

Creating Foundations for Secure Microarchitectures With Data-Oblivious ISA Extensions

Jiyong Yu

University of Illinois at Urbana–Champaign

Lucas Hsiung

SciFive

Mohamad El Hajj and Christopher W. Fletcher

University of Illinois at Urbana–Champaign

Abstract—It is not possible to write microarchitectural side channel-free code on commercial processors today. Even when we try, the resulting code is low performance. This article’s goal is to lay an ISA-level foundation, called a Data-Oblivious ISA (OISA) extension, to address these problems. The key idea with an OISA is to explicitly but abstractly specify security policy, so that the policy can be decoupled from the microarchitecture and even the threat model. Analogous to a traditional ISA, this enables an OISA to serve as a portable security-centric abstraction for software while enabling security-aware implementation and optimization flexibility for hardware. The article starts by giving a deep-dive in OISA principles and formal definitions underpinning OISA security. We also provide a concrete OISA built on top of RISC-V, an implementation prototype on the RISC-V BOOM microarchitecture, a formal analysis and security argument, and finally extensive performance evaluation on a range of data-oblivious benchmarks.

Digital Object Identifier 10.1109/MM.2020.2985366

Date of publication 6 April 2020; date of current version 22

May 2020.

■ **A, ARGUABLY THE,** central problem in secure computer architecture today is how to reason about security amid the sea of different microarchitectural side channel attacks. The prevailing approach to stop these attacks is to block leakage stemming from one hardware structure at a time. For example, by partitioning or randomizing the cache layout, we block (or at least aggravate) cache timing attacks. Yet, many hardware structures have been shown to leak secrets—from the cache to the branch predictors,⁵ speculative execution,⁸ port contention,¹ arithmetic unit timing,² etc. Given the many avenues to leak a secret, it is paramount to explore *holistic defenses* that provide a basis to block leakage through *all* hardware structures.

To our knowledge, this is the first proposal that provides a basis to block all traditional side channel and speculative execution attacks on commercial-class microarchitectures.

In this direction, the article proposes ISA design principles for what we call data-oblivious ISAs (OISAs). The key idea with an OISA is to explicitly but abstractly specify security policy, so that the policy can be decoupled from the microarchitecture and even the threat model. Analogous to a traditional ISA, this enables an OISA to serve as a portable security-centric abstraction for software while enabling security-aware implementation and optimization flexibility for hardware.

The OISA proposed in the article annotates what data is *confidential* and what instruction operands are *safe*. Inspired by information flow policies (in particular, the classic policy *High* \leftrightarrow *Low*), the hardware dynamically enforces that confidential data is never passed to unsafe operands, i.e., *Confidential data* \leftrightarrow *Unsafe operands*. Informally, “safe” in the article means “does not create a microarchitectural side channel as a function of the operand” (we also provide formal definitions), but other notions of safety can be retrofitted into the implementation without changing the OISA or the programs that sit on top of it.

OISAs enable high security, portability, and efficiency. Consider a simple example OISA instruction: a load with a *safe* address operand. Security-wise and portability-wise, the OISA guarantees that when the load executes, the address will not leak through microarchitectural

side channels. Depending on the microarchitecture, this may require closing side channels through the cache, translation lookaside buffer (TLB), etc. That is, security is not tied to closing a specific side channel and the programmer works with a simple, portable guarantee across microarchitectures. On the efficiency side, each microarchitecture can choose how to implement the safe load operation in whatever way maximizes performance while preserving security (e.g., by microcoding the load into simpler safe operations,² or using hardware partitioning,¹¹ or using cryptographic techniques⁹).

Safe loads are just one example. More generally, deciding which instruction operands to designate as safe opens a new, rich ISA design space which trades-off performance and hardware complexity.

Beyond formulating design principles for OISAs, the article proposes a concrete OISA extension built on top of RISC-V, implements (and open sources) that OISA extension on the BOOM out-of-order (OoO) speculative RISC-V core,³ and provides a formal analysis showing how the OISA provides a basis to achieve noninterference (“zero privacy leakage”) on an abstract OoO speculative machine. Crucially, the security analysis and principles are robust to modern attacks. Case in point, the article’s formal analysis shows how the OISA soundly defeats speculative execution attacks (such as Spectre⁸) without introducing special case reasoning.

To our knowledge, this is the first proposal that provides a basis to block all traditional side channel *and* speculative execution attacks on commercial-class microarchitectures.

MOTIVATION: SECURE AND EFFICIENT DATA-OBLIVIOUS PROGRAMMING

The OISA project came about by asking the following question: Is it possible today to write microarchitectural side channel-free programs on modern microarchitectures?

The answer is no. Consider the most conservative approach used by practitioners, called

data-oblivious programming.^{*} In a nutshell, a data-oblivious program is one whose hardware resource usage is independent of the program's inputs. To write such programs, the guidelines are to use only simple instructions, or otherwise ensure that complex instructions do not receive Confidential data as operands. For example, simple bitwise math is allowed, but memory operations/branches with Confidential data as addresses/predicates are not (out of fear of, e.g., cache-based/control flow-related side channels).

Despite being extremely conservative, the abovementioned guidelines fail in light of ISA-invisible microarchitecture-specific optimizations. For example, on one microarchitecture, a simple integer addition might be safe (e.g., implemented as a single-cycle operation whose timing is independent of its inputs) while on another it might be unsafe (e.g., implemented as a bit-serial operation that skips runs of 0s or zeros to save time). The article describes 11 like optimizations, which have been proposed in the literature, or are otherwise known to be implemented already, which break data-oblivious program security. These include data-in-use optimizations (such as data-dependent arithmetic) and data-at-rest optimizations (such as cache compression).

In particular, the article points out for the first time that speculative execution breaks data-oblivious program security, by steering execution so that Confidential data is consumed by an instruction whose execution can leak privacy. This is nontrivial to see for realistic programs, given the conservative guidelines used to write data-oblivious code. For example, consider data-oblivious decryption

```
1 for (i = 0; i < NUM_ROUNDS; i++)
2   state = OblDecryptRound
   (state, rkey [i])
3 leak(state)
```

That is, perform a fixed number of decryption rounds, where each round works on a part of the secret key (*rkey*) and incrementally updates the round state (*state*). Here, we assume that *OblDecryptRound*, the round logic, is data oblivious. *leak()* is a proxy for an instruction that

reveals its argument over a microarchitectural side channel.

This program is legal data-oblivious code: The branch outcome in each iteration is public information, the round logic is data oblivious, and only the plaintext is meant to be revealed after decryption is complete. Yet, unwanted privacy leaks because benign mispredictions can cause the round logic to exit early. In this example, an early mispredict of “not taken” allows the attacker to see state before all rounds complete, which allows it to perform cryptanalysis and recover the secret key *rkey*.

Core Issue: No Abstraction for Security

To summarize, data-oblivious programming today is insecure and slow. It is insecure because of ISA-invisible microarchitecture-specific optimizations. It is slow because, out of fear of leaking privacy, programmers are forced into using only the simplest of instructions.

The article sets out to address these issues by introducing new ISA-level abstractions for reasoning about security and enabling higher performance. A new ISA abstraction addresses the security problem by defining how instructions leak privacy *across all compliant microarchitectures*. It further enables higher performance by allowing data-oblivious programs to take advantage of *higher performance instructions*, as long as those instructions are deemed safe by the ISA, and gives microarchitects the ability to optimize those instructions subject to the ISA-prescribed security policies.

FORMAL DEFINITIONS FOR MICROARCHITECTURAL SIDE CHANNELS

To start, the article develops a security definition for microarchitectural side channel-free execution. There are two challenges. First, how to write the definition to account for any possible microarchitectural side channel. Second, how to write the definition so that it sheds insight on which instructions are “safe” from a microarchitectural side channel perspective.

To define privacy, we adopt a trace-based indistinguishability style definition inspired by the oblivious RAM (ORAM)⁷ literature. We

^{*}Data-oblivious programming goes by several other names, e.g., “constant-time programming” and “programming in the circuit abstraction,” depending on the community.

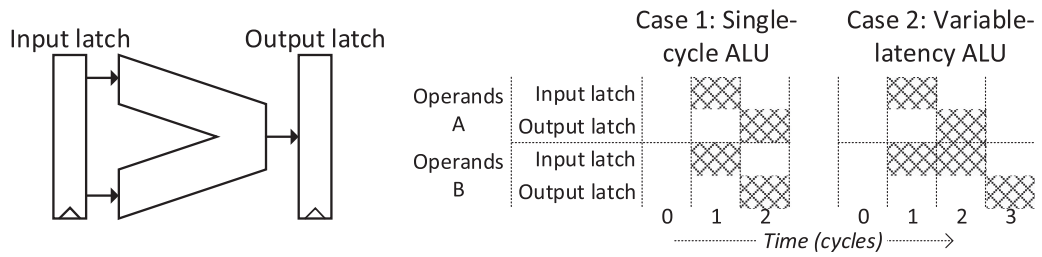


Figure 1. Changes in resource usage, as a function of data, create microarchitectural side channels. If a latch is shaded in a given clock cycle, then it means that there is (explicit) information flow from the operands to that latch in that cycle. Assume operands A and B are two sets of distinct data values, meant to induce different ALU timings.

consider a program λ , which takes public data x and confidential data y as input. That program’s execution trace, on a microarchitecture μArch , i.e., “all the atoms in the universe that are perturbed as a result of running $\lambda(x, y)$ on μArch ,” is denoted $\mu\text{Arch}(\lambda(x, y))$. The subset of this trace that the attacker can see (called the view) is denoted $\text{View}(\mu\text{Arch}(\lambda(x, y)))$. For privacy, we require that the information in the View does not depend on confidential information, i.e., that $\text{View}(\mu\text{Arch}(\lambda(x, y))) \simeq \text{View}(\mu\text{Arch}(\lambda(x, y')))$ for all confidential data y and y' . In this setting, \simeq informally means “equal, given the capabilities of any computationally bounded adversary.” For example, in ORAM schemes the view is the “memory access pattern” and ORAM seeks to make the memory access pattern independent of confidential data.

Next, we must define a view that captures any possible microarchitectural side channel that an arbitrary software-based attacker can monitor. This is nontrivial as the attacker can monitor many aspects of the program’s execution. For example, its execution time, use of the cache, arithmetic units, etc. The article makes a key observation that *all of these leakages can be modeled as confidential data-dependent changes in the program’s hardware resource usage over time*. For instance, both arithmetic units and cache sets are hardware resources and the fact that they are used at confidential data-dependent times is the crux of the attacks.

Then, the question is how to determine whether a hardware resource is currently being “used” by a program. (Note that whether a hardware resource is being “used” is independent of the logic values currently stored in that structure.) For this, we rely on an explicit gate-level

information flow abstraction similar to GLIFT.¹⁰ Figure 1 shows an example using an arithmetic logic unit (ALU) with operand-independent and then operand-dependent timing.

First assume a single-cycle ALU (see Figure 1, Case 1). Suppose the input arrives and is stored in the input latches at the rising edge of cycle 1. Using terminology from information flow, we say the input latch is tainted in cycle 1. Now, regardless of the logic values of the input, the same latches are tainted in each cycle thereafter. That is, the output latches are tainted in cycle 2, etc. Because *which latches are tainted when* is independent of the operands, we say the single-cycle ALU does not form a microarchitectural side channel.

Next assume an ALU with operand-dependent timing (see Figure 1, Case 2). For example, a multiply operation that takes one or two cycles, depending on whether an operand is 0. In this case, depending on the input, the output latch is either tainted in cycle 2 or cycle 3. Because *which latches are tainted when* is dependent on the operands, we say this ALU can form a microarchitectural side channel.

Putting everything together, we model the processor as a state machine composed of combinational logic and latches.** The subset of latches that store the Confidential input are denoted tainted at the start. Then, the View outputs a trace that indicates which subset of latches are tainted in each cycle. That is, hardware resource usage as a function of time. If the microarchitecture ensures that the View is independent of Confidential data, the microarchitecture does not leak privacy. Conversely, if the definition is not satisfied, we can pinpoint which

**W.I.o.g., we treat any state element (flip-flop, SRAM cell, etc.) as a latch.

		Input Data	
		Public	Confidential
Instruction Input Operand	Safe	Execute w/o protection	Execute w/ protection
	Unsafe	Execute w/o protection	Stop/delay instruction execution

Figure 2. Protection policies, checked before each instruction executes.

instruction caused the problem by looking at where the Views diverged. (To note, the article defines taint propagation in a nonstandard way to model only explicit information flows. This prevents taint explosion, which would render the definition not useful. Implicit information flows are modeled by quantifying over all y' .)

PRINCIPLES OF OISA DESIGN

The design principles for OISAs are twofold. First, the OISA should expose security guarantees in a microarchitecture-independent way. That is, programs written using an OISA should maintain the same security guarantees across all OISA-enabled microarchitectures. Second, OISAs should not preclude modern hardware performance optimizations, except when those optimizations have a chance to leak privacy.

To address these goals, the OISA abstraction proposed in the article has two parts. First, the OISA labels *data* to be confidential/public, to capture whether that data is a function of user secrets (i.e., the sensitive program inputs from the “Formal Definitions for Microarchitectural Side Channels” section). Second, the OISA specifies, for each instruction, whether each *instruction operand* is safe or unsafe.

Finally, compliant microarchitectures must monitor and take different actions based on what data is consumed by what instruction operands at runtime. Specifically, hardware must enforce the following rules (shown in Figure 2).

- (*Confidential* \leftrightarrow *Unsafe*) When confidential data is presented to an unsafe operand: The hardware must delay that instruction’s execution until it is nonspeculative. If the rule still applies when the instruction is nonspeculative, the program terminates with a fault (as continuing would constitute an information leak).

- (*Confidential* \rightarrow *Safe*) When Confidential data is sent to a safe operand: The hardware designer must add mechanisms to enforce the security definition given that instruction’s execution (see the “Formal Definitions for Microarchitectural Side Channels” section), for a specified view. For example, by disabling performance optimizations, scrubbing side effects and masking exceptions that occur as a function of confidential operands.
- (*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.

Despite these rules’ simplicity, they provide both security and efficiency benefits. As we will see in the “Formal Analysis,” they provide a uniform handling for both traditional- and speculative-execution-based attacks.⁸ Case in point, the only mention of speculation is a detail in the rule for *Confidential* \leftrightarrow *Unsafe*, where we say such an information flow delays the instruction’s execution until it is nonspeculative. This removes false-positive violations due to benign misspeculations. At the same time, the rules enable blocking attacks with low overhead. Case in point, the rules encode some intuitive optimizations such as “Public data does not need protection” and “it is safe to compute on confidential data with safe instructions.” The only situation where instruction execution is impeded is if confidential data is consumed by an unsafe operand.[†]

Key Idea: Abstract security policies facilitate programming simplicity, implementation flexibility, and performance optimizations. Specifying instruction operand security policy *abstractly*, i.e., as safe/unsafe, 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 security viz., the “Formal Definitions for Microarchitectural

[†]This principle directly inspired our follow-on work to block, specifically, speculative execution attacks.¹²

Base Data Oblivious ISA Extension:		
Arithmetic	operand ₁ (Safe),	operand ₂ (Safe)
Classify	operand (Safe)	
Declassify [serialized]	operand (Safe)	
Branch	predicate (Unsafe),	target (Unsafe)
Load	addr (Unsafe)	
Store	addr (Unsafe),	data (Safe)
Oblivious Memory Extension:		
Oblivious Load	addr (Safe)	
Oblivious Store	addr (Safe),	data (Safe)

Figure 3. Data-oblivious ISA policy when data is passed to an instruction operand.

Side Channels” section 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 instructions² hardware partitioning¹¹ or cryptographic techniques⁹—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.

Information flow policy and implementation. The abovementioned OISA framework describes a relatively simple security lattice (similar to $\{High, Low\}$),⁴ policy ($High \leftrightarrow Low$) and information flow propagation rule (as written, data should be marked Confidential if its value is interferent with the program’s Confidential inputs, which implies a taint algebra similar to GLIFT¹⁰). This reflects the article’s goal: to provide a comprehensive, but simple, privacy guarantee for data-oblivious programming while granting implementation flexibility to tradeoff design cost and performance. Given other use cases these parameters, and how they are concretely enforced by an implementation, can be changed for a family of OISAs, a particular OISA, or a microarchitecture that implements that OISA. For example, to support richer security lattices.⁶

DESIGN OF A CONCRETE OISA

With the principles in the “Principles for OISA Design,” section, we now propose a baseline concrete OISA that can be easily implemented on top of common existing ISAs (e.g. x86, ARM, RISC-V).

Figure 3 highlights the instructions included in the OISA, and which operands are safe/unsafe. Programmers that write data-oblivious code

will recognize this as a formalization of the guidelines used by data-oblivious programs today (see the “Core Issue: No Abstraction for Security”) section. *Arithmetic* represents all binary arithmetic operations with safe input operands. *Classify* promotes its operand from public to confidential. *Declassify* is opposite to classify, which demotes its operand from confidential to public. *Branch* performs a conditional branch, but is only allowed to specify a public destination or check a public predicate. *Load* and *store* are only permitted to take public addresses. An important detail is that since declassify has the potential to make previously protected data vulnerable, the OISA requires that the declassify instruction be verified as corresponding to program semantics. For example, on a speculative microarchitecture, this would entail delaying such an instruction until it is nonspeculative.

Extension: Memory obliviousness via safe-address loads. A common bottleneck in existing data-oblivious code is the inability to use confidential data as a load address. Therefore, we propose a new set of instructions (an oblivious-memory extension) that enable *memory-oblivious* computation.⁹ Given the OISA design principles, enabling memory-oblivious instructions is conceptually simple. Instead of emulating memory obliviousness with dummy memory operations, we designate new load/store instructions whose address operand is *safe*. This gives hardware designers the ability to build secure and efficient implementations, e.g., using partitioning¹¹ or oblivious RAM,⁷ for that specific operation.

Load instructions with safe address operands are just one example of how to accelerate secure computation with an OISA, and the article leaves extending our concrete OISA with additional safe instructions as future work. A key insight that motivates this direction is that many data-oblivious codes share common kernels (e.g., sorting) that become performance bottlenecks because the only available safe operations are simple instructions. By encapsulating these larger operations into new instructions with safe operands, a future OISA can potentially achieve constant factor or even asymptotic performance improvements. For example, a sort implemented data obliviously with simple safe instructions may cost $O(n * \log^2 n)$ operations if implemented as a

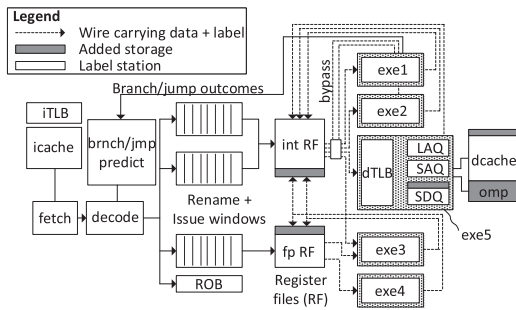


Figure 4. Microarchitectural changes needed to support the OISA from the “Design of A Concrete OISA” section (including the oblivious-memory extension, denoted “omp”). Label stations check and enforce the transition rules from Figure 2.

bitonic sort. On the other hand, if sort is specified as a single safe instruction in the OISA, an implementation based on hardware partitioning can achieve $O(n * \log n)$ time if implemented as a constant-time merge sort.

HARDWARE PROTOTYPE ON RISC-V BOOM

We prototype all hardware changes needed to support our OISA on top of the RISC-V BOOM processor (for “Berkeley OoO Machine”).³ BOOM is the most sophisticated open RISC-V processor, featuring modern performance optimizations such as speculative and OoO execution, and is similar to commercial machines that run data-oblivious code today.

Microarchitectural changes to support the OISA are shown in Figure 4. The main changes are logic at instruction issue/execute to enforce the rules from the “Principles of OISA Design” section, storage/logic to implement the oblivious-memory extension, and logic to track and denote data as confidential/public. For the latter, we implement a hardware information flow tracking mechanism similar to hardware dynamic information flow tracking, but capable of checking and updating whether data is confidential/public (the data’s label) at any stage in the pipeline.

FORMAL ANALYSIS

In parallel to our hardware prototype, we develop a formal analysis that models an abstract BOOM-class processor (OoO, speculative, superscalar), and describe how to map the abstract

BOOM to our concrete BOOM prototype. Through this model, we prove that the OISA provides a basis to satisfy strong security definitions such as those we defined in the “Formal Definitions for Microarchitectural Side Channels” section. Our security analysis is general, and applies given any implementation of several important processor structures (e.g., it models the branch predictor as an arbitrary function that takes previous branch resolutions as input).

Importantly, we are able to prove security *while* allowing high-performance hardware optimizations (e.g., OoO, speculative execution) to remain enabled in the common case and *without* ever requiring hardware flushes to structures such as the cache or branch predictors.

Security intuition. Informally, to argue security, we need to show the following.

- Each instruction’s resource usage/side-effects are independent of Confidential data.
- The sequence of instructions that are executed, i.e., the processor program counter (PC), is independent of Confidential data.[‡]

Condition (a) follows by definition by applying the rules in Figure 2 to each instruction as it executes. A more subtle point is that condition (b) also follows from applying the same rules. To see why, first consider a simple unipipelined, in-order processor with no speculation. In this case, it is clear condition (b) holds because the only instruction type from Figure 3 that changes the PC as a function of data is a branch, and the OISA requires that the branch predicate and target be Public data. What happens when we consider more advanced pipelines, e.g., with prediction and speculation? In that case, *microarchitectural state* outside of program semantics, e.g., the branch predictor state, influences the PC. To extend our security argument to these machines, we must extend what we mean by “resource usage/side-effect” to include these structures. Then, using induction, one can show that if conditions (a) and (b) hold up to fetching the i th instruction, the branch predictor state when fetching the

[‡]Similar requirements on “not tainting” the PC also govern prior work.¹⁰

$i + 1$ th instruction is independent of Confidential data, and security follows.[§]

Example: Security against speculative execution attacks. The abovementioned reasoning shows how OISAs enable security against both nonspeculative and speculative attacks. Consider the example speculative execution attack on data-oblivious decryption from the “Motivation: Secure and Efficient Data-oblivious Programming” section. This attack does not go through when using an OISA by invoking aforementioned conditions (a) and (b). That the branch predictor misspeculates and executes `leak()` prematurely is not a function of Confidential data due to condition (b). Furthermore, when `leak()` executes, it will be unconditionally stopped by the *Confidential data \leftrightarrow Unsafe operands* rule due to condition (a). Extending the analysis to attacks like Spectre,⁸ where the branch predictor is intentionally mistrained, reuses the same logic. That is, if conditions (a) and (b) hold, then the attacker’s strategy for how to mistrain the branch predictor cannot be a function of Confidential data because the program has not leaked Confidential data up to this point. In that case, intentional mistraining looks the same to the analysis as accidental misspeculation, and security follows.

EVALUATION

We evaluate OISAs 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, the OISA and hardware implementation provides an $8.8\times/1.7\times$ speedup on small/large datasets, 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 the OISA speeds up constant-time AES by $4.4\times$ and the memory oblivious ZeroTrace⁹ library by $4.6\times$ to several orders of magnitude, depending on parameters.

[§]This idea to keep the predictors a function of Public data directly inspired the mechanism to block “implicit channels” in our follow-on work.¹²

We have open-sourced our prototype design on the RISC-V BOOM processor at <https://github.com/cwfletcher/oisa>.

DISCUSSION AND FUTURE DIRECTIONS

OISAs can be extended in numerous directions, in particular as a way to compose existing hardware/software defensive mechanisms and as a novel backend for the data-oblivious stack.

Simplifying and Composing the Hardware Trusted Computing Base (TCB)

A major impediment to progress is that many hardware structures create side channels, and it is not clear whether the program is “secure” until all channels are blocked. OISAs dramatically simplify this problem, enabling a new incremental methodology for designing secure hardware and software.

In their most basic deployment, an OISA might opt to only support very basic safe instructions (e.g., bitwise operations and basic arithmetic). Such an OISA can likely be implemented with minimal changes to modern processors and already improves the state of security today. For example, by increasing our confidence that conservatively-written codes such as constant-time codes are really “constant time.”

Beyond this basic deployment, however, OISAs provide a way for computer architects to plug-and-play their high-performance “point” defenses in a compositional way. For example, a safe load can be implemented by previously proposed partitioned or randomized cache architectures.¹¹ Importantly, architects need only worry about how to implement the safe load. The generic OISA rules, e.g., *Confidential data \leftrightarrow Unsafe operands*, take care of the rest.

Composing With the Data-Oblivious Stack

Beyond writing side channel-free code for today’s hardware, there is a rich literature in the applied cryptography community—ranging from algorithm/data structure design to language and compiler support—for performing secure multiparty and encrypted computation.

We make a key observation that the underlying programming abstraction assumed for those works is the same abstraction provided by an OISA. For example, a homomorphic encryption operation is akin to a safe instruction, just using a different implementation suitable for a different threat model. This enables a new, large-scale research agenda to port insights and advances made in the applied cryptography community to/from the microarchitectural side channel community. For example, we can enable high-level programming abstractions for writing OISA-secure code by adding a new OISA backend to existing data-oblivious compiler frameworks. At the same time, the notion of safe instructions provides a new theory to explore in applied cryptography. In particular, algorithm design in encrypted computation assumes only extremely simple safe operations (e.g., bit add or multiply). With an OISA, however, we can choose which operations support safe operands, and co-design algorithms with this in mind to improve performance.

REFERENCES

1. A. C. Aldaya, B. B. Brumley, S. U. Hassan, C. P. García, and N. Tuveri, "Port contention for fun and profit," in *Proc. IEEE Symp. Secur. Privacy*, 2019, pp. 870–887.
2. M. Andryscio, D. Kohlbrenner, K. Mowery, R. Jhala, S. Lerner, and H. Shacham, "On subnormal floating point and abnormal timing," in *Proc. IEEE Symp. Secur. Privacy*, 2015, pp. 623–639.
3. C. Celio, P.-F. Chiu, B. Nikolic, D. A. Patterson, and K. Asanović, "BOOM v2: An open-source out-of-order RISC-V core," Tech. Rep. UCB/EECS-2017-157, EECS Dept., Univ. California, Berkeley, CA, USA, 2017.
4. D. E. Denning, "A lattice model of secure information flow," *Commun. ACM*, vol. 19, pp. 236–243, May 1976.
5. D. Evtvushkin, R. Riley, N. C. Abu-Ghazaleh, ECE, and D. Ponomarev, "BranchScope: A new side-channel attack on directional branch predictor," in *Proc. 23rd Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, 2018, pp. 693–707.
6. A. Ferraiuolo, M. Zhao, A. C. Myers, and G. E. Suh, "HyperFlow: A processor architecture for nonmalleable, timing-safe information flow security," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2018, pp. 1583–1600.
7. O. Goldreich and R. Ostrovsky, "Software protection and simulation on oblivious rams," *J. ACM*, vol. 43, pp. 431–473, May 1996.
8. P. Kocher et al., "Spectre attacks: Exploiting speculative execution," in *Proc. IEEE Symp. Secur. Privacy*, 2019, pp. 1–19.
9. S. Sasy, S. Gorbunov, and C. W. Fletcher, "ZeroTrace: Oblivious memory primitives from Intel SGX," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, San Diego, CA, USA, Feb. 18–21, 2018. Available: <http://dx.doi.org/10.14722/ndss.2018.23239>
10. M. Tiwari, H. M. Wassel, B. Mazloom, S. Mysore, F. T. Chong, and T. Sherwood, "Complete information flow tracking from the gates up," in *Proc. 14th Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, 2009, pp. 109–120.
11. Z. Wang and R. B. Lee, "New cache designs for thwarting software cache-based side channel attacks," in *Proc. 34th Annu. Int. Symp. Comput. Archit.*, 2007, pp. 494–505.
12. J. Yu, M. Yan, A. Khyzha, A. Morrison, J. Torrellas, and C. Fletcher, "Speculative taint tracking (STT): A comprehensive protection for speculatively accessed data," in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2019, pp. 954–968.

Jiyong Yu is currently working toward the Ph.D. degree at the University of Illinois at Urbana–Champaign. His research interests are in processor security. Contact him at jiyongy2@illinois.edu.

Lucas Hsiung received the B.S. degree from the University of Illinois at Urbana–Champaign in 2019 and now works as a Security Verification Engineer at SciFive. Contact him at lucas.hsiung@sifive.com.

Mohamad El Hajj is currently working toward the M.S. degree at the University of Illinois at Urbana–Champaign with research interests in hardware security. Contact him at melhajj2@illinois.edu.

Christopher W. Fletcher is an Assistant Professor in computer science at the University of Illinois at Urbana–Champaign. He has interests ranging from computer architecture to security to high-performance computing (ranging from theory to practice, algorithm to software to hardware). Fletcher received the Ph.D. degree from Massachusetts Institute of Technology in 2016. Contact him at cwfletch@illinois.edu.