

Data Oblivious ISA Extensions for Side Channel-Resistant and High Performance Computing

Jiyong Yu, Lucas Hsiung, Mohamad El Hajj†, Christopher W. Fletcher
University of Illinois at Urbana-Champaign, †American University of Beirut
{jiyongy2, ljhsion2, cwfletch}@illinois.edu, mae77@mail.aub.edu

Abstract—Blocking microarchitectural (digital) side channels is one of the most pressing challenges in hardware security today. Recently, there has been a surge of effort that attempts to block these leakages by writing programs *data obliviously*. In this model, programs are written to avoid placing sensitive data-dependent pressure on shared resources. Despite recent efforts, however, running data oblivious programs on modern machines today is insecure and low performance. First, writing programs obliviously assumes certain instructions in today’s ISAs will not leak privacy, whereas today’s ISAs and hardware provide no such guarantees. Second, writing programs to avoid data-dependent behavior is inherently high performance overhead.

This paper tackles both the security and performance aspects of this problem by proposing a *Data Oblivious ISA extension*. On the security side, we present ISA design principles to block microarchitectural side channels, and embody these ideas in a concrete ISA capable of safely executing existing data oblivious programs. On the performance side, we design our ISA with support for efficient memory oblivious computation, and with safety features that allow modern hardware optimizations, e.g., out-of-order speculative execution, to remain enabled in the common case.

We provide a complete hardware prototype of our ideas, built on top of the RISC-V out-of-order, speculative BOOM processor, and prove that the ISA can provide the advertised security through formal analysis of an abstract BOOM-style machine. We evaluate area overhead of hardware mechanisms needed to support our prototype, and provide performance experiments showing how the ISA speeds up a variety of existing data oblivious codes (including “constant time” cryptography and memory oblivious data structures), in addition to improving their security and portability.

I. INTRODUCTION

With the rise of cloud computing and internet services, *digital or microarchitectural side channel attacks* [1] have emerged as a central privacy threat. These attacks exploit how victim and adversarial programs share hardware/virtual resources on shared remote servers (e.g., an amazon EC2 cloud). Simply by co-locating to the same platform, researchers have shown how attackers can learn victim program secrets through the victim’s virtual memory accesses [2], hardware memory accesses [3], branch predictor usage [4], [5], arithmetic pipeline usage [6], [7], speculative execution [8], [9] and more. Given the many avenues to launch an attack, it is paramount for researchers to explore holistic and efficient defensive strategies.

Recently, there has been a surge of work that attempts to block *all* digital side channels, on commercial machines, by writing and compiling programs in a *data oblivious* fashion (e.g., [10], [11], [12], [1], [13], [14], [15], [16], [17], [18],

```
1 x = 0, y = 64
2 if (secret)
3   x = y
4 z = Memory[x]

1 x = 0, y = 64
2 z = Memory[x]
3 tmp = Memory[y]
4 z = (secret) ? tmp : z
```

(a) Insecure code. (b) Equivalent data oblivious code.

Fig. 1: Non-oblivious (1a) and equivalent data oblivious codes (1b). The word `secret` denotes private data.

[19], [20], [21], [6], [22]). Data oblivious code, a.k.a. “constant time” or “running programs in the software circuit model,” blocks side channels by disallowing private data-dependent control flow. Figure 1 gives an example. Figure 1a leaks private information over microarchitectural side channels—namely, program execution time (the ‘if-taken’ case executes more instructions) and memory footprint (if `x` and `y` touch different lines in cache). To block these leakages, a data oblivious program will evaluate both sides of the branch as shown in Figure 1b. A ternary operator—e.g., implemented as the x86 `cmov` instruction or bitwise operations—chooses the correct final result (Figure 1b, Line 4). Since executing each side of the branch is independent of the secret, and the ternary operator does work independent of the secret, running the code data obliviously does not leak the secret.

A. Challenges

Despite the promise of data oblivious programs to block side channels, future progress faces two key challenges.

Security: Existing Instruction Set Architectures (ISAs) provide no guarantees that instructions used in data oblivious codes can block leakages over microarchitectural side channels. For example, if `cmov` (used as the ternary operator in [1], [20], [14], [16]) was ever implemented as the microcode sequence `branch+mov`, the secret condition can leak based on branch predictor state and whether hardware speculation results in a squash. Being ISA-invisible, these changes can occur at any time. Case in point, Intel has stated that `cmov`’s behavior w.r.t. speculation may change in future processors ([23], Section 3.2).

Beyond `cmov`, the larger problem is that commercial ISAs such as x86 give engineers significant rope to perform *data-dependent* optimizations during program execution. For example, it is well known that arithmetic units can sometimes take data-dependent time [6], [7]. We provide a comprehensive background on related vulnerabilities in Section III-B. Any of these software-invisible optimizations can undermine the security of prior and future work that attempts to write data oblivious programs.

Performance: Data oblivious codes can incur large performance overheads. The reason, once again, is that data obliviousness does not have ISA-level support. As a result, programmers are forced to use only the simplest instructions to achieve data oblivious execution, out of fear that other instructions will leak privacy. For example, data oblivious codes must make two memory accesses in Figure 1b out of fear that a single access will reveal the address through the processor cache, or other, side channel. This overhead scales with deeper data-dependent control flow and larger data sizes.

B. This Paper

In this paper, we tackle both the security and performance aspects of this problem by developing a novel type of ISA extension which we call a *Data Oblivious ISA* extension. To our knowledge, this represents the first foundation for writing and executing secure, portable and performant data oblivious code on commercial-class (out-of-order, speculative) processors. To this end we make the following contributions:

1.) Design principles for data oblivious ISA design. Our key idea is to explicitly specify security guarantees at the ISA level, while decoupling those guarantees from the implementation details of a particular processor. Specifically, each operand of each instruction is given an ISA-level attribute specifying whether that operand is *Safe* to receive private data. If marked *Safe*, processor implementations (*microarchitectures*) using that ISA must block operand-dependent behaviors (which create side channels) from other parts of the system while that instruction is executing. Importantly, how protecting *Safe* operands is implemented is left to the hardware designer, who can devise efficient protections depending on each microarchitecture (e.g., by breaking the instruction into simpler data oblivious instructions [6] or using hardware partitioning [24]/cryptographic techniques [20]). In all cases, the programmer works with a simple, portable guarantee.

2.) Design of a concrete data oblivious ISA. With these principles, we define a set of instructions that can serve as the foundation for the rich line of ongoing work in data oblivious programming [11], [10], [14], [15], [16], [17], [21], [20], [19], [1], [18], [6], [22]. Beyond Turing completeness and security, we also want to reduce the performance overhead common with data oblivious code. To that end, we provide additional instructions that implement efficient *memory oblivious* computation [20], [18] (featuring loads/stores with private addresses). Given the principles above, this extension is conceptually simple: instead of emulating memory obliviousness with dummy memory operations (Figure 1b), we designate a new load instruction whose address operand is *Safe*, which gives hardware designers the ability to build secure and efficient implementations for that specific operation.

3.) Hardware prototype on an out-of-order speculative processor with accompanying formal analysis. To show that our ideas are practical, we prototype all hardware changes needed to support our ISA on top of the RISC-V BOOM processor (for “Berkeley Out-of-Order Machine”) [25]. BOOM is the most sophisticated open RISC-V processor, featuring modern performance optimizations such as speculative and out-of-order execution, and is similar to commercial machines

that run data oblivious code today.¹ In parallel, we develop a formal analysis that models an abstract BOOM-class processor (out-of-order, speculative, superscalar), and describe how to map the abstract BOOM to our hardware prototype. Through our prototype and formal analysis, we prove how our ISA not only provides strong security, but does so while allowing high performance hardware optimizations (e.g., speculative execution) to remain enabled in the common case.

We evaluate our proposal in terms of hardware area and performance over a range of existing data oblivious programs (including linear algebra, data structures, and graph traversal). Area-wise, our proposal takes $< 5\%$ the area of the unmodified BOOM processor. Performance-wise, our ISA and hardware implementation provides an $8.8\times/1.7\times$ speedup on small/large data sets, respectively, relative to data oblivious code running on commodity machines (and with the security and portability benefits stated before). We also show case studies, where our ISA speeds up constant time AES [29], [30] by $4.4\times$ and the memory oblivious ZeroTrace [20] library by $4.6\times$ to several orders of magnitude, depending on parameters.

II. BACKGROUND AND THREAT MODEL

A. Hardware Terminology

1) Out-of-order execution: Modern commercial processors such as the RISC-V BOOM [25] dynamically schedule and execute data-independent instructions in parallel and out of program order to improve performance. Instructions are *fetched* and *issued* (enter the scheduling system) in *program order*, *execute* (perform their operations and produce their results) possibly out of program order, and finally *retire* (make their operation externally visible by irrevocably modifying the architected system state) in program order.

2) Speculative execution: Speculative execution improves performance by executing instructions whose validity is uncertain instead of waiting to determine their validity. If such a speculative instruction turns out to be valid, it is eventually retired; otherwise, it is *squashed* and the processor’s state is rolled back to a valid state. (As a byproduct, all following instructions also get squashed.) That is, a squash causes a large pipeline disturbance. There are multiple ways an instruction stream can be speculative—e.g., due to branches, memory accesses [31], or even arithmetic instructions [32]—discussed further in Section III-B.

More details on BOOM are given in Section V-A.

B. Threat Model

We consider the setting where a victim program runs on a shared machine in the presence of adversarial software. The adversary’s goal is to learn private data in the victim program through digital side channels. For example, private inputs contributed by another party or secret program state (e.g., a cryptographic key). The program itself is considered public. We trust the processor hardware and that the victim program is correctly using the data oblivious ISA. We do not focus

¹We note that prior work [26], [27], [28] requires the use of discrete co-processors with simple microarchitecture. To match modern cloud deployments, our goal is to support *concurrent* execution of many processes on advanced microarchitectures.

on integrity of computation, and leave this issue to traditional process isolation.

We defend against two classes of adversary: supervisor-level (Ring-0) or user-level (Ring-3) software. In both cases, we strive to block digital side channels that could be exploited by the standard Intel SGX adversary used in prior work on data oblivious programming [16], [1], [20], [21], [13], [14], [15], [18]. This adversary is supervisor-level software that controls when victim threads run, and therefore can monitor/influence the victim’s hardware resource utilization (e.g., monitor/prime the cache [3]) at near-perfect resolution (e.g., via [33], [34], [2]). By extension, this adversary can monitor the victim’s termination time, and determine when a precise exception [1], [35] or system call [36] occurs. We don’t make assumptions on where the victim runs relative to adversarial code (e.g., as an adjacent SMT context, adjacent core, etc.). If the adversary is actually user-level software, our threat model is strictly conservative.²

In the case of a supervisor-level adversary, we assume the victim is running within a virtual shielding system, such as an SGX enclave [38], [39], to prevent direct inspection/tampering on victim data. The data oblivious ISA is orthogonal to which virtual shielding system is used, in the sense that shielded programs can execute oblivious instructions regardless of the exact shielding system implementation. We will therefore only discuss the data oblivious ISA, independent of the shielding system, for the rest of the paper.

Non-goals. Physical side channels (e.g., power [40] or EM [41]) are out of scope. We also do not consider Rowhammer attacks through main memory [42], but remark that modern shielding systems such as SGX can mitigate these attacks via integrity checks in memory [43].

III. DATA OBLIVIOUS EXECUTION

We now give background on data oblivious execution and give examples for where prior work on commercial ISAs (e.g., x86) and modern machines (e.g., speculative, out-of-order) is vulnerable to attack.

A. Security Definition

Data oblivious execution satisfies computational indistinguishability of program traces, once the trace is projected by an appropriate observability function.

Definition III.1. (*Confidential input privacy*). Given a program λ with Public (non-sensitive) input x and Confidential (sensitive) input y , $O(\mu\text{Arch}(\lambda(x,y))) = X = \{X_0, X_1, \dots, X_M\}$ represents the program’s observable execution trace (projected through function O) when running on a processor μArch . What information is contained in each X_i (for each time step i) depends on the observability function O . W.l.o.g. we will treat x and y as fixed-size arrays, thus λ can accept an arbitrary number of Public and Confidential inputs. Privacy for the Confidential inputs then requires:

$$\forall x \in \text{Data}_P, \forall y, y' \in \text{Data}_C : \\ O(\mu\text{Arch}(\lambda(x,y))) \simeq O(\mu\text{Arch}(\lambda(x,y')))$$

²Note that even user-level adversaries have been shown to be surprisingly powerful in their ability to monitor digital side channels [37].

where \simeq denotes computational indistinguishability, and Data_P and Data_C denote the space of Public and Confidential inputs, respectively.

We denote Definition III.1 parameterized by an observability function O and a specific microarchitecture μArch as $\text{Oblivious}[O, \mu\text{Arch}]$, dropping μArch when it is clear which microarchitecture we are referring to.

Existing data oblivious programs written for commodity machines demand a rich observability function that reveals fine-grain details about processor state [11], [10], [14], [15], [16], [17], [21], [20], [19], [1], [18], [22]. The reason is that machines today are shared, and adversaries from Section II-B can monitor internal activity such as caches and pipeline behavior. It is therefore useful to define the most conservative observability function that could apply to adversaries from Section II-B:

Definition III.2. (*BitCycle observability: Security labels at bit-level spatial granularity, cycle-level temporal granularity*). Let $S_t = \{0, 1\}^N$ denote the processor state during clock cycle t , where state includes all on-chip storage (e.g., flip-flops, SRAM). S_t^i denotes the value of the i -th bit in cycle t . Given a program execution $\lambda(x, y)$, $\text{BitCycle}(\mu\text{Arch}(\lambda(x, y))) = X = \{X_0, X_1, \dots, X_M\}$ where $X_t = \{0, 1\}^N$ and $X_t^i = 1$ indicates S_t^i contains an explicit flow³ of Confidential data in cycle t ($X_t^i = 0$ otherwise).

For example, writing data d to the processor cache at address a in cycle t sets bits in X_t , corresponding to cache memory cells at a , if either d or a were computed based on Confidential data. More generally, the definition implies the adversary can monitor every possible hardware resource pressure (e.g., flip-flop level pipeline utilization, cache footprint, etc.) every cycle. This paper’s goal is to provide a basis for programs to achieve $\text{Oblivious}[\text{BitCycle}]$ on advanced commercial-class machines (implying no timing flows [45] at bit and cycle granularity).

B. Security Issues in Existing Data Oblivious Code

Existing data oblivious codes are written extremely conservatively to remove code constructs that blatantly violate $\text{Oblivious}[\text{BitCycle}]$. For example, prior works rely solely on a carefully chosen subset of arithmetic operations (e.g., bitwise operations), conditional moves, branches with data-independent outcomes, jumps with public destinations, and memory instructions with data-independent addresses [10], [11], [12], [1], [13], [14], [15], [16], [17], [18], [19], [20], [21], [6], [22].

It is important to understand when this isn’t sufficient for security. To that end, we now detail 11 possible attack vectors on today’s data oblivious code. Importantly, we do not list many popular attacks (e.g., prime+probe in the cache [3]) as these are defeated by writing programs in the style described above. Yet, attacks can still occur because the hardware can apply invisible

³Formally: Let each memory cell S^i take two inputs: data (in^i) and write enable (we^i) where both are functions (combinational logic) taking a subset of bits in S as input. For time $t = 0$: S_0 (i.e., at $t = 0$) is initialized with starting program state, $X_0^i = 1$ iff S_0^i is Confidential data. For time $t > 0$: $X_t^i = 1$ if (a) we^i outputs 1 in cycle t and $X_{t-1}^i = 1$ or (b) we^i outputs 1 in cycle t and $X_{t-1}^j = 1$ for some j in the inputs to S^i (in^i or we^i). We note that implicit flows [44] are accounted for once BitCycle is applied to Definition III.1.

optimizations to undermine software-level transformations. In the following, we describe attack vectors known to be implemented today, and also proposals whose implementation status is unknown. However, importantly, each optimization could be implemented at any time, breaking existing codes.

Vectors 1, 2, 3: branch, jump, memory speculation: While Spectre-style attacks [8] are known to impact general purpose code, their impact on data oblivious code has not been adequately studied. We make an important observation that data oblivious code security is undermined even by ‘honest’ speculative execution. By ‘honest’, we mean the speculation is not intentionally being controlled in a malicious way, e.g., as in [8]. The root problem is that modern ISAs have limited resources (e.g., ISA-level registers) and executing unintentional instructions can cause secrets stored in *aliased resources* to be exposed accidentally.

Consider a toy example for data oblivious decryption, exploiting conditional branch misprediction (Spectre variant 1, denoted Vector 1):

```

1 for (i = 0; i < NUM_ROUNDS; i++)
2   state = OblDecRound(state, rkey[i])
3   declassify(state)

```

A legal data oblivious code can implement decryption round logic data obliviously, with the round keys `rkey` considered Confidential (Definition III.1), and wrap the round in a data-independent branch to reduce code footprint. Once decryption is complete, the program may use the plaintext in a non-oblivious way, e.g., by using it as an address to lookup a record in cache (denoted `declassify(state)`). Such a non-oblivious operation can reveal information related to the *decryption key* on a speculative machine. Specifically, if the branch mispredicts “not taken” (e.g., while the predictor is training), `state` is prematurely exposed before all rounds complete, allowing an attacker to perform cryptanalysis on encryption round intermediate state.

Removing branches or disabling branch speculation is not sufficient to fix this issue, as other forms of speculation (e.g., unconditional branches/jumps, memory disambiguation [8], [31], [46]—denoted Vectors 2 and 3) cause similar issues on legal data oblivious code.

Vectors 4, 5: sub-address optimizations: Numerous data oblivious codes, e.g., “constant time” cryptography [47], [48], make an assumption that modulating certain bits in a memory address (e.g., the bits indicating offset within a cache line) does not create observable behaviors. This assumption doesn’t hold on some microarchitectures due to hardware optimizations such as speculative store forwarding (Vector 4) and cache banking (Vector 5), and attacks exploiting these features have been shown to lead to full cryptographic breaks [49], [50].

Vector 6: input-dependent arithmetic: It is well known that complex arithmetic operations (e.g., multiply/divide, floating point square root) exhibit observable data-dependent timing based on their operands [7], [6]. While prior work can mitigate these threats by re-writing complex arithmetic using bitwise operations, this can incur over an order of magnitude performance overhead depending on the operation [6].

Vector 7: microcode: Even simple instructions may be decomposed into simpler instructions, called micro-ops, before being executed. In some cases, micro-op conversion can create data-dependent behavior. For example, `cmov` (which allows conditionals based on Confidential values [1], [20], [14], [16]) can be broken into a `branch+mov`. There is evidence to suggest that this transformation will be applied in future Intel processors ([23], Section 3.2). This breaks privacy: the branch direction will be speculatively guessed and whether a misprediction occurs changes program timing due to the squash (Section II-A).

Vectors 8, 9, 10, 11: data-based compression, data-based speculation, silent stores: Finally, there are a number of proposals whose implementation status on commercial machines is unknown. In register file [51] and cache [52] compression (analogous to OS-level page de-duplication [53]), register file and cache pressure is a function of program data (Vectors 8 and 9, respectively). Value prediction [32] (Vector 10) speculates on the result of a memory load or long-running arithmetic operation, causing a squash if the prediction is incorrect (Section II-A). Finally, silent stores [54] (Vector 11) remove redundant store operations (impacting cache pressure) when the hardware detects the memory already contains the same value at the same address. What all of the above have in common is that they are *program data-centric* optimizations that don’t discriminate between Public and Confidential data. Thus, they can undermine any data oblivious code written in any style.

Takeaway: Not only is writing data oblivious code difficult, it is fraught with danger due to subtle ISA-invisible optimizations such as those given above. Our proposed data oblivious ISA extension and hardware prototype gives hardware the visibility it needs to decide when and when not to apply leaky performance optimizations (such as those above) and enables richer hardware support for data oblivious code to speedup core operations such as oblivious memory.

IV. DATA OBLIVIOUS ISAS

We now describe data oblivious ISA design principles and give an example concrete ISA that we will later implement on top of the RISC-V BOOM.

A. Design Principles

We had two primary goals in designing the Data Oblivious ISA. First, the ISA should expose security guarantees in a microarchitecture-independent way. A single ISA may be embodied in many different microarchitectures (within and across processor generations), each with different organizations and optimizations. It isn’t reasonable to ask software to reason about each microarchitecture: a developer who writes a data oblivious code correctly once should have confidence that security will hold on each microarchitecture. Second, the ISA should not preclude modern hardware performance techniques, except when those techniques have a chance to leak privacy. Specifically, we want to be compatible with wide (multiple instructions fetched per cycle), speculative, out-of-order commercial-class machines, e.g., those described in Section II-A, and also point optimizations (e.g., banked caches, data-dependent arithmetic; c.f. Section III-B) that, left unchecked, cause security problems.

To achieve these goals, a data oblivious ISA has the following main components (which require hardware support).

1) *Dynamic tracking for Confidential (sensitive) data:*

We use hardware-based dynamic information flow tracking techniques (DIFT, similar to [55], [56]) to track how Confidential data propagates through the processor as the program executes. Conceptually, all data in the processor is *labeled* Confidential/Public at some granularity (e.g., word-level). This gives hardware the ability to decide when to apply optimizations to data in use (e.g., attack Vectors 6-7, 10-11; c.f. Section III-B) and at rest (e.g., Vectors 8-9).

Prior work does not specify precise rules for when data labeled Confidential can be processed relative to when its label is resolved. A conservative strategy is to require all {data, label} state to correspond to program order, which would preclude speculative, out-of-order execution. A more aggressive strategy is to allow speculation, and to further allow data to be used before its label is resolved.⁴ Based on our use of DIFT, it will be clear the latter approach is not secure. Instead, we adopt (and prove secure in Section VI) a middle ground which we call *coherent labels*.

Definition IV.1. (*Coherent labels*) *When reading an operand, its label must be resolved with respect to the dynamic sequence of speculative/non-speculative instructions (which does not necessarily follow program order) that have executed so far to generate that operand.*

A simple implementation that satisfies Definition IV.1 is to physically extend each data word with a label bit, which allows normal processor dependency tracking to ensure labels are resolved on time. We use this strategy for our implementation in Section V.

2) *Instruction operand-level security specifications:* In a data oblivious ISA, instruction definitions specify, for each operand, whether that operand can accept Public or both Public/Confidential data. We call the former an *Unsafe* operand and the latter a *Safe* operand. Once specified, the hardware designer must handle the following cases.

Definition IV.2. (*Confidential* \rightarrow *Safe*) *When Confidential data is sent to a Safe operand: the hardware designer must add mechanisms to enforce Definition III.1, for a specified observability function, during that instruction’s execution. For example, by disabling performance optimizations based on the Safe operand. This includes hiding exceptions that occur due to Confidential operands (see Section IV-C).*

Definition IV.3. (*Confidential* \rightarrow *Unsafe*) *When Confidential data is presented to an Unsafe operand: the hardware must stop (squash) that instruction’s execution as soon as the label is resolved. This event is called a label violation #LV. Due to Definition IV.1, #LV will be signaled immediately after register/memory read, and before the execute stage begins. If the violating instruction is the next instruction to retire (i.e., is non-speculative), terminate the program. This event is called a label fault #LF.*

⁴For example, [56] proposes storing labels in the page table. If the processor supports speculative store-forwarding [49] (Vector 4), data will be used before the label lookup completes.

That is, Definition IV.3 is similar to rules that handle badly typed programs, extended to speculative execution. Label violations (#LV) are caused by transient conditions, e.g., imperfect prediction (Section III-B, Vector 1), and are correctable. Label faults (#LF) indicate a program bug or illegal typing. Fixing bugs is outside of our scope, so we will focus on #LV.

An important question is whether #LV creates a side channel based on *when* it is triggered. We prove in Section VI-A that it does not, and further prove that #LV signals enable the ISA to block multiple additional attacks (Vectors 1-5; c.f. Section III-B), e.g., speculation on Confidential data or speculation that can reveal Confidential data, on top of the vectors blocked from Section IV-A1. Finally, Public data is handled as:

Definition IV.4. (*Public* \rightarrow *Safe/Unsafe*) *When Public data is sent to Safe or Unsafe operands, no special treatment is needed and execution can proceed without protection.*

As the above definitions apply at operand granularity, the ISA permits optimizations that are functions of individual operands. For example, zero-skip multiply can be enabled if a Public operand is 0, regardless of whether other operands are Confidential (common in, e.g., secure deep learning [57], [58]).

Specifying each instruction operand as Safe/Unsafe at the ISA level is a key design feature, and provides significant flexibility to both the ISA and hardware designer while simplifying programmer-level reasoning about security. At the ISA level, an ISA designer can decide which instructions are sufficiently important to warrant Safe operands. These choices should be made carefully: On one hand, Safe operands impose a burden on hardware designers as the processor must support mechanisms to uphold Definition III.1 for those operands. On the other hand, Safe operands do not specify an implementation strategy. Hardware designers can implement a given operation using simpler data oblivious instructions (e.g., [6]), hardware partitioning (e.g., [24]) or cryptographic techniques (e.g., [20])—depending on what is efficient given public parameters and the specific microarchitecture. In either case, programmers work with a simple guarantee: Confidential values will not be at risk when consumed by Safe operands, and dynamic execution will be terminated when violations to this policy are detected.

B. Concrete Data Oblivious ISA Specification

Using the principles from the previous section, we now present a concrete data oblivious ISA extension that we will implement on top of the RISC-V BOOM processor. Figure 2 shows data oblivious instruction encodings, supported instruction types, and the Safe/Unsafe characteristics for each operand (Section IV-A2).

1) *Label propagation:* Our ISA requires word-granularity labels, tracked in the register file and memory. In most cases, label update logic follows standard taint tracking rules, given the 2-level security lattice {Public, Confidential} [60], as shown in Figure 2. When the result is fully determined by Public operands, regardless of other operands (e.g., zero-skip multiply), the result label is set to Public (similar to GLIFT [61], but not shown in Figure 2 for simplicity).

Base Data Oblivious ISA:	Operand label constraints (S = Safe, U = Unsafe)		Instruction functionality	Label propagation	Notation (assembly)
Arithmetic (R-type)	rs2 (S)	rs1 (S)	$R[rd] \leftarrow R[rs1] \text{ op } R[rs2]$	$Lr[rd] \leftarrow Lr[rs1] \mid Lr[rs2]$	
Arithmetic (I-type)		rs1 (S)	$R[rd] \leftarrow R[rs1] \text{ op ext(imm)}$	$Lr[rd] \leftarrow Lr[rs1]$	
Declassify (I-type) (serializing)		rs1 (S)	$R[rd] \leftarrow R[rs1]$	$Lr[rd] \leftarrow 0$	ounseal %rd, %rs1
Classify (I-type)		rs1 (S)	$R[rd] \leftarrow R[rs1]$	$Lr[rd] \leftarrow 1$	oseal %rd, %rs1
Conditional move (R-type)	rs2 (S)	rs1 (S)	$R[rd] \leftarrow (R[rs1]) ? R[rs2] : R[rd]$	$Lr[rd] \leftarrow Lr[rs1] \mid Lr[rs2] \mid Lr[rd]$	ocmov %rd, %rs1, %rs2
Branch (B-type)	rs2 (U)	rs1 (U)	if ($R[rs1] \text{ op } R[rs2]$) PC = PC + imm	$Lr[rd] \leftarrow 0$	
Jump register (I-type)		rs2 (U)	$R[rd] = PC + 4$; PC = PC + imm	$Lr[rd] \leftarrow 1$	
RNG (J-type)			$R[rd] \leftarrow \text{rand}()$	-	orng %rd
Load (I-type)		rs1 (U)	$R[rd] \leftarrow M[R[rs1] + \text{ext(imm)}]$	$Lr[rd] \leftarrow Lm[R[rs1] + \text{ext(imm)}]$	orld %rd, imm(%rs1)
Store (S-type)	rs2 (S)	rs1 (U)	$M[R[rs1] + \text{ext(imm)}] \leftarrow R[rs2]$	$Lm[R[rs1] + \text{ext(imm)}] \leftarrow Lr[rs2]$	orst %rs2, imm(%rs1)
Oblivious Memory extension:			let addr := $R[rs1] + \text{ext(imm)} \% \text{OSZ}$		
Oblivious Load (I-type)		rs1 (S)	$R[rd] \leftarrow M[\text{addr}]$	$Lr[rd] \leftarrow 1$	ocld %rd, imm(%rs1)
Oblivious Store (I-type)	rs2 (S)	rs1 (S)	$M[\text{addr}] \leftarrow R[rs2]$	-	ocst %rs2, imm(%rs1)
CPUID (J-type)			$R[rd] \leftarrow \text{OSZ}$	-	ocpuid %rd

Fig. 2: Data Oblivious ISA extension. **R/Lr**, **M/Lm** denote register file data/labels, memory data/labels, respectively. The label Public is denoted logic 0, Confidential logic 1. rs1 and rs2 denote operand registers in RISC-V instructions while rd denotes destination register. R, I, B, J, S-type refers to standard RISC-V instruction formats [59]. ext extends the immediate to the word width. If assembly notation is unspecified, it follows RISC-V with an ‘o’ prefix (e.g., add becomes oadd). OSZ refers to the microarchitecture-specific oblivious memory partition size (Section IV-B6).

2) *Label declassification*: Declassification—downgrading data marked Confidential to Public—is a rare but necessary task needed to, e.g., return results. Our ISA supports a single **serializing** declassification instruction called ounseal. Serializing instructions are not executed until all other in-flight instructions (e.g., executing speculatively and out-of-order; c.f. Section II-A) retire. This is necessary for security: declassification is the only mechanism to demote Confidential to Public, and this action under malicious speculative execution could be used to bypass label checking.

3) *Instruction set*: Our ISA supports the following instruction types, which we chose to maximize compatibility with existing data oblivious codes and minimize hardware changes. First, all RISC-V integer and floating point arithmetic with Safe operands. This means programmers can implement floating point directly, without invoking bitwise libraries [6]. Second, random number generation, as many randomized data oblivious codes (e.g., [62], [63], [64]) require private random numbers (e.g., to support Oblivious RAM [18], [20]; c.f. Section IV-B6). Third, a cmov-style ternary/conditional move operator with a Safe predicate for implementing conditionals, and branches/jumps with Unsafe operands to reduce code footprint. Fourth, load/store operations (orld and orst) with Unsafe address operands.

Lastly, we support a second flavor of load/stores (with Safe address operands) which can be used to implement oblivious memory (Section IV-B6) using Confidential addresses.

4) *Mixing in non-oblivious instructions*: Oftentimes, only a small program region should be made data oblivious (e.g., the inner branch in modular exponentiation) to prevent unnecessary performance overheads. To support these situations, we support mixing data oblivious instructions with instructions from the original ISA. All operands for all original instructions are considered Unsafe. All data oblivious instructions are encoded

on top of the normal RISC-V ISA by modifying existing instruction fields (e.g., the opcode and func [59]).

1 oaddi %x1, %x0, 0	1 oaddi %x1, %x0, 0
2 oaddi %x2, %x0, 64	2 oaddi %x2, %x0, 64
3 oseal %x3, secret	3 oseal %x3, secret
4 orld %x1, 0(%x1) //Mem	4 ocmov %x1, %x3, %x2
5 orld %x2, 0(%x2) //Mem	5 ocld %x1, 0(%x1) //Mem
6 ocmov %x1, %x3, %x2	

(a) Data obl. Fig. 1b. (b) Data obl. Fig. 1b w/ OMP.

Fig. 3: Data oblivious code, using the data oblivious ISA, implementing Figure 1b. The word **secret** denotes Confidential data. %x... are RISC-V general purpose registers. %x0 is a RISC-V idiom for constant 0.

5) *Putting it all together*: To summarize the section, we show a version of Figure 1b written using our data oblivious ISA in Figure 3a. The programmer need only specify what data is Confidential via oseal. The ISA and hardware will prevent %x3 from being processed by subsequent speculative/non-speculative Unsafe operands. For example, specifying %x3 as an address to a speculative/non-speculative orld triggers a #LV/#LF, respectively.

6) *Oblivious memory extension*: A common bottleneck in existing data oblivious code is the inability to use Confidential data as memory addresses [20], [1], [24], [18]. For example, Figure 3a needed to execute two orld instructions. More generally, looking up an array with a Confidential address requires a memory scan.

To accelerate these operations, the ISA exposes two new instructions ocld and ocst, which are analogous to orld/orst (Section IV-B3) except with Safe address operands, and a new variant of CPUID ocpuid which returns a microarchitecture-specific constant OSZ (“oblivious memory partition size”).

Each microarchitecture is responsible for providing OSZ bytes of “fast” oblivious storage, called the *oblivious memory partition* (OMP), which only `ocld` and `ocst` can read/write. This storage can be used to speedup data oblivious code. For example, if `x` and `y` in Figure 1b both fall within the OMP, then Figure 3a can be rewritten as Figure 3b (saving a memory access).

How much storage is provided (the value of OSZ) and how that storage is implemented—e.g., a dedicated scratchpad, flexible cache partition, etc.—is left to hardware designers and can be decided on an implementation-by-implementation basis. (Our prototype in Section V-B uses ways in a cache.) We note that the hardware constrains addresses sent to `ocld/ocst` to fall within bounds 0 to `OSZ-1`.

To make data oblivious code portable across machines (each of which can specify a different OSZ), we provide the following software/programmer-level functions:

- `Unsafe OblObj* obl_alloc(Unsafe int size)`
- `void obl_free(Unsafe OblObj* o)`
- `Safe int obl_read(Unsafe OblObj* o, Safe int addr)`
- `void obl_write(Unsafe OblObj* o, Safe int addr, Safe int data)`

`Safe/Unsafe` qualifiers are implied based on how these functions are implemented (Section IV-A2). That is, size must be `Public`. `obl_alloc/free` dynamically allocate/free an oblivious memory object `OblObj` which exposes `type`, `base` and `bound` fields. `type = {OMP, ORAM, SCAN}` and is determined by `obl_alloc` under the hood using the following rules:

- 1) If the new object will completely fit into the OMP, based on the size argument, previous allocations, and OSZ: set `type = OMP`.
- 2) Else: depending on remaining space in the OMP and the size argument, set the type as ORAM or SCAN. Heuristics to select which are described below.

Post-allocation, users perform reads and writes to `OblObj`s through `obl_read` and `obl_write`, which instrument each operation based on the allocator’s prescribed type, as shown in Figure 4. We describe the ORAM type below.

`obl_alloc` decides on each allocation’s type based on information returned by `ocpuid`. In the current design, `ocpuid` returns OSZ, the implementation-specific size of the OMP. Future implementations may also return richer information, such as machine cache sizes/etc to make more informed decisions. Since size and branches/jumps in our ISA are `Unsafe`, the strategy selected for each allocation depends only on the program (which is `Public`) and the machine architecture. Lastly, we note that since the allocator makes decisions based on the order of previous allocations, more performance-sensitive objects should be allocated first.

ORAM and SCAN types. When the oblivious object does not fit into the OMP, the allocator may implement it as an Oblivious RAM [65] (ORAM) or memory scan. ORAMs are randomized algorithms which implement oblivious memory in poly-logarithmic time. For ORAM, we use the ZeroTrace library [20] which is a data oblivious ORAM client written in our threat model. Depending on remaining OMP space, ZeroTrace’s internal sub-structures (e.g., the ORAM stash and position map [20]) can be placed in the OMP, which we show

```

1 int obl_read(OblObj* o, int addr) {
2 #oblivious {
3   int ret; int tmp;
4   switch (o->type)
5   case OMP:
6     asm ("oaddi %0, %1, %2":
7         "=r" (tmp): "r" (addr), "r" (o->base));
8     asm ("ocld %0, 0(%1)":
9         "=r" (ret): "r" (tmp));
10    break;
11   case ORAM:
12     ret = oram("read", o, addr); break;
13   case SCAN:
14     for (int j = o->base, j < o->bound; j+=4) {
15       asm ("orld %0, 0(%1)":
16           "=r" (tmp): "r" (j));
17       asm ("ocmov %0, %1, %2":
18           "+r" (ret): "r" (j==addr), "r" (tmp));
19     } break;
20   return ret; } }

```

Fig. 4: `obl_read` implementation (`obl_write` is analogous). `#oblivious` is short-hand to indicate that the body consists only of data oblivious instructions. `oram`’s implementation is discussed in Section IV-B6. “=r”, “+r” denotes output register; “r” denotes input.

can speedup the original ZeroTrace by $> 4\times$ (Section VII-C6). SCAN is a fallback that emulates oblivious memory using normal memory, and is implemented as a sequence of `orld` and `ocmov` instructions (Figure 4).

Pointed out by [1], when scan vs. ORAM is more efficient depends on the memory size and the allocator should take this into account based on the allocation size parameter.

C. Process-OS Interface

Processes interact with the OS through exception handling, context switching and system calls. We design our ISA to cause minimal friction with the existing OS-process interface.

1) *Exceptions:* Exceptions leak data-dependent conditions (e.g., when a divide by zero occurs) in programs [35], [1]. When an exception occurs on instructions with all `Public` operands, it is handled like a normal exception. When an exception occurs on an instruction with a `Confidential` operand, the hardware must mask that exception (e.g., by replacing the result with a canonical value and leaving the label unchanged). In this design, the adversary may learn an exception has occurred only if resulting data is explicitly declassified with `ounseal`.

2) *Context switching:* Register file labels are added as part of the thread’s state on a context switch, and labels in memory are mapped to pages in a region of virtual memory that cannot be accessed directly by the program (Section V-C1). In the current design, the OMP (Section IV-B6) is counted as thread state for simplicity. While this doesn’t make context switching performance-prohibitive for the OMP sizes we consider in Section VII, it will for sufficiently large OMPs. We leave integrating the OMP into normal process virtual memory (e.g., by using the RISC-V VLS technique [66]), which would alleviate this issue, as future work. Finally, if the adversary is supervisor-level (Section II-B), we rely on the shielding system, e.g., SGX, to protect program data during context

switches. For example, in an SGX setup [38], all data (Public and Confidential) would be stored within the SGX ELRANGE.

3) *System calls*: We rely on orthogonal software techniques to sanitize system call arguments [67], [36].

V. IMPLEMENTATION

This section describes how we prototyped our data oblivious ISA on the RISC-V BOOM microarchitecture. Our design augments BOOM ‘v2,’ which is the most recent iteration of the BOOM design [25]. We give the exact parameters used for the architecture in Table II, which corresponds to the block diagram in Figure 5 and is a default BOOM configuration.

A. RISC-V BOOM Summary

We first summarize unmodified BOOM (referencing Figure 5). These details will be used for our implementation (this section) and formal analysis (Section VI).

First, multiple instructions are *fetch*ed each cycle ❶. Based on the current program counter and decoded instructions, multiple levels of branch/jump predictors issue predictions for fetched branches/jumps. Mispredicted branches/jumps are discovered in the execute stage, and cause subsequent speculatively decoded instructions to squash (Section II-A). Once *decoded*, instructions are added to the issue windows ❷ where they wait for their operands to be ready, at which point they are *scheduled* (possibly out-of-order) to execution units. Operands become ready when they are written (or written back) to one of two register files (RFs, for floats and integers) ❸, or when an execution unit finishes early and *bypass*es the result directly to the consumer instruction. RFs contain speculative and non-speculative data.

BOOM supports a configurable number of execution units ❹, each of which contains a configurable number of primitive arithmetic/branch/etc. units, shown in Table II. Each execution unit receives dedicated read/write ports to the RFs. Primitive arithmetic blocks may be *pipelined* (have input-independent latency) or *un-pipelined* (have input-dependent latency). Lastly, a load/store unit interfaces to the cache and decides whether load data should be read from the cache or store data queue which contains speculative stores (store-load forwarding). Loads may speculatively execute after stores whose address has not resolved [31]; address alias violations are caught and squashed at retire time. Finally, a reorder buffer (ROB) ❺ tracks in-flight instructions in-order to facilitate in-order commit (Section II-A).

The current BOOM does not currently support SMT/hyper-threading. We note that our ISA is compatible with an SMT-enabled machine and that the hardware mechanisms discussed below need not change to support SMT.

B. Support for New Instructions

Discussed in Section IV-B, most instructions in the data oblivious ISA have exact counterparts in RISC-V, but with additional semantics/dynamic checks for Safe/Unsafe operands. These instructions reuse existing RISC-V encodings and have altered opcode/func fields to be identified during the decode stage. Several exceptions are osea, unsea, orng, ocmov, oclld/ocst/ocpuid which don’t have RISC-V counterparts (Figure 2).

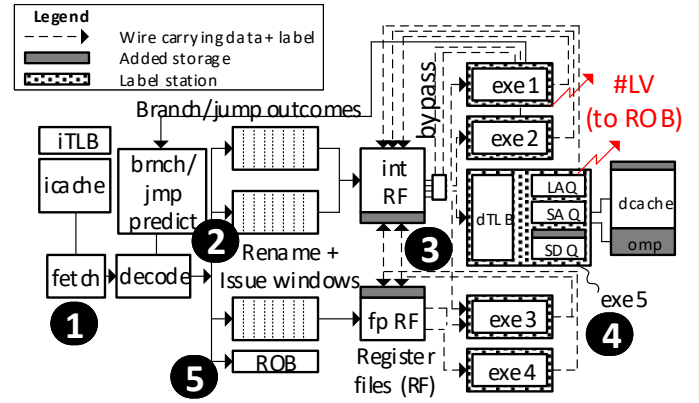


Fig. 5: RISC-V ‘BOOM v2’ pipeline [25]. ‘exeXX’ are execution units, and contain arithmetic/branch/etc units stated in Table II. Hardware modifications needed to support the data oblivious ISA (Figure 2) are shown in the legend. No modifications are needed before the int/fp register files. Label stations are discussed in Section V-C2. ‘omp’ is the oblivious memory partition (Section V-B).

We implement osea and unsea as the RISC-V addi instruction with the immediate field set to 0 (functionally a move operation), but with modified logic to set/clear label bits. As discussed in Section IV-B, unsea must also *serialize* (execute non-speculatively) to prevent malicious declassification. Since BOOM already implements serializing instructions, that modify control status registers, we reuse that functionality for unsea. Our prototype implements orng as a cryptographic PRNG (iterative AES core), although a hardware TRNG [68] may be used for a production design.

ocmov presents a challenge, as conditional move requires three operands (predicate, new value and old value) whereas no RISC-V integer instruction requires three input operands. To minimize ISA-level changes, we design a single ALU (in one execution unit) to serve ocmov instructions, and add a new RF port for that execution unit. We design this ALU to support bypassing. This design is low overhead and efficient. Having one execution unit support ocmov means we only need to add a single read port to the RF (not +1 per execution unit). Through bypassing, our design can execute back-to-back dependent ocmov, one per cycle.

Finally, our current implementation implements the oblivious memory partition (OMP) for oclld/ocst as a quarantined region of the first-level data cache. We isolate a region of the cache using way partitioning techniques [69], which are a low-complexity mechanism to divide the cache into non-interfering regions as long as the region size is a multiple of the associativity (our first-level cache is 16-way; Table II). This design has low hardware overhead. If no process has allocated oblivious objects (Section IV-B6), OMP storage can be used as normal cache memory. While an oclld/ocst instruction is looking up the OMP, all concurrent cache lookups are stalled to avoid cache bank contention [50].

C. Tracking and Checking Labels

A central feature in our data oblivious ISA is checking and tracking Public/Confidential labels as data flows through the pipeline and signalling #LV when violations occur. Noted in Section IV-B, we track labels at word granularity.

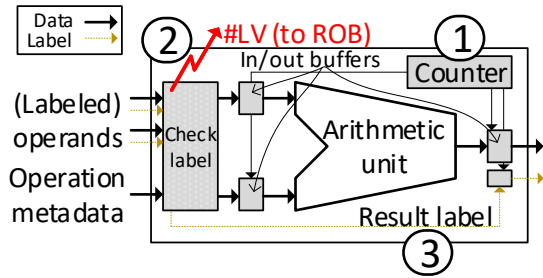


Fig. 6: Label station (Section V-C2) for an execution unit with one internal arithmetic unit. A real execution unit may contain multiple arithmetic units (Table II), in which case this logic is replicated as needed. Added hardware is shaded.

1) *Label storage*: Labels must be stored alongside each word, where-ever each word resides in the processor. This includes the RF, the SDQ, the data cache hierarchy, and intermediate pipeline registers. In all of the above structures, we treat data label as an extra bit in each word itself. This makes it easy to satisfy Definition IV.1: whenever a speculative or non-speculative instruction reads an operand, normal out-of-order processor dependency checking ensures the label is resolved.

Unfortunately, this strategy would require large changes to the DRAM/below memory levels because wider words would require wider DRAM lines and larger page tables. Thus, at the DRAM level, we store data and labels in separate disjoint pages and modify the hardware DRAM controller to join data and label into a widened cache line when on-chip (a similar scheme was used in [55]). This means any DRAM access in our system turns into two DRAM accesses.

2) *Label checks*: To satisfy Definitions IV.2 and IV.3: once a consumer instruction indicates its intent to use an operand, that operand’s label must be checked against the instruction opcode/func fields, before the use occurs. We design a parameterizable hardware module called a *label station*, which wraps each BOOM execution unit, to administer these checks. The main observation enabling the label station design is that in BOOM, all operand-dependent processor state updates are signalled from the execution units. This makes it possible to implement a shim at the input of each execution unit to perform label checks, handle label violations/faults, and disable hardware optimizations on Confidential inputs.

Specifically, the label station (visualized in Figure 6):

- ① (Definition IV.2: Confidential \rightarrow Safe) Blocks access to/from arithmetic units so that any operation processing Safe operands takes the worst case time. This is implemented using input/output buffers (e.g., flip-flops), a timer (counter), and operand/label decode logic (“Check label” in the figure). Variable-time arithmetic units and their worst-case times are given in Table II.
- ② (Definition IV.3: Confidential \rightarrow Unsafe) Checks each incoming operation for illegal label-operand violations, and signals #LV when violations are detected. All checks are performed before operands are forwarded to the execution unit. If any violation is detected, the execution unit does not receive the operation and an #LV signal is sent to the

TABLE I: Notations and simple helper functions.

$ T $	Returns number of elements in T
$T[i:j]$	Returns items with index i to j (inclusive)
λ	Public program
Fetch, Execute, Retire	Instruction stages
Arithmetic, Branch MemLoad/Store	Instruction types
stage, pc, squash, update	Trace entry format
Write(addr, data, label)	Token denoting write to program memory
$Proj(T)$	Trace with updates removed
$arg_i(pc, \lambda), dest(pc, \lambda)$	Returns instruction operand/dest fields
$op(pc, \lambda)$	Returns instruction’s implied arithmetic op
$T.append(e)$	Append e to end of T
$type(pc, \lambda)$	Return instruction at pc’s type (Branch, etc)
$done(e, \lambda)$	Returns true if $e.stage = Retire$ and $e.pc$ is the stop PC given λ
SCHEDULE, PREDICT	Instruction scheduler and predictor functions

ROB, where it is interpreted as a violation (squash) or a fault (termination, #LF), respectively.

- ③ (Label propagation) Computes the result label based on operand labels and stages the label to travel with the result when it writes back to the RF or exits early via bypass.

The module is parameterized at design-time based on what functionality is actually needed. For example, Execution unit 2 (Table II) only supports Safe-operand arithmetic and therefore doesn’t need logic to enforce Definition IV.3 (Confidential \rightarrow Unsafe). Hence, this logic is pruned away at hardware synthesis time.

VI. SECURITY ANALYSIS

We will show that the data oblivious ISA provides a basis for satisfying *Oblivious*[BitCycle] (Section III-A) by proving its security over an *abstract* out-of-order, speculative machine (AOOM), and arguing that this abstract machine can be reduced to real hardware such as the BOOM.

A. ISA Level

The following analysis assumes the data oblivious ISA disables the ounseal instruction (Section IV-B). Given our analysis, the security of data oblivious randomized algorithms which use ounseal (e.g., ZeroTrace [20]) will reduce to the security of their underlying algorithm (e.g., Oblivious RAM [65]).

1) *Abstract machine basics*: The functional model for AOOM is given in Algorithm 2, with notations/helper functions explained in Table I and Algorithm 1. Our goal was to keep the model as simple as possible, while capturing core features. Specifically, the abstract machine: (1) has a 3-stage pipeline {Fetch, Execute, Retire} where each stage is atomic and takes one unit of time, (2) has four instruction types {Arithmetic, Branch, MemLoad, MemStore}, (3) has infinite fetch bandwidth and execution units, (4) can be parameterized as an in-order or out-of-order/speculative machine. Which instruction types support Safe/Unsafe operands are encoded as conditionals checking operands for label violations (#LV). We explain how to extend the model (e.g., to account for variable latency instructions, cache, limited execution units, more pipeline stages, etc.) in Section VI-B.

2) *Execution traces*: The abstract machine AOOM takes as input a program λ , Public input x and Confidential input y and generates a trace T where each entry T_i tracks a stage of each instruction as it executes on the machine. That is, the t -th element in T is a 4-tuple:

$$T_i = (\text{stage}_i, \text{pc}_i, \text{squash}_i, \text{update}_i).$$

stage_i denotes the instruction's stage $\{\text{Fetch}, \text{Execute}, \text{Retire}\}$. pc_i denotes the instruction address. Different stages for the same logical instruction share the same pc . If $\text{stage}_i = \text{Execute}$, $\text{squash}_i = \{\text{true}, \text{false}\}$ denotes whether the instruction caused a squash during speculation (Section II-A) or due to a label violation #LV (Section IV-A2). If $\text{stage}_i \neq \text{Execute}$, $\text{squash}_i = \text{false}$. $\text{update}_i = \text{Write}(\text{addr}, \text{data}, \text{label})$ where Write is a token denoting whether program memory was written, and with what addr , data and label . The Public label is logic 0, Confidential is logic 1. If no write occurs, $\text{addr} = \perp$.

3) *Modeling time*: In our abstraction, entries in T are ordered in time as $\text{time}(T_i) \leq \text{time}(T_{i+j})$ for $i, j \geq 0$ where time is a metric for real time (e.g., clock cycles). That is, multiple events may occur in the same clock cycle (as in a real processor) or be separated far apart. Therefore, stage_i and $\text{type}(\text{pc}_i, \lambda)$ allows us to model contention in different pipeline stages for different instruction types.

4) *Modeling out-of-order and speculative execution*: A key feature in our analysis is that AOOM is parameterized by two functions, SCHEDULE and PREDICT. SCHEDULE represents control logic in a real processor and decides which stage of which instruction should be evaluated next. It takes as input the program λ and $\text{Proj}(T)$, a projection of T that removes update from each entry, i.e.,

$$\text{Proj}(T) = \{e.\text{stage}, e.\text{pc}, e.\text{squash} \text{ for } e \in T\}$$

Importantly, $\text{Proj}(T)$ constrains scheduling to not be a function of program data (i.e., $e.\text{update}$) beyond the sequence of present/past fetched instructions ($e.\text{stage}$, $e.\text{pc}$) and whether those instructions result in a squash ($e.\text{squash}$). This represents real processors as we discuss in Section VI-B. SCHEDULE outputs an index $\text{idx} \in [0, |T|)$ or \perp . If $\text{idx} = \perp$, the machine will fetch the next instruction. If $\text{idx} \neq \perp$, the machine will evaluate the next stage for the instruction at $T[\text{idx}]$. PREDICT represents branch/jump predictor logic, takes the same inputs as SCHEDULE and outputs the predicted next PC. W.l.o.g. we assume SCHEDULE and PREDICT are deterministic.⁵

Importantly, SCHEDULE and PREDICT are representative of modern processors and allow us to model simple in-order processors to advanced out-of-order speculative processors (details on this claim related to BOOM are in Section VI-B). The only assumption we will make is that SCHEDULE respects in-order Fetch and Retire, as done by machines today.

5) *Modeling machine state*: The current machine state at some point idx in the trace is computed based on the trace prefix from 0 to idx . This includes program state (register file, cache, etc.) and intermediate pipeline/machine state. Program state is calculated based on mem (Algorithm 1). We merge the register file and other memory into a single memory for simplicity. Data always travels with label, which models

⁵Heuristics based on randomness can be modeled with an additional seed input.

Definition IV.1. As mentioned in Section VI-A2, pipeline state (e.g., flip-flops/SRAM not included in program state) is modeled by the sequence of PCs and stages in the trace.

6) *Proof of Security*: We now prove that the abstract model AOOM satisfies Definition III.1 with respect to the following observability function WordStage .

Definition VI.1. (*WordStage observability: Public data and labels at Word spatial granularity, instruction stage-level temporal granularity*) Given $T = \text{AOOM}(\lambda, x, y)$,

$$\text{WordStage}(T) = \{e.\text{stage}, e.\text{pc}, e.\text{squash}, h(e) \text{ for } e \in T\}$$

where $h(e)$ returns $e.\text{update}$ (unmodified) if $e.\text{update}.\text{label} = \text{false}$, and returns $\text{Write}(e.\text{addr}, \perp, \text{true})$ otherwise.

Algorithm 1: Helper functions memin and mem .

```

/* fill memory w/ Public x, Confidential y */
function: meminit(x,y)
1 T := [];
2 for xi ∈ x do
3   T.append((Execute, ⊥, false, Write(i, xi, false)))
4 for yi ∈ y do
5   T.append((Execute, ⊥, false, Write(|x| + i, yi, true)))
6 return T;
7
/* return coherent memory snapshot, given T.
Note, an instruction that is squashed by
another instruction may still create visible
state changes in the window of time before
the other instruction reaches Execute. */
function: mem(T)
8 T' = T with all squashed instructions (trace entries) removed.
That is, remove from T any entry that occurs in between the
Fetch and Execute stage of an instruction I if I satisfies
I.stage = Execute ∧ I.squash (inclusive);
9 mem := [⊥ for t ∈ T']; // |T'| upper-bounds mem size
10 for xi ∈ T' do
11   up := xi.update;
12   if up.addr ≠ ⊥ then
13     mem[up.addr] = up.data, up.label;
14 return mem;

```

That is, WordStage only removes write data from the trace if the label corresponding to that data is Confidential. Satisfying Definition III.1 with the WordStage function implies the strongest level of privacy with respect to our abstract machine, and implies that the machine's pipeline utilization, PC sequence, set of squash events, and state w.r.t. Public data is independent of Confidential data. This is equivalent to non-interference [44] and privacy over functional/timing flows [45].

Theorem 1. $\text{Oblivious}[\text{WordStage}, \text{AOOM}]$ holds.

Proof: Let $T = \text{AOOM}(\lambda, x, y)$ and $T' = \text{AOOM}(\lambda, x, y')$ where x is an array of Public data and y, y' are arrays of Confidential data. We must show $\text{WordStage}(T) \simeq \text{WordStage}(T')$. We proceed using induction. Line numbers refer to Algorithm 2. The base case holds throughout memin (Line 1, defined in Algorithm 1) since $|y| = |y'|$. Now assume $\text{WordStage}(T[0 : n]) \simeq \text{WordStage}(T'[0 : n])$. We know SCHEDULE (Line 3) will return the same idx for both executions because $\text{Proj}(T[0 :$

Algorithm 2: Abstract machine definition. As in Figure 2, the Public label is logic 0, Confidential is logic 1.

```

function: AOOM( $\lambda, x, y$ )
1  $T := \text{meminit}(x, y);$  // initialize memory
2 while !done( $T[|T| - 1], \lambda$ ) do
3    $\text{idx} := \text{SCHEDULE}(\text{Proj}(T), \lambda);$ 
4   if  $\text{idx} = \perp$  then // Fetch new instr
5      $\text{pc} := \text{PREDICT}(\text{Proj}(T), \lambda);$ 
6      $T.\text{append}(\text{Fetch}, \text{pc}, \text{false}, \text{Write}(\perp, \perp, \text{false}));$ 
7   else
8      $\text{pc} := T[\text{idx}].\text{pc};$ 
9      $\text{stage} := T[\text{idx}].\text{stage};$ 
10    if  $\text{stage} = \text{Fetch}$  then // Execute instr
11       $T.\text{append}(\text{execute}(\text{Execute}, \text{pc}, T, \lambda));$ 
12    else if  $\text{stage} = \text{Execute}$  then // Retire instr
13       $T.\text{append}(\text{Retire}, \text{pc}, \text{false}, \text{Write}(\perp, \perp, \text{false}));$ 
14 return  $T;$ 
15
function: execute(stage, pc,  $T, \lambda$ )
16  $\text{update} := \text{Write}(\perp, \perp, \text{false});$   $\text{squash} := \text{false};$ 
17  $\text{arg}_{0,\text{data}}, \text{arg}_{0,\text{label}} := \text{mem}(T)[\text{arg}_0(\text{pc}, \lambda)];$ 
18  $\text{arg}_{1,\text{data}}, \text{arg}_{1,\text{label}} := \text{mem}(T)[\text{arg}_1(\text{pc}, \lambda)];$ 
19 if  $\text{type}(\text{pc}, \lambda) = \text{Arithmetic}$  then
20    $\text{data} := \text{arg}_{0,\text{data}} \text{op}(\text{pc}, \lambda) \text{arg}_{1,\text{data}};$ 
21    $\text{label} := \text{arg}_{0,\text{label}} \vee \text{arg}_{1,\text{label}};$ 
22    $\text{update} := \text{Write}(\text{dest}(\text{pc}, \lambda), \text{data}, \text{label});$ 
23 else if  $\text{type}(\text{pc}, \lambda) = \text{Branch}$  then
24   if  $\text{arg}_{0,\text{label}} \vee \text{arg}_{1,\text{label}}$  then
25      $\text{squash} := \text{true};$  // #LV: Confidential->Unsafe
26   else
27      $\text{fidx} := \text{index of Fetch for current instr in } T;$ 
28      $\text{guess} := \text{direction for PREDICT}(\text{Proj}(T[0 : \text{fidx}], \lambda));$ 
29      $\text{actual} := \text{arg}_{0,\text{data}} \text{op}(\text{pc}, \lambda) \text{arg}_{1,\text{data}};$ 
30      $\text{squash} := \text{guess} \neq \text{actual};$  // mispredict
31 else
32   if  $\text{arg}_{0,\text{label}}$  then
33      $\text{squash} := \text{true};$  // #LV: Confidential->Unsafe
34   else
35     if  $\text{type}(\text{pc}, \lambda) = \text{MemLoad}$  then
36        $\text{data}, \text{label} := \text{mem}(T)[\text{arg}_{0,\text{data}}];$ 
37        $\text{addr} := \text{dest}(\text{pc}, \lambda)$ 
38     else if  $\text{type}(\text{pc}, \lambda) = \text{MemStore}$  then
39        $\text{data}, \text{label} := \text{arg}_{1,\text{data}}, \text{arg}_{1,\text{label}};$ 
40        $\text{addr} := \text{arg}_{0,\text{data}}$ 
41      $\text{update} := \text{Write}(\text{addr}, \text{data}, \text{label})$ 
42 return  $\text{stage}, \text{pc}, \text{squash}, \text{update};$ 

```

$n]) = \text{Proj}(T'[0 : n])$ by the induction hypothesis and λ is fixed (Section VI-A4). We now proceed by cases depending on idx :

Case 1 ($\text{idx} = \perp$). A new instruction will be fetched for both executions. We know PREDICT (Line 5) will return the same PC across both executions, following the same logic as SCHEDULE. By Line 6, it is clear the induction step holds.

Case 2 ($T[\text{idx}].\text{stage} = \text{Fetch}$). The instructions at idx in T and T' will be executed. There are 4 cases in Execute, depending on the instruction type $\text{type}(\text{pc}, \lambda)$, and a careful inspection shows

that each satisfies the induction step. Of note, the only possible deviation between T and T' is over update.data , in the case of Arithmetic and MemStore instructions when $\text{update.label} = \text{true}$, which is allowed by Definition VI.1.

Case 3 ($T[\text{idx}].\text{stage} = \text{Execute}$). The instructions at idx in T and T' will be retired. By Line 13, it is clear the induction step holds. ■

Takeaway: security holds with hardware optimizations enabled. An important takeaway is that security holds despite speculative and out-of-order execution, meaning that these optimizations can be safely enabled with our data oblivious ISA. This is underpinned by three important assumptions. First, the SCHEDULE and PREDICT are constrained to be functions of the program λ and Proj of the current trace (discussed in Section VI-B). Second, we perform local checks to determine label violations #LV (Lines 25 and 33), implementing Definition IV.3. This formalizes why label stations are secure (Section V-C2). Third, instructions with Safe operands don't leak secrets, implementing Definition IV.2.

B. Implementation Level

We now map our ISA-level security analysis (Sections IV-B and VI-A) to our prototype on BOOM (Section V), referred to as BOOM.

1) *Threat vectors in unmodified BOOM:* Unmodified BOOM hardware (Section V-A) supports speculation over branches, jumps and unresolved store instructions (Vectors 1-3; c.f. Section III-B) as well as arithmetic units with input-dependent timing (Vector 6, Table II).⁶ Our implementation of the OMP (Section V-B) is also susceptible to cache bank contention (Vector 5) because it uses space in the data cache.

2) *Securing BOOM:* Recall, the primary hardware mechanisms we added to get security are dynamic information flow tracking (Section V-C1), label stations per execution unit (Section V-C2) with logic to handle #LV, and logic to isolate the OMP (Section V-B).

In Section VI-A, we proved *Oblivious*[WordStage, AOOM]. We now show how to use the proof to argue *Oblivious*[BitCycle, BOOM]—i.e., cycle-level security of our implementation—which implies that Vectors 1-3 and 5-6 are blocked. There are two steps: (1) mapping AOOM to BOOM and (2) mapping WordStage to BitCycle.

Step 1: mapping AOOM to BOOM. We kept AOOM simple to illustrate key points for presentation, but can make the model more sophisticated to better represent our actual implementation. In particular: relative to BOOM, AOOM has (a) an idealized SCHEDULE/PREDICT, (b) an unlimited number of execution units without dependencies, (c) less instruction types, (d) less pipeline stages, (e) fixed latency (“atomic”) arithmetic/memory operations, (e) no support for jump/memory speculation, and (f) no support for the OMP.

For (a): we note that BOOM's actual scheduler/predictors operates over a subset of the inputs given to SCHEDULE and

⁶We note BOOM also supports load/store forwarding but is not susceptible to Vector 4 because the data TLB is accessed sequentially before checking the SAQ (Section V-A).

PREDICT. For example, BOOM’s branch/jump predictors take as input the sequence of fetched PCs (speculative and non-speculative) and branch outcomes for retired instructions [25]. The sequence of PCs is contained in $Proj(T)$ and branch outcomes can be extrapolated from $Proj(T)$ given the program λ (Section VI-A4). The issue/scheduler logic is similar: employing data-independent information such as instruction age in the window and public instruction dependencies. For (b): limited execution units (structural hazards) and instruction dependencies (data hazards) can be modeled as additional rules in the SCHEDULE function. For example, SCHEDULE may not consider executing an in-flight instruction until its dependencies given in-flight instructions (recent trace entries) have executed. For (c) and (d): it is straightforward to add extra stages (e.g., BOOM’s issue stage) and instruction types. We note that the 4 types (Arithmetic, etc.) in the analysis cover the major cases in the data oblivious ISA (Section IV-B). For (e): variable arithmetic can be modeled by creating additional execute stages $Execute^i$ for $i < L$. To match our label station hardware (Section V-C2), the limit L can be set to a constant based on the operand label and the instruction type. Caches can similarly be modeled by allowing L to become a function of the sequence of previous memory addresses. Additional types of speculation (e.g., jump and memory speculation) can be modeled using Branch type instructions. For (f): it is easy to model the OMP (Section IV-B6) in a fashion analogous to MemLoad and MemStore, as hardware mechanisms moderating access to the OMP ensure that the OMP satisfies non-interference w.r.t. other cache lookups.

Note, we also don’t model RF or cache compression (Vectors 8-9) as these optimizations don’t appear in BOOM. They can similarly be modeled with addition #LV checks on data at rest.

Step 2: mapping WordStage to BitCycle. It is straightforward to extend Theorem 1 to include time (Section VI-A3), meaning the analysis can satisfy a “WordCycle” observability function as opposed to WordStage. We have performed a careful audit of our BOOM hardware prototype (analyzing logic paths through major flows in the processor at cycle level) to ensure intermediate storage elements (e.g., pipeline flip-flops) correspondingly carry Confidential data in data-independent cycles. We have not found violations to this rule outside of one exception (given below).

The main insight that enables the correspondence is that stages (e.g., Fetch, Execute) in our abstract model (Section VI-A2) correspond to intermediate storage elements. That is, each instruction stage corresponds to storage elements, e.g. pipeline registers in functional units, and the proof shows how the sequence of stages in the trace is invariant to Confidential input. This means that modeling AOOM with additional stages/sub-stages (e.g., rename, issue, $Execute^i$, etc.) provides a basis for modeling complex machines at the flip-flop and cycle level. We leave automated analyses on the design (e.g., using [60]) as future work.

The one exception (noted above) is label stations. We assume label stations are implemented correctly, meaning a software adversary cannot view flip flop-level label state within a label station, while that unit is processing Confidential data, or within the OMP. This is a minor detail, handled in the analysis by removing state bits in BitCycle when units are processing

TABLE II: RISC-V BOOM parameters we use for our prototype and evaluation. Arithmetic units with a ‘(xx)’ next to their name are un-pipelined (variable latency), where ‘xx’ denotes the worst-case latency. The prefix ‘i’ denotes integer, ‘f’ denotes floating point. **CondMove** and **Omp** denote logic for ocmov and the oblivious memory partition (Section IV-B), respectively, and are only present on our modified BOOM.

Core μ arch	out-of-order, speculative
Fetch/issue width	4 instructions fetched/issued per cycle
Execution unit 1	iALU, Branch, iMul, iDiv (6-66)
Execution unit 2	iALU, CondMove
Execution unit 3	IntToFP casting
Execution unit 4	fAdd, fMul, fDiv (5-21), fSqrt (5-29), FPToInt casting
Execution unit 5	Load/store + Omp (memory unit)
L1 I/D cache	32 KB, 4 way/64 KB, 16 way; 64 B cache lines
I/D TLB	16/32 entries

Confidential data. A similar principle is applied in showing security for Execution Leases [24].

Lastly, we note that WordCycle is equivalent to BitCycle in modern processors, including the BOOM. In modern ISAs like RISC-V, operations occur at word granularity meaning that labels only need to be tracked at word granularity.

VII. EVALUATION

We now evaluate our data oblivious ISA in terms of area overhead (given our prototype on RISC-V BOOM) and performance over data oblivious workloads. We also show two case studies, showing how the data oblivious ISA secures and accelerates constant time cryptographic code and memory oblivious libraries.

A. Methodology

We evaluate our system through hardware prototyping to show area overheads and software simulation to show performance.

1) *Hardware prototyping:* We build on top of the open-source BOOM design [25] which is written in the Chisel hardware description language [70]. We parameterized the prototype according to Table II and synthesized the design using a 32 nm commercial process and the Synopsys flow. We report standard cell (logic cell) area for logic and flip-flops post-synthesis, and report SRAM area using the widely used Cacti tool [71]. BOOM maps the instruction/data caches/TLBs and branch predictor tables to SRAM. Remaining storage structures (e.g., the SDQ, register files) are mapped to flip-flops. The BOOM word width is 64 bits.

2) *Software simulation:* The BOOM hardware only features a single-level cache, whereas commercial machines feature two- or three-level caches to reduce traffic to DRAM. Thus, to measure more realistic performance figures for our system we use Multi2Sim [72], parameterized to match Table II as closely as possible. For all experiments, we use a 256 KB 4-way level 2 cache (that is shared by data and instructions) and a 2 MB 16-way level 3 cache. This configuration is similar to a single slice on an Intel Skylake machine.

TABLE III: Area for baseline and modified BOOM cores. Units are μm^2 .

	BOOM	BOOM + data oblivious ISA	Overhead
Logic	363,900	388,658	6.80%
SRAM	384,232	391,291	1.84%
Total	748,132	779,949	4.25%

3) *OMP usage*: We use a 32 KB OMP (Section IV-B6) that is built into level 1 data cache. This is sufficient to store ORAM sub-structures (Section IV-B6) and also big enough to fit tables for constant time cryptographic routines (e.g., AES T-tables and RSA multiplier tables). Some workloads do not benefit from the OMP (e.g., some do not have data-dependent memory access patterns). In this case, a bit in thread state disables the OMP to recover cache space.

B. Hardware Prototyping and Area Results

We show area results for unmodified BOOM and BOOM extended to support our data oblivious ISA in Table III. Our prototype supports all instructions in Section IV-B and Figure 2. The main hardware components needed to support the data oblivious ISA are storage for DIFT, logic/storage for label stations, logic to partition the OMP, and a random number generator for *orng* (Section V). For structures that need to store labels, we store those labels alongside the data in whatever medium the data was stored in. That is, labels in data cache are stored in SRAM, labels in the SDQ and register files are stored in flip-flops. The largest single area overhead comes from an iterative AES core that we downloaded from OpenCores [73] to implement *orng*. This unit has area $10,935 \mu\text{m}^2$ (3% of the logic area for the unmodified BOOM), and can be replaced by a hardware TRNG (whose area is negligibly small [68]) in a production design.

The takeaway is that hardware overheads are tolerable, both on the logic and SRAM side, showing the practicality of the proposal on advanced commercial-class machines.

C. Performance Results

We now perform studies to evaluate the performance overhead of running data oblivious code securely, with and without the oblivious memory partition.

1) *Comparison systems*: We compare two systems—*doisa* and *doisa_omp*—to a baseline insecure system. All three systems use the same microarchitecture (Table II). Benchmarks run on insecure are written in a non-data oblivious fashion (i.e., without the constraints in Section III-B). Benchmarks run on *doisa* are data oblivious, and written using all instructions in Figure 2 except *ocld/ocst* (the oblivious memory extension; c.f. Section IV-B6). Thus, *doisa* will be similar performance-wise to existing data oblivious codes, e.g., *Raccoon* [1], which don’t have access to an OMP. Benchmarks run on *doisa_omp* use all instructions in Figure 2.

2) *Workloads*: We evaluate a suite of common workloads (Table IV) which have previously been written and evaluated data obliviously [1], [64], [74], [62] on existing x86 ISAs. These codes are divided into three categories. First, codes that are nearly data oblivious in their default form (*mat. mult.*,

TABLE IV: Benchmarks and input data sizes for comparing insecure, *doisa* and *doisa_omp*.

Name	Implementation	Data size (small / large)
<i>mat. mult</i>	data oblivious by default	256x256 / 1024x1024
<i>neural network</i>	“”	64-1K-8 / 1024-32K-256 (2 layers)
<i>findmax</i>	“”	8K / 1M integers
<i>sort</i>	bitonic-sort (<i>doisa</i>), data obl. merge-sort (<i>doisa_omp</i>)	4K / 256K integers
<i>pagerank</i>	GraphSC [74]	1K / 16K nodes
<i>binary search</i>	memory scan (<i>doisa</i>), obl. memory (<i>doisa_omp</i>)	8K / 16M integers
<i>kmeans</i>	obl. memory for histogram	64/256 clusters, 4K/32K points
<i>heap push</i>	ODS [64]	8K / 32M integers in heap
<i>heap pop</i>	ODS [64]	8K / 32M integers in heap
<i>sparse dijkstra</i>	ObliVM [62]	256 / 4K vertices

neural network, *findmax*). Second, codes that rely heavily on data oblivious sort as a subroutine (*sort* and *pagerank*). Third, codes that rely heavily on oblivious memory (*binary search*, *kmeans*, *heap*, *dijkstra*). We will also perform case studies showing our proposal’s applicability in two additional important settings—constant time cryptography and oblivious memory—in Sections VII-C5 and VII-C6.

3) *Data set sizes*: For each benchmark, we evaluate ‘small’ and ‘large’ data sizes. ‘small’ indicates the largest input size that wholly fits into the 32 KB OMP (Section VII-A). We use this configuration for two reasons. First, to show the benefit of having an OMP. Second, to performance compare against prior work (*Raccoon* [1], which uses similar data sizes). Finally, we show the ‘large’ data size to illustrate overheads where program data does not completely fit into the OMP. In that case, we fallback to ORAM or SCAN as described in Section IV-B6.

4) *Results*: Figure 7 shows the overhead of $\{\text{doisa}, \text{doisa_omp}\} \times \{\text{small}, \text{large}\}$ relative to insecure. The main takeaway is that *doisa_omp* achieves significant ($8.8 \times / 1.7 \times$ for small/large data sizes) speedup over *doisa*. Furthermore, *doisa_omp* has only $3.2 \times / 40.4 \times$ slowdown on relative to insecure on the same data sizes. This shows that our data oblivious ISA makes data oblivious computing practical in cases where data fits in the OMP.

There are two avenues for future work. First, enhance the OMP to support larger sizes (e.g., beyond the level 1 data cache, see Section IV-C2). As we see on the large data set size, overhead for both *doisa* and *doisa_omp* can be large for workloads that depend on oblivious memory, as large data sizes cannot fit into the OMP. Second, engineer more sophisticated instructions supporting Safe operands. For example, *sort* is an important kernel in multiple data oblivious codes [75], [74], [16], [62]. A data oblivious ISA can support an *osort* instruction with Safe operands directly, and use techniques such as hardware partitioning to speedup that operation.

5) *Case study: constant time AES*: An important commercial use-case for data oblivious code today is “constant time” cryptography. Many papers have demonstrated how unprotected codes—e.g., T-table AES [29] and naive modular exponentiation for RSA—leak privacy over microarchitectural side channels (specifically, through the cache [3]). As a result, practitioners use slower codes to improve security—e.g., S-box or bitslice AES [30] and montgomery ladder exponentiation




Fig. 7: Performance comparison between doisa and doisa_omp relative to insecure for small/large data sets.

for RSA.⁷

Our data oblivious ISA provides a basis for running high-performance cryptography securely. To demonstrate the benefit, we compare the performance of T-table AES [29] (high performance, low security) vs. bitslice AES [30] (low performance, high security). On one hand, we retrofit T-table AES using our ISA and store the T-tables in the OMP to prevent cache attacks (the rest of the code is naturally data oblivious). This gives us a high performance, high security code. Our ISA can securely run both the fully unrolled code or a variant with a loop over the number of rounds, regardless of branch prediction accuracy (Section III-B). On the other hand, we argue that on commodity machines today, highly sensitive applications will have to resort to codes like bitslice AES.

Both codes are compiled with gcc using `-O3` optimizations. Relative to an insecure T-table AES code (insecure), our data oblivious T-table AES (doisa_omp) has a $2.17\times$ slowdown, while bitslice AES has a $9.6\times$ slowdown against the same baseline. Our slowdown relative to insecure is caused by the compiler not optimizing code around `ocld` instructions. Thus, doisa_omp can achieve even lower slowdown with better compiler support.

6) *Case Study: ZeroTrace* [20]: Beyond cryptography, there is a rich literature to accelerate data structure operations data obliviously [64], [62], [18]. These schemes typically use oblivious memory as a subroutine. We now demonstrate how the data oblivious ISA can speedup this subroutine by comparing our oblivious memory API to the original ZeroTrace [20] proposal. Discussed in Section IV-B6, our library combines ZeroTrace with the OMP to achieve speedup for different oblivious memory sizes.

Results are shown in Figure 8. doisa_omp provides significant speedup in all size regimes. For small data, doisa_omp places the entire memory in the OMP, providing $O(1)$ ($> 1000\times$ speedup) time access to that data. For larger data, doisa_omp uses SCAN or ORAM, depending on which strategy yields best performance, and places the ORAM stash in the OMP in the latter case. An important finding in the ZeroTrace paper is that stash management, written data obliviously, creates a

⁷Discussed in Section III-B, even hardened codes may be insecure due to subtle hardware optimizations.




Fig. 8: Comparison between oblivious memory primitives. Scan is the SCAN code from Figure 4, shown for completeness. non-recursive/recursive Path ORAM are baseline ZeroTrace [20].

performance bottleneck.⁸ Since the stash does not grow as a function of the ORAM capacity, we can use the OMP to store the stash and manage it more efficiently, which allows us to improve over baseline ZeroTrace by $\geq 4.6\times$ in all regimes.

VIII. RELATED WORK

Beyond data oblivious code written for today’s ISAs, there is a rich literature to improve algorithm/data structure [65], [77], [75], [74], [78], [79], [62], [80], [81], [65], [64] performance in the *software circuit* abstraction. Additionally, there is rich literature to write (e.g., [63], [82]) and compile (e.g., [62], [83], [80]) programs to software circuits. An important observation is that, although many of these works target cryptographic backends such as garbled circuits, their underlying programming abstraction (software circuits) is very similar to the data oblivious abstraction. For example, bitwise crypto can be easily mapped to integer-wide operations. Thus, our proposal can be used as a secure hardware backend for these works.

Also working in the software circuit abstraction, TinyGarble [80] advocates the use of *sequential logic* to reduce code/data footprint to, e.g., reduce cache footprint. While optimizing data oblivious code to have better footprints is beyond our scope, our ISA provides mechanisms—e.g., conditional moves to reuse locations in memory, data-independent branches—to allow programs to reduce their own footprint.

Secure co-processor proposals Ghost rider [27] and Ascend [26] have the same security goal (Definition III.1) as this paper, but assume a coarse-grain observability function that only captures the processor’s external pin activity (whereas this paper considers fine-grain observability; c.f. Section III-A). These proposals also assume simple processor pipelines and scheduling (e.g., one process per chip at a time). Relative

⁸We note that an alternate ORAM, Circuit ORAM [76], was designed to avoid stash management overheads. Unfortunately, Circuit ORAM has worse bandwidth— $12 * \log n$ vs. $8 * \log n$ and $3.5 * \log n$ for data size n —than Path ORAM, which relies on a stash. Since our oblivious memory extensions make stash management essentially free, our scheme based on Path ORAM will outperform Circuit ORAM.

to these works, our goal is to show how to retrofit *existing* high-performance machines to concurrently run sensitive and non-sensitive programs, which matches how programmers are writing data oblivious code today.

There is a rich literature to block cache side channel attacks [84], [85], [86], which is related to our oblivious memory partition. Our proposal is complementary, as our goal is to expose a common, portable programmer-level interface for schemes like these. We note, however, that none of these works satisfy Definition III.1 at cycle granularity as is. For example, PLCache [84] was designed to mitigate cache attacks and therefore doesn't try to hide whether there was a hit or miss *within* a secure partition.

IX. CONCLUSION

This paper proposes an Oblivious ISA extension to enable secure and high-performance data oblivious computing in the data oblivious abstraction. We propose ISA principles, a concrete ISA, a complete prototype on an advanced microprocessor, and accompanying formal analyses for all of the above. Long term, we hope this paper serves as a step for writing and running safe, portable and performant data oblivious code for sensitive applications.

REFERENCES

- [1] A. Rane, C. Lin, and M. Tiwari, "Raccoon: Closing digital side-channels through obfuscated execution," in *USENIX Security'15*.
- [2] W. Wang, G. Chen, X. Pan, Y. Zhang, X. Wang, V. Bindschaedler, H. Tang, and C. A. Gunter, "Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX," *CoRR'17*.
- [3] D. A. Osvik, A. Shamir, and E. Tromer, "Cache attacks and countermeasures: The case of aes," in *CT-RSA'06*.
- [4] O. Aciicmez, J.-P. Seifert, and C. K. Koc, "Predicting secret keys via branch prediction," *IACR'06*.
- [5] D. Evtushkin, R. Riley, N. C. Abu-Ghazaleh, ECE, and D. Ponomarev, "Branchscope: A new side-channel attack on directional branch predictor," in *ASPLOS'18*.
- [6] M. Andryscio, D. Kohlbrenner, K. Mowery, R. Jhala, S. Lerner, and H. Shacham, "On subnormal floating point and abnormal timing," in *S&P'15*.
- [7] J. Großschädl, E. Oswald, D. Page, and M. Tunstall, "Side-channel analysis of cryptographic software via early-terminating multiplications," *IACR'09*.
- [8] P. Kocher, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, "Spectre attacks: Exploiting speculative execution," *ArXiv'18*.
- [9] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, "Meltdown," *ArXiv'18*.
- [10] D. J. Bernstein, "Curve25519: New diffie-hellman speed records," in *PKC'06*.
- [11] D. J. Bernstein, "The poly1305-aes message-authentication code," in *FSE'05*.
- [12] D. Molnar, M. Piotrowski, D. Schultz, and D. Wagner, "The program counter security model: Automatic detection and removal of control-flow side channel attacks," *IACR'05*.
- [13] D. B. S. G. Ben A. Fisch, Dhinakaran Vinayagamurthy, "Iron: Functional encryption using intel sgx," in *CCS'17*.
- [14] O. Ohrimenko, F. Schuster, C. Fournet, A. Mehta, S. Nowozin, K. Vaswani, and M. Costa, "Oblivious multi-party machine learning on trusted processors," in *USENIX Security'16*.
- [15] Z. L. L. K. Fahad Shaon, Murat Kantarcioglu, "Sgx-bigmatrix: A practical encrypted data analytic framework with trusted processors," in *CCS'17*.
- [16] W. Zheng, A. Dave, J. G. Beekman, R. A. Popa, J. E. Gonzalez, and I. Stoica, "Opaque: An oblivious and encrypted distributed analytics platform," in *NSDI'17*.
- [17] S. Eskandarian and M. Zaharia, "An oblivious general-purpose SQL database for the cloud," *CoRR'17*.
- [18] P. Mishra, R. Poddar, J. Chen, A. Chiesa, and R. A. Popa, "Oblix: An efficient oblivious search index," in *S&P'18*.
- [19] S. Tople and P. Saxena, "On the trade-offs in oblivious execution techniques," in *Detection of Intrusions and Malware, and Vulnerability Assessment* (M. Polychronakis and M. Meier, eds.), Springer'17.
- [20] S. Sasy, S. Gorbunov, and C. W. Fletcher, "ZeroTRACE: Oblivious memory primitives from intel sgx," in *NDSS'18*.
- [21] A. Ahmad, K. Kim, M. I. Sarfaraz, and B. Lee, "Obliviate: A data oblivious filesystem for intel sgx," in *NDSS'18*.
- [22] B. Coppens, I. Verbauwhede, K. D. Bosschere, and B. D. Sutter, "Practical mitigations for timing-based side-channel attacks on modern x86 processors," in *S&P'09*.
- [23] "Speculative execution side channel mitigations." <https://software.intel.com/sites/default/files/managed/c5/63/336996-Speculative-Execution-Side-Channel-Mitigations.pdf>. Revision 1.0, January 2018.
- [24] M. Tiwari, X. Li, H. M. G. Wassel, F. T. Chong, and T. Sherwood, "Execution leases: A hardware-supported mechanism for enforcing strong non-interference," in *MICRO'09*.
- [25] C. Celio, P.-F. Chiu, B. Nikolic, D. A. Patterson, and K. Asanovi, "Boom v2: an open-source out-of-order risc-v core," tech. rep., EECS Department, University of California, Berkeley, Sep 2017. <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-157.html>, <https://ccelio.github.io/riscv-boom-doc/>.
- [26] C. Fletcher, M. Van Dijk, and S. Devadas, "A secure processor architecture for encrypted computation on untrusted programs," in *STC'12*.
- [27] C. Liu, A. Harris, M. Maas, M. Hicks, M. Tiwari, and E. Shi, "Ghostrider: A hardware-software system for memory trace oblivious computation," *SIGPLAN Not.*, vol. 50, pp. 87–101, Mar. 2015.
- [28] K. Nayak, C. W. Fletcher, L. Ren, N. Chandran, S. Lokam, E. Shik, and V. Goyal, "Hop: Hardware makes obfuscation practical," in *NDSS'17*.
- [29] "T-table AES (OpenSSL)." https://github.com/openssl/openssl/blob/master/crypto/aes/aes_core.c.
- [30] "Bitslice AES (Bitcoin)." <https://github.com/bitcoin-core/ctaes>.
- [31] G. B. Bell and M. H. Lipasti, "Deconstructing commit," in *ISPASS'04*.
- [32] M. H. Lipasti, C. B. Wilkerson, and J. P. Shen, "Value locality and load value prediction," *SIGPLAN Not.*, vol. 31, pp. 138–147, Sept. 1996.
- [33] A. Moghimi, G. Irazoqui, and T. Eisenbarth, "Cachezoom: How SGX amplifies the power of cache attacks," *CoRR'17*.
- [34] Y. Xu, W. Cui, and M. Peinado, "Controlled-channel attacks: Deterministic side channels for untrusted operating systems," in *S&P'15*.
- [35] M.-W. Shih, S. Lee, T. Kim, and M. Peinado, "T-sgx: Eradicating controlled-channel attacks against enclave programs," February 2017.
- [36] T. Hunt, Z. Zhu, Y. Xu, S. Peter, and E. Witchel, "Ryoan: A distributed sandbox for untrusted computation on secret data," in *OSDI'16*.
- [37] D. Gullasch, E. Bangerter, and S. Krenn, "Cache games – bringing access-based cache attacks on aes to practice," in *S&P'11*.
- [38] Intel, "Intel Software Guard Extensions Programming Reference." software.intel.com/sites/default/files/329298-001.pdf, 2013.
- [39] P. Subramanyan, R. Sinha, I. Lebedev, S. Devadas, and S. A. Seshia, "A formal foundation for secure remote execution of enclaves," in *CCS'17*.
- [40] P. C. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *CRYPTO'99*.
- [41] A. Nazari, N. Sehatbakhsh, M. Alam, A. Zajic, and M. Prvulovic, "Eddie: Em-based detection of deviations in program execution," in *ISCA'17*.
- [42] Y. Kim, R. Daly, J. Kim, C. Fallin, J. H. Lee, D. Lee, C. Wilkerson, K. Lai, and O. Mutlu, "Flipping bits in memory without accessing them: An experimental study of dram disturbance errors," in *ISCA'14*.
- [43] S. Gueron, "A memory encryption engine suitable for general purpose processors," *IACR'16*.

- [44] G. Smith, "Principles of secure information flow analysis," in *Malware Detection*, 2007.
- [45] J. Oberg, S. Meiklejohn, T. Sherwood, and R. Kastner, "Leveraging gate-level properties to identify hardware timing channels," *IEEE TCAD'14*.
- [46] "Spectre variants." [https://en.wikipedia.org/wiki/Spectre_\(security_vulnerability\)](https://en.wikipedia.org/wiki/Spectre_(security_vulnerability)).
- [47] S. Gueron, "Efficient software implementations of modular exponentiation," *IACR'11*.
- [48] Intel, "Intel Software Guard Extensions Software Development Kit." <https://software.intel.com/en-us/sgx-sdk>.
- [49] A. Moghimi, T. Eisenbarth, and B. Sunar, "Memjam: A false dependency attack against constant-time crypto implementations," *CoRR'17*.
- [50] Y. Yarom, D. Genkin, and N. Heninger, "Cachebleed: A timing attack on openssl constant time rsa," *IACR'16*.
- [51] S. Jourdan, R. Ronen, M. Bekerman, B. Shomar, and A. Yoaz, "A novel renaming scheme to exploit value temporal locality through physical register reuse and unification," in *MICRO'98*.
- [52] A. R. Alameldeen and D. A. Wood, "Adaptive cache compression for high-performance processors," *SIGARCH Comput. Archit. News*, vol. 32, pp. 212–, Mar. 2004.
- [53] C. A. Waldspurger, "Memory resource management in vmware esx server," *SIGOPS Oper. Syst. Rev.*, vol. 36, pp. 181–194, Dec. 2002.
- [54] K. M. Lepak and M. H. Lipasti, "Silent stores for free," in *MICRO'00*.
- [55] M. Dalton, H. Kannan, and C. Kozyrakis, "Raksha: A flexible information flow architecture for software security," *SIGARCH Comput. Archit. News*, vol. 35, pp. 482–493, June 2007.
- [56] G. E. Suh, J. W. Lee, D. Zhang, and S. Devadas, "Secure program execution via dynamic information flow tracking," *SIGARCH Comput. Archit. News*, vol. 32, pp. 85–96, Oct. 2004.
- [57] B. D. Rouhani, M. S. Riazi, and F. Koushanfar, "Deepsecure: Scalable provably-secure deep learning," *CoRR'17*.
- [58] C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan, "Gazelle: A low latency framework for secure neural network inference," *IACR'18*.
- [59] A. Waterman, Y. Lee, D. A. Patterson, and K. Asanovi, "The risc-v instruction set manual, volume i: User-level isa, version 2.0," Tech. Rep. UCB/Eecs-2014-54, EECS Department, University of California, Berkeley, May 2014.
- [60] D. Zhang, Y. Wang, G. E. Suh, and A. C. Myers, "A hardware design language for timing-sensitive information-flow security," *SIGPLAN Not.*, vol. 50, pp. 503–516, Mar. 2015.
- [61] M. Tiwari, H. M. Wassel, B. Mazloom, S. Mysore, F. T. Chong, and T. Sherwood, "Complete information flow tracking from the gates up," *SIGARCH Comput. Archit. News*, vol. 37, pp. 109–120, Mar. 2009.
- [62] C. Liu, X. S. Wang, K. Nayak, Y. Huang, and E. Shi, "Oblivm: A programming framework for secure computation," in *S&P'15*.
- [63] D. Darais, C. Liu, I. Sweet, and M. Hicks, "A language for probabilistically oblivious computation," *CoRR'17*.
- [64] X. S. Wang, K. Nayak, C. Liu, T.-H. H. Chan, E. Shi, E. Stefanov, and Y. Huang, "Oblivious data structures," *IACR'14*.
- [65] E. Stefanov, M. van Dijk, E. Shi, T.-H. H. Chan, C. Fletcher, L. Ren, X. Yu, and S. Devadas, "Path oram: An extremely simple oblivious ram protocol," *CCS'13*.
- [66] H. Cook, K. Asanovi, and D. A. Patterson, "Virtual local stores: Enabling software-managed memory hierarchies in mainstream computing environments," tech. rep., 2009.
- [67] B. Yee, D. Sehr, G. Dardyk, J. B. Chen, R. Muth, T. Ormandy, S. Okasaka, N. Narula, and N. Fullagar, "Native client: A sandbox for portable, untrusted x86 native code," in *S&P'09*.
- [68] V. Fischer, "Random number generators for cryptography (design and evaluation)." <https://summerschool-croatia.cs.ru.nl/2014/slides/Random%20Number%20Generators%20for%20Cryptography.pdf>.
- [69] L. Domnitser, A. Jaleel, J. Loew, N. Abu-Ghazaleh, and D. Ponomarev, "Non-monopolizable caches: Low-complexity mitigation of cache side channel attacks," *TACO'12*.
- [70] J. Bachrach, H. Vo, B. Richards, Y. Lee, A. Waterman, R. Avizienis, J. Wawrzyniek, and K. Asanovic, "Chisel: constructing hardware in a scala embedded language," in *DAC'12*.
- [71] N. Muralimanohar and R. Balasubramonian, "Cacti 6.0: A tool to understand large caches."
- [72] R. Ubal, B. Jang, P. Mistry, D. Schaa, and D. Kaeli, "Multi2Sim: A Simulation Framework for CPU-GPU Computing," in *PACT'12*.
- [73] "Open cores." <https://opencores.org/>.
- [74] K. Nayak, X. S. Wang, S. Ioannidis, U. Weinsberg, N. Taft, and E. Shi, "Graphsc: Parallel secure computation made easy," in *S&P'15*.
- [75] M. Blanton, A. Steele, and M. Alisagari, "Data-oblivious graph algorithms for secure computation and outsourcing," in *ASIA CCS'13*.
- [76] X. Wang, H. Chan, and E. Shi, "Circuit oram: On tightness of the goldreich-ostrovsky lower bound," *IACR'14*.
- [77] "Bitonic sort." https://en.wikipedia.org/wiki/Bitonic_sorter.
- [78] J. Doerner, D. Evans, and abhi shelat, "Secure stable matching at scale," *IACR'16*.
- [79] T.-H. H. Chan, Y. Guo, W.-K. Lin, and E. Shi, "Cache-oblivious and data-oblivious sorting and applications," *IACR'17*.
- [80] E. M. Songhori, S. U. Hussain, A. Sadeghi, T. Schneider, and F. Koushanfar, "Tinygarble: Highly compressed and scalable sequential garbled circuits," in *S&P'15*.
- [81] S. Zahur and D. Evans, "Circuit structures for improving efficiency of security and privacy tools," in *S&P'13*.
- [82] S. Cauligi, G. Soeller, F. Brown, B. Johannesmeyer, Y. Huang, R. Jhala, and D. Stefan, "Fact: A flexible, constant-time programming language," *SecDev'17*.
- [83] S. Zahur and D. Evans, "Obliv-c: A language for extensible data-oblivious computation," *IACR'15*.
- [84] Z. Wang and R. B. Lee, "New cache designs for thwarting software cache-based side channel attacks," in *ISCA'07*.
- [85] F. Liu, Q. Ge, Y. Yarom, F. Mckeen, C. Rozas, G. Heiser, and R. B. Lee, "Catalyst: Defeating last-level cache side channel attacks in cloud computing," in *HPCA'16*.
- [86] D. Gruss, J. Lettner, F. Schuster, O. Ohrimenko, I. Haller, and M. Costa, "Strong and efficient cache side-channel protection using hardware transactional memory," in *USENIX Security'17*.