t = 1$ in Figure 4, line 4. The required randomness is generated pseudorandomly based on (K_2, H_S, H_R) for ORAM and (K_3, H_S, H_R) for sstore. As mentioned in Section IV-A, during initialization, ORAM in $\mathcal{F}_{token}$ reads mem_0 word by word (not shown in figure). For each word read, sstore performs Merkle tree verification with $H_S := \text{digest}(\text{mem}_0)$. Similarly, when the input is read, sstore verifies it with $H_R := \text{digest}(\text{inp})$. sstorestate and oramstate uniquely determine the initialization state. Hence, if the receiver rewinds, the execution trace remains the same. The commitment H_R ensures that the receiver cannot change his input after invoking "initialize". This completes the formal scheme description of the UC functionality $\mathcal{F}_{token}$.

D. Proof of Security

Theorem 1. *Assuming that Enc is an INT-CTXT + IND-CPA authenticated encryption scheme,* ORAM *satisfies obliviousness (Section IV-A),* sstore *adopts a semantically secure encryption scheme and a collision resistant Merkle hash tree scheme and the security of* PRF*, the protocol described in Figures 3 and 4 UC realizes $\mathcal{F}_{obf}^{\mathcal{RAM}}$ (Figure 2) in the $\mathcal{F}_{token}$-hybrid model.*
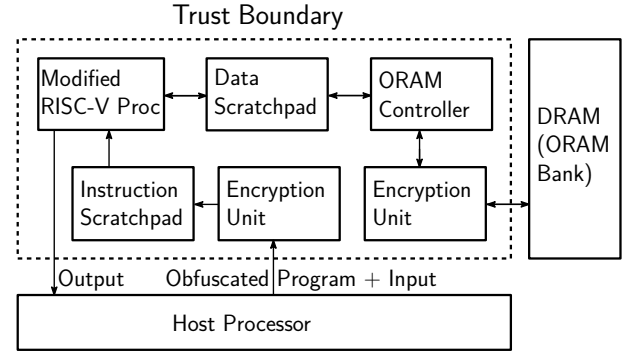

Fig. 6: HOP Architecture

Proof. We refer the reader to the full version of the paper for a detailed proof of security.

V. IMPLEMENTATION

The final architecture of HOP (with the optimizations from Section III) is shown in Figure 6. We now describe implementation-specific details for each major component.

A. Modified RISC-V Processor and Scratchpad

We built HOP with a RISC-V processor which implements a single stage 32bit integer base user-level ISA developed at UC Berkeley [12]. A RISC-V C cross-compiler is used to compile C programs to be run on the processor. The RISC-V processor is modified to include a 16 KB instruction scratchpad and a 512 KB data scratchpad (Section III-D). The RISC-V processor and the compiler are modified accordingly to accommodate the new scratchpad load/unload instructions (described below). While HOP uses a single stage RISC-V processor, our system does not preclude additional hardware optimizations in commodity processors such as multi-issue, branch predictor, etc. Our only requirement to support such processor structures is the ability to calculate, for that program over all inputs, a suitably conservative maximum runtime T.

New scratchpad instructions. For our prototype, we load the scratchpad using a new instruction called **spld**, which is specified as follows:

$$\textbf{spld } addr, \#mem, spaddr$$

In particular, $addr$ is used to specify the starting address of the memory that needs to be loaded in scratchpad. $\#mem$ is the number of memory locations to be loaded on the scratchpad starting at $addr$ and $spaddr$ is the location in scratchpad to store the loaded data. When the processor intercepts an **spld** instruction, it performs two operations: 1. It writes back the data stored in this scratchpad location to the appropriate address in main memory (ORAM). 2. It reads $\#mem$ memory locations starting at main memory address $addr$ into scratchpad locations starting at $spaddr$. Of course, **spld**'s precise design is not fundamental: we need a way to load an on-chip memory such that it is still feasible to statically determine T.

Example scratchpad use. Figure 7 shows an example scenario where **spld** is used. The program shows a part of the code

```
 1: int decompress(char *chunk) {
 2:    int compLen = 0;
 3:    // initial processing
 4:    burrowsWheeler(chunk, compLen);
 5:    // more processing
 6:    writeOutput(chunk);
 7:    return compLen;
 8: }
 9: void main() {
10:    char *inp = readInput();
11:    for (i = 0; i < len(inp); i += len) {
12:      spld(inp + i, CSIZE, 0);
13:       len = decompress(inp + i);
14:    }
15: }
```

Fig. 7: Example program using **spld**: `bzip2`

used for decompressing data using the `bzip2` compression algorithm. The algorithm decompresses blocks of compressed data and outputs data of size CSIZE independently. Each block of data may be read and processed multiple times during different steps of compression (run-length encoding, Burrows-Wheeler transform, etc.). Hence, each such block is loaded into the scratchpad (line 12) before processing. This ensures that every subsequent access to this data is served by the scratchpad instead of memory (thereby reducing expensive ORAM accesses). After decompressing the block, **spld** is executed for the next block of compressed data.

B. ORAM Controller

We use a hardware ORAM controller called 'Tiny ORAM' from [19], [20]. The ORAM controller implements an ORAM tree with 25 levels, having 4 blocks per bucket. Each block is 512 bits (64 Bytes) to match modern processor cache line size. This corresponds to a total memory of 4 GB. The ORAM controller uses a stash of size 128 blocks and an on-chip position map of 256 KB. For integrity and freshness, Tiny ORAM uses the PosMap MAC (PMMAC) scheme [19]. We note that PMMAC protects data integrity but does not achieve malicious security. We estimate the cost of malicious security using a hardware Merkle-tree on ORAM in Table II. We disable the PosMap Lookaside Buffer (PLB) in Freecursive ORAM to avoid leakage through the total number of ORAM accesses.

C. Encryption Units

For all encryption units, we use tinyaes from Open-Cores [2]. The encryption units communicate with the external DRAM (bandwidth of 64 Bytes/cycle) as well as the host processor. Data is encrypted before writing to the DRAM. Similarly, all data read from the DRAM is decrypted first before processed by the ORAM controller. Another encryption unit is used to decrypt the obfuscated program before loading it into the instruction scratchpad.

VI. EVALUATION

We now present a detailed evaluation of HOP for some commonly used programs, and compare HOP to prior work.

TABLE II: **Resource allocation and utilization of HOP on Xilinx Virtex V7485t FPGA.** For each row, first line indicates the estimate. % utilization is mentioned in parentheses. LUT: Slice LookUp Table, FFs: Flip-flops or slice registers, BRAM: Block RAM.

	LUT	FFs	LUT-Mem	BRAM
Total Estimate	169472	51870	81112	566.5
(% Utilization)	(55.8%)	(8.5%)	(62.0%)	(55.0%)
HOP Estimate	103462	39803	38725	437
(% Utilization)	(34.0%)	(6.6%)	(47.7%)	(42.4%)
(HOP− ORAM) Estimate	21626	6579	1	83
(% Utilization)	(7.1%)	(1.1%)	(∼0%)	(8.1%)
Estimate with Merkle tree	221041	81410	81126	566.5
(% Utilization)	(72.8%)	(13.4%)	(62.0%)	(55.0%)

A. Methodology

We measure program execution time in processor cycles, and compare with our own baseline scheme (to show the effectiveness of our optimizations), an insecure processor as well as related prior work. For each program, we choose parameters so that our baseline scheme requires about 100 million cycles to execute. We also report processor idle time, the time spent on dummy arithmetic instructions and dummy memory accesses to adhere to an $A^N M$ schedule (Section III-C).

For the programs we evaluate (except *bzip2*; c.f., Section VI-D), we calculate T manually. We remark that the average input completion time and worst case time are very similar for these programs. To find T for larger programs, one may use established techniques in determining worst case execution time (e.g., a tool from [54]).

In our prototype, evaluating an arithmetic instruction takes 1 cycle while reading/writing a word from the scratchpad takes 3 cycles. Given the parameters in Section V-B, an ORAM access takes 3000 cycles. For our HOP configurations with a scratchpad, we require both scratchpad read/writes and arithmetic instructions to take 3 cycles in order to hide which is occurring. Following Section III-C, we set $N = 3000$ when not using a scratchpad; with a scratchpad, we use $N = 1000$. For our evaluation, we consider programs ranging from those with high locality (e.g., `bwt-rle`) to those that show no locality (e.g., `binsearch`).

B. Area Results

We synthesized, placed and routed HOP on a Xilinx Virtex V7485t FPGA for parameters described in Section V. HOP operates at 79.3 MHz on this FPGA. The resource allocation and utilization figures are mentioned in Table II. The first three rows represent the total estimate, estimate for HOP (i.e. excluding RISC-Vprocessor, and the scratchpad) and an estimate for HOP that does not account for ORAM. The last row shows the total overhead including an estimate for a Merkle tree scheme. Excluding the processor, scratchpad and ORAM, HOP consumes < 9% of the FPGA resources. We see that the total area overhead of HOP is small and can be built on a single FPGA chip.

Fig. 8: Execution time for different programs with (i) baseline scheme, (ii) $A^N M$ schedule and (iii) Scratchpad + $A^N M$.

C. Main Results

Figure 8 shows the execution time of HOP variants relative to an insecure processor. For each program, there are three bars shown. The first bar is for the baseline HOP scheme (i.e., Section III-B only); the second bar only uses an $A^N M$ schedule without a scratchpad (adds Section III-C); and the third bar is our final scheme that uses a scratchpad and the $A^N M$ schedule (adds Section III-D). All schemes are relative to an insecure processor that does not use ORAM or hide what instruction it is executing. We assume this processor uses a scratchpad that has the same capacity as HOP in Section V-A. The time required to insert the program and data is not shown.

Comparison of HOP variants. As can been seen in the figure, the $A^N M$ schedule without a scratchpad gives a $1.5\times \sim 18\times$ improvement. Adhering to an $A^N M$ schedule requires some dummy arithmetic or memory instructions during which the processor is essentially idle. We observe that for our programs, the idle time ranges between 43% and 49.9% of the execution time, consistent with the claim in Section III-C.

Effect of a scratchpad. The effect of a scratchpad largely depends on program locality. We thus classify programs in our evaluation into four classes:

1) Programs such as `binsearch, heappop` do not show locality. Thus, a scratchpad does not improve performance.
2) Programs such as `sum, findmax` stream (linear scan) over the input data. Given that an ORAM block is larger than a word size (512 bits vs 32 bits in our case), a scratchpad in these streaming applications can serve the next few (7 with our parameters) memory accesses after **spld**. A larger ORAM block size can slightly benefit these applications while severely penalize programs with no locality, and therefore is not a good trade-off.
3) Programs that maintain a small working set at all times will greatly benefit from a scratchpad. We evaluate one such program `bwt-rle`, which performs Burrows-Wheeler transform and run length encoding, and is used in compression algorithms.
4) Lastly, some programs are a mix of the above cases — some data structures can be entirely loaded into the scratchpad whereas some cannot (e.g. a Radix sort program).

Comparison to insecure processor. The remaining performance overhead of the optimized HOP (the third bar) comes from several sources. First, the performance of ORAM: The number of cycles to perform a memory access using ORAM is much higher than a regular DRAM. In HOP, an ORAM access is $40\times$ more expensive than an insecure access. Second, dummy accesses to adhere to a schedule: As shown in Section III-C, the performance overhead due to dummy accesses $\leq 2\times$. For programs such as `bwt-rle`, HOP has a slowdown as low as $8\times$. This is primarily due to the reduction in ORAM accesses by maintaining a small working set in the scratchpad.

D. Case Study: bzip2

To show readers how our system performs on a realistic and complex benchmark, we evaluate HOP on the open-source algorithm `bzip2` (re-written for a scratchpad, cf. Figure 7). We evaluate the decompression algorithm only, as the decompression algorithm's performance does not heavily depend on the input if one fixes the input size [1]. This allows us to run an average case input and use its performance to approximate the effect of running other inputs. To give a better sense for how the optimizations are impacted by different inputs, we don't terminate at a worst-case time T but rather terminate as soon as the program completes.

We run tests on two inputs, both highly compressible strings. For the first input, HOP achieves $106\times$ speedup over the baseline scheme and $17\times$ slowdown over the insecure version. For the second input, HOP achieves $234\times$ speedup over the baseline and $8\times$ slowdown over the insecure version. Thus, the gains and slowdowns we see from the prior studies extend to this more sophisticated benchmark.

E. Comparison with Related Work

We now compare against prior work on obfuscation with hardware (these prior works were not implemented) and several works with related threat models.

1) Comparison to prior obfuscation from trusted hardware proposals [15], [17], [27]: We now compare against [15], [17], [27] which describe obfuscation using trusted hardware. Note that none of these schemes were implemented.

Part of the proposals in [15], [17] require programs to be run as *universal circuits* under FHE while [27] evaluates programs as universal circuits directly on hardware (i.e., by feeding the encrypted inputs of each gate into a stateless hardware unit: where it decrypts the inputs, evaluates the gate, and re-encrypts the output). We will now compare HOP to these circuit-based approaches. Again, we stress that all of [15], [17], [27] require the use of trusted hardware for their complete scheme and thus can be viewed similarly to HOP from a security perspective.

Table III shows the speedup achieved by HOP relative to universal circuits run under FHE (left) and bare hardware (right). We assume the cost of a universal circuit capable of evaluating any c gate circuit is $18 * c * \log c$ gates [39]. We compare the approaches on the `findmax` and `binsearch` benchmarks, using a dataset size of 1 GB for each. We show `findmax` as it yields a very efficient circuit and a best-case situation for the circuit approach (relative to the corresponding RAM program); `binsearch` shows the other extreme.

For [15], [17], we assume a BGV-style FHE scheme [9], using the NTRU cryptosystem, with polynomial dimension and ciphertext space parameters chosen using [24], to achieve 80 bits of security.[6] For [27], we assume each NAND gate takes 20 cycles to evaluate (10 cycles for input decryption with AES, 0 cycles for evaluation, 10 cycles for re-encryption). For HOP, we assume the parameters from Section V.

| | FHE [15], [17] | | Hardware [27] | |
	On + Off	On	On + Off	On
findmax	$1 * 10^9$	$2 * 10^9$	$4 * 10^3$	$1 * 10^4$
binsearch	$4 * 10^9$	$4 * 10^{15}$	$6 * 10^3$	$1 * 10^{10}$

TABLE III: HOP speedup ($\times$) relative to universal circuit approaches. findmax and binsearch are over 1 GB datasets.

In the Table, On+Off ('online and offline') assumes one search query is run: in that case, HOP's performance is reduced due to the time needed to initially load the ORAM. The On ('online only') column shows the amortized speedup when many search queries are made without changing the underlying search database (i.e., without re-loading the ORAM each time). This shows an inherent difference to works based on universal circuits: those works represent programs as circuits, where optimized algorithms such as binsearch do not see speedup. In all cases, HOP shows orders of magnitude improvement to the prior schemes.

We note that our comparison to [15], [17] is conservative: we only include FHE's time to perform AND/OR gate operations and not the cost of auxiliary FHE operations (re-linearization, modulus switching, bootstrapping, etc). Lastly, FHE is only one part of [15], [17]: we don't include the cost of NIZK protocols, etc. which those schemes also require.

2) Comparison with $i\mathcal{O}$ [37]: We compare HOP with an implementation of indistinguishability obfuscation ($i\mathcal{O}$) that does not assume a trusted hardware token. Note that while VBB obfuscation is not achievable in general, $i\mathcal{O}$ is a weaker notion of obfuscation. With [37], evaluating an 80-bit point function (a simple function that is 0 everywhere except at one point) takes about 180 seconds while HOP takes less than a msec, which is about 5-6 orders of magnitude faster.

3) Comparison with GhostRider [40]: Recall from Section II that GhostRider protects input data to the program *but not the program*. Since our privacy guarantee is strictly greater than GhostRider, we now compare to that work to show the cost of extra security. Note: we compare to the GhostRider compiler and not the implementation in [40] which uses a different parameterization for the ORAM scheme. This comparison shows the additional cost that is incurred by HOP to hide the program. We don't show the full comparison for lack of space, but point out the following extreme points: For programs with unpredictable access patterns (binsearch, heappop), GhostRider outperforms HOP by $\sim 2\times$. HOP's additional overhead is from executing dummy instructions to

adhere to a particular schedule. For programs with predictable access patterns (sum, findmax, hist), GhostRider's performance is similar to that of an insecure processor.

F. Time for Context Switch

Since it was not required for our performance evaluation, we have not yet implemented context switching (Section III-E) in our prototype. Recall, context switching means the receiver interrupts the processor, which encrypts and writes out all the processor state (including CPU state, instruction scratchpad, data scratchpad, ORAM position map and stash) to DRAM. We estimate the time of a context switch as follows. The total amount of data stored by our token is ~ 800 KB (Section V). Assuming a DRAM bandwidth of 10 GB/s and a matching encryption bandwidth, it would take $\sim 160\mu s$ to perform a context switch to run another program. Note that this assumes all data for a swapped-out context is stored in DRAM (i.e., the ORAM data already in the DRAM need not be moved). If it must be swapped out to disk because the DRAM must make room for the new context, the context switch time grows proportional to the ORAM size.

VII. PRACTICAL DEPLOYMENT AND APPLICATIONS

In this section, we present practical deployment considerations and potential applications for a HOP processor.

Parties involved in the system. In a practical deployment, there would be three parties involved in this system: The sender is a software provider (e.g. Microsoft), the receiver is the end user and the manufacturer is a hardware company (e.g., Intel, TSMC). Software providers are incentivized to use this framework to hide the IP of their proprietary programs so as to sell those programs to customers without being pirated. Hardware manufacturers are incentivized to provide security in order to retain the software providers as customers (e.g., Intel SGX was initially envisioned as a buy-in service).

Potential applications. The focus of this paper is to build hardware that satisfies the definition of VBB obfuscation. Thus, our model assumes obfuscation of 'batch programs' – those which take inputs, and compute non-interactively until a result is produced. Some examples of such programs are compilers, compression algorithms, machine learning algorithms, etc. In our setting, the program itself should contain some sensitive IP (to warrant obfuscation). Given the pervasive nature of batch programs, we see HOP being applicable in both commercial and military settings. We note that even the military outsources its fabrication (and therefore its trust) to external foundries (e.g., global foundries handles runs for the NSA).

Beyond batch programs, it is possible to support streaming applications with little change to the model. In particular, while HOP runs, it can accept streams of public data (e.g., a video feed) in a similar fashion to Stream Ascend [57]. Importantly, this has no impact on security, as long as HOP doesn't change its observable behavior given the data in the stream and the program accepts this data at fixed intervals of time.

---

[6]When represented as circuits, both findmax and binsearch look like a linear PIR. Over a 1 GByte dataset, we evaluate this function with a 10-level FHE circuit, which gives an FHE polynomial dimension (n) of ~ 8192 and ciphertext space q of $\sim 2^{128}$ (using terminology from [9]). With these parameters, a single polynomial multiplication/addition using NTL [50] costs 14 ms / .4 ms on a 3 GHz machine.

VIII. CONCLUSION

This paper makes two main contributions. First, we construct an optimized hardware architecture - called HOP - for running obfuscated RAM programs. We give a matching theoretic model for our optimized architecture and prove it secure. A by-product of our analysis shows the first obfuscation for RAM programs using 'stateless' tokens. Second, we present a complete implementation of our optimized architecture and evaluate it on real-world programs. The complete design requires 72% the area of a V7485t Field Programmable Gate Array (FPGA) chip. Run on a variety of benchmarks, HOP achieves an average overhead of $8\times \sim 76\times$ relative to an insecure system. To the best of our knowledge, this effort represents the *first* implementation of a provably secure VBB obfuscation scheme in any model under any assumptions.

REFERENCES

[1] "bzip2 man pages," http://www.bzip.org/1.0.5/bzip2.txt.

[2] "Open cores," http://opencores.org/.

[3] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi, "The em sidechannel (s)," in *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 2002, pp. 29–45.

[4] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang, "On the (im)possibility of obfuscating programs," in *CRYPTO*, 2001.

[5] N. Bitansky, R. Canetti, S. Goldwasser, S. Halevi, Y. T. Kalai, and G. N. Rothblum, "Program obfuscation with leaky hardware," in *Advances in Cryptology–ASIACRYPT 2011*. Springer, 2011, pp. 722–739.

[6] N. Bitansky, S. Garg, H. Lin, R. Pass, and S. Telang, "Succinct randomized encodings and their applications," in *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*. ACM, 2015, pp. 439–448.

[7] A. Bittau, A. Belay, A. Mashtizadeh, D. Mazières, and D. Boneh, "Hacking blind," in *IEEE S&P*, 2014.

[8] D. Boneh, R. A. DeMillo, and R. J. Lipton, "On the importance of checking cryptographic protocols for faults," in *International Conference on the Theory and Applications of Cryptographic Techniques*, 1997.

[9] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, "(Leveled) Fully Homomorphic Encryption without Bootstrapping," in *ITCS*, 2012.

[10] R. Canetti, "Universally composable security: A new paradigm for cryptographic protocols," in *FOCS*, 2001.

[11] R. Canetti, J. Holmgren, A. Jain, and V. Vaikuntanathan, "Succinct garbling and indistinguishability obfuscation for RAM programs," in *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*. ACM, 2015, pp. 429–437.

[12] C. Celio and E. Love, "The sodor processor collection," http://riscv.org/download.html#tab_sodor.

[13] D. Champagne and R. B. Lee, "Scalable architectural support for trusted software," in *HPCA*, 2010.

[14] J. H. Cheon, K. Han, C. Lee, H. Ryu, and D. Stehlé, "Cryptanalysis of the multilinear map over the integers," in *EUROCRYPT*, 2015.

[15] K.-M. Chung, J. Katz, and H.-S. Zhou, "Functional encryption from (small) hardware tokens," in *ASIACRYPT*, 2013.

[16] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Communications of the ACM*, 2008.

[17] N. Döttling, T. Mie, J. Müller-Quade, and T. Nilges, "Basing obfuscation on simple tamper-proof hardware assumptions." *IACR Cryptology ePrint Archive*, 2011.

[18] C. W. Fletcher, M. v. Dijk, and S. Devadas, "A secure processor architecture for encrypted computation on untrusted programs," in *STC*, 2012.

[19] C. W. Fletcher, L. Ren, A. Kwon, M. van Dijk, and S. Devadas, "Freecursive oram:[nearly] free recursion and integrity verification for position-based oblivious ram," in *ASPLOS*, 2015.

[20] C. W. Fletcher, L. Ren, A. Kwon, M. Van Dijk, E. Stefanov, D. Serpanos, and S. Devadas, "A low-latency, low-area hardware oblivious ram controller," in *FCCM, 2015*. IEEE, 2015.

[21] S. Garg, "Program obfuscation via multilinear maps," in *Security and Cryptography for Networks*, 2014.

[22] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters, "Candidate indistinguishability obfuscation and functional encryption for all circuits," in *FOCS*, 2013.

[23] C. Gentry, S. Halevi, S. Lu, R. Ostrovsky, M. Raykova, and D. Wichs, "Garbled ram revisited," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2014.

[24] C. Gentry, S. Halevi, and N. P. Smart, "Homomorphic evaluation of the aes circuit," in *Advances in Cryptology–CRYPTO 2012*, 2012.

[25] O. Goldreich, "Towards a theory of software protection and simulation by oblivious rams," in *ACM symposium on Theory of computing (STOC)*, 1987.

[26] O. Goldreich and R. Ostrovsky, "Software protection and simulation on oblivious rams," *J. ACM*, 1996.

[27] V. Goyal, Y. Ishai, A. Sahai, R. Venkatesan, and A. Wadia, "Founding cryptography on tamper-proof hardware tokens," in *TCC*, 2010.

[28] D. Grawrock, *Dynamics of a Trusted Platform: A Building Block Approach*, 1st ed. Intel Press, 2009.

[29] S. Hada, "Zero-knowledge and code obfuscation," in *ASIACRYPT 2000*, 2000.

[30] K. Järvinen, V. Kolesnikov, A.-R. Sadeghi, and T. Schneider, "Garbled circuits for leakage-resilience: Hardware implementation and evaluation of one-time programs," in *CHES 2010*, 2010.

[31] M. Kainth, L. Krishnan, C. Narayana, S. G. Virupaksha, and R. Tessier, "Hardware-assisted code obfuscation for fpga soft microprocessors," in *Design, Automation & Test in Europe Conference & Exhibition*, 2015.

[32] J. Katz, "Universally composable multi-party computation using tamper-proof hardware," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2007.

[33] Á. Kiss and T. Schneider, "Valiant's universal circuit is practical," in *EUROCRYPT*, 2016.

[34] P. C. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *CRYPTO'99*, 1999.

[35] V. Kolesnikov and T. Schneider, "A practical universal circuit construction and secure evaluation of private functions," in *FCDS*, 2008.

[36] V. Koppula, A. B. Lewko, and B. Waters, "Indistinguishability obfuscation for turing machines with unbounded memory," in *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*. ACM, 2015, pp. 419–428.

[37] K. Lewi, A. J. Malozemoff, D. Apon, B. Carmer, A. Foltzer, D. Wagner, D. W. Archer, D. Boneh, J. Katz, and M. Raykova, "5gen: A framework for prototyping applications using multilinear maps and matrix branching programs," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 981–992.

[38] D. Lie, C. Thekkath, M. Mitchell, P. Lincoln, D. Boneh, J. Mitchell, and M. Horowitz, "Architectural support for copy and tamper resistant software," *ACM SIGPLAN Notices*, 2000.

[39] H. Lipmaa, P. Mohassel, and S. Sadeghian, "Valiant's universal circuit:

Improvements, implementation, and applications," Cryptology ePrint Archive, Report 2016/017, 2016, http://eprint.iacr.org/2016/017.

[40] C. Liu, A. Harris, M. Maas, M. Hicks, M. Tiwari, and E. Shi, "Ghostrider: A hardware-software system for memory trace oblivious computation," in *ASPLOS*, 2015.

[41] M. Maas, E. Love, E. Stefanov, M. Tiwari, E. Shi, K. Asanovic, J. Kubiatowicz, and D. Song, "Phantom: Practical oblivious computation in a secure processor," in *CCS*, 2013.

[42] F. McKeen, I. Alexandrovich, A. Berenzon, C. V. Rozas, H. Shafi, V. Shanbhogue, and U. R. Savagaonkar, "Innovative instructions and software model for isolated execution." in *HASP@ ISCA*, 2013.

[43] O. Ohrimenko, M. Costa, C. Fournet, C. Gkantsidis, M. Kohlweiss, and D. Sharma, "Observing and preventing leakage in mapreduce," in *CCS*, 2015.

[44] R. Pass, E. Shi, and F. Tramr, "Formal abstractions for attested execution secure processors."

[45] A. Rane, C. Lin, and M. Tiwari, "Raccoon: Closing digital side-channels through obfuscated execution," in *USENIX Security*, 2015.

[46] L. Ren, X. Yu, C. Fletcher, M. van Dijk, and S. Devadas, "Design space exploration and optimization of path oblivious ram in secure processors," in *Symposium on Computer Architecture*, 2013.

[47] S. Schrittwieser and S. Katzenbeisser, "Code obfuscation against static and dynamic reverse engineering," in *Information Hiding*, 2011.

[48] F. Schuster, M. Costa, C. Fournet, C. Gkantsidis, M. Peinado, G. Mainar-Ruiz, and M. Russinovich, "Vc3: Trustworthy data analytics in the cloud using sgx," in *IEEE S&P*, 2015.

[49] E. Shi, T.-H. H. Chan, E. Stefanov, and M. Li, "Oblivious ram with $O(\log^3 N)$ worst-case cost," in *ASIACRYPT*, 2011.

[50] V. Shoup *et al.*, "Ntl: A library for doing number theory," 2001.

[51] E. M. Songhori, S. Zeitouni, G. Dessouky, T. Schneider, A.-R. Sadeghi, and F. Koushanfar, "Garbledcpu: A mips processor for secure computation in hardware," in *DAC*, 2016.

[52] E. Stefanov, M. van Dijk, E. Shi, C. Fletcher, L. Ren, X. Yu, and S. Devadas, "Path ORAM – an extremely simple oblivious ram protocol," in *CCS*, 2013.

[53] G. E. Suh, D. Clarke, B. Gassend, M. Van Dijk, and S. Devadas, "Aegis: architecture for tamper-evident and tamper-resistant processing," in *Conference on Supercomputing*, 2003.

[54] L. Tan, "The worst case execution time tool challenge 2006: The external test," in *Leveraging Applications of Formal Methods, Verification and Validation, 2006. ISoLA 2006.*, 2006.

[55] L. G. Valiant, "Universal circuits (preliminary report)," in *STOC*, 1976.

[56] X. S. Wang, S. D. Gordon, A. McIntosh, and J. Katz, "Secure computation of mips machine code," ePrint 2015/547, Tech. Rep., 2015.

[57] X. Yu, C. W. Fletcher, L. Ren, M. v. Dijk, and S. Devadas, "Generalized external interaction with tamper-resistant hardware with bounded information leakage," in *Proceedings of the 2013 ACM Workshop on Cloud Computing Security Workshop*, ser. CCSW '13. New York, NY, USA: ACM, 2013, pp. 23–34. [Online]. Available: http://doi.acm.org/10.1145/2517488.2517498

[58] X. Zhuang, T. Zhang, and S. Pande, "Hide: an infrastructure for efficiently protecting information leakage on the address bus," in *ACM SIGPLAN Notices*, 2004.

APPENDIX A
PROOF OF SCHEDULE OVERHEAD

Claim: For *any* program and input, the setting of N from Section III-C results in $\leq 50\%$ of processor cycles performing dummy work. In other words, the schedule incurs $\leq 2\times$ performance overhead relative to the best possible A-M schedule (which is insecure over the timing channel) and incurs no dummy work.

Proof: Without loss of generality, we break up a program into a sequence of instruction *epochs*, where each epoch consists of a continuous run of arithmetic instructions followed by a continuous run of memory instructions. Denote the i-th epoch as $A^{n_i} M^{p_i}$. For example, the program

A A A A M A A M M

has 2 epochs, with $n_1 = 4, p_1 = 1, n_2 = 2, p_2 = 3$.

Without loss of generality, we align the start of each epoch with the beginning of an $A^N M$ schedule. Given our choice of N, we examine the number of processor cycles spent doing dummy operations in each epoch. For the rest of the analysis, we abbreviate $|M| = $ ORAM latency and $|A| = $ Arithmetic latency.

Consider the start of epoch i (i.e., the first A instruction). To progress from the start of the epoch to the first M instruction (excluded) in the epoch, we perform $|A| * N * \lfloor \frac{n_i}{N} \rfloor + |A| * (n_i \mod N)$ real cycles and $|M| * \lfloor \frac{n_i}{N} \rfloor + |A| * (N - (n_i \mod N))$ dummy cycles worth of work. To progress from the first M instruction (including) to the end of the epoch, we perform $|M| * p_i$ real cycles and $|A| * N * (p_i - 1)$ dummy cycles worth of work. Note that by our definitions of epochs, we have that $p_i \geq 1$.

Also note that $|M| = |A| * N$ by our choice of N. Combining these two time periods, we spend $|M| * (\lfloor \frac{n_i}{N} \rfloor + p_i) + |A| * (n_i \mod N)$ real cycles and $|M| * (\lfloor \frac{n_i}{N} \rfloor + p_i - 1) + |A| * (N - (n_i \mod N))$ dummy cycles worth of work. ∎

APPENDIX B
OBFUSCATION IN THE PUBLIC-KEY SETTING

For the sake of simplicity, we describe our construction and proof in the model where a single sender embeds a symmetric key into a secure processor and provides this to the receiver along with the obfuscated program to execute. However, we note that we can extend our results to reuse the token and allow multiple senders to obfuscated the program for a receiver. For example, suppose two senders S_1 and S_2 would like to both send encrypted programs to be executed by a receiver R on a hardware token (provided by a trusted hardware manufacturer). The hardware would then be initialized with a secret key sk_{enc} of a public-key CCA secure encryption scheme (with public key pk_{enc}) along with a verification key vk_{sig} of a signature scheme (with signing key sk_{sig}). The signing key sk_{sig} would be owned by a trusted certificate authority and would also be stored in the token. Now, in our construction, we would replace the symmetric key CCA secure authenticated encryption with a public key CCA secure encryption, where all ciphertext are authenticated with a signature scheme. When S_1 wishes to send an obfuscated program P_1 to a receiver R, S_1 would pick a signing key/verification key pair (sk_{S_1}, vk_{S_1}). S_1 will obtain a signature of vk_{S_1} from the trusted certificate authority (denote this signature by σ and note that this signature will verify under the verification key vk_{sig}). Now, S_1 will encrypt P_1 with pk_{enc} and authenticate all ciphertexts with sk_{S_1} and provide these ciphertexts along with σ to the receiver. The receiver will feed in encrypted ciphertexts along with σ to the token. The token, when decrypting ciphertexts, will first check the validity of vk_{S_1} by verifying σ and the signatures of all the ciphertexts. If all the checks pass, the token will decrypt the ciphertexts using sk_{enc}. When encrypting state to be sent back to the receiver, the token will encrypt it with pk_{enc} and sign it with sk_{sig}. This will mimic the symmetric key CCA secure authenticated encryption scheme that we use in our single sender/receiver scheme.