

Speculative Taint Tracking (STT): A Comprehensive Protection for Speculatively Accessed Data

Jiyong Yu

University of Illinois at Urbana–Champaign

Mengjia Yan

Massachusetts Institute of Technology

Artem Khyzha

Tel Aviv University

Adam Morrison

Tel Aviv University

Josep Torrellas and Christopher W. Fletcher

University of Illinois at Urbana–Champaign

Abstract—Speculative execution attacks present an enormous security threat, capable of reading arbitrary program data under malicious speculation, and later exfiltrating that data over microarchitectural covert channels. This article proposes speculative taint tracking (STT), a high-security and high-performance hardware mechanism to block these attacks. The main idea is that it is safe to execute and selectively forward the results of speculative instructions that read secrets, as long as we can prove that the forwarded results do not reach potential covert channels. The technical core of the article is a new abstraction to help identify all covert channels, and an architecture to quickly identify when a covert channel is no longer a threat. We further conduct a detailed formal analysis on the scheme and prove security in a companion document. When evaluated on SPEC06 workloads, STT incurs 8.5% or 14.5% performance overhead relative to an insecure machine.

Digital Object Identifier 10.1109/MM.2020.2985359

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

May 2020.

■ **SPECULATIVE EXECUTION ATTACKS** such as Spectre⁵ have opened a new chapter in hardware security. In these attacks, malicious speculative execution causes doomed-to-squash instructions to *access* and later *transmit* secrets over covert channels such as the cache.⁹ For

example, in Spectre V1 (see Figure 1) a branch misprediction enables the attacker to access and leak/transmit arbitrary program data by controlling the out-of-bounds address `&array1[off]`. We refer to such data, which is brought into the pipeline by a speculative instruction, as *secret*.

A secure, but conservative, way to block *all* speculative execution attacks—regardless of covert channel—is to delay executing all instructions that can access a secret until such instructions become nonspeculative. In nearly all attacks today, this would imply blocking all loads until they are nonspeculative, which would be tantamount to disabling speculative execution.

This article proposes a principled, high-performance mechanism that achieves the same security guarantee as the above conservative scheme. The key idea is that speculative execution is safe unless speculatively accessed data (secrets) reaches a covert channel. In many cases, speculative instructions either do not leak secrets or do not form covert channels, and so can execute freely under speculation. For example, the first load in Spectre V1 forms a covert channel, but it only leaks the attacker-selected address `&array1[off]`—not the secret data in that address. Likewise, many instructions (e.g., simple arithmetic) do not form covert channels even if their operands are secret values.

This article presents speculative taint tracking (STT), a framework that tracks the flow of speculatively accessed data through in-flight instructions (similar to dynamic information flow tracking/DIFT⁷) until it is about to reach an instruction that may form a covert channel. STT then delays the forwarding of the data until the instruction becomes nonspeculative or the execution squashes due to mis-speculation.

```

1  if ( off < array1_size ) { // mispredicts
2    x = array1[ off ]; // secret accessed
3    y = array2[x]; } // secret transmitted

```

Figure 1. Spectre variant 1.

To be secure and efficient, we address two key challenges.

- First, we develop an abstraction that indicates how and when instructions can form covert channels, so as to delay data forwarding to the latest safe time.
- Second, we identify and develop a microarchitecture to indicate exactly when data should be considered secret, so as to re-enable data forwarding at the earliest safe time.

This article proposes a new abstraction through which to view covert channels on speculative microarchitectures, discovers new points where instructions can create covert channels, and discovers a new class of covert channels

Challenge #1: New Abstractions for Describing *All* Microarchitectural Covert Channels

Covert channels come in different shapes and sizes. For example, attackers can monitor how loads interact with the cache,⁵ the timing of SIMD units,⁶

execution pipeline port contention,² branch predictor state,¹ and more. To comprehensively block leakage through these different channels, it is necessary to understand their common characteristics.

To address this challenge, this article proposes a new abstraction through which to view covert channels on speculative microarchitectures, discovers new points where instructions can create covert channels, and discovers a new class of covert channels. We find that all covert channels are one of two flavors, which we call explicit and implicit channels (which are related to explicit and implicit information flow⁸). In an *explicit channel*, data is directly passed to an instruction whose execution creates operand-dependent hardware resource usage, and that resource usage reveals the data. For example, how a load impacts the cache depends on the load address.⁵ In an *implicit channel*, data indirectly influences how (or that) an instruction(s) execute, and these changes in resource usage reveal the data. For example, the instructions executed after a branch reveal the branch predicate.^{2,6} This article further defines subclasses of the implicit channel, based on when the leakage occurs and based on the nature of the secret-dependent condition that forms the channel.

Key Advance: Safe Prediction. Through its investigation of implicit channels, this article makes a key advance by showing how to use hardware predictors safely. Spectre attacks were born from attackers mistraining predictors to leak secrets. Through its abstraction for implicit channels, STT enforces a policy that prevents arbitrary predictor mistraining from leaking any secret data over any covert channel. The article shows how this enables existing predictors to stay enabled without leaking privacy, dramatically improving performance. In the future, we expect the idea of safe prediction to enable further innovation, i.e., by enabling the design of new predictors without fear of opening new security holes. Indeed, our follow-on work uses this idea to safely improve the performance of instructions that create explicit channels.¹¹

Challenge #2: Mechanisms to Quickly and Safely Disable Protection

Once we have mechanisms to block secret data from reaching covert channels, the next question is when and how to disable that protection, if speculation turns out to be correct. This is crucial for performance, as delaying data forwarding longer than necessary increases the chance that delayed instructions reach the head of the reorder buffer (ROB) and block retirement.

STT tackles this problem with a safe but aggressive approach, by re-enabling data forwarding as soon as data becomes a function of retired register file state. This represents the earliest safe point, but is nontrivial to implement in hardware. For example, a delayed instruction's operand(s) may be the result of a complex dependence chain across many control flow and speculative operations. Intuitively, determining that data is a function of nonspeculative information would require retracing a backwards slice of the program's execution, which is costly to do quickly.

Despite the above challenges, STT proposes a simple hardware mechanism that can disable protection/re-enable forwarding for an arbitrary instruction in a single cycle, using hardware similar to traditional instruction wake-up logic. The key idea is that to determine whether data is a function of retired state, it is sufficient to determine whether the *youngest* load, whose return value influences

the data, has become nonspeculative. Checking this condition is akin to tracking a single extra dependence for each instruction, as opposed to performing complex backwards slice tracking.

Security Guarantees and Formal Analysis

Alongside the main paper, we formally prove that STT enforces a novel form of noninterference³ with respect to speculatively accessed data. In a nutshell, we show that hardware resource usage patterns over time are independent of data that eventually squashes (covering microarchitectural interference- and timing-based attacks). We released a companion technical report¹² with detailed formal analysis and a security proof for this property.

Putting It All Together

Putting everything together, STT provides both high security and high performance. It does not require partitioning or flushing microarchitectural resources, and does not require changes to the cache/memory subsystem or the software stack. When evaluated on SPEC06 workloads, STT incurs 8.5% or 14.5% performance overhead (depending on the threat model) relative to an insecure machine.

ATTACKER MODEL AND PROTECTION SCOPE

Attacker Model. STT assumes a powerful adversary that can monitor any microarchitectural covert channel from anywhere in the system, and induce arbitrarily speculative execution to access secrets and create covert channels. For example, the attacker can monitor covert channels through the cache/memory system,⁵ data-dependent arithmetic,⁴ port contention,² branch predictors,¹ etc.

Scope: Protecting Speculatively Accessed Data. We distinguish attacks based on whether the access instruction is doomed-to-squash (*transient*) or bound to retire (*nontransient*). STT's goal is to block attacks involving doomed-to-squash access instructions, shown in Figure 2. These attacks can access data that a correct (not miss-speculated) execution would never access, which often results in being able to read from any location in memory. Attacks involving bound-to-retire access instructions are out of

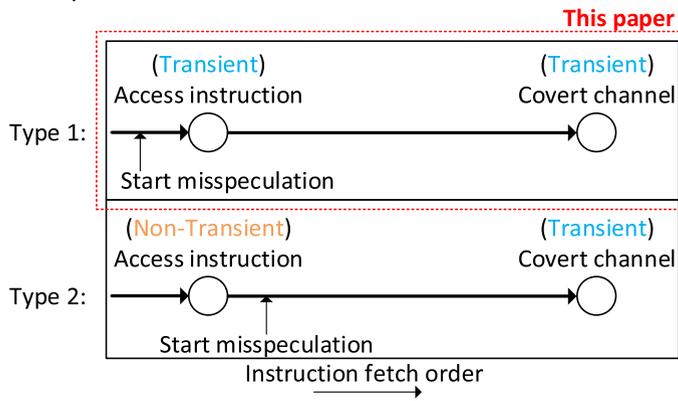


Figure 2. STT’s scope is to protect speculatively accessed data from leaking over any microarchitectural covert channel. Protecting values which have retired is outside of scope.

scope. They can only leak *retired (or bound to retire) register file state*, not arbitrary memory, and their leakage can be reasoned about by programmers or compilers and blocked using complementary techniques (e.g., Yu *et al.*¹⁰).

ABSTRACTION FOR COVERT CHANNELS

STT proposes a novel abstraction for covert channels (see Figure 3). In our abstraction, covert channels are broken into two classes: *explicit* and *implicit* channels. An *explicit channel*, related to explicit flow in information flow,⁸ is one where data (e.g., a secret) is *directly* passed to an instruction whose execution creates operand-dependent hardware resource usage, and that resource usage reveals the data. An example is a load instruction’s changes to the cache state. An *implicit channel*, related to implicit flow,⁸ is one where data *indirectly* influences how (or whether) an instruction or several instructions execute, and these changes in resource usage reveal the data. An example is a branch instruction, whose outcome determines subsequent instructions and thus whether some functional unit is used.

We further distinguish between implicit channels, depending on when they leak secrets and what type of branch they feature. First, we find that implicit channels can leak at two points: when a prediction is made (e.g., a branch prediction) and when a resolution occurs (e.g., a branch resolves). Second, we find that implicit channels

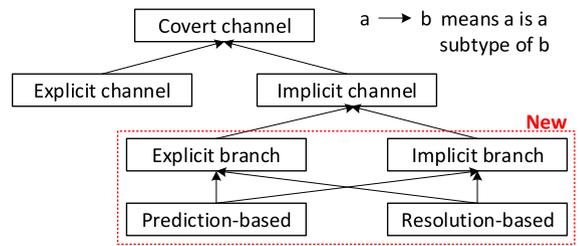


Figure 3. STT’s new classification schema for covert channels.

can feature either an *explicit* or an *implicit* branch. An explicit branch is a control-flow instruction, while an implicit branch is a conceptual branch formed in the hardware due to an optimization. For example, store-to-load forwarding between a store and a younger load can be viewed as an implicit branch that checks for an address alias as shown in Figure 4. Written this way, it is clear that store-to-load forwarding can create a covert channel: depending on whether there is an alias, the processor either looks up the cache or forwards from the local store queue.

Insights From Analysis of Implicit Channels

Since it was proposed in the STT paper, the classification for implicit channels has proven to be a robust and useful way to represent and pinpoint the root cause of hardware security vulnerabilities. For example, in the NetSpectre attack,⁶ it might be said that the attack root cause is SIMD unit power-on time. STT’s abstraction shows, however, that the root cause is an explicit branch, and that “fixing” the SIMD unit does not prevent the attack.

Even more subtly, the abstraction demonstrates and provides cases where implicit flow and privacy leakage *do occur*, despite not occurring according to program semantics. Consider a simple example where a load is control- and data-independent of a sensitive branch, e.g., “if (`secret == rV`) { `rX <- rW`; } load `rZ <- (rY)`;”. How this load executes is important, to understand whether a potential cache-based covert channel exists. Traditional software-level analysis would indicate that the execution of the load is independent of the secret (the branch outcome). Yet, on a speculative microarchitecture, STT’s abstraction would classify this code as a resolution-based implicit channel.

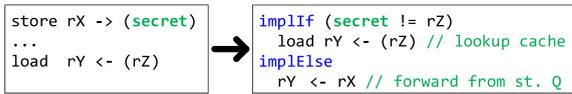


Figure 4. Rewriting a store-load pair as an implicit branch. `implIf` reveals a potential covert channel as a function of memory aliasing to the older store. This occurs if the microarchitecture supports store-to-load forwarding or memory-dependence speculation.

As shown in Figure 5, if the branch mis-speculates and subsequently squashes, the load may execute either once or twice depending on the value of `secret`.

Finally, the abstraction applies to a large set of microarchitectural optimizations. For example, the representation of store-to-load forwarding (see Figure 4) also captures the behavior of memory-dependence speculation with a store set predictor. Here, the store set predictor is modeled as a *prediction* on the implicit branch (`implIf` in the figure). As we will see, being able to represent different optimizations as *predictions on implicit branches* will enable STT to apply a uniform mechanism to block leakage through a variety of structures (e.g., branch, store set, etc., predictors).

STT: DESIGN

Framework and Concepts

STT requires that the microarchitect define what instructions write secrets into registers (*access instructions*, mainly loads), what instructions can form explicit channels (*transmitters*), and what instructions form implicit channel branch predicates (for both explicit and implicit branches). Finally, the architect must define the *Visibility Point*, after which speculation is considered safe (e.g., at the point of the oldest unresolved branch, or at the head of the ROB). If the visibility point refers to an instruction older than an access instruction, we call the access instruction *unsafe*; otherwise it is considered *safe*.

We provide guidelines for microarchitects on identifying access and transmit instructions. An instruction should be classified as an access instruction if it has the potential to read a secret. Except for loads, there are only

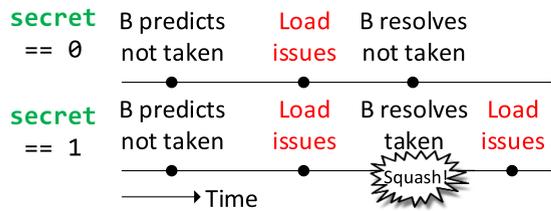


Figure 5. Resolution-based implicit channel due to secret-dependent pipeline squashes. When the branch (B) resolves, it leaks the secret based on whether a squash occurs, as this causes the younger load to execute once or twice. There is an analogous case when the (public) predictor state takes the branch.

a handful of such instructions, which can be identified manually.

An instruction should be classified as a transmit instruction if its execution creates operand-dependent resource usage that can reveal the operand (partially or fully). Identifying implicit branches is similar: the architect must analyze whether the resource usage of some in-flight instruction changes as a function of *some other* instruction’s operand. This definition can be formalized by analyzing (offline) how information flows in each functional unit at the SRAM-bit and flip-flop levels to determine whether resource usage depends on the input value, in the style of the OISA¹⁰ or GLIFT⁸ formal frameworks. Automatically performing such analysis is important future work.

Taint and Untaint Propagation

Conceptually, in each clock cycle, STT applies the following taint rules to instructions in the ROB:

The Output Register of an Access Instruction is tainted if and only if the access instruction is unsafe.

The output register of a non-access instruction is tainted if and only if at least one of its input operands is tainted.

In the implementation, taint propagation is piggybacked on the existing register renaming logic in an out-of-order core. Tainting is therefore fast. In contrast, it is difficult to propagate “untaint” to

all dependencies of an access instruction that becomes safe in a single cycle. We address this with a single-cycle implementation for untaint in the “STT: Implementation” section.

Unlike prior DIFT schemes,⁷ STT does not require tracking taint in any part of the memory system or across store-to-load forwarding. The reason is that because loads are access instructions, the taint of their output is determined only based on whether they have reached the visibility point. That is, the output of an unsafe load is always tainted.

Blocking Covert Channels

Given STT’s rules for tainting/untainting data and its abstraction for covert channels, STT blocks all covert channels by applying a uniform rule across each type.

Blocking Explicit Channels STT blocks explicit channels by delaying the execution of any transmit instruction whose operands are tainted until they become untainted. This scheme imposes relatively low overhead because it only delays the execution of transmit instructions if they have tainted operands. For example, a load that only *reads* a (potential) secret but does not transmit one—such as the load on line 2 in Figure 1—executes without delay. The load on line 3, however, will be delayed and eventually squashed, thereby defeating the attack.

Blocking Implicit Channels STT blocks implicit channels by enforcing an invariant that the sequence of instructions fetched/executed/squashed never depends on tainted data. That is, *STT makes the program counter independent of tainted data*. To enforce this invariant efficiently, without needing to delay execution of instructions following a tainted branch, we introduce two general principles to neutralize the sources of implicit channels:

Prediction-Based Implicit Channels are eliminated by preventing tainted data from affecting the state of any predictor structure.

Resolution-based Implicit Channels are eliminated by delaying the effects of branch resolution until the branch’s predicate becomes untainted.

STT’s principles can be applied to efficiently make *any* hardware predictor impossible to exploit as a covert channel for leaking speculatively accessed data.

Conceptually, the protection mechanism does not need to reason about whether an implicit channel is caused by an explicit or implicit branch: both types have a predicate and the policy with respect to the predicate is the same in both cases. The implementation, however, must identify the predicate. We illustrate this by showing how the STT microarchitecture handles explicit branches.

Applying Principle #1 (Prediction-Based Channels). STT requires that every frontend predictor structure be updated based only on untainted data. This makes the execution path fetched by the frontend unaffected by the output of unsafe access instructions. STT passes a branch’s resolution results to the direct/indirect branch predictors only after the branch’s predicate and target address become untainted; if the branch gets squashed before this, the predictor will not be updated.

Figure 6(c) demonstrates the effect of STT on a speculative execution of the code snippet in Figure 6(a), in which the branch B0 is mispredicted as taken. No matter how many experiments the attacker runs, the predicted direction of the branch B will not be a function of secret, because the branch predictor is not updated when B resolves. As a result, the execution path does not depend on secret (top versus bottom)—it only depends on the predicted branch direction (left versus right).

Applying Principle #2 (Resolution-Based Channels). STT delays squashing a branch that resolves as mispredicted until the branch’s predicate becomes untainted. As a result, a doomed-to-squash branch with a tainted predicate (such as the branch B in Figure 6(c)) will never be squashed and re-executed, preventing the implicit channel leak discussed in the “Insights From Analysis of Implicit Channels” section. As Figure 6(c) shows, the doomed-to-squash branch B is eventually squashed once an older (mispredicted) branch with an untainted predicate squashes. Thus, the squash does not leak any information about the branch’s resolution. Importantly, it is safe to

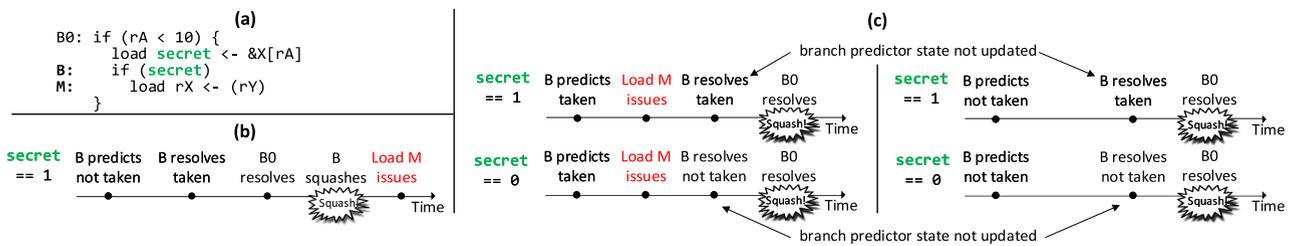


Figure 6. STT executing the code in (a), which includes an untainted branch B0, an access instruction reading `secret`, and an implicit channel (due to branch B). (a) Implicit channel formed through the squash/control dependency on B. (b) When earlier branch B0 is predicted correctly. (c) When earlier branch B0 is predicted incorrectly (left: B predicts taken, right: B predicts not taken).

resolve a branch *as soon as* its predicate becomes untainted, even if an *older branch with a tainted predicate* has not yet resolved.

STT only increases the latency of *recovering* from a *tainted branch* misprediction. For example, in Figure 6(b), the load does not execute immediately after B resolves. Fortunately, tainted branch mispredictions are only a small fraction of overall branch mispredictions, which are infrequent in the first place because successful speculation requires accurate branch prediction.

Implicit Branches. The STT paper applies the above principles to secure several common microarchitectural optimizations that can be formulated as implicit branches, namely: store-to-load forwarding, memory dependence speculation, and memory consistency speculation. In the process, the paper details various optimizations and cases which arise when dealing with implicit channels. In particular: whether the explicit/implicit branch has a prediction step, can be resolved early or can be optimized in some other way. For example, because store-to-load forwarding can only result in two observable outcomes (issue the load or forward from a prior store), we hide which one occurs by unconditionally accessing the cache.

STT: IMPLEMENTATION

We previously assumed untaint information propagated along data dependencies instantly. This is difficult to implement in hardware because a word of tainted data may be a function of complex dependence chains involving many access instructions.

A tainted register needs to be untainted once all access instructions on which it depends reach the visibility point, i.e., become safe. Our key observation is that it suffices to track only when the *youngest access instruction* becomes safe, because instructions become non-speculative in program order in the processor ROB. We call this youngest access instruction the youngest root of taint (YRoT).

Determining the YRoT is done through modifications to rename logic in the processor frontend. Specifically, the YRoT for an instruction *X* being renamed is given by the max of 1) the YRoT(s) of the instruction(s) producing the arguments for *X*, if those instructions are not access instructions; or 2) the ROB index of the instruction(s) producing the arguments for *X*, otherwise. (By convention, we assume the ROB index increases from ROB head to tail.) After rename, the YRoT is stored alongside the instruction in its reservation station and is conceptually an extra dependence for that instruction. When the visibility point changes, its new position is broadcast to in-flight instructions, akin to a normal writeback broadcast, and instructions whose YRoT is less than the visibility point's new position are allowed to execute (assuming their other dependencies are satisfied). The entire architecture requires modest changes to the frontend rename logic, storage in reservation stations for the YRoT, and logic to compare the YRoT to the visibility point which is comparable to normal instruction wakeup logic.

Figure 7 shows an example. Assume the Spectre attack model, i.e., that the visibility point will be set to the ROB index of the oldest unresolved branch. The ROB contains three unresolved

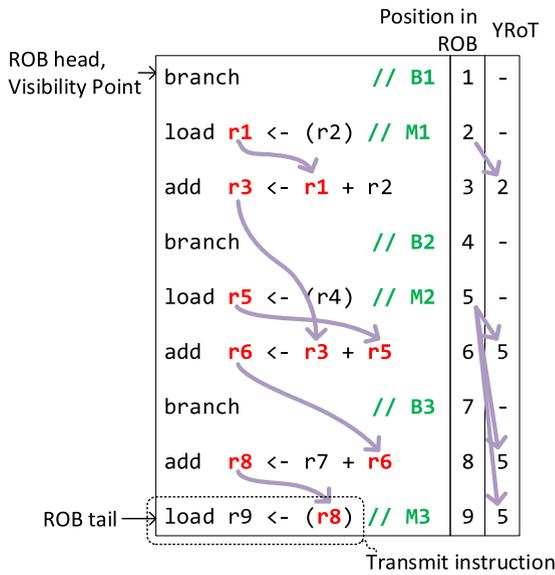


Figure 7. Example of YRoT tracking showing a snapshot of ROB state. Addition (add) instructions are used to represent arithmetic (non-loads). If the YRoT is set to “-,” that means the instruction’s youngest dependent access instruction is a part of retired state.

branches (B1–B3) and a transmit instruction (M3) whose operand/address r8 is a function of the return value of two access instructions (M1 and M2). M3 is a transmit instruction (because it is a load) and can potentially leak secrets because mis-speculations on branches B1 and B2 can influence the data returned by loads M1 and M2, which in turn contribute to the address of M3 through data dependencies.

On one hand, the data dependence chain from load M1 all the way to load M3 is quite complex. That is, the instruction at ROB index 6 depends

on index 5 and index 3, index 8 depends on 6, etc. Re-traversing this dataflow graph to propagate untaint, akin to tracing backwards slices, would be expensive. On the other hand, the YRoT dependence chain is relatively simple. Each instruction just tracks whichever is the youngest load that contributes to its dependence chain (e.g., load M2 for instructions 6, 8, and 9). When branches B1 and B2 resolve, the visibility point advances to branch B3 (ROB index 7). Since 7 is greater than 5 (the YRoT for the transmit instruction M3), M3 is now allowed to execute. Note, the dependence chain could have been more complex, with additional branches and arithmetic dependencies separating load M2 and load M3, but this would not change the moment that it is safe to execute load M3.

Importantly, the above scheme is only secure after applying STT’s mechanisms to block both explicit and implicit channels (see the “STT: Design” section). That is, the scheme requires that r8 is not a function of speculative data at the exact moment load M2 becomes nonspeculative. This requires that branch B3 not be influenced by speculative data (achieved by protections for implicit channels) and that other intervening instructions that can cause explicit channels not execute until they are likewise safe (achieved by protections for explicit channels).

FORMAL ANALYSIS/SECURITY PROOF

We formally prove¹² that STT enforces a novel notion of noninterference: at each step of the execution, the value of a *doomed* register—a register written to by a bound-to-squash access instruction—does not influence future visible events in the execution. This applies to all micro-architectural timing- and interference-based attacks. For instance, the property ensures that the program’s completion time and hardware resource usage—for all hardware structures including cache, branch predictor, etc.—is completely independent of doomed values.

The key challenge in the analysis is how to avoid “looking into the future” to determine if an instruction is doomed to squash. We address this by running the STT machine alongside a nonspeculative in-order processor, which allows

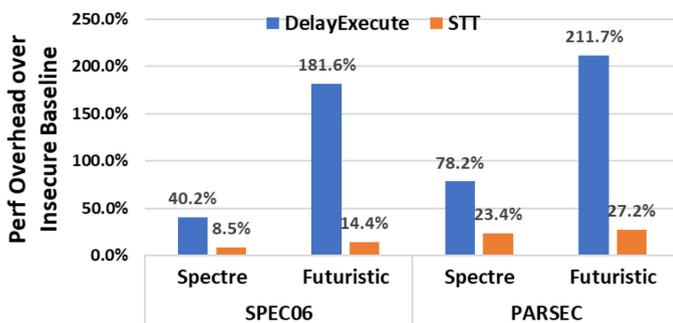


Figure 8. Performance evaluation on SPEC06 and PARSEC benchmark suites. STT outperforms the baseline secure scheme (DelayExecute) with much smaller performance overhead, for both Spectre and Futuristic attacker models.

us to verify the STT machine's branch predictions and determine whether a prediction leads to mis-speculation.

EVALUATION RESULTS

We evaluate STT on 21 SPEC and 9 PARSEC workloads. The results are shown in Figure 8. Relative to an insecure machine, STT adds only 13.0%/18.2% overhead (averaged across both SPEC and PARSEC benchmarks), depending on whether the attack model considers only control-flow speculation (Spectre) or all types of speculation (Futuristic). This indicates that defending against stronger attack models is viable with STT without sacrificing much performance. Compared to the baseline secure scheme (DelayExecute) described in the introduction, STT reduces overhead by 4.0× in the Spectre model and 10.5× in the Futuristic model, on average.

ACKNOWLEDGMENTS

We thank J. Emer, S. Adve, and S. Mukherjee for very helpful discussions. We would especially like to thank our colleagues at Intel who contributed significant feedback throughout the project's development, in particular, F. Liu, M. Fernandez, F. McKeen, and C. Rozas. This work was supported in part by NSF under Grant CNS-1816226; in part by Blavatnik ICRC at TAU, ISF under Grant 2005/17; and in part by an Intel Strategic Research Alliance Grant.

REFERENCES

1. O. Aciicmez, J.-P. Seifert, and C. K. Koc, "Predicting secret keys via branch prediction," *Proc. 7th Cryptographers' Track RSA Conf. Topics Cryptology*, 2007, pp. 225–242, doi: [10.1007/11967668_15](https://doi.org/10.1007/11967668_15).
2. A. Bhattacharyya *et al.*, "SMoTherSpectre: Exploiting speculative execution through port contention," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2019, pp. 785–800.
3. J. A. Goguen and J. Meseguer, "Security policies and security models," in *Proc. IEEE Symp. Secur. Privacy*, 1982, Art. no. 11.
4. J. Groschädl, E. Oswald, D. Page, and M. Tunstall, "Side-channel analysis of cryptographic software via early-terminating multiplications," in *Proc. Int. Conf. Inf. Secur. Cryptology*, 2009, pp. 176–192.
5. P. Kocher *et al.*, "Spectre attacks: Exploiting speculative execution," in *Proc. IEEE Symp. Secur. Privacy*, 2019, pp. 1–19.
6. M. Schwarz, M. Schwarzl, M. Lipp, and D. Gruss, "Netspectre: Read arbitrary memory over network," in *Proc. Eur. Symp. Res. Comput. Secur.*, 2019, pp. 279–299.
7. G. E. Suh, J. W. Lee, D. Zhang, and S. Devadas, "Secure program execution via dynamic information flow tracking," in *Proc. 11th Int. Conf. Archit. Support Program. Lang. Operating Syst.*, 2004, pp. 85–96.
8. M. Tiwari, H. M. Wassel, B. Mazloom, S. Mysore, F. T. Chong, and T. Sherwood, "Complete information flow tracking from the gates up," in *Proc. 14th Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, 2009, pp. 109–120.
9. Y. Yarom and K. Falkner, "Flush+Reload: A high resolution, low noise, L3 cache side-channel attack," in *Proc. Usenix Secur. Symp.*, 2014, pp. 719–732.
10. J. Yu, L. Hsiung, M. E. Hajj, and C. W. Fletcher, "Data oblivious ISA extensions for side channel-resistant and high performance computing," in *Proc. 26th Netw. Distrib. Syst. Secur. Symp.* [Online]. Available: <https://eprint.iacr.org/2018/808>.
11. J. Yu, N. Mantri, J. Torrellas, A. Morrison, and C. W. Fletcher, "Speculative data-oblivious execution (SDO): Mobilizing safe prediction for safe and efficient speculative execution," in *Proc. Int. Symp. Comput. Archit.*, 2020.
12. J. Yu, M. Yan, A. Khyzha, A. Morrison, J. Torrellas, and C. W. Fletcher, "Speculative taint tracking (STT): A formal analysis," Univ. Illinois Urbana–Champaign and Tel Aviv Univ., Tech. Rep., 2019. [Online]. Available: http://cwffletcher.net/Content/Publications/Academics/TechReport/stt-for-mal-tr_micro19.pdf.

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

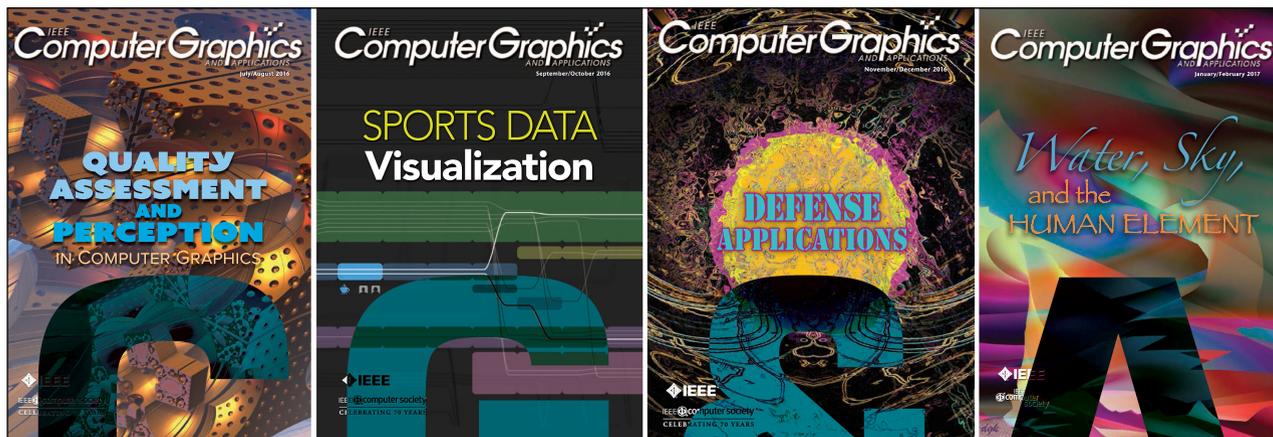
Mengjia Yan is an Assistant Professor in the Electrical Engineering and Computer Science Department, Massachusetts Institute of Technology. Her research interest lies in the areas of computer architecture and hardware security, with a focus on side channel attacks and defenses. Yan received the Ph.D. degree from the University of Illinois at Urbana–Champaign (UIUC). Contact her at mengjia@csail.mit.edu.

Artem Khyzha is a Postdoctoral Fellow in the School of Computer Science, Tel Aviv University. His research interests include formal methods for software and hardware systems. Khyzha received a joint Ph.D. degree in computer science from Technical University of Madrid and IMDEA Software Institute. Contact him at artkhyzha@mail.tau.ac.il.

Josep Torrellas is the Saburo Muroga Professor of Computer Science at the University of Illinois at Urbana–Champaign (UIUC). He is the Director of the Center for Programmable Extreme Scale Computing and past Director of the Illinois-Intel Parallelism Center. His research interests include computer architecture and parallel processing. Torrellas received the Ph.D. degree from Stanford University. Contact him at torrellas@cs.uiuc.edu.

Adam Morrison is an Assistant Professor in the School of Computer Science, Tel Aviv University. His research interests include high-performance trustworthy computer and systems architecture. Morrison received the Ph.D. degree in computer science from Tel Aviv University. Contact him at mad@cs.tau.ac.il.

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



IEEE Computer Graphics and Applications bridges the theory and practice of computer graphics. Subscribe to CG&A and

- stay current on the latest tools and applications and gain invaluable practical and research knowledge,
- discover cutting-edge applications and learn more about the latest techniques, and
- benefit from CG&A's active and connected editorial board.