

# Designing Secure TLBs

**Shuwen Deng**

Tsinghua University

Beijing 100190, China

**Jakub Szefer**

Yale University

New Haven, CT 06510 USA

**Wenjie Xiong**

Virginia Polytechnic Institute and State University

Blacksburg, VA 24061 USA

*Editor's notes:*

This article reveals timing-based side-channel and covert-channel attacks from the translation look-aside buffers (TLBs) and discusses how to design secure TLBs.

—Gang Qu, University of Maryland, USA

timing-based vulnerabilities exhaustively. Rather than modeling software attacks, the three-step approach analyzes all possible victim or attacker behaviors that

**RESEARCH ON TIMING-BASED ATTACKS** (and defenses) in processors has a long history. To date, researchers have focused mainly on the memory subsystem when showing the different timing-based attacks, and have, for example, demonstrated a plethora of timing-based channels in caches [6]. All the attacks have shown the possibility to extract sensitive information via timing-based channels, and often the focus is on extracting cryptographic keys.

Timing-based channels in translation look-aside buffers (TLBs), which is the focus of this work, are distinct from caches in that they are triggered by memory translation requests, not by direct access to data. They also have a different granularity (pages versus cache lines for data or instruction caches), and, in commercial processors, TLBs have more complicated logic, compared to caches, due to support for various memory page sizes. Furthermore, defending cache attacks does not protect against TLB attacks [4]. Moreover, there has not been a systematic security analysis of the TLB vulnerabilities, nor concrete proposals for secure TLB design. This article provides both.

This work starts by providing a novel three-step modeling approach to enumerate all possible TLB

affect the TLB state. In total, 24 possible vulnerabilities were found, of which only eight map to existing attacks [4], [5]. We believe that the other 16 are new attack types not previously considered. Based on the three-step model, microsecurity benchmarks are then semiautomatically generated.

Armed with the three-step model and the security benchmarks, the security of different typical configurations of TLBs is tested using Rocket core implementation of the reduced instruction set computer (RISC)-V processor architecture. *Standard TLBs*, that is, fully-associative (FA) or set-associative (SA) TLBs, which include process identifiers (IDs), for example, address space identifier (ASID) in RISC-V architecture, are shown to be vulnerable to many of the attacks. Consequently, this work presents new defenses. Especially, we present two new secure TLB designs in hardware: a static-partition (SP) TLB and an RF TLB, the latter of which is more complex but can defend against all attacks. These are the first hardware defenses to TLB attacks. To help understand the impact of the new secure TLBs on the system performance, a RISC-V Rocket core-based processor with the new secure TLBs is synthesized on the Zynq ZC706 and ZedBoard FPGAs.

## Background

This section reviews existing work on caches (most closely related to TLBs) and the few existing works on TLBs.

*Digital Object Identifier 10.1109/MDAT.2023.3287938*

*Date of publication: 27 June 2023; date of current version:*

*21 February 2024.*

## Timing-based attacks and caches

In modern processor caches, there are timing differences between cache hits (fast) and cache misses (slow), and these variations in timing have been exploited to leak sensitive information. Especially, a large number of different cache timing-based side-channel and covert-channel attacks have been presented in the literature [6]. And, many secure hardware cache designs aim to prevent these different attacks [9]. However, even if the cache-based attacks are mitigated, TLB-based attacks are the next attack vector that malicious attackers might use—and hence are the focus of this work.

## Timing-based channels in TLBs

Compared with caches, there are two published TLB-based timing attacks.<sup>1</sup> TLBleed attack [4] uses timing-based channels combined with machine learning to create an attack that can leak bits of secret keys from the Rivest–Shamir–Adleman encryption (RSA) algorithm (they also show an attack for the Edwards-curve digital signature algorithm (EdDSA) algorithm). They leverage the prime + probe attack strategy previously applied in processor caches.

Before TLBleed, the double page fault attack [5] leverages the cache collision attack strategy previously applied in processor caches. It requires the victim to access some kernel memory pages twice and uses the fact that access to previously allocated kernel virtual pages will bring in TLB entries, even if a page fault is generated and accesses permission checks failed. The timing of the second access thus reveals information on whether an inherent TLB hit happened.

Beyond these individual attacks, there are neither exhaustive categorizations nor models of possible TLB timing-based attacks—as are proposed in this work.

## Existing approaches to securing TLBs

Currently, we are only aware of five approaches (mostly software-based) that can help mitigate some TLB attacks, but are not as effective as our hardware-secure TLBs.

First, today's Linux system makes use of virtual addresses and process IDs, for example, ASID on RISC-V, to identify different processes in the SA

TLBs to do the partition. Second, in the Sanctum [1] secure processor design, the per-core SA TLBs are flushed by a *security monitor* software whenever a core switches between the enclave and nonenclave codes. Third, Intel SGX also flushes SA TLBs during switching between the enclave and non-enclave codes. Fourth, the InvisiSpec [11] work proposes to prevent observable changes to data translation lookaside buffers (D-TLBs) only for speculative attacks. Fifth, some processors employ FA TLBs, which, by design, do not have different TLB sets (there is only one set).

Unlike all the existing work, this work presents two new hardware secure TLB designs, including the RF TLB, which can prevent all types of timing-based attacks according to the three-step model, and has about the same performance as an SA TLB.

## Framework

Our article's key idea focused on the new attack vector in modern processors: the timing-based channel attacks due to the TLBs. We provide a systematic approach to analyze the full set of vulnerabilities and provide corresponding hardware defenses to mitigate them.

## Modeling TLB timing-based vulnerabilities

We first presented a novel three-step modeling approach that was used to exhaustively enumerate all possible TLB timing-based vulnerabilities.

## Threat model and assumptions

A TLB timing-based attack involves an attacker and a victim. In many cases, they are executing on the same processor core, a set of cores, or a set of hyperthreads that share the same physical TLB, but this is not required for all types of attacks. In this article, we use  $A$  and  $V$  to denote the attacker and the victim with different process IDs. For the attacks where the attacker and the victim are in the same address space, the attacker can trigger some known address memory operation as if it were the victim, for example, states  $V_a$  and  $V_{a\text{alias}}$  in Table 1 can be actually attackers.

We assume, in hardware, all memory operations are identified by the virtual memory address,  $vaddr$  (including null address in the case of certain TLB flush-related operations) and the process ID (including null process ID in the case of certain TLB flush-related operations), for example, ASID in RISC-V.

<sup>1</sup>The Leaky Cauldron [8] attack is also related to TLB and targets Intel SGX. However, it does not depend on hits and misses in the TLB; instead, it relies on the assumption that the attacker can evict the enclave entries in the TLB, so an enclave's memory access will trigger a page table walk, and the malicious OS can get the page access pattern trace.

The victim is assumed to have some security-critical memory range,  $x$ , within which the access pattern depends on the secret the attacker wants to learn. An example of a security-critical region is the set of page entries accessed during the execution of the RSA functions of *libgcrypt*, where the value of the key bit (either 0 or 1) determines which specific memory pages are accessed. The timing of the accesses to the security-critical memory range is affected by the timing of TLB-related operations, and it can reveal information such as cryptographic keys.

The attacker is assumed to know the victim software, for example, what implementation of a cryptographic algorithm it uses, but not the secret cryptographic keys. He or she is assumed to know the size, *ssize*, and the location, *sbase* (in virtual memory) of the security-critical memory range  $x$ . The attacker can measure the timing of its own memory operations or the operations of the victim, but cannot access the actual sensitive data being processed by the victim.

### Introduction of the three-step model

One observation we make is that all existing TLB timing-based attacks take three steps. In Step 1, a memory operation is performed, placing the TLB block (also called TLB slot or TLB entry) in a known initial state (e.g., a new translation is put into the block or block is invalidated).

Then, in Step 2, a second memory operation alters the state of the TLB block from the initial state. Finally, in Step 3, a final memory operation is performed, and the timing of the final operation reveals some information about the relationship among the addresses from Step 1, Step 2, and Step 3. Attacks with more than three steps can be reduced to a three-step attack (details are in our article [2]). Table 1 lists all the 10 possible states of the TLB block for each step of our three-step model. Each step in the model represents a state of a TLB block.

### Derivation of all TLB vulnerabilities

Based on the states possible in each step, there are in total  $10 * 10 * 10 = 1,000$  combinations of possible three steps. We developed an algorithm that can process the list of all three steps and eliminates ones that cannot lead to an attack based on a list of derived rules (details mentioned in our article [2]).

After applying the script that implements our simplification algorithm, 34 three-step access patterns remain as candidates for possible timing-based TLB

**Table 1. Ten possible states for a single TLB block in our three-step vulnerability modeling procedure.**

States	Description
$V_u$	The TLB block contains translation for a memory address $u$ , translation which is placed in the TLB block due to a memory access by the victim. Attacker does not know $u$ , but $u$ is from a range $x$ of memory locations, range which is known to the attacker. The address $u$ may have same page index as $A_a$ or $V_a$ and thus conflict with them in the TLB block. The goal of the attacker is to learn the page address or index of $V_u$ .
$A_a$ or $V_a$	The TLB block contains translation for a memory address $a$ . The translation is placed in the TLB block due to a memory access by the attacker, $A_a$ , or the victim, $V_a$ . The attacker knows the address $a$ , independent of whether the access was by the victim or the attacker themselves. The address $a$ is from within the range of sensitive locations $x$ . The address $a$ may or may not be the same as the address $u$ .
$A_{a\text{alias}}$ or $V_{a\text{alias}}$	The TLB block contains translation for a memory address $a^{\text{alias}}$ . The translation is placed in the TLB block due to a memory access by the attacker, $A_{a\text{alias}}$ , or the victim, $V_{a\text{alias}}$ . The address $a^{\text{alias}}$ is within the range $x$ . It is not the same as $a$ , but it has same page index and can map to the same TLB block, i.e. it “aliases” to the same block.
$A_{inv}$ or $V_{inv}$	The TLB block previously containing translation for a memory address is now invalid. The translation is “removed” from the TLB block by the attacker $A_{inv}$ or the victim $V_{inv}$ as the result of TLB block being invalidated, e.g. due to synchronization updates to in-memory memory-management data structures or due to context switch between processes which causes OS to flush per-core TLB entries.
$A_d$ or $V_d$	The TLB block contains translation for a memory address $d$ . The translation is placed in the TLB block due to a memory access by the attacker, $A_d$ , or the victim, $V_d$ . The address $d$ is not within the range $x$ .
★	Any data, or no data, can be in the TLB block. The attacker has no knowledge of page address in this TLB block.

attacks. These 34 access patterns are further manually reduced to a list of 24 types of timing-based TLB vulnerabilities, listed in Table 2. Due to space limitations, details on why the 10 patterns cannot form vulnerabilities are not included in the article.

To summarize all the vulnerability types, Table 2 shows the list of all 24 vulnerability types, along with more coarse-grained attack strategies, which cover one or more vulnerability types. The list of vulnerability types can be further collected into four simple macrotypes: 1) internal interference miss-based (IM); 2) internal interference hit-based (IH); 3) external interference miss-based (EM); and 4) external interference hit-based (EH). Most of the vulnerability types have not been explored before, except some mapping to existing double-page fault attacks [5] and the TLBleed attack [4].

**Table 2. Timing-based TLB vulnerabilities.** The *Attack strategy* column gives our common name for each set of one or more specific vulnerabilities that would be exploited in an attack in a similar manner (many of the names are borrowed from cache timing-based attacks in literature). The *Vulnerability type* column gives the three steps that define each vulnerability. For *Step 3, fast* indicates that a TLB hit must be observed, while *slow* indicates that a TLB miss must be observed. The *Macro type* column proposes the categorization of the vulnerability belongs to. “E” is for external interference vulnerabilities. “I” is for internal interference vulnerabilities. “M” is for miss-based vulnerabilities. “H” is for hit-based vulnerabilities. The *Attack* column shows if a type of vulnerability has been previously presented in the literature.

Attack Strategy	Vulnerability Type			Macro Type	Attack
	Step 1	Step 2	Step 3		
TLB Internal Collision	$A_{inv}$	$V_u$	$V_a$ (fast)	IH	(1)
	$V_{inv}$	$V_u$	$V_a$ (fast)	IH	(1)
	$A_d$	$V_u$	$V_a$ (fast)	IH	(1)
	$V_d$	$V_u$	$V_a$ (fast)	IH	(1)
	$A_{qalias}$	$V_u$	$V_a$ (fast)	IH	(1)
	$V_{qalias}$	$V_u$	$V_a$ (fast)	IH	(1)
TLB Flush + Reload	$A_{inv}$	$V_u$	$A_a$ (fast)	EH	<b>new</b>
	$V_{inv}$	$V_u$	$A_a$ (fast)	EH	<b>new</b>
	$A_d$	$V_u$	$A_a$ (fast)	EH	<b>new</b>
	$V_d$	$V_u$	$A_a$ (fast)	EH	<b>new</b>
	$A_{qalias}$	$V_u$	$A_a$ (fast)	EH	<b>new</b>
	$V_{qalias}$	$V_u$	$A_a$ (fast)	EH	<b>new</b>
TLB Evict + Time	$V_u$	$A_d$	$V_u$ (slow)	EM	<b>new</b>
	$V_u$	$A_a$	$V_u$ (slow)	EM	<b>new</b>
TLB Prime + Probe	$A_d$	$V_u$	$A_d$ (slow)	EM	(2)
	$A_a$	$V_u$	$A_a$ (slow)	EM	(2)
TLB version of Bernstein’s Attack	$V_u$	$V_a$	$V_u$ (slow)	IM	<b>new</b>
	$V_u$	$V_d$	$V_u$ (slow)	IM	<b>new</b>
	$V_d$	$V_u$	$V_d$ (slow)	IM	<b>new</b>
	$V_a$	$V_u$	$V_a$ (slow)	IM	<b>new</b>
TLB Evict + Probe	$V_d$	$V_u$	$A_d$ (slow)	EM	<b>new</b>
	$V_a$	$V_u$	$A_a$ (slow)	EM	<b>new</b>
TLB Prime + Time	$A_d$	$V_u$	$V_d$ (slow)	IM	<b>new</b>
	$A_a$	$V_u$	$V_a$ (slow)	IM	<b>new</b>

(1) Double Page Fault attack [5].

(2) TLBleed attack [4].

## Key contribution

The *key contribution* of this article was to show the first systematic modeling approach that can be used to reason about all timing-based attacks on TLBs. Our work developed a novel model of the attacker and victim behaviors in relation to the TLB states. Rather than modeling software attacks, the three-step approach analyzed all possible victim or attacker behaviors that affected the TLB states. All possible combinations of the attacker and victim behaviors were evaluated and systematically reduced to only three-step behaviors that can result

in timing-based attacks. In total, 24 possible vulnerabilities were found, including 16 new attack types not previously considered.

## Microsecurity benchmarks

Building on the three-step model, our article then showed how to automatically generate microsecurity benchmarks to test TLBs to check if they are vulnerable to each of the attack types. To generate the microsecurity benchmarks, we leverage a Python script that follows a three-step template to generate assembly code of all the types of vulnerabilities shown in Table 2.

We use channel capacity [3] to quantify the amount of information about the secret address translation that the attack gains from a specific timing-based attack.

## Secure TLB designs

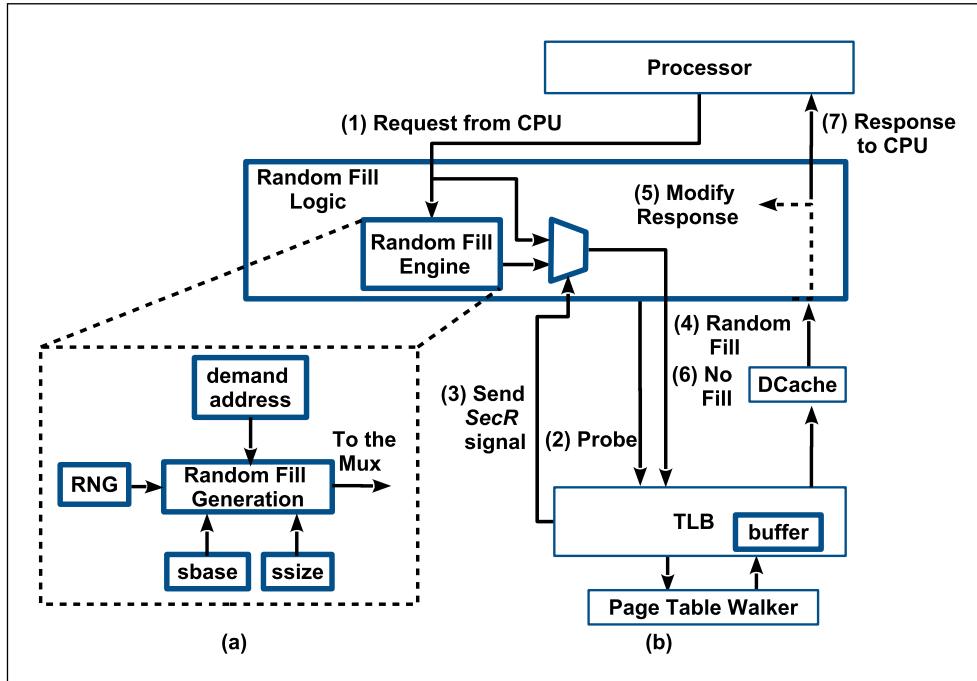
After showing the insecurity of standard TLBs, our work also proposed the first hardware defenses for TLB attacks: the new SP TLB and the new RF TLB and realize them in a Rocket core implementation of the RISC-V processor. Especially, RF TLB was more complex in logic but could defend against all of the attacks compared to SP TLB.

The first type of TLB, SP TLB, is an SA TLB where certain ways are assigned to a victim process and other ways are assigned to all remaining processes, which by default are assumed to be potential attacker processes. The process ID, for example, ASID in RISC-V, is used to differentiate the victim and the attacker. The number of ways assigned to each is set at design time but could be further extended to be dynamic at run time.

To protect all the vulnerabilities, we propose RFTLB, which can decorrelate the requested memory access from actual TLB entries that are brought into the TLB, making the attacker’s observations nondeterministic. For TLB hits, the behavior is the same as the SA TLB. For TLB misses, depending on the memory address region, a random address translation will be fetched into the TLB (“random fill”), while the originally requested address is directly sent back to the central processing unit (CPU) without filling the TLB (“no fill”).

The RF TLB also introduces the *Sec* bit that is used to identify certain memory translation entries are belonging to secure data.

RF TLB block diagram is shown in Figure 1b. All the bold lines and blocks are the added hardware and logic extension. In the TLB array, an extra field (a



**Figure 1. (a) RFE. (b) RF TLB block diagram.**

secure bit  $Sec$ , either 0 or 1) is added to each of the TLB entries to indicate whether it contains an address translation within the secure region. In addition, the existing process ID field (e.g., ASID in RISC-V) in each TLB entry is used to differentiate the victim and the attacker process. By default, we set specific process ID 1 for the victim program and all other ASIDs to be attackers.

An extra set of registers is added to store the process ID of the victim process and the start address  $sbase$  and the size  $ssize$  of the secure region (the base and size are defined in terms of pages, usually 4KiB). The registers can be managed by a trusted operating system (OS) to change the victim process ID and secure regions when different victim programs need protection.

An extra *buffer* is added which stores the equivalent of one TLB entry. It is used as temporary storage for translation data that is returned to the CPU, but which should not be placed in the TLB. It will be cleaned up after the address is returned.

The RF engine (RFE), shown in Figure 1a, is used to generate addresses that should be used for TLB updates.<sup>2</sup> In Figure 1b (1 and 2), the “no fill” *fill\_type* will first be sent to TLB. On a TLB miss, the

TLB will *probe* the page address without filling TLB entries to see if the chosen entry has a valid secure page address translation. Then, (3) the  $Sec_R$  bit is set and sent back. Next, (4) if it is a request to the secure region or the  $Sec_R$  bit is one, an RF request will be triggered. If the original request is in the secure region, a random virtual page address is derived from RFE within the secure region [ $sbase$ ,  $sbase + ssize$ ], and a translation will be put into the TLB entry. If the original request comes from the nonsecure region, most of the higher bits of the requested address remain the same while the bits that correspond to the TLB set index<sup>3</sup> will be randomized to make the eviction indeterministic. Next, (5) the RF logic will modify the response and prevent the RF result from being sent to the processor. Then, (6) the original page address is finally requested, and “no fill” *fill\_type* will be sent to the TLB to obtain the translation. Finally, (7) this address will be stored in the *buffer*, without modifying TLB entries, and be sent back to the processor.

## Security evaluation

The three-step model and the security benchmarks were used to analyze the security of the

<sup>2</sup>We assume the OS has pregenerated page table entries that may correspond to the random virtual address generated by the RFE, which may not be actually used by the original program, to prevent OS- or software-based timing attacks due to page faults when a page entry for a random address is looked up by the TLB.

<sup>3</sup>The TLB set index to be randomized has bit size  $S_n = \log_2[\min(ssize, nsets)]$ , where  $nsets$  is the number of sets in TLB. A random set index will be generated within the region  $[sbase[S_n - 1, 0], sbase[S_n - 1, 0] + \min(ssize, nsets)]$  for RF.

new designs in simulation. Based on the analysis, we showed that the proposed secure TLBs could defend not only against the previously publicized attacks, but also against other new timing-based attacks in TLBs found using our new three-step modeling approach.

As demonstrated in Table 3, while evaluating the security of the standard and secure TLBs using the microsecurity benchmarks running on RISC-V simulation, we showed that results matched with theoretical mutual information calculation. For the specific security effectiveness of different TLBs, our security evaluation showed that standard SA TLBs could defend 10 types of external hit-based related attacks. For the new hardware defenses, our evaluation showed that SP TLB was able to further prevent four more external miss-based vulnerabilities, in total defending 14 out of 24 vulnerabilities. Meanwhile, the RF TLB was able to prevent all of the 24 possible timing-based vulnerabilities in TLBs. The proposed secure TLBs could defend not only against the previously publicized attacks, but also against other possible timing-based attacks in TLBs found using our new three-step modeling approach.

### Performance evaluation

Finally, we were able to maintain the performance overhead to a small number while protecting the full security. We tested performance by synthesizing the hardware on FPGAs and running RSA decryption tests alongside SPEC 2006 benchmarks under Linux on the FPGAs. To help understand the impact of the new secure TLBs on the system performance, a RISC-V Rocket Core-based processor with the new secure TLBs was synthesized on the Zynq ZC706 and

ZedBoard FPGAs. This allowed for running security software alongside SPEC 2006 benchmarks and a full Linux system. Based on our evaluation, for example, the SP TLB had 3x misses per kiloinstructions (MPKI) compared to the standard SA TLB, while the RF TLB had 9% more MPKI than the standard TLB. For the RF TLB, the hardware cost of the defenses was about 8% more logic.

### Discussion

Research on timing-based attacks (and defenses) in processors has a long history. Most mitigations of timing-based attacks in the memory subsystem have focused on the design of secure caches. Meanwhile, our work focused on preventing timing-based attacks due to TLBs and presented the first hardware defenses for TLBs. In our securing TLB work, both the systematic approach of finding all possible TLB timing side-channel vulnerabilities and the hardware defenses of TLBs regarding timing side-channel vulnerabilities were presented as promising methodologies and solutions, respectively, to examine and tackle TLB timing side channels. Especially, our hardware defenses were tested in real hardware (FPGA), while much of the existing work is evaluated in simulation only. Therefore, our work can give more confidence and possibly be more easily adapted to be used in industrial and commercial products.

The methodology and solution presented in our work have already influenced other research projects on timing side channels and TLBs. Our hardware defenses for TLBs have been shown to be effective even with newly developed attacks. For example, PThammer [12] developed new attacks focusing on eviction-based cache and TLB attacks. It exploited the fact that cache and TLB were shared between sensitive data and public data. It also confirmed our hardware defenses are effective and can readily mitigate PThammer by partitioning or randomizing the TLB. Without the use of our secure TLBs, the PThammer could easily compromise TLBs and lead to information leaks in real systems.

Meanwhile, our TLB work has also been shown in work [10] to be able to prevent transient execution attacks. The recent Spectre and Meltdown attacks have shown that timing-based channels are more dangerous than previously thought. Whether by

**Table 3. Comparison of SA TLB, SP TLB, and RF TLB simulation and theoretical results.  $C^*$  and  $C$  represent mutual information based on simulation and theoretical calculation, respectively. Bold  $C^*$  and  $C$  are the ones with a value of 0 or about 0, indicating that this TLB can prevent the corresponding vulnerability. Small numbers are rounded up. Some vulnerability types are not shown to save space.**

Attack Category	Vulnerability Type	SA TLB		SP TLB		RF TLB	
		$C^*$	$C$	$C^*$	$C$	$C^*$	$C$
TLB Evict+Probe	$V_d \rightsquigarrow V_u \rightsquigarrow A_d$ (slow)	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>
TLB Prime+Time	$A_d \rightsquigarrow V_u \rightsquigarrow V_d$ (slow)	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>
TLB Flush+ Reload	$A_d \rightsquigarrow V_u \rightsquigarrow A_d$ (fast)	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>
TLB Prime+Probe	$A_d \rightsquigarrow V_u \rightsquigarrow A_d$ (slow)	0.99	1	<b>0.02</b>	<b>0</b>	<b>0.01</b>	<b>0</b>
TLB Evict+Time	$V_u \rightsquigarrow A_d \rightsquigarrow V_d$ (slow)	1	1	<b>0.03</b>	<b>0</b>	<b>0</b>	<b>0</b>
TLB Internal Collision	$A_d \rightsquigarrow V_u \rightsquigarrow V_d$ (fast)	1	1	0.98	1	<b>0.01</b>	<b>0</b>
TLB Bernstein's Attack	$V_u \rightsquigarrow V_a \rightsquigarrow V_u$ (slow)	0.99	1	0.99	1	<b>0.01</b>	<b>0</b>

themselves or in combination with speculative execution such as the Spectre and Meltdown attacks, the timing-based channels in microarchitecture pose threats to system security and should be mitigated. Our work proposed hardware defenses to mitigate timing-based channels in TLBs, and therefore, it is effective even for the new transient execution attacks which utilize TLB microarchitecture timing-based channels.

Research based on our work has clearly acknowledged the effectiveness of our modeling approach and the hardware defense methods for the TLB timing side-channel vulnerabilities, for example, [7]. We believe that our work on securing the TLBs from the timing-based channels has already created an impact and will continue to serve as a catalyst for higher security, especially in future computer architectures.

**THIS ARTICLE PROPOSED** a novel three-step modeling approach that exhaustively enumerates all possible TLB timing-based vulnerabilities. It showed how to automatically generate microsecurity benchmarks that test for TLB vulnerabilities. It gave details of two new hardware-secure TLB designs: an SP TLB and an RF TLB. The simulations confirmed the theoretical channel capacity calculations and full system performance on FPGA showed that the new secure TLBs are as good as regular TLBs, while protecting against various attacks. The proposed secure TLBs can defend not only against the previously publicized attacks, but also against other possible timing-based attacks in TLBs found using our new three-step modeling approach. ■

## References

- [1] V. Costan, I. A. Lebedev, and S. Devadas, “Sanctum: Minimal hardware extensions for strong software isolation,” in *Proc. USENIX Secur. Symp.*, 2016, pp. 857–874.
- [2] D. Shuwen, W. Xiong, and J. Szefer, “Secure TLBs,” in *Proc. 46th Int. Symp. Comput. Archit.*, 2019, pp. 346–359.
- [3] A. J. Goldsmith and P. P. Varaiya, “Capacity of fading channels with channel side information,” *IEEE Trans. Inf. Theory*, vol. 43, no. 6, pp. 1986–1992, Nov. 1997.
- [4] B. Gras, K. Razavi, H. Bos, and C. Giuffrida, “Translation leak-aside buffer: Defeating cache side-channel protections with TLB attacks,” in *Proc. USENIX Secur. Symp. (USENIX)*, 2018, pp. 955–972.

- [5] R. Hund, C. Willems, and T. Holz, “Practical timing side channel attacks against kernel space ASLR,” in *Proc. IEEE Symp. Secur. Privacy*, May 2013, pp. 191–205.
- [6] C. Percival, “Cache missing for fun and profit,” 2005. [Online]. Available: <http://www.daemonology.net/papers/htt.pdf>
- [7] D. Skarlatos, Z. N. Zhao, R. Paccagnella, C. W. Fletcher, and J. Torrellas, “Jamais vu: Thwarting microarchitectural replay attacks,” in *Proc. 26th ACM Int. Conf. Architectural Support Program. Lang. Operating Syst.*, 2021, pp. 1061–1076.
- [8] 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,” in *Proc. Conf. Comput. Commun. Secur.*, 2017, pp. 2421–2434.
- [9] Z. Wang and R. B. Lee, “New cache designs for thwarting software cache-based side channel attacks,” in *Proc. ACM SIGARCH Comput. Archit. News*, Vol. 35, 2007, pp. 494–505.
- [10] W. Xiong and J. Szefer, “Survey of transient execution attacks,” 2020, arXiv:2005.13435.
- [11] M. Yan, J. Choi, D. Skarlatos, A. Morrison, C. W. Fletcher, and J. Torrellas, “InvisiSpec: Making speculative execution invisible in the cache hierarchy,” in *Proc. Int. Symp. Microarchit. (MICRO)*, 2018, pp. 428–441.
- [12] Z. Zhang, Y. Cheng, D. Liu, S. Nepal, Z. Wang, and Y. Yarom, “PThammer: Cross-user-kernel-boundary rowhammer through implicit accesses,” 2020, arXiv:2007.08707.

**Shuwen Deng** is an assistant professor with the Department of Electronic Engineering, Tsinghua University, Beijing 100190, China. She works on computer architecture, hardware, and system security. Deng has a BS in microelectronics from Shanghai Jiao Tong University, Shanghai, China, and a PhD from the Department of Electrical Engineering, Yale University, New Haven, CT, USA. She is a Member of IEEE.

**Wenjie Xiong** is an assistant professor with the Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA 24061 USA. Her research interests comprise physically unclonable functions and side-channel attacks and defenses. Xiong has a BS in microelectronics and psychology from Peking University, Beijing, China, and a PhD from the Department of Electrical

Engineering, Yale University, New Haven, CT, USA.  
She is a Member of IEEE.

**Jakub Szefer** is an associate professor with the Electrical Engineering Department, Yale University, New Haven, CT 06510 USA, where he leads the Computer Architecture and Security Laboratory (CASLAB). His research interests are at the intersection of computer architecture, hardware security, and FPGA security. Szefer has a BS with highest honors in electrical and computer engineering from the University of Illinois at Urbana–Champaign, Champaign, IL, USA, and an MA and a PhD in electrical engineering

from Princeton University, Princeton, NJ, USA, where he worked with Prof. Ruby B. Lee on secure hardware architectures. He is a Senior Member of IEEE.

■ Direct questions and comments about this article to Shuwen Deng, Tsinghua University, Beijing 100190, China; shuwend@tsinghua.edu.cn.