

Research Report

Let's Not Speculate: Discovering and Analyzing Speculative Execution Attacks

Andrea Mambretti¹, Matthias Neugschwandtner², Alessandro Sorniotti², Engin Kirda¹,
William Robertson¹, Anil Kurmus²

¹Northeastern University
Boston, MA
USA

²IBM Research – Zurich
8803 Rüschlikon
Switzerland

LIMITED DISTRIBUTION NOTICE

This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies (e.g., payment of royalties). Some reports are available at <http://domino.watson.ibm.com/library/Cyberdig.nsf/home>.



Research

Africa • Almaden • Austin • Australia • Brazil • China • Haifa • India • Ireland • Tokyo • Watson • Zurich

Let's Not Speculate: Discovering and Analyzing Speculative Execution Attacks

Andrea Mambretti*, Matthias Neugschwandtner†, Alessandro Sorniotti†, Engin Kirda*,

William Robertson*, Anil Kurmus†

*Northeastern University

{mbr, ek, wkr}@ccs.neu.edu

†IBM Research

{eug, aso, kur}@zurich.ibm.com

Abstract—Speculative execution attacks exploit vulnerabilities at a CPU’s microarchitectural level, which, until recently, remained hidden below the instruction set architecture, largely undocumented by CPU vendors. New speculative execution attacks are released on a monthly basis, showing how aspects of the so-far unexplored microarchitectural attack surface can be exploited. In this paper, we generalize speculative execution related attacks and identify common components. The structured approach that we employed helps us to identify potential new variants of speculative execution attacks. We explore one such variant, SPLIT SPECTRE, in depth and demonstrate its applicability to a real-world scenario with the SpiderMonkey JavaScript engine. Further, we introduce SPECULATOR, a novel tool to investigate speculative execution behavior critical to these new microarchitectural attacks. We also present our findings on multiple CPU platforms.

I. INTRODUCTION

A developer’s view of the CPU when writing a low-level program is defined by the CPU’s instruction set architecture (ISA). The ISA is a well-defined, stable interface the developer can use to access, and change the architectural state of a CPU. The software is in full control over memory, registers, interrupts and I/O. At the same time, the CPU has a lower-level state of its own – the extra-architectural state of the microarchitecture, commonly referred to as the *microarchitectural state*. In general, the ISA provides no direct access to the CPU microarchitecture, allowing the microarchitecture to evolve independently, while keeping the programming interface stable. The microarchitecture of a CPU is subject to frequent changes and is different among vendors. A CPU’s microarchitecture typically also implements security controls, such as process isolation.

Recent works [46], [25], [29] have shown how security controls can be bypassed by submitting carefully crafted inputs at the level of the ISA interface. These attacks exploit undocumented behavior at the microarchitectural level, and have been discovered through reverse engineering and trial-and-error. The full breadth of this class of attacks is not entirely understood, owing to the fact that details about the microarchitectural level of modern commercial CPUs are not publicly available. The research community cannot provide complete answers to questions about the existence of new attacks and the effectiveness of defenses.

We propose a two-pronged strategy towards this objective: the first prong is a technique to classify this family of attacks

that we refer to as *Speculative Execution Attacks* (SEAs). The classification serves to highlight the different phases of the attacks, what instructions or sequence thereof are involved at each phase, which privilege level is requested/involved, and which principal requests its execution (victim or attacker). A Cartesian product of the possible variants of all constituent parts can then be used to construct a set of candidate attacks. Verifying whether a candidate attack is viable is not straightforward. Owing to the aforementioned lack of documentation, the verification step must be an empirical one – the attack must be prototyped and its effectiveness validated against a test machine. A negative test result does not constitute a sufficient reason to discard the candidate, as the failure might be attributable to a different set of undocumented microarchitectural aspects. This phase can be aided if the appropriate tooling and automation is available. Unfortunately, none is available to date. This observation justifies the second prong of our strategy, a tool that we call SPECULATOR. This tool supports the validation of the candidate attack set, and the discovery of undocumented microarchitectural features that influence their outcome.

In this paper, we describe our classification approach and SPECULATOR in detail. We show the effectiveness of these approaches by describing their usage to reveal a new variant of Spectre v1 [25], which we call SPLIT SPECTRE. SPLIT SPECTRE requires a smaller piece of vulnerable code available in the victim’s attack surface compared to the original attack, making it a potentially pernicious vulnerability.

Our paper makes the following contributions:

- A new approach of classifying and decomposing speculative execution attacks to understand existing SEAs and identify new ones.
- A new performance counter-based method and tool, SPECULATOR, to empirically verify candidate SEAs.
- A novel variation on Spectre v1, SPLIT SPECTRE, requiring a smaller piece of vulnerable code available in the victim’s attack surface.

II. DISSECTING SPECULATIVE EXECUTION ATTACKS

Speculative execution attacks (SEAs) exploit a new class of vulnerabilities, targeting a particular microarchitectural CPU design with specially crafted software. These attacks leverage known attack vectors such as side channels, but go much

further by combining them with vulnerabilities at the micro-architectural level. Numerous variants of SEAs have been disclosed since the beginning of 2018. In this section, we propose a general definition and analysis of SEAs with the aim of clearly distinguishing SEA variants in order to motivate and guide the analysis of new attacks and defenses in this area.

Before delving into the dissection of SEAs, we need to distinguish SEAs from the more general category of out-of-order execution attacks. Spectre v1 and v2 [46], [25] are the first discovered SEAs, with Spectre v1.1 [24], Spectre v4 [21], NetSpectre [7] and Netspectre-AVX being follow-ups. In contrast, attacks such as Meltdown [29], Spectre v3a [8], Foreshadow [40] and Foreshadow-NG [43] do not rely on speculative execution behavior, and may be classified in the more general category of out-of-order execution attacks. Since this paper focuses on analyzing speculative execution, we have opted to leave them out.

A. Attack scenarios, Privilege boundaries

SEAs, much like side channel attacks, can be performed in a variety of scenarios involving one victim and one attacker thread. The notion of *thread* here is in the general, hardware-related sense (e.g. VMM thread, guest thread, (un)-sandboxed thread, or user/kernel thread). These two threads run with different privileges, with the attacker thread typically running with a lower privilege. There can also be scenarios where both threads are at the same privilege level, but have access to different data. In all cases, however, a boundary separating attacker and victim contexts resides between the two threads.

In addition, the two threads are in temporal or spatial co-residence, as well as in spatial and temporal proximity. An example of spatial co-residence is two threads running on the same hardware thread, one closely after the other (i.e. temporal proximity) – even on a machine not supporting simultaneous multi-threading (SMT). An example of temporal co-residence is two SMT threads (hyperthreads in Intel nomenclature) running at the same time, influencing each other. These are in spatial proximity because they run on the same core.

B. SEA Phases

For our generalization of SEAs, we decompose them into four phases, and describe how existing and new attacks fit into these categories.

- ❶ Prepare side channel: In this phase, the CPU performs operations that will increase the chances of the attack succeeding. For instance, the attacker can prime caches to prepare for a prime-and-probe [38] cache side channel measurement, make sure important target data is flushed, or ensure that the attacking thread and victim thread are co-located.
- ❷ Prepare speculative execution: In this phase, the CPU executes code that will allow speculative execution to start. This is code that is typically executed within the context of the victim.
- ❸ Speculative execution start: In this phase, the CPU executes an instruction whose outcome decides the next instruction to be executed, such as a conditional branch

instruction. Between the time window where this instruction is issued and when it is retired, modern CPUs guess the outcome of the branch to avoid stalling the pipeline, and execute code speculatively. This is known as speculative execution [27].

- ❹ Speculative execution, side channel send: In this phase, the CPU executes (but not necessarily retires) instructions that will result in a micro-architectural state change.
- ❺ Side channel receive: In this phase, the CPU executes instructions that transform the micro-architectural state change that occurred in the previous step into an architectural state change.

C. Privilege boundaries and attack impact

The core element that turns speculative execution into an attack is the breach of a privilege boundary that is established through hardware isolation support by the CPU. These privilege boundaries typically aim to provide confidentiality and integrity of the data residing within the boundary (i.e. preventing data from being read or modified directly from outside the boundary). All accesses to such data are mediated by code running within the privilege boundary, and that code may only be invoked from a lower privilege through well-defined entry points.

In the case of currently known SEAs, the attacker’s aim is limited to breaching confidentiality of data residing beyond the privilege boundary by either accessing arbitrary data or leaking specific metadata, such as pointer values, of the running program. In addition, the currently known privilege boundaries that can be bypassed by speculative execution are:

- kernel vs. user-mode code
- hardware enclave (SGX) vs. user-mode or kernel-mode code
- sandboxed code in the same process, for example JavaScript JIT code
- processes-to-process boundary
- remote node to local node boundary

We note that code at each SEA phase previously described can potentially be run either in the higher privileged mode (victim-provided code) or lower privileged one (attacker-provided code). We show later in the paper that this insight leads to a new Spectre v1 variant.

D. Classification of vulnerabilities

We now qualify existing attacks according to our SEA categorization. In Table I, we specify for each phase of each attack the type of instruction processed by the CPU, and whether that instruction is running in high-privilege (⊙) or low-privilege mode (⊕). For instance, the Spectre v2 BPF-based exploit [46] uses a prime-and-probe side channel (phases ❶ and ❺) and trains the branch target buffer (phase ❷). All these phases are performed by attacker-provided, low-privilege code.

This way of categorizing SEA attacks can be used to reveal potential new variants, either by finding another type of instruction for a given phase, by combining two different

variants, or by switching the required privilege level. To illustrate: Netspectre uses an evict-and-reload [7] strategy for its side channel, which could well be adapted to v1, v2, v1.1, RSB and for NetSpectre-AVX. With this in mind, there may be little benefit in doing so, as evict-and-reload side channels tend to be noisier than prime-and-probe side channels. Similarly, one can think of other types of side channels that can be used in SEAs.

Beyond simply adapting a new side channel for existing attacks, we find that the privilege level for ④ in Spectre v1 can be switched, resulting in a new variant which we call SPLITSPECTRE and examine in more detail in the following section. We focus on a v1-based example in this paper as we found it to be the most promising. However the same idea can be equally applied to Spectre v2, Spectre v1.1, and Spectre RSB attacks.

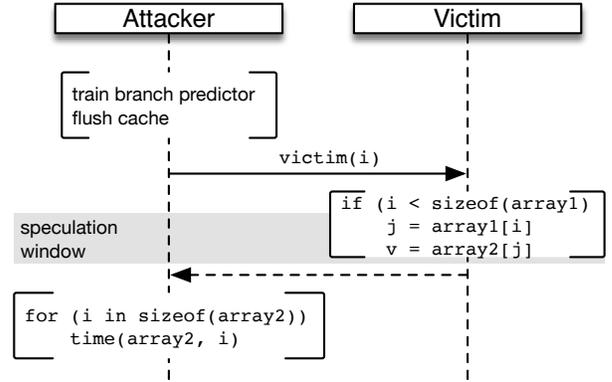
E. Running example and new attack: SPLITSPECTRE

In Spectre v1, the victim code that is executed speculatively (“gadget”) consists of three components: *i*) a conditional branch on a variable, typically a length check, *ii*) a first array access that uses the variable from the conditional branch as an offset, and *iii*) a second array access that uses the result of the first array access as an offset. If the conditional branch triggers speculative execution of the following array accesses (phase ③), the first array access may access an out-of-bounds memory region, revealing the contents of this region through a side channel (phase ④) by measuring the access time to the second array after executing the gadget (phase ⑤).

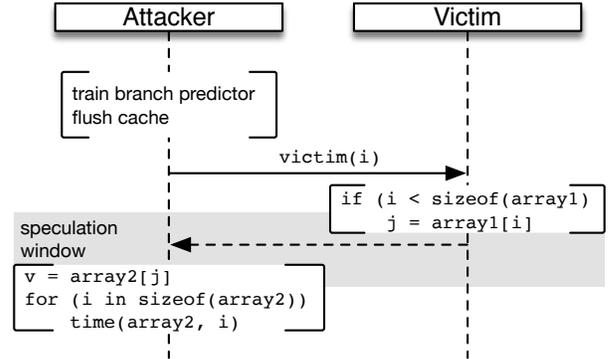
Although Spectre v1 is powerful and does not rely on SMT, it requires such a gadget to be present in the victim’s attack surface. Google Project Zero writes in their original blog post on Spectre v1 [46] that they could not identify such a vulnerable code pattern in the kernel, and instead relied on eBPF to place one there themselves.

In this point lies the strength of our new Spectre v1 variant, SPLITSPECTRE. As its name implies, it splits the Spectre v1 gadget into two parts: one consisting of the conditional branch and the array access (phase ③), and the other one consisting of the second array access that constitutes the sending part of the side channel (phase ④). This has the advantage that the second part, phase ④, can now be placed into the attacker-controlled code. It is more likely that an attacker finds such gadgets, thereby alleviating one of the main difficulties of performing a v1 attack. Furthermore, the attacker can choose to employ amplification of a v1 attack by reading multiple indices of the second array to deal with imprecise time sources.

Figure 1 compares the regular Spectre v1 with our split version. As shown in the figure, the speculation window needs to be sufficiently large such that it still covers the second part. We define the *speculation window* (short for speculative execution window) as the time interval between the event that triggers speculative execution, e.g. a branch condition, and the point in time when it is resolved and the speculatively executed instructions are either retired or rolled back. The speculation window is measured in cycles and determines how many instructions of a given sequence are speculatively



(a) Regular Spectre v1. The gadget requires two dependent array accesses in the victim’s attack surface.



(b) Split Spectre v1. The second, dependent array access from a regular v1 gadget moves to the attacker code.

Fig. 1: A comparison of regular Spectre v1 and SPLITSPECTRE. While SPLITSPECTRE only requires a simple array access, the speculation window needs to be sufficiently large to contain both the gadget and the second array access exercised by the attacker.

executed. The number of instructions of a given sequence that can be speculatively executed at a given time also depends on the CPU’s microarchitecture. For example, some instructions are more “expensive” in the sense that they are split into a number of μ ops, and thus take a long time to execute. Also, the combination of instructions in a sequence affects how fast they execute: similar instructions might lead to congestion on the execution ports, as they require similar execution units.

The speculation window caps the maximum number of instructions executed between the two parts. Extending the length of the speculation window is an instrumental part in extending the capabilities of a speculative execution attacker and the reach of a SPLITSPECTRE attack. In the course of the paper, we show how we use SPECULATOR to evaluate SPLITSPECTRE and speculative execution aspects relevant to its feasibility.

III. SPECULATOR

Speculative execution is not well-documented compared to other features of modern CPUs. Being part of the microarchitecture, its implementation details are hidden behind the ISA and subject to optimization, which manufacturers keep to themselves.

Attack	① Prepare SC	② Prepare SE	③ SE start	④ SE SC send	⑤ SC receive
Spectre v1 [25], [46]	Prime caches ●	Train branch predictor ⊙	Compare ⊙	Load ⊙	Probe caches ●
Spectre v2 [25], [46]	Prime caches ●	Branch target injection ●	Indirect branch ⊙	Load ⊙	Probe caches ●
Spectre v1.1 [24]	Prime caches ●	Train branch predictor ⊙	Compare ⊙	Load† ⊙	Probe caches ●
Spectre RSB [26]	Prime caches ●	Poison RSB ●	Return ⊙	Load ⊙	Probe caches ●
NetSpectre [7]	Evict caches ⊙	Train branch predictor ⊙	Compare ⊙	Load ⊙	Transmit gadget: reload ⊙
NetSpectre-AVX [7]	Reset AVX ⊙	Train branch predictor ⊙	Compare ⊙	AVX ⊙	Transmit gadget: AVX ⊙
SPLITSPECTRE (new)	Prime caches ●	Train branch predictor ⊙	Compare ⊙	Load ●	Probe caches ●

Table I: Classification of Speculative Execution Attacks. ●: code provided by the attacker, running in low privilege; ⊙: code provided by the victim, running in high privilege; SE: speculative execution; SC: side channel; † reached through speculative buffer overflow, attacker chosen code.

However, understanding the internals of speculative execution is key to comprehending the limits of SEAs, and to designing adequate mitigations and defenses against SEAs. For this reason, we have designed and implemented SPECULATOR, a tool whose purpose is to reverse-engineer the behavior of different CPUs in order to build a deeper understanding of speculative execution. SPECULATOR aggregates the relevant sources of information available to an observer of speculative execution, chief among them CPU performance counters and model-specific registers, so that the behavior of different code snippets can be observed from a speculative execution standpoint. In this section, we describe the design and implementation of SPECULATOR.

A. Performance Monitor Capabilities

Modern CPUs provide relevant information through the performance counter interface. This interface is offered by most manufacturers, and it exposes a set of registers (some fixed and some programmable) that can be used to retrieve information on various aspects of the execution. Through these registers, counters for events or duration related to microarchitectural state changes such as cache accesses, retired instruction, and mispredicted branches, are made available to the developer. Events are manufacturer- and architecture-specific. This interface was originally made available to provide a method for developers to improve the performance of their code. The interface is typically used as follows: through a setup step, developers can choose which events will be measured by programmable counters out of a wide set of supported ones. Measurements can be started and stopped programmatically in order to carefully control the events of which precise sequence of instructions is being measured. Setting up, starting, and stopping measurements often requires supervisor mode (ring 0 in x86 nomenclature) instructions, whereas accessing counters is usually available in user mode.

SPECULATOR builds on top of performance counters to observe the nature and effects of speculative execution. One challenge with this approach is that the performance counters interface was not designed with this objective in mind. One of the contributions of this paper is the identification of effective ways of using the interface, and a useful set of counters to accurately infer the behavior of speculative execution.

B. Objectives

The main objective of SPECULATOR is to accurately measure microarchitectural state attributes associated to the speculative portion of the execution of user-supplied snippets of code. Accuracy refers to the degree with which the tool is capable of isolating the changes to the microarchitectural state caused by the snippet being analyzed from that of the tool itself and the rest of the system (e.g. the OS or other processes). An incomplete list of SPECULATOR observables are 1) which parts of the snippet are speculatively executed, 2) what causes speculative execution to start and stop, 3) what parameters affect the amount of speculative execution, 4) how do specific instruction affect the behavior of speculative execution, 5) which security boundaries are effective in the prevention of speculative execution, and 6) how consistently CPUs behave within the same architecture and across architectures and vendors. The creation of a new tool is justified because none of the existing ones, such as perf_events [14] or Likwid [35], provide the required information with sufficient accuracy.

Perf_events has two modes of operations, sampling and counting. During sampling, there is no way to have precise quantitative information about code execution, and therefore it is not suitable for our purpose. When evaluating perf_events' counting mode, we experienced for very small snippets a certain level of overhead (in the order of 500 μ ops). This overhead was caused by the perf_event design decision of integrating all its operations (e.g. start counters, stop counters) in the kernel. Since the test snippets are 20-30 instructions long on average, this overhead completely prevents inferring any kind of relevant behavior.

Likwid operates instead in user space just as SPECULATOR, instrumenting the counters through the MSR register. However, its design only allows system-wide measurements and does not provide the same flexibility of handling the counter as the snippet progresses in its execution.

We also considered other tools and libraries such as Oprofile [28], Perfmon2 [16], Perfctl [33], and PAPI [36]. Unfortunately, all of these possess either the same issues of measure inaccuracy or lack of flexibility, or otherwise are outdated and unmaintained. Performance comparisons among some of these interfaces are provided by Zapanuks et al. [45] and Weaver [42].

Another SPECULATOR objective is to provide tooling for the generation and manipulation of code snippets. The ability to inspect individual snippets and snippet groups during

speculative execution gives the user the ability to focus on combinations of instructions that are relevant for specific use-cases. For instance, during our tests, there were cases such as inspecting the behavior of `clflush` where we were interested mainly in a single snippet behavior. Meanwhile, in other tests, we were interested in how the behavior of a sequence of instructions changes with variations in the back-end load. In this scenario, we were interested in how the measures varied between a snippet and the following.

Additionally, support for multiple platforms enables the inference of general facts about speculative execution.

C. Design and Implementation

Figure 2 describes the architecture of SPECULATOR and its three main components: a pre-processing unit, a runtime unit called the Monitor, and a post-processing unit.

The task of the pre-processing unit is to compile the provided input into the appropriate execution format, and to introduce the instrumentation required by the performance monitor interface to be able to observe the value of the selected set of hardware counters. Input can be provided as a snippet of C or assembly code, or as a template for the generation of code snippets. Code snippets are generated from templates in an incremental fashion, resulting in the output of multiple snippets with an increasing number of instructions taken from a pre-compiled JSON list. Each instruction is inserted by the SPECULATOR snippet generator in the specific location defined in the source template (Step 1 in Figure 2). The introduction of such “incremental” snippets is justified by the fact that the addition of a single assembly instruction may trigger optimizations that – while preserving the expected program semantics – alter the behavior of the CPU at a microarchitectural level and affect the nature of speculative execution. Having incremental snippets helps to verify when optimizations are triggered and take them into account during the analysis of the results.

After the generation of the executable (also referred to as the test application), the SPECULATOR runtime is invoked on each of the generated outputs (Step 2). To ensure that the Monitor does not perturb the measurements, the process executing the snippet and the monitor are pinned on different cores. Monitor is responsible to configure the counters on the core used by the test application (Step 3). As previously mentioned, there are many programmable counters that can be used so we provide a configuration file that can be loaded into SPECULATOR to easily switch among them.

Once the Monitor has set up the environment, it loads and executes the snippet in a separate process, and waits for it to complete (Step 4). The test application prologue and epilogue will interact with the environment created by the Monitor, resetting, starting and stopping the counters as needed. The counters related to the core where the test application runs are stopped by the test application just before termination. When the test application terminates, the Monitor will be signaled by the Operating System. At this point, the Monitor can retrieve the values of the counters from the core where the test application runs (Step 6) and store them in a result file.

The Monitor can be configured to run a specific test N times. In this case, the result file will contain the values of each run.

Once the tests results are collected from the Monitor, they are handed to the post-processing unit (Step 7). This unit aggregates the results from multiple runs by computing statistics (e.g. mean and standard deviation) and by removing clear outliers.

D. Triggering Speculative Execution

The user of SPECULATOR supplies as input a code snippet to determine how the CPU behaves when speculative execution takes place. We note that in the absence of branch misprediction, instructions that are speculatively executed will eventually retire and there should be no undesired microarchitectural side-effects. The more interesting case for the SPECULATOR user is a snippet containing a branch, or other speculative execution trigger, that the CPU does not predict accurately, leading to the speculative execution of instructions that will not retire. In this scenario, SPECULATOR helps the user detect which instructions the CPU executed and how they influenced the microarchitectural state.

In order to automate the generation of test cases, SPECULATOR provides the user with a template, described in Figure 3. The template is used as follows: the user supplies a snippet, expecting i) it to be speculatively executed, ii) that none of its instructions will retire, and iii) that SPECULATOR will report counters relating to its execution. In order to achieve this, the template prefixes the snippet supplied by the user with a branch instruction. The template begins with a setup step that aims to train the branch predictor not to take that branch. After the branch predictor is trained not to take the branch, the program state is set to require the branch to be taken to ensure that the snippet will be speculatively executed and that none of its instructions will retire. The template then starts the performance counters that were previously setup by the Monitor and executes the branch, after which it stops performance counters. In order to prolong or shorten the speculative execution of the user snippet, the condition variable of the branch can be placed in registers or memory. On the microarchitectural level, a variable placed in memory can also be cached in one of the levels of the cache hierarchy.

E. Speculative Execution Markers

In the context of SPECULATOR we are mostly interested in determining the behavior of the CPU when instructions that are speculatively executed do not retire. A first natural question is whether non-retired instructions were speculatively executed at all and, if so, how many of them. An accurate detection of these events is (perhaps surprisingly) not trivial. Indeed, the CPU strives to undo most observable architectural side-effects from non-retired speculatively executed instructions. However, as we know from the Spectre and Meltdown works [25], [29], not all side effects are undone. One possible approach to detect non-retired speculative execution would be to rely on the side-channels exploited in these works. This approach has several shortcomings: it has a relatively low single-run

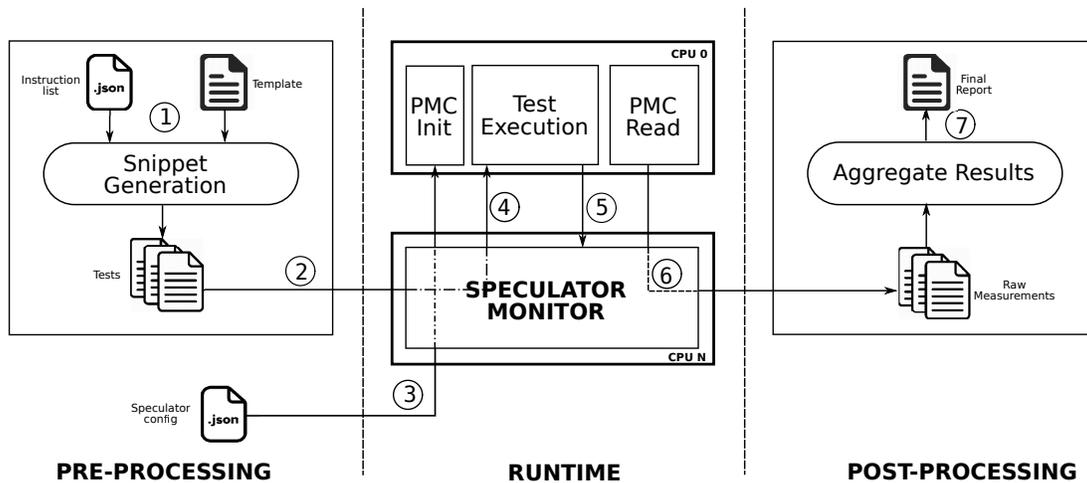


Fig. 2: The architecture of SPECULATOR. A template with the speculative execution trigger and a list of instructions to be speculatively executed are the input to the code generation. The code snippets are run repeatedly under supervision of the speculator monitor, which captures the event specified in the configuration file. Finally, the measurements are post-processed to present a final report on speculative execution behavior.

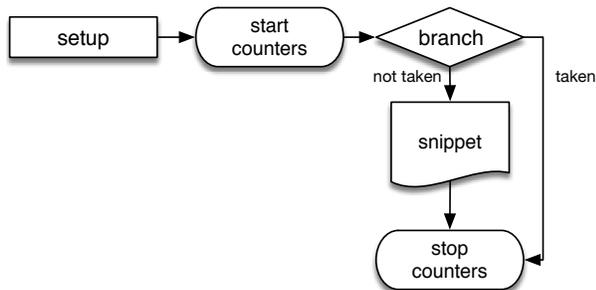


Fig. 3: Flow chart of the experiment template that is used in SPECULATOR. The setup code brings the branch predictor in a specific state that will cause the later branch to mispredict and speculatively execute the code snippet consisting of the instructions. The speculative execution of the instructions is measured by the PMC infrastructure, which is triggered by the corresponding start/stop instructions indicated in the flow chart.

detection accuracy, it is costly to setup and read, and it requires otherwise unnecessary changes to program observables.

A more effective approach is based on markers of speculative execution, that is, special instructions or sequences thereof (which we will refer to as markers) that are detectable by performance counters even when they do not retire. The approach requires appending the marker to the snippet which is fed as input to SPECULATOR, and ensuring that there is no other occurrence of the marker in the snippet. If SPECULATOR detects the marker, the detection can be used as proof that the CPU executed the snippet.

The choice of which markers to use is manufacturer- and architecture-specific, given that not all CPUs expose the same set of counters. In general, the marker must cause a microarchitectural event that is detectable by a performance counter irrespective of its retired status. For example, counters that measure *issued* or *executed* instructions of a specific type irrespective of their retired status constitute a good marker. The selection of which counter to use on a given architecture

requires manual inspection of the CPU architecture programmer’s manual. In what follows, we report our findings on the available markers for Intel processors:

`UOPS_EXECUTED.CORE/THREAD` counts the number of μ ops executed by the CPU. It can be used to report the exact number of μ ops that were executed out of the user-supplied snippet by subtracting the number of μ ops that retire in the template (the branch and the instrumentation to stop performance counters) from the output value of the counter. This counter is subject to μ -fusion of instructions and does not count instructions that do not require execution such as NOP. An exception to that rule is `FNOP`, which is tracked by this counter as well.

`UOPS_ISSUED.SINGLE_MUL` belongs to a group of counters triggered only by a specific set of instructions. This counter is fired whenever a single-precision floating-point instruction that operates on the `XMM` register is issued. This means that such an operation can be inserted at the end of the user-supplied snippet to verify whether this counter is incremented or not. This counter has been dropped by Intel on most recent CPUs (e.g. Skylake) and therefore its usage is limited across platforms.

Similarly to `UOPS_ISSUED.SINGLE_MUL`, `UOPS_ISSUED.SLOW_LEA` is triggered by only a specific set of instructions. It counts `LEA` instructions with three source operands (e.g. `lea rax, [array+rax*2]`). Unfortunately, certain operations such as `clflush` are considered by the CPU as `SLOW_LEA` operations, so extra care must be taken to subtract any number of those present outside of the user-supplied snippet.

`LD_BLOCKS.STORE_FORWARD` is incremented for each store forward that result in a failure. An example of a sequence that triggers this kind of situation is shown in Listing 1.

The following markers are available on the AMD Zen architecture:

`DIV_OP_COUNT`, counting the number of executed `div`

Listing 1: Failed store forward example

```

1 mov DWORD[array], eax
2 mov DWORD[array+4], edx
3 movq xmm0, QWORD[array]

```

Architecture	CPU	Design
Intel Haswell	i5-4300U	tock
Intel Broadwell	i5-5250U	tick
Intel Skylake	i7-6700K	tock
AMD Zen	Ryzen 1700	

Table II: The CPUs per architecture we use SPECULATOR on. While Haswell and Skylake are new designs – “tocks” in Intel nomenclature – Broadwell is a “tick”, a die-shrink of Haswell.

instructions.

NUMBER_OF_MOVE_ELIMINATION_AND_SCALAR_OP_OPTIMIZATION, like LD_BLOCKS_STORE_FORWARD, does not track the execution of an instruction, but rather the effect of a certain instruction sequence. In this case, it tracks in how many cases move elimination was successful.

IV. USING SPECULATOR: THE EXAMPLE OF SPLIT SPECTRE

To find out whether the SPLIT SPECTRE attack can be exploited in practice, we use SPECULATOR to investigate several speculative execution properties related to it. This also serves as an example of the study of an attack aided by the SPECULATOR tool. The results that we uncover are applicable to SPLIT SPECTRE, and are also of independent interest. Since some of our findings are hardware-dependent, we also show the differences based on the underlying CPU architecture (Table II).

A. Out-of-order execution bandwidth

Speculative execution is no different in how it uses the resources available in both the front- and the back-end of a CPU compared to regular execution. On Intel platforms, instructions that have been fetched and decoded into μ ops by the frontend are entered in the reorder buffer of the backend. This buffer contains all μ ops that are currently “in flight”, which means they are either ready for execution, are currently being executed, or have finished execution. The buffer’s name derives from the fact that on modern CPUs μ ops are executed *out-of-order*. This means they are dispatched to execution units based on their data flow dependencies, rather than the control flow of the program. After being executed, they remain in the reorder buffer until they are retired. Retirement of μ ops happens at an assembly-instruction granularity and *in-order*, honoring the control flow of the program. When μ ops are retired, the outcome of their computation is committed to the program’s state.

The size of the reorder buffer is a natural upper bound on the length of a sequence of instructions that can be speculatively executed. That is, the reorder buffer would hold the branch instruction that triggered speculative execution plus

the instructions of the code path being speculatively executed. The branch instruction is the first one that is retired in-order, potentially causing all other μ ops in the buffer to be canceled in case of misprediction. If the branch instruction takes time to retire, e.g. because it depends on a compare that requires a slow memory access, chances are higher that the reorder buffer is filled with μ ops that are speculatively executed than for a branch that retires quickly. If the reorder buffer is full, the whole CPU back-end stalls.

A large reorder buffer is beneficial for SPLIT SPECTRE and most attacks that exploit speculative execution because it lets a larger amount of instructions be speculatively executed, enhancing the capabilities of a speculative execution attacker. While the size of the reorder buffer is typically a known attribute of a CPU, we decided to empirically verify this number to show how precise measurements taken by SPECULATOR are. In our experiment, we use the UOPS_EXECUTED_CORE counter (see Section III-E). Since the counter operates at the granularity of a core, we disable SMT to reduce the noise caused by Hyperthreads that are scheduled on the same core. We also use the BR_MISP_RETIRED counter, which counts the number of mispredicted, retired branch instructions.

When relying on the count of executed μ ops to measure the reorder buffer size, we need to keep in mind that the μ ops actually need to execute before the branch that triggered speculative execution is retired. This means we need instructions that execute quickly to achieve maximum throughput. Since “regular” instructions would easily saturate the available execution ports and units, we pick the NOP instruction. NOP is decoded into a single μ op, which occupies a single slot in the reorder buffer. It does not actually execute and thus neither requires an execution unit nor is it captured by the counter that measures executed μ ops. We thus put an arbitrary regular instruction as a marker at the end of the NOP-sled, increasing the latter in size for each test generated. When running this test with SPECULATOR, we expect to measure a constant amount of μ ops executed up to the point, where the NOP-sled takes up all slots in the reorder buffer and the terminating instruction is no longer speculatively executed. Indeed, the results match our expectation: as can be seen in Figure 4, the number of executed μ ops is constant up until 188 NOPs on Broadwell and 220 NOPs on Skylake. In addition to the NOPs we also need to account for the branch instruction, taking up two slots in the reorder buffer as well as the marker instruction, taking up yet another two entries. In total, this is in line with the specifications published by Intel, which state a reorder buffer size of 192 entries for Broadwell and 224 entries for Skylake.

Interestingly, the number of executed instructions differs for the architectures: it is 34 and 32 for Broadwell and 32 to 30 for Skylake, in spite of the code being exactly the same. Presumably, this is caused by extended μ op-fusion introduced as optimization on Skylake. Fused μ ops count as a single μ op.

AMD’s Zen platform has a construct similar to Intel’s reorder buffer: the retire queue. Every μ op that has entered the backend and not been either retired or canceled takes a slot in this queue. Our Ryzen CPU does not feature a counter for executed μ ops, so we can only provide a measurement based on our marker instruction in this case. The marker instruction,

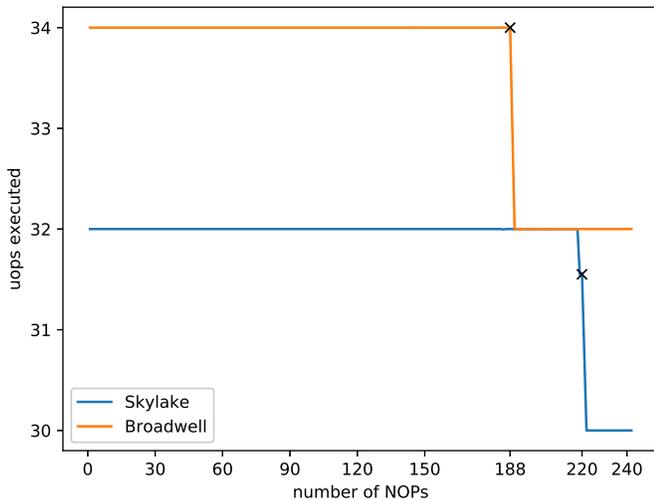


Fig. 4: Reorder buffer size test results on Broadwell and Skylake. Since the marker instruction is no longer executed for a sufficiently large number of NOPs, the number of executed μ ops drops at the size of the reorder buffer.

which takes up four μ ops in this case, is executed up until 186 NOPs. This is in line with the size of the retire queue, which is specified to have 192 entries ($= 186 + 2 + 4$). Interestingly, the speculation window seems to be halved when we switch off SMT: we recognize execution of the marker instruction only up to 91 NOPs.

B. Nesting Speculative Execution

So far, our approach to extend the speculation window was based on placing the value used to evaluate the conditional branch in a memory region that has been flushed from the cache. Next, we use not just one, but multiple conditional branches to investigate whether this results in a longer speculation window. We use SPECULATOR to evaluate the effectiveness of the approach. This experiment has multiple potential outcomes: given two nested branches, an outer and an inner one, either *i*) the inner branch is not speculatively executed until the branch condition on the outer branch is resolved, or *ii*) speculative execution continues to the inner branch and beyond. In the second case, we are interested in the speculative execution behavior if the inner branch is resolved while the result of the outer one is still pending.

We design our experiment with three nested conditional branches, outermost to innermost, with the branch conditions being independent of one another. The conditions are set up with decreasing complexity, such that the outermost will take longest to resolve. We achieve this by involving an uncached value that is subject to multiple expensive operations (`divs`) in the outermost branch condition, a simple uncached value in the middle branch condition, and a cached value in the innermost branch condition. As usual, we train the branch predictor for all branches in the setup phase such that it is going to mispredict all targets in the measurement phase. To evaluate which code paths are (speculatively) executed, we

repeat the experiment multiple times with marker instructions placed in the opposite branch target paths.

We performed this experiment on both Broadwell and Skylake, yielding identical results: in both cases, nested speculative execution takes place, i.e. speculative execution continues along the trained branch targets for all branches. Second, if a nested branch condition is resolved before its parent branch and a misprediction has occurred, speculative execution picks up the opposite branch target. If a parent branch is resolved, all mispredicted code paths, including nested speculative execution, is canceled. For SPLITSPPECTRE this means that while nested speculative execution takes place, it does not widen the speculation window.

C. Speculative execution across system calls

A fundamental aspect of SPLITSPPECTRE is that its two parts are situated on either side of a privilege boundary. As mentioned in Section II-C, one such boundary isolates user from kernel mode. We thus investigate whether speculative execution continues across the context switch from user- to kernel mode. To this end we design a simple test scenario, where the speculatively executed snippet issues a system call. For the system call itself we picked `sys_getppid` because of its low complexity – an execution only amounts to 47 instructions. We use the counter for executed μ ops and tune it to capture either just μ ops executed in user mode or kernel mode.

We performed the experiment on the Broadwell and Skylake microarchitectures with identical results:

- The number of μ ops executed in user mode corresponds to the instructions before the system call and does not increase with additional instructions added after the system call.
- The number of μ ops executed in kernel mode does not increase compared to a baseline measurement taken without speculative execution of the code snippet.

Thus, we conclude that a system call effectively stops speculative execution: it stops after the system call returns from kernel mode. We further conclude that a SPLITSPPECTRE attack across the system call boundary is not feasible on the tested Intel CPUs.

D. SPLITSPPECTRE in SpiderMonkey

Based on the results of our analysis using SPECULATOR, we mounted a SPLITSPPECTRE attack in a real-world setting. We chose a browser-like setting, where untrusted JavaScript is executed in a trusted runtime environment, establishing a privilege boundary. Recall that a V1 gadget consists of a bounds check and two array accesses, the first one using the provided index and the second one using the content of the first array at that position as an index into the second array. In order to mount a regular Spectre V1 attack, we would require a complete Spectre V1 gadget available in the in the JavaScript engine. The intuition behind SPLITSPPECTRE permits us to relax this requirement and only require the first half of a V1 gadget, i.e. the bounds check and the first

array access. The second half of this gadget is provided by attacker-controlled JavaScript code (Figure 5). The attack can only work if speculative execution spans across the privilege boundary from the bounds check in the runtime environment to the second array access in the attacker-controlled, unprivileged code.

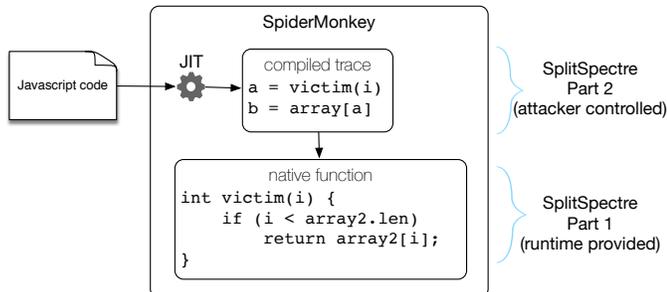


Fig. 5: A conceptual view of a SPLITSPECTRE attack instance with JavaScript.

We implemented SPLITSPECTRE on SpiderMonkey 52.7.4, Firefox’s JavaScript engine. We use the standard configuration parameters and conducted experiments on our Haswell, Skylake, and Ryzen CPUs.

We start our experiments by introducing a built-in native JavaScript accessor function to SpiderMonkey’s source code that returns the content of a pre-allocated array at a given index. This function is the first part of the speculative execution gadget that needs to be part of the victim’s attack surface. To simplify the code, we explicitly flush the bounds of the array. Our attacker code is an adapted regular V1 PoC code for JavaScript JIT engines, with just the first array access replaced by the call to the victim function. The time measurement is done using the SharedArrayBuffer technique, which reads the content of such a buffer while it is being incremented in the background by a web worker that is running in parallel.

The attack works: we leak a string of ten characters with a success rate of over 80% Table III, and we leak the full string with a success rate of 10%. Investigating the distance between the two parts of the speculation gadget, we measure the distance after 50 training runs of the JavaScript code that causes Spidermonkey’s tracing JIT to compile an optimized IonJIT trace implementing the JavaScript code in assembly. The distance between the bounds check and the second array access is 43 instructions, which is small enough for the attack to produce reliable results.

We proceed with our experiments by replacing our native built-in function with code already present in the SpiderMonkey source. Our scan for a suitable gadget reveals the built-in `string.charCodeAt()` function, which returns the character code of a string at a given index and is implemented in native code. Internally, `string.charCodeAt()` calls `string.charCodeAt_impl()`, which includes the bounds check and actual access. Unfortunately, the speculation window is not large enough for the attack to work with `string.charCodeAt()`: It turns out that after 50 training runs, the distance between the compare in `string.charCodeAt_impl()` and the dereference of the

second array in the JIT trace is 90 instructions. An examination of the extracted execution trace with SPECULATOR shows that the number of speculatively executed ops reaches a plateau at around 40 instructions into the trace for Skylake and 27 for Broadwell (Figure 6).

We also examine the execution trace on an AMD Ryzen CPU using a marker instruction, since the Zen performance counters do not feature a generic counter for executed instructions. We see the marker instruction being executed for the full length of the trace. However, the granularity of time measurement is too coarse-grained to permit a successful read of the cache side channel. Amplifying the attack by adding multiple dependent array accesses would extend the trace so that it no longer fits into the speculation window.

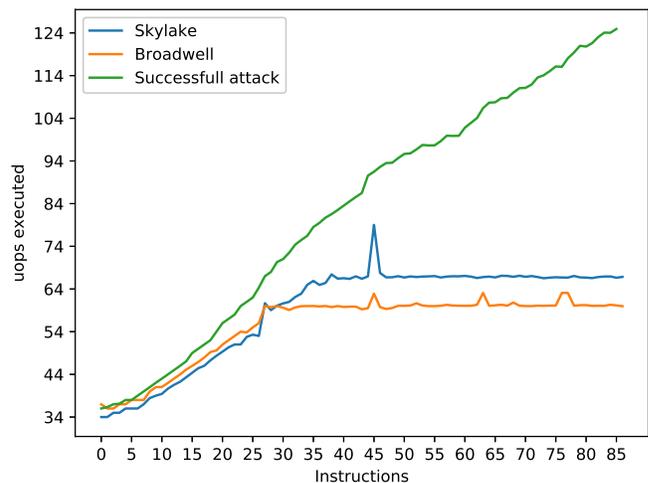


Fig. 6: An examination of the SPLITSPECTRE execution trace between the length check of `string.charCodeAt_at_impl()` and the second array access using SPECULATOR. The plateau of executed μops at around 27 (Broadwell) respectively 40 (Skylake) instructions shows that we are not reaching the second array access in speculative execution despite the total number of μops in the trace being lower than the capacity of the reorder buffer on both architectures. The spikes in the plateaus are caused by mispredicted branches in the trace itself, which lead to nested speculative execution of fast-executing code paths.

We further optimize the attack by reducing the amount of code that is executed between the bounds check and the second access. This was achieved by implementing the second access and the call to the victim function in web assembly, which allows even more attacker control over the compiled JIT trace. However, using WebAssembly actually increases the number of instructions between the compare and the second access to 107. The reason is that the native call is not made directly from within the WebAssembly, but additional JavaScript glue code is invoked.

JIT engine authors have already reacted with countermeasures [41], [11] in order to mitigate Spectre V1 in the context of browsers. These countermeasures mostly address sources for high-precision timers. Diluting the timing and disabling homebrew sources such as SharedArrayBuffers mitigate this version of JavaScript SPLITSPECTRE. However, it remains to be seen if amplification of the attack’s timing properties make

Runs	100
Only highest scoring char	76.6%
1st and 2nd highest scoring char	80.7%
Full string leaked	10%

Table III: Success rates for the SPLIT SPECTRE attack on JavaScript. We perform 100 runs, each run trying to leak a string of 10 consecutive characters. We provide numbers on both the highest and the second highest scoring characters.

it feasible if only coarse-grained time sources are available.

On top of timing-related countermeasures, the V8 engine also masks addresses and array indices in JITted code before dereferences. While this mitigates a standard Spectre V1 attack, it does not help with SPLIT SPECTRE, where the bounds check is actually not exercised in JITted code, but the engine code itself.

All things considered, our analyses lead us to conclude that the attack is viable, and that the ability to trigger it in practice depends on the identified microarchitectural properties of individual CPU families. We leave a comprehensive analysis of these properties for the various CPU architectures/models as an item of future work, which can be aided by SPECULATOR.

V. USING SPECULATOR: MICROARCHITECTURAL INSIGHTS BEYOND SPLIT SPECTRE

We also used SPECULATOR to investigate microarchitectural aspects beyond the ones directly related to SPLIT SPECTRE.

A. Speculation window size

In Section II-E, we defined that the speculation window size is determined by the clock cycles that it takes until a speculation trigger is resolved. In this section, we provide our measurements of the speculation window for the different triggers used in the Spectre v1, v2, and v4 attacks. For measuring clock cycles, we again leverage the facilities provided by the PMC of the respective platform. On Intel, a predefined counter tracks elapsed clock cycles according to the same settings as the configurable counters. On AMD, the APERF counter tracks elapsed clock cycles in general.

The theoretical upper limit of instructions that can be executed during speculative execution is given by the size of the reorder buffer, which we evaluated in Section IV-A. In practice, it is also limited by the execution ports and units available for executing those transactions. Thus, we also investigate instruction sequences that do not lead to a bottleneck on those resources during speculative execution.

Conditional branches. Conditional branches are the speculative execution triggers used in Spectre v1 to check for an out-of-bounds access to an array. The speculation window size depends on how fast the CPU determines that the actual branch target differs from the information provided by the branch target buffer. We place the conditional value that determines the actual branch target in different locations and involve it in additional computation to investigate how this affects the size of the speculation window. As a baseline, we measure how long the execution of the additional instructions takes. We then

Conditional branch	Broadwell	Skylake	Zen
Register access	14	16	7
Access to cached memory	19	17	9
Access to uncached memory	144	280	321
Mul with register	19	19	2
Mul with cached memory	33	33	8
Mul with uncached memory	154	290	362
Div with register	35	41	17
Div with cached memory	34	39	30
Div with uncached memory	164	306	353

Table IV: Speculation window of a conditional branch depending on the type of instructions needed to resolve the branch as well as the placement of the value involved in the condition, measured in cycles.

Indirect branch target location	Broadwell	Skylake	Zen
Register	81	85	25
Cached memory	87	85	31
Uncached memory	248	349	351

Table V: Speculation window of an indirect control flow transfer, measured in cycles. The speculation window size depends on where the target of the indirect control flow transfer is stored.

measure how long the execution of the instructions together with the conditional branch takes. Any difference in the time it takes the conditional branch to retire reflects the placement of the variable and the effect of the additional instructions involved. All measurements are performed a thousand times. Note that controlling the performance counters involves a system call. Since system calls stop speculation, we can only measure how long the retirement of an instruction sequence takes. Since the measurement technique differs between the two CPU vendors, results for Intel and AMD cannot be compared.

Table IV shows the results of this experiment. We see that complex instructions such as `div`, which translates to multiple μ ops, widen the speculation window. The same is true for a cache miss, when the CPU needs to fetch the data from main memory.

At the same time, access to cached memory contributes little to the speculation window compared to a register access. Measuring a range from four to twelve cycles, the results for Broadwell and Skylake are in accordance with Intel’s performance analysis guide [1] which states four cycles as the average for an access to L1 and ten cycles for L2.

On AMD, we see even less impact between register and cached accesses. In addition, adding a complex instruction on top of an access has a negligible effect on the speculation window size.

Indirect control flow transfer. Indirect control flow transfers are the speculative execution triggers used in Spectre v2. The speculation window size depends on how fast the CPU determines that the target in the branch history buffer does not match the actual target. Table V shows the speculation window sizes depending on the location of the indirect branch target.

Store to load forwarding. Modern CPU designs feature store and load queues, which capture the effects and dependencies of corresponding load and store operations before the data is

even written to or read from the cache. This infrastructure allows for efficient store to load forwarding: if an instruction writes to a certain memory address and a following instruction reads from that very address, the CPU can leverage the result of the first instruction, which is written to the store queue, for executing the second instruction. This avoids unnecessarily stalling the execution of the second instruction until the first is retired. In a recent attack, this behavior has been used for a “speculative buffer overflow” [24].

We are interested in the behavior a failed store to load forwarding causes. In this case, we deviate from our default SPECULATOR template and remove the branch instruction. Instead, we create a snippet with a data dependency that is not detected by the CPU in a combination with a sequence of store and load operations that triggers store-to-load forwarding.

Running the snippet in SPECULATOR reveals that store-to-load forwarding fails and the load instruction is in fact executed twice. This means that a failed store-to-load forwarding also creates a situation similar to speculative execution results being discarded because of a mispredicted branch, although it provides a significantly smaller speculation window.

Spectre v4 (a speculative store bypass) makes use of speculative execution through store-to-load forwarding. For this trigger we measure a speculation window of 55 cycles on average on Broadwell. We also measure the speculatively executed instructions using `FNOP`, which provides us with an upper bound for the speculation window in terms of instructions. We measure an average of 15 μ ops with a maximum of 23 μ ops (Figure 7).

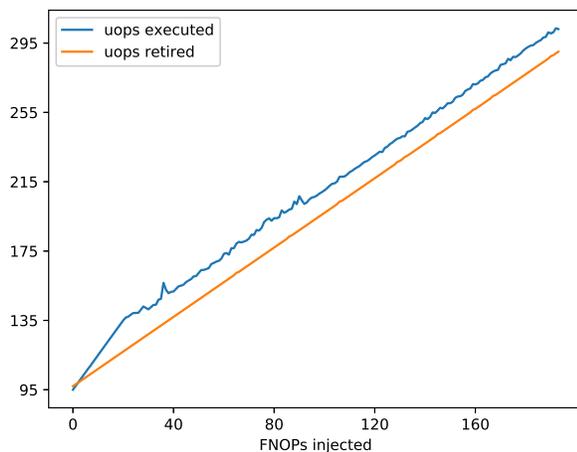


Fig. 7: Speculation window of a store-to-load forward failure, measured in executed `FNOPs` on Broadwell.

Max speculation with optimized instruction sequence. During our experiments, we observed multiple situations in which the CPU back-end stalled. For instance, the CPU could stall due to exhaustion of execution units for a certain operation (e.g. `MOV`, `MUL`) or, for instance, data dependencies of multiple operations where one or more data loads caused cache misses. In a hypothetical scenario, we wanted to verify how many non-`NOP` executed μ ops the CPU speculates within the maximum time window (e.g. access to uncached memory in combination with a `DIV` instruction). Based on the layout of the back-end

of our Broadwell CPU under test, to the best of our abilities, we crafted an optimized sequence of instructions to account for the delay of each operation and the available execution unit. Our tests show that the maximum number of non-trivial speculatively executed instructions we could achieve was 160, with 187 being the maximum for `FNOP`.

B. Stopping Speculative Execution

Many instruction set architectures feature an instruction that stops speculative execution in the sense that no following instruction is speculatively executed. On x86 (and x86_64), one such instruction is `lfence`, short for “load fence”, the name reflecting its initial purpose of serializing all memory load operations issued prior to this instruction. In addition to this behavior, it also works as a barrier for speculative execution: the operational description in Intel’s manual [5] specifies that `lfence` waits on following instructions until preceding instructions complete.

We verify this behavior using SPECULATOR by creating a snippet with an `lfence` instruction followed by an increasing sequence of regular instructions. As expected, the counter for executed μ ops remains constant among the test runs irrespective of the number of instructions following `lfence`.

C. Flushing the Cache

The x86 instruction set provides a convenient, dedicated instruction to cause the CPU to flush the cache line indicated by a memory address from all caches, `clflush`. It is very useful in settings where an attacker can execute assembly instructions, as it allows easy eviction of data from the cache.

We use SPECULATOR to investigate how `clflush` behaves when executed speculatively. To this end we create a snippet that first flushes the cache line corresponding to a value stored in memory and then loads the value. This is shown at line 3 and line 19 respectively in Listing 2. We perform two runs, one where the setup code warms up the cache by loading the value from memory (line 7) and one where the value is left uncached. In both tests, within the speculatively executed sequence, we place a `clflush` followed by an `lfence` instruction to stop the speculation, making sure that the final load is not executed during speculation as well (line 15). We measure the execution cycles on both runs, which shows a difference of over 160 clock cycles between the two settings (Figure 8). This is a clear indication that while `clflush` is speculatively executed, it does not affect the cache until retired.

```

1  setup
2  .loop:
3  clflush[counter]
4  clflush[var]
5  lfence
6
7  mov eax, DWORD[var] ;cached version
8  lfence ;only
9
10 start_counter
11
12 cmp 12, DWORD[counter]
13 je .else
14
15 clflush[var]

```

```

16 | lfence
17 |
18 | .else:
19 |     mov eax, DWORD[var]    ;final load
20 |     lfence
21 |
22 |     stop_counter
23 |
24 |     inc DWORD[i]
25 |     cmp DWORD[i], 13
26 |     jl loop

```

Listing 2: Clflush test snippet structure

Another conclusion we can draw from this experiment is that in order to make sure `clflush` is effective, it needs to be combined with an instruction that stops speculative execution, such as `lfence`.

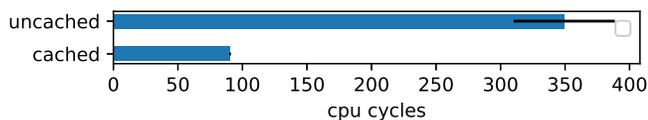


Fig. 8: Execution time measurements of a speculative access to a value in memory, once cached, once uncached. The difference between the two measurements demonstrates that the speculatively executed `clflush` instruction before the access does not actually affect the cache.

D. Executable Page Permission

Memory page permissions control access to memory regions at page-level granularity. As we have seen with Meltdown and Foreshadow, such permission checks might be lazily evaluated after an instruction is already executed, but before it is retired. Related work has so far focused on data read or write access to memory pages. We are interested in execute permissions enforced by the NX bit, a hardware extension introduced by modern processors to mitigate the classic textbook stack-based code injection exploits. If the control flow of a program is diverted to a page without execute permissions, the processor will trap into the kernel to handle the fault. This raises the question whether during speculative execution the permission is honored or it is possible to execute instructions from a page without such a permission set.

Our corresponding experiment sets up a branch misprediction with a following control flow transfer to a memory region we control the access bits to, essentially testing whether the data in this memory region is executed. We ensure that the data from the page is in the L2 cache during speculative execution and the addresses are in the TLB. The result of the experiment is that the execute page table permission is honored during speculative execution by all architectures we examined. This is even true if an instruction spans over two pages: it will not be executed if the second page is set non-executable.

E. Memory Protection Extensions

Instead of performing bounds checks purely in software, Intel’s MPX instruction set extension [34] available on the

SkyLake platform provides hardware support for both efficiently keeping track of bounds information associated with pointers and corresponding spatial memory checks before dereferencing pointers. Pointer bounds information is stored in memory and loaded to dedicated registers before it can be used to check the upper bound using the `bndcu` and the lower bound using the `bndcl` instruction. If a bound check fails, a `#BR` exception is raised and the CPU traps into the kernel.

We used SPECULATOR to measure if and how much code following a bounds check instruction is speculatively executed. The setup executes the regular code path without the bounds violation for ten iterations and then fails on a `bndcu` twice. To measure the speculative execution window size, we first used an increasing run of `NOPs` in conjunction with a terminating slow `LEA` marker instruction. In this experiment, we measured that we speculatively execute the marker instruction for a sled of up to 122 `NOPs`. In our second experiment, we used `FNOP` instead of regular `NOP`, which is tracked by the `UOPS_EXECUTED` counter. As is shown in Figure 9, in this case, the number of executed μ ops increases up to a sled of 22 `FNOPs` and remains constant beyond.

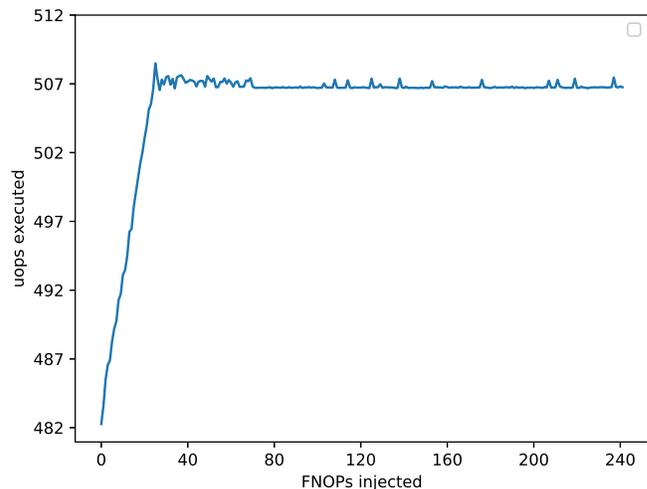


Fig. 9: Speculative execution after an MPX bounds violation.

F. Issued vs. Executed μ ops

All performance counters that address a certain μ op group such that they are suited as isolated markers for speculative execution count issued μ ops. Since issued μ ops are not necessarily executed, as is the case for the `NOP` instruction, we performed a dedicated experiment. We use the template introduced in Section III-D and generated tests where the code snippet just contains an increasing number of `RIP`-relative load instructions. As Figure 10 demonstrates, the number of executed μ ops increases at the same rate as the counter for slow load effective address instructions, which are load μ ops with three sources.

VI. DISCUSSION

SPECULATOR depends on a CPU’s infrastructure to expose microarchitectural state, in Intel’s nomenclature known as

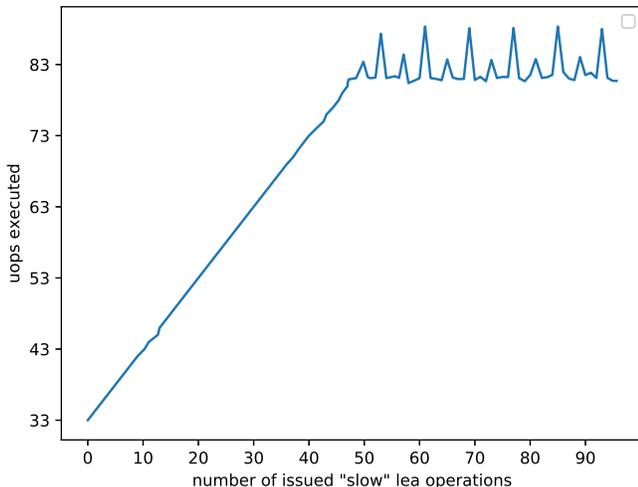


Fig. 10: Performance counter numbers for an increasing number of speculatively executed relative load instructions. The graph shows that the number of issued instructions corresponds to the number of executed instructions, justifying the use of such instructions as markers.

PMC. In fact, it depends on the correctness of the information provided by this infrastructure, as well as the availability to be able to monitor certain groups of events. While the monitoring infrastructure itself as well as the ability to track specific events is different for every CPU architecture, the possibility to track instructions that have been executed, yet not retired, is available on all the platforms we investigated. Thus, while the exact implementation of SPLIT SPECTRE might not be directly transferable without changes to platforms we have not looked at, the method is.

Besides, SPECULATOR also depends on a properly functioning CPU. Results should be similar among multiple CPUs of the same architecture, with inconsistencies indicating a defective unit.

VII. RELATED WORK

Speculative Execution. Optimizing CPU instruction throughput through speculative execution has been proposed and implemented in the 1990s [27], [37]. For information about the microarchitecture of CPUs with respect to out-of-order and speculative execution, we mostly have to rely on the material provided by the CPU manufacturers [5], [4], [2]. Unfortunately this material often just hints at important aspects, not providing detail on how mechanisms such as the branch predictor actually work. Agner Fog’s work [17] sheds light on those details, providing detailed information backed by a substantial amount of experimental research on the microarchitectural aspects of CPUs. This information is leveraged in processor simulators such as gem5 [10].

Cache Side Channels. All Spectre variants including our new SPLIT SPECTRE rely on cache side channels to infer the memory contents accessed by speculative execution. Cache side channels have been extensively studied: First, Tromer et al. introduced both the “evict-and-time” and “prime-and-probe” techniques to efficiently perform a cache attack on

AES [38]. Prime and probe is a popular technique, which was also used for certain Spectre variants. “Flush-and-reload” [44] is a technique that allows for higher precision and is used in NetSpectre. Recently, other techniques such as “flush-and-flush” [20] and “prime-and-abort” [15] were presented. Flush and flush leverages the fact that `clflush` executes faster in case of a cache hit. Prime and abort makes use of Intel’s transactional memory mechanism to detect when an eviction has happened without the need to probe the cache.

Security Issues. Since the beginning of 2018, three security issues related to speculative execution known as Spectre and Meltdown were revealed [46], [29], [25]. CPU vendors reacted with reports on those issues [22], [19] and how they affected their CPU architectures. These initial reports were followed by more security issues, involving further speculative execution triggers [21], [9], [3] and side channels [18], even affecting Intel’s virtualization and secure enclave technology SGX [40]. In addition to that, research groups have established remote Spectre attack vectors over the network [7]. The classic buffer overflow to overwrite the return address on the stack also has a speculative execution context twist, as shown in [24], [26].

Mitigations. Apart from the microcode updates shipped by CPU vendors, certain mitigations against SEAs can be implemented in software. Especially JavaScript engines deployed mitigations against Spectre v1 such as diluting timing precision, disabling concurrent threads to prevent homebrew-timers and masking pointer accesses to prevent speculative out-of-bounds accesses [11], [41], [6]. Linux has deployed retpoline [39] in the kernel to mitigate Spectre v2 by trapping mispredicted indirect branches and the KAISER patches [13] to protect against Meltdown by separating page tables organization for user- and kernel space. Also compiler tool chains have picked up the topic, with LLVM working on introducing data dependencies on loads that might be speculatively executed [12], [31] and MSVC adding speculation barrier instructions such as `lfence` to the compiled binary code [32]. At the same time, research groups have proposed to address the issue in silicon, such as adding microarchitectural shadow structures to the CPU for leakage-free speculation [23] or exposing the microarchitectural state in the ISA [30].

VIII. CONCLUSION

In this paper, we shed light on security-relevant speculative execution and microarchitectural behavior. We presented SPECULATOR, a novel tool that allow targeted and precise measures of microarchitectural characteristics. Using SPECULATOR, we then investigate speculative execution. We study aspects such as the speculation window for various speculative execution triggers, which is an important factor for the payload of a speculative execution attack. We also show which events stop speculative execution and that some security controls such as NX are still in effect during speculative execution, while others do not act as a barrier such as Intel’s MPX bounds checks.

Based on these findings, we then verified the feasibility of a new variant of SEA that we call SPLIT SPECTRE. We motivated its importance with new upcoming more powerful families of

processors showing how the gap for having a successful real world attack decreases the longer the CPU is able to speculate.

We plan to release our tool, SPECULATOR, which we used to investigate speculative execution behavior, as open source.

REFERENCES

- [1] Performance Analysis Guide for Intel Core i7 Processor and Intel Xeon Processors. https://software.intel.com/sites/products/collateral/hpc/vtune/performance_analysis_guide.pdf.
- [2] Preliminary Processor Programming Reference (PPR) for AMD Family 17h Models 00h-0Fh Processors. http://support.amd.com/TechDocs/54945_PPR_Family_17h_Models_00h-0Fh.pdf, 2017.
- [3] Analysis and mitigation of speculative store bypass. <https://blogs.technet.microsoft.com/srd/2018/05/21/analysis-and-mitigation-of-speculative-store-bypass-cve-2018-3639/>, 2018.
- [4] Intel Architectures Optimization Reference Manual. <https://www.intel.com/content/dam/www/public/us/en/documents/manuals/64-ia-32-architectures-optimization-manual.pdf>, 2018.
- [5] Intel Software Developer Manual. <https://software.intel.com/en-us/articles/intel-sdm>, 2018.
- [6] JIT mitigations for Spectre. <https://github.com/Microsoft/ChakraCore/commit/08b82b8d33e9b36c0d6628b856f280234c87ba13>, 2018.
- [7] Netspectre: Read arbitrary memory over network. <https://misc0110.net/web/files/netspectre.pdf>, 2018.
- [8] Rogue system register read. <https://software.intel.com/security-software-guidance/software-guidance/rogue-system-register-read>, 2018.
- [9] Speculative store bypass disable, 2018.
- [10] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill, and D. A. Wood. The gem5 simulator. *SIGARCH Computer Architecture News*, 39(2), Aug. 2011.
- [11] M. Bynens. V8 Untrusted code mitigations. <https://github.com/v8/v8/wiki/Untrusted-code-mitigations>, 2018.
- [12] C. Carruth. Speculative Load Hardening. <https://lists.llvm.org/pipermail/llvm-dev/2018-March/122085.html>, 2018.
- [13] J. Corbet. Kaiser: hiding the kernel from user space. <https://lwn.net/Articles/738975/>.
- [14] A. C. de Melo. The New Linux perf tools. <http://www.linux-kongress.org/2010/slides/1k2010-perf-acme.pdf>, 2010.
- [15] C. Disselkoe, D. Kohlbrenner, L. Porter, and D. Tullsen. Prime+abort: A timer-free high-precision l3 cache attack using intel TSX. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 51–67, Vancouver, BC, 2017. USENIX Association.
- [16] S. Eranian. Perfmon2: a flexible performance monitoring interface for linux. In *Proc. of the 2006 Ottawa Linux Symposium*, pages 269–288, 2006.
- [17] A. Fog. The microarchitecture of Intel, AMD and VIA CPUs: An optimization guide for assembly programmers and compiler makers. <https://www.agner.org/optimize/microarchitecture.pdf>, 2018.
- [18] B. Gras, K. Razavi, H. Bos, and C. Giuffrida. Translation leak-aside buffer: Defeating cache side-channel protections with TLB attacks. In *USENIX Security Symposium*, 2018.
- [19] R. Grisenthwaite. Cache Speculation Side-channels. https://developer.arm.com/-/media/Files/pdf/Cache_Speculation_Side-channels.pdf, 2018.
- [20] D. Gruss, C. Maurice, K. Wagner, and S. Mangard. Flush+flush: A fast and stealthy cache attack. In J. Caballero, U. Zurutuza, and R. J. Rodríguez, editors, *Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 279–299, Cham, 2016. Springer International Publishing.
- [21] J. Horn. Spectre v4. <https://bugs.chromium.org/p/project-zero/issues/detail?id=1528>, 2018.
- [22] Intel. Analysis of speculative execution side channels. <https://newsroom.intel.com/wp-content/uploads/sites/11/2018/01/Intel-Analysis-of-Speculative-Execution-Side-Channels.pdf>, 2018.
- [24] V. Kiriansky and C. Waldspurger. Speculative Buffer Overflows: Attacks and Defenses. <https://people.csail.mit.edu/vlk/spectre11.pdf>, 2018.
- [23] K. N. Khasawneh, E. M. Koruyeh, C. Song, D. Evtushkin, D. Ponomarev, and N. B. Abu-Ghazaleh. SafeSpec: Banishing the Spectre of a Meltdown with Leakage-Free Speculation. *CoRR*, 2018.
- [25] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom. Spectre attacks: Exploiting speculative execution. In *IEEE Symposium on Security and Privacy*, 2019.
- [26] E. M. Koruyeh, K. N. Khasawneh, C. Song, and N. B. Abu-Ghazaleh. Spectre Returns! Speculation Attacks using the Return Stack Buffer. *CoRR*, 2018.
- [27] B. W. Lampson. Lazy and speculative execution in computer systems. In *ACM SIGPLAN Conference on Functional Programming*, 2008.
- [28] J. Levon. Oprofile. <http://oprofile.sourceforge.net>.
- [29] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, A. Fogh, J. Horn, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg. Meltdown: Reading kernel memory from user space. In *USENIX Security Symposium*, 2018.
- [30] J. Lowe-Power, V. Akella, M. K. Farrens, S. T. King, and C. J. Nitta. Position paper: A case for exposing extra-architectural state in the isa. In *Proceedings of the 7th International Workshop on Hardware and Architectural Support for Security and Privacy*, 2018.
- [31] O. Oleksenko, B. Trach, T. Reiher, M. Silberstein, and C. Fetzer. You shall not bypass: Employing data dependencies to prevent bounds check bypass. *CoRR*, 2018.
- [32] A. Pardoe. Spectre mitigations in MSVC. <https://blogs.msdn.microsoft.com/vcblog/2018/01/15/spectre-mitigations-in-msvc/>, 2018.
- [33] M. Pettersson. Perfctr. <http://user.it.uu.se/~mikpe/linux/perfctr/>.
- [34] S. Ramakesavan and J. Rodriguez. Intel Memory Protection Extensions Enabling Guide. <https://software.intel.com/en-us/articles/intel-memory-protection-extensions-enabling-guide>, 2016.
- [35] T. Rhl, J. Eitzinger, G. Hager, and G. Wellein. LIKWID Monitoring Stack: A Flexible Framework Enabling Job Specific Performance monitoring for the masses. In *IEEE International Conference on Cluster Computing (CLUSTER)*, 2017.
- [36] D. Terpstra, H. Jagode, H. You, and J. Dongarra. Collecting performance data with papi-c. In *Tools for High Performance Computing 2009*, pages 157–173. Springer, 2010.
- [37] K. B. Theobald, G. R. Gao, and L. J. Hendren. Speculative execution and branch prediction on parallel machines. In *International Conference on Supercomputing*, 1993.
- [38] E. Tromer, D. A. Osvik, and A. Shamir. Efficient cache attacks on aes, and countermeasures. *Journal of Cryptology*, 23(1):37–71, Jan 2010.
- [39] P. Turner. Retpoline: a software construct for preventing branch-target injection. <https://support.google.com/faqs/answer/7625886>.
- [40] J. Van Bulck, M. Minkin, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx. Foreshadow: Extracting the keys to the Intel SGX kingdom with transient out-of-order execution. In *USENIX Security Symposium*, 2018.
- [41] L. Wagner. Mitigations landing for new class of timing attack. <https://blog.mozilla.org/security/2018/01/03/mitigations-landing-new-class-timing-attack/>, 2018.
- [42] V. M. Weaver. Linux perf_event features and overhead. In *The 2nd International Workshop on Performance Analysis of Workload Optimized Systems, FastPath*, volume 13, 2013.
- [43] O. Weisse, J. V. Bulck, M. Minkin, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, R. Strackx, T. F. Wenisch, and Y. Yarom. <https://foreshadowattack.eu/foreshadow-NG.pdf>, 2018.
- [44] Y. Yarom and K. Falkner. Flush+reload: A high resolution, low noise, l3 cache side-channel attack. In *Proceedings of the 23rd USENIX Conference on Security Symposium, SEC'14*, pages 719–732, Berkeley, CA, USA, 2014. USENIX Association.
- [45] D. Zapparanuks, M. Jovic, and M. Hauswirth. Accuracy of performance counter measurements. In *Performance Analysis of Systems and Software, 2009. ISPASS 2009. IEEE International Symposium on*, pages 23–32. IEEE, 2009.
- [46] G. P. Zero. Reading privileged memory with a side-channel. <https://googleprojectzero.blogspot.ch/2018/01/reading-privileged-memory-with-side.html>, 2018.