

STBPU: A Reasonably Secure Branch Prediction Unit

Tao Zhang
William & Mary
tzhang06@wm.edu

Timothy Lesch
William & Mary
tjlesch@wm.edu

Kenneth Koltermann
William & Mary
khkoltermann@wm.edu

Dmitry Evtyushkin
William & Mary
devtyushkin@wm.edu

Abstract—Modern processors have suffered a deluge of threats exploiting branch instruction collisions inside the branch prediction unit (BPU), from eavesdropping on secret-related branch operations to triggering malicious speculative executions. Protecting branch predictors tends to be challenging from both security and performance perspectives. For example, partitioning or flushing BPU can stop certain collision-based exploits but only to a limited extent. Meanwhile, such mitigations negatively affect branch prediction accuracy and further CPU performance. This paper proposes Secret Token Branch Prediction Unit (STBPU), a secure BPU design to defend against collision-based transient execution attacks and BPU side channels while incurring minimal performance overhead. STBPU resolves the challenges above by customizing data representation inside BPU for each software entity requiring isolation. In addition, to prevent an attacker from using brute force techniques to trigger malicious branch instruction collisions, STBPU actively monitors the prediction-related events and preemptively changes BPU data representation.

I. INTRODUCTION

Although hardware attacks such as microarchitectural side channels [1, 3, 6, 22, 27, 41, 54, 55], covert channels [20, 25, 46, 52], and power analysis [2, 33, 45, 48, 53] attacks have been known for a long time, only recently did researchers demonstrate the true power of microarchitectural attacks with newly discovered transient execution attacks, such as Meltdown [39, 70] and Spectre [15, 31, 32, 34, 44, 66]. These attacks are based on speculative (transient) execution, a performance optimization technique present in nearly all of today’s processors. While this technique improves CPU performance, with a carefully crafted exploit, it completely undermines memory protection, giving unauthorized users the ability to read arbitrary memory [32, 39], bypass crucial protections [31] or even perform arbitrary computations [18].

Microarchitectural attacks are possible because performance optimizations such as caches, prefetchers, and various predictors were not traditionally designed with security in mind. For example, data structures used to implement these mechanisms are often shared, making various conflicts possible. Some of these conflicts result in the leakage of sensitive data. One such mechanism is the branch prediction unit (BPU). To maximize BPU’s utilization, it is typically shared between hardware threads; it is not flushed on mode and context switches while addresses are truncated, making it prone to various branch collisions [17, 62]. This enables attacks such as side channels [2, 19, 21] that are capable of leaking encryption keys or bypassing address space layout randomization (ASLR), and

the recently introduced speculative execution attacks [31, 32]. At the same time, shared BPUs are beneficial for performance. They allow high utilization of hardware structures to reduce the cost and enable efficient branch history accumulation [50]. Therefore, naïve protections which disable sharing or flushing BPU structures have high performance overhead. Recently Intel introduced microcode updates implementing countermeasures against Spectre attacks [28]. While being effective at mitigating attacks, they can impose the performance overhead as high as 440% [56, 68].

Despite significant efforts directed towards designing various secure microarchitectural components e.g., caches [16, 30, 40, 61, 71, 73, 74] and memory buses [4, 38, 65], secure BPU designs remain a handful of attempts [23, 36, 72, 81]. When designing a microarchitecture security mechanism, it is important to correctly estimate the attacker’s capabilities. Otherwise, it risks to be defeated by more advanced attack algorithms as was recently demonstrated with randomized caches [9, 13, 57, 59, 69].

In this paper, we propose Secret-Token Branch Prediction Unit (STBPU), a secure BPU design aimed to protect against collision-based BPU attacks and eliminate BPU side channels. STBPU prevents attacks by disallowing software entities from creating controlled branch instruction collisions and thus affecting each other in an unsafe way. This is done by customizing the branch instruction representation for each software entity in the form of address mappings and by encrypting data stored in BPU. In STBPU, each software entity is provided with a unique, randomly-generated secret token (ST) that customizes the data representations. STBPU detects active attacks by monitoring related hardware events and automatically re-randomizes the ST to prevent attackers from reverse-engineering the ST value and creating collisions.

II. BACKGROUND

A. BPU Baseline Model

A typical ISA permits the following types of branch instructions: i) Direct jumps/calls where target addresses are encoded as an offset from the current instruction pointer and stored as an immediate value. ii) Conditional jumps that are only taken if a certain flag in the flag register is set. The target of this branch is encoded similarly to direct jumps. iii) Indirect jumps/calls where targets are stored in a register or in memory, and can change throughout program execution. iv) Return instructions

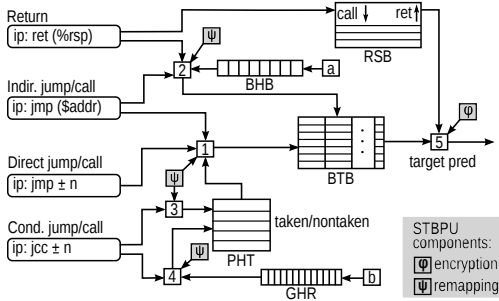


Fig. 1: BPU with STBPU components highlighted

are a special type of indirect jumps where the target is stored on top of the call stack.

Below we describe a BPU baseline model that we will utilize as a foundation to build STBPU. The baseline reflects the branch predictor (including structure sizes) used in Intel Skylake microarchitecture. Derived from recent reverse-engineering works [19, 21, 32, 34, 43, 78], it represents a generalization of mechanisms used in modern Intel processors. STBPU can be applied to other branch predictor configurations and designs. This is possible because STBPU does not interfere with underlying prediction mechanisms and only changes the branch instruction representation inside BPU data structures. We demonstrate this by adapting the STBPU to protect several advanced predictors such as TAGE-SC-L [67] and Perceptron [29].

The BPU consists of the following main structures: shift registers such as the global history register (GHR) and branch history buffer (BHB), branch target buffer (BTB), pattern history table (PHT), and return stack buffer (RSB). Figure 1 depicts how these structures are utilized during a BPU lookup with highlighted components that are modified by STBPU. The figure also shows several important functions which are referenced later.

Shift registers such as GHR and BHB are used in the BPU as a low-cost way of retaining complex branch history. GHR stores the global history of taken/not-taken branches and is used in the prediction of conditional branches. BHB is used by the indirect branch predictor. Its purpose is to accumulate the branch context. When a direct branch (or a call) is executed, its virtual address is folded using XOR and mixed with the current state of BHB [32]. This context is used as part of BTB lookup, enabling BPU to predict the target of an indirect branch when it depends on the sequence of previously executed branches.

BTB serves the purpose of caching target addresses of branch instructions. It is implemented as an 8-way, 4096-entry table. Each entry stores a truncated address of the 32 least significant address bits of the target. Function [5] is then utilized to convert a 32-bit entry into a 48-bit virtual address during prediction by combining 16 higher bits from the branch instruction pointer with 32 lower bits from BTB. While the BTB is used to store targets for all branch types, it has two addressing modes. In mode one, the virtual address of a branch instruction is used to compute an index and tag. In mode two, in addition to virtual address, the BHB is used to perform a lookup. Mode two is only used when predicting indirect branches, and serves as a

fall-back mechanism for predicting returns. This addressing enables storing multiple targets for a single indirect branch depending on the context [19, 32, 78].

PHT is a large (16k entry) table consisting of n-bit (e.g. 2-bit) saturating counters; each counter implements a simple finite-state machine with states ranging from strongly non-taken to strongly-taken. This structure is used as a base predictor to predict the direction of conditional branches. Previous studies [7, 21, 26, 78] indicated the presence of a mechanism similar to gshare [76] with two distinct modes of addressing: i) a simple 1-level mode where the virtual memory address of a branch is used to find a PHT entry, and ii) a more complex 2-level mode where the branch virtual memory address is hashed with global history register (GHR), enabling the accurate prediction of complex patterns.

RSB is used to predict return instructions. The RSB is implemented as a fixed size (16-entry) hardware stack [34, 43]. A call instruction pushes a return address on the RSB and a return instruction pops it. Similarly to the BTB, RSB stores only 32 bits of the target. Due to limited capacity, the RSB can underflow. In this case, returns are treated as indirect branches, and the indirect predictor is utilized for prediction.

Microcode BPU Protections. Intel has proposed a set of microcode-based protections which aim to mitigate speculative execution attacks on legacy CPUs by limiting BPU structure sharing. These protections include Indirect Branch Restricted Speculation (IBRS), Indirect Branch Prediction Barrier (IBPB), and Single Threaded Indirect Branch Prediction (STIBP) [28]. IBRS prevents higher privilege processes from speculating with BPU data placed by lower privilege processes. This is done by flushing BPU structures when entering the kernel. IBPB provides protection by flushing the contents of the BPU on context switches. While effectively stopping BPU interference, flushing BPU structures results in a loss of useful history, causing significant performance reduction [56, 68]. Additionally, recent research demonstrated exploitable branch collisions within same address space [63, 78]. Thus enforcing security only during context and mode switch is not complete. STIBP logically segments the BPU such that the threads on the same physical core become isolated.

B. BPU Attacks

BPU can be manipulated to enable attacks of different types. For example, an attacker can attempt to passively observe and recover branch instruction patterns. This happens during side and covert channel attacks. On the other hand, an active attacker can manipulate the BPU state by executing branch instructions. Such a state triggers a malicious speculative execution causing data leakage. Moreover, attacks range based on what BPU property they utilize. First, there are attacks that exploit the most fundamental principle of BPU to make predictions based on the previous behavior of a branch. E.g., if a conditional branch was taken 100 times in a row, it is likely to be taken the next time. An example of such an attack is Spectre-v1 [32]. Second, there are attacks that exploit branch collisions (aliasing). Collisions appear when two

	Reuse-based (RB)		Eviction-based (EB)	
	Home effect (HE)	Away effect(AE)	Home effect (HE)	Away effect (AE)
Attack steps	BTB: 1. V: $\text{jmp } s \rightarrow d$; BTB $\leftarrow (s, d)$ 2. A: $\text{jmp } s \rightarrow d'$; (s, d) reused 3. A sees misprediction PHT: 1. V: $\text{jt } s \rightarrow d$; PHT $\leftarrow (s, t)$ 2. A: $\text{jnt } s \rightarrow s + 1$; (s, t) reused 3. A sees misprediction RSB: 1. V: $\text{call } s \rightarrow d$; RSB $\leftarrow (s + 1)$ 2. A: $\text{ret } \rightarrow s'$; $(s + 1)$ reused 3. A sees misprediction	BTB: 1. A: $\text{jmp } s \rightarrow d$ 2. V: $\text{jmp } s \rightarrow d'$ 3. V speculatively executes d PHT: 1. A: $\text{jnt } s \rightarrow d$; PHT $\leftarrow (s, t)$ 2. V: $\text{jt } s \rightarrow d$; (s, nt) reused 3. V speculatively executes $s + 1$ RSB: 1. A: $\text{call } s \rightarrow d$; RSB $\leftarrow (s + 1)$ 2. V: $\text{ret } \rightarrow s'$; $(s + 1)$ reused 3. V speculatively executes $s + 1$	BTB: 1. A: $\text{jmp } s \rightarrow d$; BTB $\leftarrow (s, d)$ 2. V: $\text{jmp } s' \rightarrow d'$; BTB $\leftarrow (s', d')$ $ H(s) = H(s'), (s, d)$ is evicted 3. A sees s mispredicted PHT: <i>PHT entries are not evicted</i> RSB: 1. A: $\text{call } s \rightarrow d$; RSB $\leftarrow (s + 1)$ then fills RSB 2. V: $\text{call } s' \rightarrow d'$; RSB $\leftarrow (s' + 1)$ evicting $(s + 1)$ 3. A sees misprediction	BTB: 1. V: $\text{jmp } s \rightarrow d$; BTB $\leftarrow (s, d)$ 2. A: $\text{jmp } s' \rightarrow d'$; BTB $\leftarrow (s', d')$ $ H(s) = H(s')$ 3. V: CPU uses static prediction PHT: <i>PHT entries are not evicted</i> RSB: 1. V: $\text{call } s \rightarrow d$; RSB $\leftarrow (s + 1)$ 2. A: overflows RSB by looping $\text{call } s' \rightarrow d'$ 3. V: CPU uses static prediction
Adversarial effects	Visible source and target of branch/call addresses, taken/nontaken patterns [3, 19, 21, 34, 37]	Timing channel due to A controlling predictions in V [3], speculative execution attacks [15, 32, 34, 43, 66, 78]	V's jmp taken/nontaken [3] and call patterns, branch instruction virtual address [32]	Timing channel due to A forcing static default predictions [3], speculatively execute gadget at static prediction address [14]

A: attacker; V: victim; $\text{jmp } s \rightarrow d$: jump from s to d ; $\text{call } s \rightarrow d$: call function d from callsite s ; BTB/PHT/RSB $\leftarrow (s, d)$: store target d for branch s in BTB/PHT/RSB; $H()$: BTB/PHT hash function; $s + 1$: next instruction after s

TABLE I: Attack surface classification for BPU collision-based attacks by event and adversarial effect types

different branch instructions map into the same BPU entry and affect one another’s behavior. In this work, we focus only on collision-based attacks. We believe that mitigating them is an important task on its own for a number of responses. i) There exist a large number of well-documented collision attacks that have truly devastating effects on security [5, 14, 15, 51, 78]. ii) Protecting from non-collision attacks requires different principles, such as delaying speculative execution [64] or limiting its observability [77]. iii) Even in systems that implement safe speculation, branch collisions can still happen, causing side channel attacks. Because of that, we believe protecting from collision and non-collision attacks are two orthogonal tasks.

There are two BPU features that are present in nearly all CPUs that make collision-based attacks possible. First, the BPU data structures are shared among all software executed on a CPU core, enabling branch collisions between different processes. Second, the BPU operates with compressed virtual addresses. For instance, out of 48 bits of branch virtual address, only 30 are utilized. Then, these bits are further compressed [32]. This allows collisions to appear within the same virtual address space, e.g., collisions between branches in kernel and user process [19]. The deterministic nature of the BPU makes it possible for an attacker to trigger collisions in a controlled way. Our proposed solution aims at eliminating such determinism.

We detail the entire collision-based attack surface in Table I. First, we classify attacks by where adversarial effect takes place, either within the attacker’s code (home effect) or in the victim’s code (away effect). Secondly, we classify by the kind of the effect. A collision in BPU structures results in either data placed by another software entity being reused, or such data evicted and replaced. We refer to these as reuse-based and eviction-based attacks correspondingly. Please note that there can be different adversarial effects enabled by same type of collision. For instance, a collision in BTB between two different branches can result in i) BTB-data reuse, ii) BTB-eviction and iii) activating malicious speculative execution. While i) and ii) results in side channel leakage of branch-related information iii) is used as part of speculative execution attack to reveal victim’s memory contents. As can be seen from

the table, there is a diverse range of dangerous collision-based attacks. By eliminating collision-based BPU attacks STBPU can stop many practical exploits and substantially improve security of microprocessors.

III. THREAT MODEL

We assume a powerful attacker that has a *complete* understanding of all hardware components and structures in the STBPU. The attacker has access to normal reverse engineering resources, such as time measurements and performance counters, and has access to a wide variety of hardware covert channels. The STBPU design calls for new special purpose registers as detailed in Section IV; the adversary is assumed to be unable to read/modify the contents of these registers. Such a role is delegated to a privileged software entity (OS, hypervisor) which attacker does not control.

We assume the attacker cannot gain access to ST for the victim process neither when it is in the special purpose register, nor when system software stores it. Former is impossible because the register can only be accessed from the privileged mode. Later happens only in the event of system software compromise. The ST can be considered as part of processes’s context that is saved and restored on context switches. The event of attacker gaining access to context data would be equal to a full compromise. In such case, there is no point for attacker to use side channels.

We consider attacks presented in Table I including both side channel attacks in which victim executes a sensitive data dependent branch branch as well as speculative execution attacks where victim is forced to speculatively execute leakage gadget code. We assume the following two attack scenarios:

Sensitive Process as Victim. In this scenario, an attacker tries to learn sensitive data from a victim process by manipulating the BPU state and recording observations. The attacker has control over user-level process co-located on the same CPU core and is capable of performing activities that are normally allowed to untrusted process such as accessing fine grain hardware timers via `rdtscp` instructions. We assume the victim and attacker can either execute on two logical cores within the same physical core or share the same logical

core with time-slicing. This scenario also includes recently introduced transient trojans [78] where collisions occurring within the same memory segments are exploited.

Kernel/VMM as Victim. The attacker takes a form of a software entity with lower privilege level, i.e. untrusted user process. The attacker tries to learn sensitive data owned by a higher privileged entity (OS kernel or VMM) by manipulating with BPU state and recording observations. Here, victim and attacker share a same contiguous virtual address space. Attacker is restricted from executing privilege instructions.

IV. STBPU DESIGN

As discussed in Section II, BPU attacks are possible due to deterministic mapping mechanisms, allowing attackers to create branch collisions. STBPU aims to stop these attacks by replacing these deterministic mechanisms with keyed remapping mechanisms which prevent branch collision construction. The design philosophy of STBPU is to create different data representations for separate software entities inside the BPU data structures. Each software entity requiring isolation is assigned a unique ST, which is a random integer that controls how branch virtual addresses are mapped into BPU structures. This ST is also used to encrypt/decrypt stored data. Compared to naïve protections based on flushing or partitioning, our approach has a number of benefits.

Consider a protection scheme where branch target poisoning is prevented by flushing the BTB on context switches. Invalidating the entire branch target history will negatively affect performance in cases where context switches are frequent. Similarly, to protect from target collisions between kernel and user branches, BTB must be flushed on mode switches (e.g. all syscalls). Partitioning hardware resources reduces the effective capacity of BPU structures resulting in a higher miss rate and lower prediction accuracy. Instead, a customized mapping approach allows separate software entities to co-exist in the BPU with minimal performance overhead; performance evaluated in Section VII. STBPU utilizes two key approaches to enable safe resource sharing.

- We make collision creation difficult by ensuring all remapping functions are dependent upon both branch address and ST.
- STBPU detects when a potential attacker process has recovered sufficient information that enables deterministic collision creation by monitoring hardware events.

A. ST re-randomization

The ST of the current process in the BPU is re-randomized once certain (OS controlled) thresholds are reached. Note that in STBPU design, the OS is trusted and is responsible for setting parameters such as the re-randomization threshold. This is a common assumption for systems protecting against microarchitectural attacks since compromising OS gives the attacker full control over the system, making such attacks non-necessary. On the other hand, such a design choice makes our mechanism more flexible and permits the OS to adjust the strength of enforcement based on factors such as whether

a certain process is considered sensitive or the attacker’s capabilities. For instance, if a more effective side channel attack is discovered after STBPU is deployed, the underlying hardware mechanism will still remain effective and will only require the OS to readjust the thresholds. Moreover, for the extreme cases of sensitive processes the OS may opt to set the threshold as low as 1, forcing re-randomization after every branch instruction, effectively disabling the BPU mechanism.

STBPU can be also adapted for systems with OS not trusted (e.g. SGX), then another system component needs to be responsible for managing tokens and thresholds. For instance, in the case of SGX, the enclave entering routine can serve this purpose. Alternatively, simple logic of ST management in STBPU should also enable hardware only implementation. Re-randomizing ST effectively resets the customization of the BPU data representation for that process. Although it leads to the loss of branch history (by making it unusable), our analysis indicates that such events are infrequent. Re-randomizing the ST of one process does not remove stored history of a process with a different ST. This is the key difference compared to flushing-based approaches. We derive the re-randomization thresholds through the analysis in Section VI.

While potentially dangerous, branch history sharing between programs benefits performance. Consider a server application that spawns a new process for each incoming connection. Since each process executes the same code, the accumulated BPU state is used by the newly spawned process. This allows the new process to avoid the lengthy period of BPU training. STBPU permits selective history sharing by allowing the OS to provide multiple copies of the same program to utilize the same ST value. However, when sharing is not desired, each thread can be given a unique ST.

B. Hardware Mechanisms and Interfaces

Since current BPU designs are highly optimized in terms of performance and hardware cost, we restrict ourselves to only modifying BPU mapping mechanisms, adding registers, and encrypting stored targets. Such changes will provide similar performance to the unprotected design and make STBPU agnostic to a particular BPU design. In STBPU, each hardware thread is provided with an extra register to store the ST of the current process. Only the OS is allowed to read/modify these registers, and these registers are inaccessible in unprivileged CPU mode. As such, the OS facilitates history retention across context/mode switches by loading the appropriate STs. We also add several model-specific registers (MSRs) that store thresholds and counters for automatic ST re-randomization. These MSRs monitor the events that indicate an active attacker process. We monitor two events: i) branch mispredictions which includes incorrectly predicted direction of conditional branches and targets of any branch, and ii) BTB evictions. In Section VI, we explain how these events are utilized to deter BPU attacks. Initially, the counter values are set to their respective threshold values. When an event is observed, the corresponding counter is decremented. When a counter reaches 0, the current ST is re-randomized, and the CPU reset the

	Baseline input	STBPU input	Output	Function
[1]	32 s	32 ψ , 48 s	9 ind, 8 tag, 5 offs	$R_1(80 \mapsto 22)$
[2]	58 BHB	32 ψ , 58 BHB	8 tag	$R_2(90 \mapsto 8)$
[3]	32 s	32 ψ , 48 s	14 ind	$R_3(80 \mapsto 14)$
[4]	18 GHR, 32 s	32 ψ , 16 GHR, 48 s	14 ind	$R_4(96 \mapsto 14)$
[5]	48 s , L(GHR)	32 ψ , 48 s , L(GHR)	10/13 ind, 8/12 tag	$R_t(80 \uparrow \mapsto 25)$
[6]	48 s	32 ψ , 48 s	10 ind	$R_t(80 \mapsto 10)$

L(GHR) — represents geometric series of global history lengths
 s — represents the source bits of branch instructions

TABLE II: I/O bits for baseline and STBPU functions

counter with the threshold value. The OS treats these registers as a part of software context saving, and recovering their values on context/mode switches. We assume re-randomization is done by fetching a value from low-latency in-chip pseudo-random number generator [42].

The ST register is a 64-bit register divided into two 32-bit chunks, ψ and φ . The first chunk ψ acts as a key for a keyed remapping functions making BPU mapping unique for each process. We replace functions [1], [2], [3] and [4] in Figure 1 with STBPU remappings $R_{1..4}$ accordingly. We add functions R_t and R_p that are used for STBPU implementation with the TAGE and Perceptron predictors. Both baseline and STBPU remapping functions reduce input data (address, BTB, GHR bits) into fixed size index, tag, and offsets used by the BPU to perform lookups. Section V-A describes how $R_{1..4,t,p}$ were selected. Additionally, these functions utilize the entire 48-bit virtual address unlike legacy functions that use truncated address bits as inputs. This is crucial to prevent the same address space attacks [78]. Table II details all input/output bit changes between the baseline and STBPU models.

We use a simple scheme based on XOR to encrypt data stored in BPU structures to stop attackers from redirecting execution to a desired speculative gadget even if collisions occur. In the case of a collision, speculative execution will be redirected to an encrypted (random) address. This will effectively stall malicious speculative execution. In STBPU, every entry stored in BTB and RSB is XORed with φ of the current process. Note that the baseline BPU stores only 32 bits of target addresses, so the 32-bit φ is sufficient for encrypting all stored bits. We use a simple XOR encryption for two reasons: i) XOR operations are extremely fast with trivial hardware implementation, and ii) automated ST re-randomization makes the simple XOR encryption sufficiently strong (discussed in Section VI). To decrypt data in BTB and RSB, we modify the function [5], which XORs target bits with φ before extending them to 48-bit address.

V. IMPLEMENTATION

In Section IV, we defined remapping functions $R_{1..4,t,p}$ which replace the methods of calculating indexes, tags, and offsets for lookup purposes in the baseline BPU model. Remapping functions $R_{1..4,t,p}$ can be thought of as *non-cryptographic* hash functions. Given the size constraints of the BPU structures, collisions between different inputs to functions $R_{1..4,t,p}$ will occur; this fact prevents functions $R_{1..4,t,p}$ from providing cryptographic security, regardless of implementation. This inherent weakness is remedied with periodic re-randomization of STs; the security of such re-randomizations are discussed in Section VI. The mapping

functions used in the baseline model are not fully reverse engineered, but we can safely assume some fast compression functions are used with delays of no more than 1 clock cycle. Using performance and security as our guides, we placed several important constraints upon functions $R_{1..4,t,p}$:

- C1 The compute delay for $R_{1..4,t,p}$ must not exceed C clock cycles, where C may vary from CPU to CPU. For our purposes, we choose C to be 1 clock cycle. We enforce this by limiting the number for transistors of each remapping function on the critical path.
- C2 The function must provide *uniformity*: outputs of $R_{1..4,t,p}$ should be uniformly distributed across their respective output spaces.
- C3 The function must demonstrate *avalanche effect* [24]: The outputs of $R_{1..4,t,p}$ must appear to be pseudo-random, and the relationship between inputs and outputs should be non-linear.

We analyzed existing hardware supported hashing mechanisms, but found none that satisfied our specific requirements. Specifically, existing multi-round hash functions exceed the single CPU cycle constraint. Later we describe a mechanism we developed to automatically generate remapping functions taking into account aforementioned constraints. In addition to remapping, STBPU requires encryption of branch addresses stored inside BPU. We found out that existing lightweight cryptographic functions are not suitable for our purposes for two main reasons: First, using strong ciphers does not directly translate into better security which are primarily designed to withstand known plaintext/ciphertext attacks. However, STBPU threat model is much different as attackers never observe encrypted addresses (ciphertext) nor partially matched plaintext/ciphertext. They only observe collisions (not knowing with their own or victim’s branch) and need to reverse-engineer the rest of the address bits. Besides, knowing their own STs does not provide immediate access to collision creation or simplifies collision-based attacks. In Section VI, we show that the number of mispredictions and evictions attackers must incur to successfully infer a ST far exceeds the thresholds that will trigger ST re-randomization. Thus, encrypting with a more advanced cipher would not increase the level of security. Secondly, more sophisticated encryption schemes introduce significant delays in CPU frontend. For instance, we explored PRINCE-64 [12] and Feistel-Network [47] to encrypt stored branch targets. While comparably fast, PRINCE-64 and Feistel-Network will still consume multiple clock cycles and consume more energy due to higher number of gates compared to a simple subsingle-cycle XOR operation.

A. Automation of Finding Remapping Functions

Automated Remap Generation Algorithm. Designing the remapping mechanisms is a multi-variable optimization problem. To solve it, we developed an algorithm that takes in a list of hardware constraints, and randomly generates remapping function candidates. The algorithm composes the function from a predetermined pool of primitives. Each remapping function is iteratively generated and tested one layer at a time,

where a layer is a block of these primitives. After a layer is added, the current function is tested against the supplied constraints. There are three possible scenarios that occur during each round of testing. i) The current design satisfies all constraints, and subsequently stored for later optimization. ii) The current design violates one or more constraints, and is discarded. iii) The current design does not outright violate the constraints, but is incomplete. In case 3, our algorithm changes the weights used for primitive selection during the creation of the next layer to improve the current design.

Constraint Selection of C1. Our algorithm requires an input of several variable constraints for the generated remapping functions to satisfy C1. These constraints are: the maximum count of transistors along the critical path, the maximum number of transistors in parallel (breadth), the maximum number of total transistors for the design, the number of input and output pins, the maximum number of functional layers (blocks) the design can have, and the maximum number of wires an arbitrary wire can cross over.

Modern processors are designed to perform 15-20 gate operations in a single cycle [58], which translates to roughly 30-45 transistors along the critical path. The delay incurred by each transistor in the critical path is relatively independent of the CPU clock cycle; therefore, the faster the CPU’s clock cycle, the smaller the number of transistors that can be completed within 1 clock cycle. Therefore, we assume 45 is the absolute maximum number of transistors we allow in the critical path with preference set for shorter critical paths.

Primitive Selection. Much research has been conducted into cryptographic hash primitives [10, 11, 35, 79, 80] that provide building blocks for hash functions with strong properties. We leverage these primitives from SPONGENT [11] and PRESENT [10] hashes. Out of those S-boxes (establishing non-linearity by substations) are perhaps most critical. To increase the simplicity of remapping function generation, we separate primitives into two categories: non-invertible compression primitives and mixing primitives.

Non-invertible primitives tend to employ XOR logic gates to obfuscate the relationship between input and output. For many such primitives, multiple inputs generate the same, smaller output which makes reverse-engineering difficult. Combining multiple non-invertible layers increases complexity of attacks aiming to pair a known output to an unknown input. These primitives *compress* input size $|m|$ to an output size $|n|$ where $|m| > |n|$. Table II shows the disparity between the input and output sizes for $R_{1..4,t,p}$ functions, and indicates the need for optimized compression primitives. Mixing primitives are primarily used to introduce *non-linearity* to a hash design which makes deterministically changing the output by varying the input difficult. These primitives are primarily composed of $|m| \mapsto |m|$ sized S-boxes and P-boxes (performing permutations). Since the hardware complexity of S-boxes increases superlinearly with the size of $|m|$, we limit our S-boxes to a maximum of 4 input/outputs. These S-boxes can be implemented efficiently with combinatorial logic or transistor/diode matrices. P-boxes are constrained by the




Fig. 2: R_1 remapping function construction

maximum wire crossover set for the algorithm.

Validation of Uniformity (C2) and Avalanche Effect (C3)

Remapping functions that satisfy the hardware constraints are then tested against constraints C2 and C3. We first employ the balls and bins analysis and compute the coefficient of variation (CV) of bins to approximate the uniformity (C2) of the output space [60]. C3 is satisfied when a remapping adheres to a strict avalanche criterion. To quantify the avalanche effect of F , for each input λ , we generate a set of unique inputs, S , where each input in S differs from λ by a single bit flip. We then compute the hamming distance between $F(\lambda)$ and $F(S_i)$, for all inputs in S . Using these hamming distances, we determine the CV of the hamming distances for a particular λ . We test each F with 1 million random inputs and compute the average hamming distance for all inputs. The ideal case occurs when: i) the average hamming distance over 1 million random inputs is roughly 50%. ii) For all inputs, the CV of the average hamming distance for each input is 0. iii) For all bit positions of an output of F , the difference between the minimum and maximum hamming distances for a bit flip in any bit position is 0.

B. Optimization and Remapping Selection

The final selection of remapping functions $R_{1..4,t,p}$ is primarily based upon the results from the previous tests. The result is a multiobjective optimization problem where the ideal state for different desired metrics may be maximized or minimized. To make all metrics comparable, we normalized each metric so that the optimal value is 0. We then considered this to be a simple weighted optimization problem where we seek functions that yield the lowest sum of all metrics recorded when testing for uniformity and the avalanche effect. Let F be a particular function in the group of potential functions G for remapping function R_i , for $i \in R_{1..4,t,p}$:

$$\min \sum_i^k w_i g(F), F \in G \quad (1)$$

All weights were set to 1 to avoid prioritizing one metric over another. Further prioritizing then can be done by hardware developers for a specific CPU design. For space reasons, we do not show the designs for all of $R_{1..4,t,p}$ since they share many similar characteristics. Instead, we show the chosen design for R_1 in Figure 2 where stages 1, 3, and 5 are substitution layers using $4 \mapsto 4$ and $3 \mapsto 3$ S-boxes. For space reasons, not all types of S-boxes are shown. Under the design of R_1 , we show the logical mappings for S-boxes used by PRESENT and SPONGENT. P-boxes are $n \mapsto n$ in size with the pin mappings

generated randomly by our remap function generator. C-S boxes are compression structures that map $|m|$ bits to an output size of $|n|$ bits where $|m| > |n|$. This design of R_1 has a critical path length of 36 transistors, so it is capable of being computed within a single clock cycle.

VI. SECURITY ANALYSIS

We assume any attackers can have complete knowledge of all STBPU remapping functions, full control of execution flow, and are capable of executing branches to/from any address within their processes. The goal is to enable malicious branch instruction collisions that allow mounting one of the collision-based attacks. STBPU makes collisions non-deterministic, forcing the attackers to rely on either brute force approaches or reverse-engineering the ST value. Further, attackers can utilize recently proposed fast attack algorithms such as GEM [59] and PPP [57] that target randomized caches [9, 13].

Parameter	Description	A : Branch in attacker(A)'s address space
W_{struct}	Number of ways	V : Branch in victim(V)'s address space
I_{struct}	Number of sets (indexes)	$\psi_{a/v}$: A/V $R()$ 32-bit token
T_{struct}	Entry tag bit entropy	$\varphi_{a/v}$: A/V target encryption token
O_{struct}	Entry offset bit entropy	τ_Q : Target of arbitrary branch Q
Ω_{struct}	Entry target bit entropy Q	E_Q : Entry stored for arbitrary branch Q

TABLE III: Parameters used in STBPU analysis

A. Analysis of Branch Predictor Attacks under STBPU

An attacker possessing knowledge of their ST (ψ/φ) voids the security provided by the STBPU because they can deterministically generate outputs with any of the remapping functions used by the STBPU. Before we discuss how STBPU affects attacks on BPU, we show the parameters for security analysis in Table III and list several important axioms below:

- A1 Attackers do not know the numerical outputs of $R_{1..4,t,p}$.
- A2 Due to A1, all the current state of the STBPU must come from detection of mispredictions and evictions.
- A3 Attacker does not have inherent knowledge or control of ST of any process.

1) *Target Injection Attacks*: Recall that we encrypt the targets stored in the BTB and RSB through the following means: $E_A = \varphi_a \oplus \tau_A$. With Spectre V2, the attacker supplies a malicious τ_A using branch A that collides with the victim's branch V causing V to speculate with τ_A . With the SpectreRSB, the attacker places a malicious return address τ_A on the stack that the victim speculates with. In both cases, the target the victim will use from the STBTB or STRSB is now $\tau_V = \varphi_a \oplus \tau_A \oplus \varphi_v$. If there is a Spectre gadget located in the victim's address space at address G , the attack is successful if $\tau_V = G$. Due to A3, the attacker does not have knowledge or control of φ_a or φ_v ; consequently, the only variable the attacker can change is the address of τ_A to make $\tau_V = G$. The probability that τ_A results in $\tau_V = G$ is $\frac{1}{\Omega_{STBTB}}$ or $\frac{1}{\Omega_{STRSB}}$. As such, the attacker must execute $\frac{\Omega_{STBTB}}{2}$ or $\frac{\Omega_{STRSB}}{2}$ different τ_A values to have a 50% chance of successfully executing their target injection attack. Each incorrect τ_A will result in the misprediction counter decrementing towards zero.

2) *Reuse-based Attacks*: Address mappings are randomized so that there is only a probability that an arbitrary A and V will collide in the STBPU. Even though A and V are mapped with $R_{1..4,t,p}$, the probability that attacker branch A collides with victim branch V in the STBTB/STPHT is not bound by birthday attack complexity because V is a static, specific address. The probability of collision is $P(A \Rightarrow V) = (\frac{1}{I})(\frac{1}{TO})$. Note, we break up the probability that A and V are in the same set vs. the probability that A and V have matching tag and offsets because tag/offset comparisons are only done if A and V are in the same set. This adds uncertainty for reuse-based side channels where the attacker wishes to determine the direction of V since a lack of misprediction by A or V could mean that A and V do not collide, or that V was not taken. To increase the probability that an arbitrary A collides with a static V , the attacker can execute a set of branches $S_B = \{b_1, \dots, b_n\}$ where n is large so that one branch in S_B might collide with V . The probability that one of the branches in S collides with V is $P(S_B \Rightarrow V) = \sum_{i=1}^n P(S_{B_i} \Rightarrow V)$. However, noise is added using this method because it is possible that branches in S_B will collide with each other. The probability that two branches in S_B collide can be approximated with birthday attack complexity because the branches in S_B are arbitrary.

In order to ensure that no branches in S_B collide with any other branch in S_B , the attacker execute the following steps: i) Choose a new branch b_{new} with a new address in attacker's address space. ii) For every branch b_i in S_B , execute b_i and b_{new} . iii) If no MISPs. between b_i and b_{new} , $S_B = S_B \cup \{b_{new}\}$. In order to achieve a 50% probability of collision between A and a branch in S_B , the size of S_B must be $\frac{ITO}{2}$. The number of MISPs M and evictions E generated whilst generating S_B of size $n = \frac{ITO}{2}$ can be approximated as follows:

$$M \approx \sum_{i=0}^n \sum_{j=0}^{j=i} \frac{1}{\sqrt{\frac{\pi}{2}I}} \cdot \frac{1}{\sqrt{\frac{\pi}{2}TO}} = \frac{n(n+1)}{2\sqrt{\frac{\pi}{2}I} \cdot \sqrt{\frac{\pi}{2}TO}} \quad (2)$$

$$E \approx \frac{ITO}{2} - IW$$

Note the reuse-based side channel attacks on PHT do not generate evictions. The size of the STBTB is IW which is significantly smaller than $\frac{ITO}{2}$, so entries in the BTB will constantly be evicted as the attacker grows S_B .

Attacks such as BranchScope [21] and BlueThunder [26] are viable against processors using hybrid directional predictors largely due to the inclusion of a base directional predictor in these hybrid BPUs. Due to the complexity of TAGE tables and Perceptron weights, it is *significantly* easier to maliciously modify the base directional predictor than the complex TAGE/Perceptron structures. Since the remapping mechanisms used in our TAGE/Perceptron structures are different than the remapping functions used for the base directional predictor, little information is gained by an attacker observing mispredictions from both the base and complex directional components. Due to A1, an attacker will not know which TAGE bank or Perceptron weight set produced a prediction. The thresholds for re-randomization stemming from mispredictions from the directional predictor are based on the *least* complex attack

on the directional predictor. More complex attacks will be affected by re-randomization to a greater extent.

3) *Same Address Space Attacks*: Recently discovered same address space attacks [78] are classified as target injection attacks, but in this case both A and V are located inside the attacker’s address space. As such, encrypting the target of A with φ_a provides no security because V will decrypt τ_A with φ_a . However, due to R_i , there is only a probability that A and V will collide; this probability is the same as for reuse-based attacks. Therefore, the number of mispredictions and evictions generated while performing a same address space attack are also approximated by Equation (2).

4) *Eviction-based Attacks*: The attacker cannot deterministically create BTB eviction sets without knowing ψ_a since address mappings change when ψ_a is re-randomized. With W_{stbtt} ways, detecting an eviction in an arbitrary set requires $W_{stbtt} + 1$ colliding branches (same index, different tag and/or offset). The attacker wants to fill STBTB sets so that if V is executed, it disturbs one of the attacker’s primed sets. To increase the chances that V will enter a primed set, the attacker must prime as many sets as possible. Assuming the ideal case when the attackers does not have conflicts between their own branches, they need to cover $P \cdot I$ sets to achieve P probability of a successful attack. For example, the probability that A enters the same set as a static V is $\frac{1}{I}$, so to have a 50% chance of priming the set V enters, the attacker must prime $\frac{I}{2}$ sets. Naïvely, the probability of randomly guessing W_{stbtt} branches to form a single set of branches S_e that enters the same STBTB set is:

$$P(S_e) = \frac{1}{I^{W_{stbtt}-1}} \quad (3)$$

Since this probability is not favorable, the attacker could apply a fast algorithm GEM [59] to construct every eviction set. The attackers uses GEM because bottom-up strategies like PPP becomes less efficient without a partitioned randomized structure [57] or specific cache conditions [9, 13, 69]. We assume the ideal scenario for the attacker is when most of the branches tested follow a perfect uniformity. In this case, given a particular branch, the probability to have W branches belonging to the same set is directly related to the total number of test entries. For instance, there is a 50% probability that in a group of $\frac{IW}{2}$ branches that at least W branches share the same index. Thus, in order to achieve P attack rate, the attacker needs to test at least PIW branches as the initial set since the total attack lines in L in GEM. (E.g., $L > 44$ for an efficient GEM in [59]). With the original setting in GEM, the attacker sets the group size $G = W + 1$ and starts to eliminate groups of branches. Although the total branch accesses will be approximately $2.3 \cdot W \cdot L$, the total eviction number will be less as the majority of the probe during each iteration will be hit. Since the probability that each group will produce an eviction is approximately equal to $1 - 1/e$. The evictions generated by testing will be negligible as $(W + 1) \cdot 1 - \frac{1}{e} \cdot n$ since the total rounds n for GEM converge on the list of conflicting lines are relatively small. However, when first placing L branches, the attacker has to trigger the same amount of evictions.

Summarizing the procedure to construct required eviction sets above, we can now approximate evictions numbers generated whilst building sets for P attack rate as follows:

$$E \approx PI \times (PIW + (W + 1) \times (1 - \frac{1}{e}) \times 3) \quad (4)$$

5) *Re-randomization Thresholds for Baseline Model*: STBPU has the same parameters as the baseline Intel Skylake BPU. The BTB has 8 ways and 512 sets. The stored entries have a compressed 8-bit tag and a 5 bit offset. The PHT has 1 way and 2^{14} sets. Using Equation 2, the number of mispredictions and evictions an attacker will trigger before a successful reuse-based side channel attack on BTB is 6.9×10^8 and $\approx 2^{21}$, respectively. Correspondingly, for a PHT reuse-based side channel, the number of triggered mispredictions is $\approx 8.38 \times 10^5$. For a BTB eviction-based side channel, the average number of triggered evictions is $\frac{I}{2}$ or 5.3×10^5 per Equation 4. For Spectre V2 and SpectreRSB, the number of triggered mispredictions is $\approx 2^{31}$. To prevent attacks, we use the lowest misprediction and eviction thresholds as the upper bounds for re-randomization of ST when evaluating the performance of STBPU.

6) *Denial-of-Service Attack on STBPU*: While the primary goal of a typical attacker is to reveal some sensitive data via a side channel or speculative execution attack, they can also attempt to perform a denial-of-service (DoS) attack. In this attack, the goal is to cause an abnormal slowdown of a victim process by triggering excessive branch mispredictions. We consider two DoS attack scenarios: i) Eviction-based: attacker attempts to evict from BPU data associated with a branch that is critical for the victim’s performance. ii) Reuse-based: attacker fills BTB with bogus data hoping to make the victim speculatively execute code at a wrong address causing a delay due to the recovery from incorrect speculative execution. On high level, STBPU makes both of these attacks more challenging because they rely on branch instruction collisions which are difficult to create in STBPU. Now we will discuss each attack in more detail.

STBPU cannot eliminate the possibility of the first attack because, in STBPU, internal data structures such as BTB remain to be shared. However, the attack becomes more difficult to carry out with STBPU. Since the victim and attacker are guaranteed to use different STs, the attacker must default to a brute force. Due to unknown branch-to-BTB mappings, finding eviction sets becomes a difficult task. Since BTB is a set-associative structure, to guarantee eviction of a certain entry, the attacker needs to find n branches mapped into the same set, where n is the number of ways in BTB. Since the attacker is blind, the attacker must rely on constantly executing a large number of branches hoping to evict the victim’s entries.

The second attack is very difficult in the case of STBPU. In order to cause a hit in BTB, the attacker’s and victim’s branches need to have the same index, tag, and offset after they are remapped by STBPU mechanisms with different STs. Based on our analysis above, such an event is unlikely to happen. Moreover, because the stored address is encrypted with the ST of a different process, the predicted address

would most likely point to an invalid address. Thus, erroneous speculative execution would not happen.

VII. EVALUATION OF STBPU DESIGN

Evaluating BPU design under realistic conditions is a challenging task. Sharing BPU resources creates various possibilities for branch conflicts that affect prediction accuracy. Moreover, some BPU protection mechanisms (e.g., Intel’s IBRS) are triggered by system events such as mode and context switches. BPU resources need to be flushed upon context/mode switches to avoid BPU training or state leakage between user and kernel processes. Workloads that involve frequent system calls and interrupts may experience performance degradation and negatively affect other programs executing on the same core. In this situation, standard benchmark suits that are typically compute-bound and do not trigger frequent library calls, mode and context switches may not accurately evaluate BPU performance effects. Thus, a good evaluation environment needs to capture system-wide events and include real applications. A trace-based simulation is a logical choice for this. Meanwhile, the complex performance side effects caused by branch mispredictions and evictions require detailed performance data (e.g., IPC) using a cycle accurate simulator. To address both the abovementioned aspects, we evaluate STBPU using two simulation frameworks.

First, we utilize the Intel processor trace (PT) technology to collect large amounts of branch instruction traces captured from different workloads within the same CPU physical core, including user applications that cause frequent mode switches and context switches and the SPEC benchmarks. These traces then will be passed through an in-house BPU simulator with the BPU baseline found in the Intel Skylake processor. The simulator also runs different secure models such as STBPU and reports prediction accuracy. Secondly, to evaluate fine-grained microarchitectural performance effects, we implemented the STBPU mechanisms inside gem5 [8] and conducted simulations in syscall-emulation (SE) mode using DerivO3CPU model with configurations that mimic similar modern processors. The detailed configuration is listed in Table IV. All gem5 simulations were performed simulating 110 million instructions with a warm-up of 10 million instructions.

ISA	Single thread: X86-64, 3.4GHz; SMT: Alpha, 3.4GHz
BPU	BTB entries: 4096, 8-way, RAS size: 16
Core	8-issue, OoO, IQ/LQ/SQ entries: 64/32/32, ROB: 192, ITLB/DTLB: 64/64
Cache	L1-I/L1-D: 32KB/32KB both 8-way, L2: 256KB 4-way, LLC: 4MB 16-way

TABLE IV: Parameters used in gem5 simulation

A. Re-randomization Threshold

In Section VI, we demonstrate the misprediction and eviction thresholds for ST re-randomization when various STBPU attacks have a P attack success rate. For BranchScope attack, to have a 50% chance of success, the number of triggered mispredictions is estimated at $\approx 8.38 \times 10^5$. For a BTB eviction-based side channel attack, the number of triggered evictions is $\approx 5.3 \times 10^5$. These are the lower-bound numbers of mispredictions and evictions triggered by any attack discussed in this paper. We aim to re-randomize ST well before the

attacker has a reasonable probability of a successful attack. To do so, we utilize results from the previously discussed security analysis and derive the re-randomization thresholds as follows. We first denote the attack complexity C as the least number of evictions or mispredictions that the attack needs to trigger to succeed with a 50% chance. Please note that we use 50% probability rather than 100% since on average the attacker will succeed with half the number of attempts needed for the fully exhaustive key search. Let the variable r be the attack difficulty factor and Γ be the re-randomization threshold. As such, $\Gamma = r \cdot C$. An attack has a 50% success rate when $r = 1$. For instance, if $r = 0.1$, then the re-randomization thresholds for mispredictions and evictions are set to 8.3×10^4 and 5.3×10^4 , while 4.15×10^4 and 2.65×10^4 when $r = 0.05$. For further experiments, we set r to 0.05 and derive the re-randomization thresholds from this value as it offers strong security guarantees with a low impact on performance.

B. STBPU Performance Evaluation

1) *Prediction Accuracy with real branch trace:* We evaluate the STBPU impact on BPU accuracy and compare it to existing naïve protections modeled after microcode protections based on flushing or partitioning BPU resources. To do so, we utilize our trace-based BPU simulator based on Intel PT technology. It avoids simulating the complex state of microarchitectural components. Instead, it is designed to allow rapid testing of BPU models using branch traces from a live system running a variety of real-world scenarios.

Each simulation instance is collected from an Intel Core i7-8550U machine that captures traces from a live physical core and includes any OS/library code executed, including naturally occurring context, mode switches, and interrupts. This allows realistically simulating complex cross-process BPU effects and assessing how BPU flushing or ST re-randomization affects performance. To evaluate single-process compute-bound scenarios, we collected 23 traces from different workloads in SPEC CPU 2017. In addition, we captured traces from user and server applications, including Apache2 workloads under different prefork settings, Google Chrome traces when running single or multiple browser workloads, MySQL server, and OBS Studio.

As previously discussed, our baseline BPU model is based on recent reverse-engineering efforts [19, 21, 32, 34, 43, 78]. To evaluate STBPU, we applied the ST mechanisms from Section IV-B to the BPU baseline model. We also created two models that mimic the baseline model with Intel’s microcode-based protections, namely μ code protection 1 and 2, modeling IPBP+IBRS protection with and without STIBP. Please note that microcode-based protections cannot prevent branch collisions from occurring within the same context. To prevent such collisions, more structural BPU changes are required. In particular, instead of storing compressed and truncated addresses in BTB, the full 48-bit address must be stored. As a result, the number of entries the BTB is capable of storing must be reduced (assuming unchanged hardware budget). We refer to such a model as conservative, which fully prevents any known




Fig. 3: Overall branch prediction accuracy: STBPU against other secure BPU models

collision-based BPU attack by flushing or partitioning. Note that STBPU achieves the same security level via customizing BPU data representations and has better performance.

The result from simulating the above five models is demonstrated in Figure 3 where we aggregate all the effective predictions into a single metric: overall accuracy effective (OAE). OAE counts a branch correctly predicted if all necessary (target and direction) predictions are correct; otherwise, it's counted as mispredicted. Figure 3 shows the overall accuracy of the various BPU models against the SPEC2017 benchmarks and user applications. STBPU demonstrates an average 1.3% overall effective prediction accuracy penalty. For comparison, the microcode and the conservative BPU models suffer at least around 12% overall accuracy loss with multiple cases of nearly 30% reduction. With this, we conclude that based on the BPU accuracy data STBPU outperforms the microcode protections that utilize flushing and partitioning.

2) *Cycle Accurate Evaluation using gem5*: Our next evaluation focuses on the comprehensive impact of STBPU on Out-of-Order (OoO) CPU in terms of cycle accurate performance, evaluating effects of STBPU on advanced branch predictors, and SMT performance. We tested three advanced BPU models: TAGE_SC_L_8KB, TAGE_SC_L_64KB [67], and PerceptronBP [29]. To demonstrate the consistency of accuracy between gem5 and our previous evaluation, we also ported and tested our baseline model from Section VII-B1. We refer to it as SKLCond. We compared the direction prediction accuracy between SKLCond in gem5 with our previous baseline model using the same workloads. We observed on average less than 5% direction prediction difference which validates our simulator consistency.

We treated the aforementioned four BPUs as baseline models and implemented four STBPU models. In single process evaluations, we simulated each pair of STBPU models and their non-ST counterparts across 18 SPEC2017 workloads. Figure 4 illustrates the reduction of direction / target predictions rate and the normalized IPC between STBPU designs and their non-secure counterparts. We observe all 4 STBPU designs can achieve less than 2% reduction on average target prediction rate and less than 1.3% reduction on average of direction prediction rate. The less than 4% average IPC reduction demonstrates the high effectiveness of STBPU designs.

We used the same eight BPU models in our gem5 SMT simulations. Instead of running a single workload at a time, we grouped the individual workloads in pairs and simulated these pairs in SMT mode. In order to accurately evaluate




Fig. 4: STBPU single workload evaluation in gem5

the STBPU impacts on overall throughput, we calculated the Harmonic means (Hmeans) [49] of IPCs since each workload is equally valued. Figure 5 displays the overall IPC and the impact on accuracy. We observed the ST_SKLcond models suffer the most in SMT mode. This is because running tasks in SMT mode introduces more frequent ST re-randomizations. However, the reduction of throughput is less than 5%. We believe this is because the ST_SKLcond model does not have a separate threshold register as TAGE models do for TAGE-table mispredictions. This causes more frequent direction mispredictions as shown in the first chart of Figure 5. This effect further affects the overall performance. On the other hand, the advanced BPU models overall retain their efficiencies with minimized accuracy reduction and throughput slowdown.

3) *Aggressive ST Re-randomization and Performance*: It is common to see a constant arms race between protection mechanisms and more advanced attacks [57, 59]. STBPU can withstand faster attack algorithms by reducing the ST re-randomization threshold to lower values. This would result in a more aggressive protection scheme but can negatively affect the performance. To measure such an effect on performance,




Fig. 5: SMT Evaluation of STBPU using workload pairs




Fig. 6: Effects on performance when using more aggressive re-randomization thresholds with the TAGE_SC_L_64KB BPU, result are averaged from 42 combinations of SPEC CPU 2017 workload pairs. The X-axis represents the r parameter.

we experimented with lowering the r parameter. This is equivalent to assuming a new attack that is faster 10 times, 100 times, and even more.

To demonstrate an extreme case, we select an advanced BPU model most sensitive to branch history loss and thus re-randomizations. We test it in the SMT setting which is more prone to trigger branch mispredictions and evictions. Figure 6 demonstrates how reducing the r parameter affects the performance of the TAGE_SC_L BPU protected with STBPU. It shows that the thresholds can be safely reduced and maintain accuracy above 95%. However, setting the threshold too low results in ST re-randomizations happening after every few hundreds of mispredictions or evictions. This practically

ceases any BPU training.

VIII. RELATED WORK

To protect against Spectre attacks, Intel processors are enhanced with Spectre-specific microcode updates, including IBRS, IBPB, and STIBP protections [28]. The implementation of these mechanisms is not well documented and varies from one microarchitecture to another as their performance cost. Since these protections are added on top of unsafe BPU designs, they usually come with a very costly performance overhead. As a result, in practice, they are not used to their full extent and are only enabled by the OS in critical cases such as selectively protecting only a handful of processes.

Several previous academic works proposed BPU modifications to protect against side channel and speculative execution attacks. BRB [72] stores and reloads the entire history of the directional predictor for each process, effectively mitigating PHT collision-based attacks such as BranchScope. BSUP [36] first encrypts the PC and then encrypts the entries of BPU, making it unsuitable for SMT processors.

Zhao et al. [81] encode branch contents (directions and destination histories) and indexes using thread-private random numbers to achieve isolation between threads or privilege levels. Their approaches re-generate random numbers upon context and mode switches, which cannot defend against the transient execution attacks from same-address-space [75, 78]. Besides, our work implements ST re-randomization based on BPU events allowing efficient branch history retention.

The BPU in the Samsung Exynos processor is also protected with XOR-based encryption as branch history data enhancement [23]. Since this mechanism aims to prevent speculative execution attacks such as Spectre variant 2, Exynos only encrypts stored branch targets of indirect branch instructions and returns. However, other forms of branch collisions may still result in side channel leakage [81]. Additionally, in Exynos, an output of the hash function serves as a key for encrypting branch target data. It is derived from a number of process and machine-specific inputs. In our work, the OS is given more flexibility for managing the ST, which allows selective branch history sharing, adjustment of re-randomization frequency, and enforcing BPU isolation for various types of software entities such as sandboxes and libraries.

IX. CONCLUSION

We presented the STBPU, a secure branch prediction design that defends against collision-based BPU side channel and speculative execution attacks. We performed a systematization of BPU-related attacks and provided a detailed security analysis against recent attacks. While providing security, STBPU demonstrates high performance for branch predictors modeled after real-world chips and utilizing advanced models.

X. ACKNOWLEDGMENTS

The work in this paper is partially supported by Intel Corporation and National Science Foundation grant #1850365. The statements made herein are solely the responsibility of the authors and do not necessarily reflect those of the sponsors.

REFERENCES

- [1] O. Aciğmez, B. B. Brumley, and P. Grabher, “New results on instruction cache attacks,” in *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 2010, pp. 110–124.
- [2] O. Aciğmez, Ç. K. Koç, and J.-P. Seifert, “On the power of simple branch prediction analysis,” in *Proceedings of the 2nd ACM symposium on Information, computer and communications security*. ACM, 2007, pp. 312–320.
- [3] O. Aciğmez, Ç. K. Koç, and J.-P. Seifert, “Predicting secret keys via branch prediction,” in *Cryptographers Track at the RSA Conference*. Springer, 2007, pp. 225–242.
- [4] S. Aga and S. Narayanasamy, “InvisiMem: Smart memory defenses for memory bus side channel,” in *ACM SIGARCH Computer Architecture News*, vol. 45, no. 2. ACM, 2017, pp. 94–106.
- [5] E. Barberis, P. Frigo, M. Muench, H. Bos, and C. Giuffrida, “Branch History Injection: On the Effectiveness of Hardware Mitigations Against Cross-Privilege Spectre-v2 Attacks,” in *USENIX Security*, 2022.
- [6] D. J. Bernstein, “Cache-timing attacks on AES,” 2005.
- [7] S. Bhattacharya, C. Maurice, S. Bhasin, and D. Mukhopadhyay, “Branch Prediction Attack on Blinded Scalar Multiplication,” *IEEE Transactions on Computers*, vol. 69, no. 5, pp. 633–648, 2019.
- [8] N. Binkert, B. Beckmann, G. Black, S. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib Bin Altaf, N. Vaish, M. Hill, and D. Wood, “The gem5 simulator,” *SIGARCH Computer Architecture News*, vol. 39, pp. 1–7, 08 2011.
- [9] R. Bodduna, V. Ganesan, P. Slpsk, K. Veezhinathan, and C. Rebeiro, “Brutus: Refuting the security claims of the cache timing randomization countermeasure proposed in ceaser,” *IEEE Computer Architecture Letters*, vol. 19, no. 1, pp. 9–12, 2020.
- [10] A. Bogdanov, L. R. Knudsen, G. Leander, C. Paar, A. Poschmann, M. J. B. Robshaw, Y. Seurin, and C. Vikkelsoe, “PRESENT: An Ultra-Lightweight Block Cipher,” in *Cryptographic Hardware and Embedded Systems - CHES 2007*, P. Paillier and I. Verbauwhede, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 450–466.
- [11] A. Bogdanov, M. Knežević, G. Leander, D. Toz, K. Varıcı, and I. Verbauwhede, “spongint: A Lightweight Hash Function,” in *Cryptographic Hardware and Embedded Systems – CHES 2011*, B. Preneel and T. Takagi, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 312–325.
- [12] J. Borghoff, A. Canteaut, T. Güneysu, E. B. Kavun, M. Knezevic, L. R. Knudsen, G. Leander, V. Nikov, C. Paar, C. Rechberger *et al.*, “PRINCE—a low-latency block cipher for pervasive computing applications,” in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2012, pp. 208–225.
- [13] T. Bourgeat, J. Drean, Y. Yang, L. Tsai, J. Emer, and M. Yan, “Casa: End-to-end quantitative security analysis of randomly mapped caches,” in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2020, pp. 1110–1123.
- [14] C. Canella, J. V. Bulck, M. Schwarz, M. Lipp, B. von Berg, P. Ortner, F. Piessens, D. Evtuyushkin, and D. Gruss, “A Systematic Evaluation of Transient Execution Attacks and Defenses,” in *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 249–266. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity19/presentation/canella>
- [15] G. Chen, S. Chen, Y. Xiao, Y. Zhang, Z. Lin, and T. H. Lai, “SgxPectre: Stealing Intel Secrets from SGX Enclaves Via Speculative Execution,” pp. 142–157, June 2019.
- [16] B. Coppens, I. Verbauwhede, K. De Bosschere, and B. De Sutter, “Practical mitigations for timing-based side-channel attacks on modern x86 processors,” in *Security and Privacy, 2009 30th IEEE Symposium on*. IEEE, 2009, pp. 45–60.
- [17] M. Evers, P.-Y. Chang, and Y. N. Patt, “Using hybrid branch predictors to improve branch prediction accuracy in the presence of context switches,” in *ACM SIGARCH Computer Architecture News*, vol. 24, no. 2. ACM, 1996, pp. 3–11.
- [18] D. Evtuyushkin, T. Benjamin, J. Elwell, J. A. Eitel, A. Sapello, and A. Ghosh, “Computing with time: Microarchitectural weird machines,” in *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2021, pp. 758–772.
- [19] D. Evtuyushkin, D. Ponomarev, and N. Abu-Ghazaleh, “Jump over ASLR: Attacking branch predictors to bypass ASLR,” in *Microarchitecture (MICRO), 2016 49th Annual IEEE/ACM International Symposium on*. IEEE, 2016, pp. 1–13.
- [20] —, “Understanding and mitigating covert channels through branch predictors,” *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 13, no. 1, p. 10, 2016.
- [21] D. Evtuyushkin, R. Riley, N. C. Abu-Ghazaleh, D. Ponomarev *et al.*, “BranchScope: A New Side-Channel Attack on Directional Branch Predictor,” in *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2018, pp. 693–707.
- [22] Q. Ge, Y. Yarom, D. Cock, and G. Heiser, “A survey of microarchitectural timing attacks and countermeasures on contemporary hardware,” *Journal of Cryptographic Engineering*, vol. 8, no. 1, pp. 1–27, 2018.
- [23] B. Grayson, J. Rupley, G. Z. Zuraski, E. Quinnell, D. A. Jiménez, T. Nakra, P. Kitchin, R. Hensley, E. Brekel-

- baum, V. Sinha *et al.*, “Evolution of the samsung exynos CPU microarchitecture,” in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2020, pp. 40–51.
- [24] N. Hua, E. Norige, S. Kumar, and B. Lynch, “Non-crypto Hardware Hash Functions for High Performance Networking ASICs,” in *2011 ACM/IEEE Seventh Symposium on Architectures for Networking and Communications Systems*, Oct 2011, pp. 156–166.
- [25] C. Hunger, M. Kazdagli, A. Rawat, A. Dimakis, S. Vishwanath, and M. Tiwari, “Understanding contention-based channels and using them for defense,” in *High Performance Computer Architecture (HPCA), 2015 IEEE 21st International Symposium on*. IEEE, 2015, pp. 639–650.
- [26] T. Huo, X. Meng, W. Wang, C. Hao, P. Zhao, J. Zhai, and M. Li, “Bluethunder: A 2-level Directional Predictor Based Side-Channel Attack against SGX,” *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2020, no. 1, pp. 321–347, Nov. 2019. [Online]. Available: <https://tches.iacr.org/index.php/TCHES/article/view/8401>
- [27] M. S. Inci, B. Gulmezoglu, G. Irazoqui, T. Eisenbarth, and B. Sunar, “Cache attacks enable bulk key recovery on the cloud,” in *International Conference on Cryptographic Hardware and Embedded Systems*. Springer, 2016, pp. 368–388.
- [28] Intel, “Intel analysis of speculative execution side channels,” January 2018.
- [29] D. A. Jiménez and C. Lin, “Dynamic Branch Prediction with Perceptrons,” in *Proceedings of the 7th International Symposium on High-Performance Computer Architecture*, ser. HPCA 01. USA: IEEE Computer Society, 2001, p. 197.
- [30] T. Kim, M. Peinado, and G. Mainar-Ruiz, “STEALTH-MEM: System-Level Protection Against Cache-Based Side Channel Attacks in the Cloud.” in *USENIX Security symposium*, 2012, pp. 189–204.
- [31] V. Kiriansky and C. Waldspurger, “Speculative Buffer Overflows: Attacks and Defenses,” *arXiv preprint arXiv:1807.03757*, 2018.
- [32] 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 preprint arXiv:1801.01203*, 2018.
- [33] P. Kocher, J. Jaffe, and B. Jun, “Differential power analysis,” in *Annual International Cryptology Conference*. Springer, 1999, pp. 388–397.
- [34] E. M. Koruyeh, K. N. Khasawneh, C. Song, and N. Abu-Ghazaleh, “Spectre returns! speculation attacks using the return stack buffer,” in *12th USENIX Workshop on Offensive Technologies (WOOT 18)*, 2018.
- [35] G. Leander and A. Poschmann, “On the Classification of 4 Bit S-Boxes,” in *Proceedings of the 1st International Workshop on Arithmetic of Finite Fields*, ser. WAIFI ’07. Berlin, Heidelberg: Springer-Verlag, 2007, pp. 159–176. [Online]. Available: https://doi.org/10.1007/978-3-540-73074-3_13
- [36] J. Lee, Y. Ishii, and D. Sunwoo, “Securing Branch Predictors with Two-Level Encryption,” vol. 17, no. 3, Aug. 2020. [Online]. Available: <https://doi.org/10.1145/3404189>
- [37] S. Lee, M.-W. Shih, P. Gera, T. Kim, H. Kim, and M. Peinado, “Inferring fine-grained control flow inside SGX enclaves with branch shadowing,” in *26th USENIX Security Symposium, USENIX Security*, 2017, pp. 16–18.
- [38] T. S. Lehman, A. D. Hilton, and B. C. Lee, “PoisonIvy: Safe speculation for secure memory,” in *The 49th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Press, 2016, p. 38.
- [39] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, “Meltdown,” *arXiv preprint arXiv:1801.01207*, 2018.
- [40] 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 *High Performance Computer Architecture (HPCA), 2016 IEEE International Symposium on*. IEEE, 2016, pp. 406–418.
- [41] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, “Last-level cache side-channel attacks are practical,” in *Security and Privacy (SP), 2015 IEEE Symposium on*. IEEE, 2015, pp. 605–622.
- [42] J. M., “Intel digital random number generator (drng) software implementation guide,” Oct 2019. [Online]. Available: <https://software.intel.com/en-us/articles/intel-digital-random-number-generator-drng-software-implementation-guide>
- [43] G. Maisuradze and C. Rossow, “ret2spec: Speculative Execution Using Return Stack Buffers,” *CoRR*, vol. abs/1807.10364, 2018. [Online]. Available: <http://arxiv.org/abs/1807.10364>
- [44] —, “Speculose: Analyzing the Security Implications of Speculative Execution in CPUs,” *arXiv preprint arXiv:1801.04084*, 2018.
- [45] S. Mangard, “A simple power-analysis (SPA) attack on implementations of the AES key expansion,” in *International Conference on Information Security and Cryptology*. Springer, 2002, pp. 343–358.
- [46] C. Maurice, C. Neumann, O. Heen, and A. Francillon, “C5: cross-cores cache covert channel,” in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 2015, pp. 46–64.
- [47] A. J. Menezes, P. C. Van Oorschot, and S. A. Vanstone, *Handbook of applied cryptography*. CRC press, 2018.
- [48] T. S. Messerges, E. A. Dabbish, and R. H. Sloan, “Examining smart-card security under the threat of power analysis attacks,” *IEEE transactions on computers*, vol. 51, no. 5, pp. 541–552, 2002.
- [49] P. Michaud, “Demystifying multicore throughput metrics,” *IEEE Computer Architecture Letters*, vol. 12, no. 2, pp. 63–66, 2012.
- [50] P. Michaud, A. Sez nec, and R. Uhlig, “Trading conflict

- and capacity aliasing in conditional branch predictors,” in *ACM SIGARCH Computer Architecture News*, vol. 25, no. 2. ACM, 1997, pp. 292–303.
- [51] A. Milburn, K. Sun, and H. Kawakami, “You Cannot Always Win the Race: Analyzing the LFENCE/JMP Mitigation for Branch Target Injection,” *arXiv preprint arXiv:2203.04277*, 2022.
- [52] H. Naghibijouybari and N. Abu-Ghazaleh, “Covert Channels on GPGPUs,” *Computer Architecture Letters*, 2016.
- [53] S. B. Ors, F. Gurkaynak, E. Oswald, and B. Preneel, “Power-Analysis Attack on an ASIC AES implementation,” in *Information Technology: Coding and Computing, 2004. Proceedings. ITCC 2004. International Conference on*, vol. 2. IEEE, 2004, pp. 546–552.
- [54] D. A. Osvik, A. Shamir, and E. Tromer, “Cache attacks and countermeasures: the case of AES,” in *Cryptographers Track at the RSA Conference*. Springer, 2006, pp. 1–20.
- [55] C. Percival, “Cache missing for fun and profit,” 2005.
- [56] A. Prout, W. Arcand, D. Bestor, B. Bergeron, C. Byun, V. Gadepally, M. Houle, M. Hubbell, M. Jones, A. Klein *et al.*, “Measuring the Impact of Spectre and Meltdown,” *arXiv preprint arXiv:1807.08703*, 2018.
- [57] A. Purnal, L. Giner, D. Gruss, and I. Verbauwhede, “Systematic Analysis of Randomization-based Protected Cache Architectures,” in *42th IEEE Symposium on Security and Privacy*.
- [58] M. K. Qureshi, “CEASER: Mitigating Conflict-Based Cache Attacks via Encrypted-Address and Remapping,” in *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Oct 2018, pp. 775–787.
- [59] M. K. Qureshi, “New attacks and defense for encrypted-address cache,” in *2019 ACM/IEEE 46th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2019, pp. 360–371.
- [60] M. Raab and A. Steger, ““Balls into Bins” - A Simple and Tight Analysis,” in *Proceedings of the Second International Workshop on Randomization and Approximation Techniques in Computer Science*, ser. RANDOM ’98. Berlin, Heidelberg: Springer-Verlag, 1998, pp. 159–170. [Online]. Available: <http://dl.acm.org/citation.cfm?id=646975.711521>
- [61] H. Raj, R. Nathuji, A. Singh, and P. England, “Resource management for isolation enhanced cloud services,” in *Proceedings of the 2009 ACM workshop on Cloud computing security*. ACM, 2009, pp. 77–84.
- [62] M. Ramsay, C. Feucht, and M. H. Lipasti, “Exploring efficient SMT branch predictor design,” in *Workshop on Complexity-Effective Design, in conjunction with ISCA*, vol. 26, 2003.
- [63] X. Ren, L. Moody, M. Taram, M. Jordan, D. M. Tullsen, and A. Venkat, “I See Dead μ ops: Leaking Secrets via Intel/AMD Micro-Op Caches.”
- [64] C. Sakalis, S. Kaxiras, A. Ros, A. Jimborean, and M. Sjalander, “Efficient invisible speculative execution through selective delay and value prediction,” in *2019 ACM/IEEE 46th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2019, pp. 723–735.
- [65] B. Saltaformaggio, D. Xu, and X. Zhang, “Busmonitor: A hypervisor-based solution for memory bus covert channels,” *Proceedings of EuroSec*, 2013.
- [66] M. Schwarz, M. Schwarzl, M. Lipp, J. Masters, and D. Gruss, “Netspectre: Read arbitrary memory over network,” pp. 279–299, 2019.
- [67] A. Seznec, “TAGE-SC-L Branch Predictors Again,” 2016.
- [68] N. A. Simakov, M. D. Innus, M. D. Jones, J. P. White, S. M. Gallo, R. L. DeLeon, and T. R. Furlani, “Effect of Meltdown and Spectre Patches on the Performance of HPC Applications,” *arXiv preprint arXiv:1801.04329*, 2018.
- [69] W. Song, B. Li, Z. Xue, Z. Li, W. Wang, and P. Liu, “Randomized last-level caches are still vulnerable to cache side-channel attacks! But we can fix it,” in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 955–969.
- [70] C. Trippel, D. Lustig, and M. Martonosi, “Meltdown-Prime and SpectrePrime: Automatically-Synthesized Attacks Exploiting Invalidation-Based Coherence Protocols,” *arXiv preprint arXiv:1802.03802*, 2018.
- [71] V. Varadarajan, T. Ristenpart, and M. M. Swift, “Scheduler-based Defenses against Cross-VM Side-channels.” in *USENIX Security Symposium*, 2014, pp. 687–702.
- [72] I. Vougioukas, N. Nikoleris, A. Sandberg, S. Diestelhorst, B. M. Al-Hashimi, and G. V. Merrett, “BRB: Mitigating Branch Predictor Side-Channels.” in *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2019, pp. 466–477.
- [73] Z. Wang and R. B. Lee, “New Cache Designs for Thwarting Software Cache-based Side Channel Attacks,” in *Proceedings of the 34th Annual International Symposium on Computer Architecture*, ser. ISCA ’07. New York, NY, USA: ACM, 2007, pp. 494–505. [Online]. Available: <http://doi.acm.org/10.1145/1250662.1250723>
- [74] —, “A novel cache architecture with enhanced performance and security,” in *Proceedings of the 41st annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Computer Society, 2008, pp. 83–93.
- [75] W. Xiong and J. Szefer, “Survey of transient execution attacks and their mitigations,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 3, pp. 1–36, 2021.
- [76] T.-Y. Yeh and Y. N. Patt, “Two-level adaptive training branch prediction,” in *Proceedings of the 24th annual international symposium on Microarchitecture*. ACM, 1991, pp. 51–61.
- [77] J. Yu, M. Yan, A. Khyzha, A. Morrison, J. Torrellas, and C. W. Fletcher, “Speculative Taint Tracking (STT): A Comprehensive Protection for Speculatively Accessed

- Data,” in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, ser. MICRO '52. New York, NY, USA: Association for Computing Machinery, 2019, p. 954968. [Online]. Available: <https://doi.org/10.1145/3352460.3358274>
- [78] T. Zhang, K. Koltermann, and D. Evtushkin, “Exploring Branch Predictors for Constructing Transient Execution Trojans,” in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS 20. New York, NY, USA: Association for Computing Machinery, 2020, p. 667682. [Online]. Available: <https://doi.org/10.1145/3373376.3378526>
- [79] W. Zhang, Z. Bao, D. Lin, V. Rijmen, B. Yang, and I. Verbauwhede, “RECTANGLE: a bit-slice lightweight block cipher suitable for multiple platforms,” *Science China Information Sciences*, vol. 58, no. 12, pp. 1–15, Dec 2015. [Online]. Available: <https://doi.org/10.1007/s11432-015-5459-7>
- [80] W. Zhang, Z. Bao, V. Rijmen, and M. Liu, “A New Classification of 4-bit Optimal S-boxes and Its Application to PRESENT, RECTANGLE and SPONGENT,” in *Fast Software Encryption*, G. Leander, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015, pp. 494–515.
- [81] L. Zhao, P. Li, R. Hou, M. C. Huang, J. Li, L. Zhang, X. Qian, and D. Meng, “A lightweight isolation mechanism for secure branch predictors,” in *2021 58th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2021, pp. 1267–1272.