

CounterPoint: Using Hardware Event Counters to Refute and Refine Microarchitectural Assumptions

Nick Lindsay
Yale University
New Haven, Connecticut, USA
nick.lindsay@yale.edu

Anurag Khandelwal
Yale University
New Haven, Connecticut, USA
anurag.khandelwal@cs.yale.edu

Caroline Trippel
Stanford University
Stanford, California, USA
trippel@stanford.edu

Abhishek Bhattacharjee
Yale University
New Haven, Connecticut, USA
abhishek@cs.yale.edu

Abstract

Hardware event counters offer the potential to reveal not only performance bottlenecks but also detailed microarchitectural behavior. In practice, this promise is undermined by their vague specifications, opaque designs, and multiplexing noise, making event counter data hard to interpret.

We introduce CounterPoint, a framework that tests user-specified microarchitectural models—expressed as μ path Decision Diagrams—for consistency with performance counter data. When mismatches occur, CounterPoint pinpoints plausible microarchitectural features that could explain them, using multi-dimensional counter confidence regions to mitigate multiplexing noise. We apply CounterPoint to the Haswell Memory Management Unit as a case study, shedding light on multiple undocumented and underdocumented microarchitectural behaviors. These include a load-store queue-side TLB prefetcher, merging page table walkers, abortable page table walks, and more.

Overall, CounterPoint helps experts reconcile noisy hardware performance counter measurements with their mental model of the microarchitecture—uncovering subtle, previously hidden hardware features along the way.

CCS Concepts: • Computing methodologies → Model development and analysis; • Software and its engineering → Virtual memory.

Keywords: hardware event counters; address translation

ACM Reference Format:

Nick Lindsay, Caroline Trippel, Anurag Khandelwal, and Abhishek Bhattacharjee. 2026. CounterPoint: Using Hardware Event Counters to Refute and Refine Microarchitectural Assumptions. In *Proceedings of the 31st ACM International Conference on Architectural*



This work is licensed under a Creative Commons Attribution 4.0 International License.

ASPLOS '26, Pittsburgh, PA, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2359-9/2026/03

<https://doi.org/10.1145/3779212.3790145>

Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '26), March 22–26, 2026, Pittsburgh, PA, USA. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3779212.3790145>


1 Introduction

Hardware event counters (HECs) are specialized registers embedded in CPUs and hardware accelerators that provide low-overhead, fine-grained insights into microarchitectural behavior during execution. First introduced in the 1980s—most notably in the DEC VAX and early RISC machines—HECs were originally designed to support performance tuning and system-level debugging.

Since then, their role has expanded. While HECs remain essential for identifying performance bottlenecks [40, 42, 63, 118], they are now also used to calibrate microarchitectural software simulators [58, 96], build analytical models of hardware [4, 8, 63], correlate microarchitectural activity with power and thermal behavior [27, 56, 59, 62, 98, 119], and more [1, 2, 9, 36, 87, 112, 120]. As their utility has grown, so has their prevalence: modern x86-64 processors now expose thousands of HECs—more than a 10× increase since 2009 (Figure 1a).

The promise of a broad set of HECs. In principle, a rich set of HECs should allow experts to gain deeper insight into their mental model of the hardware, even without access to proprietary RTL or internal documentation. These insights are crucial for building accurate performance models and calibrating architectural simulators for future hardware [4, 12, 15, 20, 21, 23, 24, 94], and go beyond traditional HEC uses that measure only broad performance metrics like CPU and memory utilization [43, 69, 82, 99, 101, 102, 116].

Address translation provides a prime example. Modern processors devote many HECs to this function—for instance, IBM's Power9 includes 96 HECs dedicated solely to address translation [55]. Researchers often attempt to leverage these counters to reverse-engineer address translation hardware, enabling accurate integration into software simulators and analytical models of system performance [4, 12, 15, 20, 21, 23, 24, 94]. These models commonly assume specific behavior



(a) The estimated number of HECs events in x86-64 systems increased over 10X between 2009 and 2019. (Y-axis plotted on a log-scale.) (b) The number of model constraints implied by a model scales with the number of HECs used to understand hardware behavior. (c) HEC measurement noise increases with active HECs; beyond a point, model constraint violations can no longer be reliably detected.

Figure 1. The rapid growth of HECs has increased manual effort to construct and compose model constraints, and amplified multiplexing noise that obscures constraint violations. (a) The blue line shows the number of documented HEC ‘names’, assuming a single core. The red line shows the number of ‘addressable’ events after accounting for per-core replication and the conservative removal of events that, while still documented (and potentially informative), have been deprecated by the vendor. Each red data point represents a microarchitecture paired with its typical core count in a server system. This graph shows only documented events and does not include the thousands of additional undocumented HECs identified in recent work [117]¹. (b) The number of model constraints grows superlinearly with the number of HECs (our x-axis shows increasing HEC count for an Intel Haswell MMU, in steps associated with all the HECs in a logical group; e.g., 10 HECs for L2 TLB events) and worsens significantly when including hypothetical HECs across all MMU caches (shown in green). (c) For a representative model constraint on the Intel Haswell MMU ((1) in Table 1), we show that as measurement noise increases—both overall and for individual HECs—it becomes impossible to determine whether the model constraint is violated with 99% confidence once 19 HECs are active. Here, noise is defined as the standard deviation in the observed HEC values.

for the Paging Directory Entry (PDE) cache, which eliminates memory accesses to non-leaf page table levels during a walk. It is assumed that the PDE cache is accessed exactly once per page-table walk, implying that the number of PDE cache misses (`load.pde$miss`) should not exceed page walks (`load.causes_walk`):

$$\text{load.pde\$miss} \leq \text{load.causes_walk}$$

Surprisingly, our measurements on Intel Haswell show that this expected relationship—a sanity check of the expert’s mental model, which we term a *model constraint*—does not always hold. This challenges a widely held assumption in address translation research and casts doubt on the validity of much simulation-based work that relies on it. This example also illustrates two benefits of having a diverse set of HECs.

First, a diverse set of HECs helps detect violations of the model constraints associated with an expert’s mental model of the hardware. Here, we can spot the violation only because Haswell exposes the `load.causes_walk` HEC—which many other processors lack. Without it, researchers rely on generic TLB miss HECs that miss such nuances.

Second, a diverse set of HECs helps explain *why* a model constraint is violated, enabling refinement toward a more accurate representation of the hardware. For example, comparing counters for retired TLB misses, PDE cache misses,

and page table walks uncovers two likely undocumented behaviors: (i) merged walks to the same virtual address occurring after PDE cache lookup, and (ii) aborted translation requests that terminate after PDE cache lookup but before a page table walk begins. Without the full set of HECs, these effects would remain hidden.

The reality of a broad set of HECs. The PDE cache example is an ideal case where a diverse set of HECs helps reveal that there is a flaw in the expert’s original mental model of the microarchitecture. In practice, however, experts are rarely able to leverage the full introspective power of HECs because they rely on manual and ad hoc approaches to doing so. In particular, they face two challenges:

First, understanding microarchitectural behavior requires identifying how HECs relate to one another—that is, determining *all* the model constraints that observed HEC data must satisfy to align with an expert’s mental model (*i.e.*,

¹We counted the total number of HEC names in the Linux `perf` counter database [53, 104] to determine the set of ‘Named’ events per microarchitecture. We estimated the number of ‘Addressable’ events by (i) conservatively removing deprecated HECs, (ii) distinguishing between core and uncore HECs, and (iii) accounting for per-core replication. We account for per-core replication by summing together the uncore counters with the number of core counters multiplied by the typical core count of server systems of that microarchitecture. We will further explore these trends in future work.

their set of assumptions about the microarchitectural implementation). Our PDE cache model constraint illustrates how surprisingly difficult it can be to reason about even simple relationships involving just two HECs. As more HECs are used to check whether observed behavior matches expert expectations, the number of model constraints grows super-linearly (Figure 1b). These model constraints become increasingly complex, often involving dozens of HECs in intricate relationships that are hard to reason about (Section 2). Manual approaches to identifying and evaluating model constraints quickly become intractable. The challenge worsens when observed HEC values violate model constraints, forcing experts to revise their mental models—and then deduce entirely new, complex sets of associated model constraints.

Second, modern architectures allow recording thousands of *logical* HECs, but these are multiplexed onto a much smaller number of *physical* HECs—typically just 4 to 8 at a time. This means that HEC measurements are approximate rather than exact, leading to measurement noise that makes it even more difficult to evaluate model constraints. Multiplexing noise typically grows rapidly with the number of HECs being measured. Beyond a point, the growing number of HECs makes it nearly impossible to determine whether a representative model constraint is truly violated (Figure 1c).

Extracting the promise of HECs with CounterPoint.

To bridge the gap between the promise and reality of HECs, we invent CounterPoint²—a framework that helps experts reconcile HEC data with their mental models of the microarchitecture. CounterPoint automates the demanding task of generating all model constraints associated with a model and checking them against noisy HEC data, enabling exploration of the accuracy of a wider range of microarchitectural models. CounterPoint is centered around three key insights.

First, experts can more naturally express their envisioned hardware as a directed acyclic graph (DAG) linking hardware components to HEC activity, rather than directly constructing model constraints. DAGs are well-suited to formal tools that can automatically derive these constraints. We introduce the μ path Decision Diagram (μ DD)—a specialized DAG for capturing an expert’s mental model of microarchitectural structures and how interactions among these structures increment HECs. A μ DD concisely describes a set of microarchitectural execution paths (μ paths) that micro-ops (μ ops) may follow. Each μ path is associated with specific microarchitectural events, including those that increment performance counters, enabling natural and complete generation of all constraints implied by the model.

Second, experts can refine microarchitectural models more effectively when model constraints are as tightly upper- and lower-bounded as possible (e.g., constraint (3) in Table 1 is

most useful when its left-hand side tightly lower-bounds the number of memory references in a page table walk). Tight constraints increase sensitivity to even minor deviations from the expert’s mental model. Such tightness is more likely when HEC relationships are expressed at the granularity of micro-ops, enabling precise attribution of events to specific hardware behaviors. μ DDs naturally exploit this by modeling micro-op flows through execution paths, yielding model constraints with inherently tight bounds.

Third, while more HECs increase multiplexing noise, they also increase intrinsic correlations (e.g., page table walks often correlate with TLB misses). These correlations allow statistical methods to build tight *counter confidence regions*—ranges of HEC values likely to occur with a given probability from noisy data. Compared to traditional methods that treat counter noise independently [8, 69], this approach substantially reduces the impact of noise, enabling CounterPoint’s automated analysis to scale well beyond the number of physically available HECs.

The CounterPoint approach. Experts begin by expressing their mental model of the microarchitecture in a domain-specific language, which CounterPoint translates into a μ DD (Figure 2). Experts also run workloads on the target hardware, collecting as many active HECs as needed for analysis.

Given a μ DD, CounterPoint applies convex geometry techniques to derive the *model cone*—all HEC value combinations producible by micro-ops traversing the μ DD. The model cone represents the values that simultaneously satisfy all model constraints, and eliminates the need for manual derivation. CounterPoint then processes noisy HEC measurements from real hardware, extracting intrinsic correlations to define tight *counter confidence regions*: ranges of HEC values inferred with high confidence despite multiplexing noise.

Finally, with feasibility testing, CounterPoint compares the counter confidence region against the model cone. If they do not intersect, the expert’s model is inconsistent with the HEC observations, implying that some model constraints are violated. CounterPoint reports these violations, guiding how the μ DD may be revised for consistency. This enables iterative exploration: the expert proposes new μ DDs, and CounterPoint tests them until a consistent model is found.

Evaluating CounterPoint via a case study. We demonstrate CounterPoint’s capabilities by applying it to the Intel Haswell Memory Management Unit (MMU), where we uncover several previously undocumented and underdocumented features. These include a load-store queue-side TLB prefetcher (as well as its trigger conditions and interaction with page hotness tracking), hardware mechanisms that merge and abort page table walks, and a cache for the root level of the page table. The Haswell MMU serves as a compelling case study: it embodies complex hardware–software interactions, and has been foundational for a decade of address translation research [4, 7, 41, 63, 72, 76, 86, 103, 108,

²CounterPoint enables using hardware event *counters* to *point* out gaps in an expert’s understanding of microarchitectures and makes it easier to explore improvements or *counterpoints* to their assumptions.




Figure 2. CounterPoint automatically determines the feasibility of a microarchitectural model against HEC data. Models are described using a DSL and transformed into a μ path Decision Diagram (μ DD), which is analyzed to determine the model cone (the set of model constraints). Counter confidence regions are constructed for each observation to handle multiplexing noise. Observations are tested against all model constraints simultaneously. CounterPoint’s counter confidence region bounds are sharper than other approaches, enabling more violations to be identified, and thereby enabling more opportunities to refine the expert’s microarchitectural assumptions. CounterPoint effortlessly supports dozens of HECs and constraints.

115, 121]. Yet, it is poorly modeled in state-of-the-art software simulators [38, 39, 65, 114], motivating recent efforts to use HECs to reverse-engineer accurate models [4, 12, 15, 20, 21, 23, 24, 94]. As a rigorous test of CounterPoint’s analysis of sophisticated microarchitectural behavior, the Haswell MMU case study provides a foundation for extension to other components and more modern microarchitectures.

Technical contributions. Overall, this work:

- Defines HEC model constraints and demonstrates their ability to expose hidden microarchitectural behavior.
- Introduces the μ DD, a compact representation that encodes both microarchitectural assumptions and HEC semantics.
- Defines the model cone, shows how it can be naturally derived from the μ DD, and proves its equivalence to the model constraints of the μ DD.
- Applies counter confidence regions to mitigate measurement noise, enabling reliable inference even when the number of counters exceeds hardware limits.
- Develops μ DD feasibility testing to automatically validate measured HEC data against a μ DD’s implied constraints.
- Reveals several likely undocumented and underdocumented features in a commercial Intel CPU—including TLB prefetching, early paging-structure cache lookups, and merged page table walks—using CounterPoint’s automated analysis.

In sum, CounterPoint uses HECs to refine expert understanding of hardware—challenging incorrect assumptions and uncovering subtle, otherwise hidden effects. Such insights are essential for building trustworthy models as computer systems grow increasingly complex and opaque—we therefore publicly release CounterPoint³.

2 CounterPoint: A Bird’s-Eye Overview

The pros and cons of model constraints. Model constraints are valuable because they let experts identify exactly

when and how their assumptions about the microarchitecture break down. The HECs involved in a violated model constraint highlight which parts of the model may be incorrect. However, to be fully effective, all (often dozens of) model constraints must be enumerated, and each must be correct and tight. By tight, we mean the bounds leave minimal slack: loose constraints can miss infeasible observations, whereas tight constraints clearly delineate what is possible versus impossible, making violations easier to detect.

Manually deriving all the constraints is onerous, even for an expert. Table 1 shows just a subset of constraints for a simple Intel Haswell MMU model. Each may involve many HECs and depend on the intersection of multiple microarchitectural assumptions.

Worse, constraints are easy to formulate either too loosely or incorrectly. For example, one might bound the number of page walker loads on Haswell, with its four-level page table: $walk_ref \leq 4 \cdot (load.causes_walk + store.causes_walk)$ This is correct but not tight, since it ignores page size and MMU cache hits (unlike Constraint 2 in Table 1).

Alternatively, one could try to exploit the fact that larger pages shorten page table walks:

$$walk_ref \leq 4 \cdot load.walk_done_4k + 4 \cdot store.walk_done_4k + 3 \cdot load.walk_done_2m + 3 \cdot store.walk_done_2m + 2 \cdot load.walk_done_1g + 2 \cdot store.walk_done_1g$$

But this version is too strong: it rejects valid cases where walks inject memory accesses but do not terminate (e.g., invalid translations). As Constraint 2 shows, the tightest correct bound is actually a far more nuanced relationship.

Even simpler constraints require tightness. For instance, Figure 3a shows an infeasible observation of HEC values detectable only with enough relevant constraints. With fewer or irrelevant counters (Figures 3b and 3c), the violation slips through. When scaling to dozens of model constraints, many of which include complex relationships among dozens of HECs each, all these problems compound.

³CounterPoint will be maintained at: <https://github.com/NicholasLindsay/counterpoint-public>

Table 1. The Haswell MMU requires reasoning about dozens of model constraints; we show three representative examples. Deriving the exact constraints for a given model is challenging because they stem from subtle microarchitectural assumptions. For example, Constraint 1 relies on expert knowledge that no retired TLB miss suffered a prior page fault. Constraint 2 relies on even more subtle knowledge that an upper bound on the number of memory references injected by a page table walker is determined by (i) PDE cache hit/miss status; (ii) the page size of the translation and (iii) the fact that every walk makes at least one memory reference. For brevity, we define: $\text{walk_ref} \triangleq \text{walk_ref}.\text{11} + \text{walk_ref}.\text{12} + \text{walk_ref}.\text{13} + \text{walk_ref}.\text{mem}$.

(1) $\text{load.ret_stlb_miss} \leq \text{load.walk_done}$	2 HECs
Every TLB-miss micro-op that retires must have obtained a valid translation from a page-table walk.	
(2) $\text{walk_ref} \leq \text{load.causes_walk} + \text{store.causes_walk} + 3 \cdot \text{load.pde\$_miss} + 3 \cdot \text{store.pde\$_miss}$ $- \text{load.walk_done_2m} - \text{store.walk_done_2m} - 2 \cdot \text{load.walk_done_1g} - 2 \cdot \text{store.walk_done_1g}$	12 HECs
The number of memory accesses made by the page table walker is upper bounded by the distribution of combinations of page sizes and PDE cache interactions.	
(3) $\text{load.causes_walk} + \text{store.causes_walk} + \text{load.walk_done_1g} + \text{store.walk_done_1g} \leq \text{walk_ref}$	8 HECs
Every page table walk must result in one or more page table walker memory accesses. Walks that complete with 1GB page emit two memory references when the MMU cache for the root page table level is absent.	




Figure 3. The ability of HECs to test assumptions depends on their number and semantics, shown here pictorially. The orange regions represent points which satisfy all model constraints; the blue dot represents an observation; the red and green boxes represent two alternative constructions of counter confidence regions. Model constraints correspond to edges in 2D or faces in 3D. (a) Consider a model cone constructed from the three HECs shown. These counters imply three constraints: $\text{load.ret_stlb_miss} \leq \text{load.walk_done}$ because each retired STLB miss must correspond to a unique, successfully completed page table walk; $\text{load.ret_stlb_miss} \leq \text{load.causes_walk}$ because each retired STLB miss must trigger exactly one page table walk; and $\text{load.walk_done} \leq \text{load.causes_walk}$ because only a subset of initiated page table walks ultimately complete. The first two inequalities rely on the assumption that STLB misses are never merged. Using all three HECs clearly exposes a violation of these constraints, indicating a flaw in the expert’s mental model. (b) All three HECs were required to detect this flaw; removing load.walk_done eliminates the second and third constraints, making the model violation undetectable. (c) Simply substituting $\text{load.pde\$_miss}$ (or any other counter) is insufficient, because the semantics of each counter matter. Using this alternative counter adds the constraint $\text{load.pde\$_miss} \leq \text{load.causes_walk}$, but this constraint still fails to reveal the model violation. (d) Counter confidence regions replace point observations with value ranges; exploiting correlations yields tighter bounds than assuming independence.

The geometry underlying model constraints. Model constraints are powerful for validating hardware assumptions but are unscalable as they are derived in an ad hoc manner. The challenge is not in their use, but their derivation. We observe that model constraints naturally arise because microarchitectural events occur in predefined groups rather than in isolation—for example, each completed walk for a 4KB page (load.walk_done_4k) involves 1 to 4 page table walker memory accesses (walk_ref). Our insight is that instead of manually deriving constraints, experts can more easily enumerate all valid groupings, letting CounterPoint automatically determine whether an observed set of events

could result from some combination of groups. This approach enables scalable feasibility checking of the model constraints.

We enable experts to specify how μ ops interact with the microarchitecture—including their effect on HECs—using μ DDs. A μ DD is a specialized DAG where each path represents a single HEC group, enabling automated testing of observed HEC values against feasibility constraints. μ DDs are centered on μ ops because they form a natural unit for grouping microarchitectural events: they are familiar to experts, fine-grained enough to capture low-level hardware interactions, and directly responsible for incrementing HECs. The DAG representation is concise; a few nodes can efficiently describe an exponential number of μ paths.

The group-matching problem naturally induced by the μ DD is fundamentally a counting problem that can be framed in terms of convex geometry. The resulting geometric object—*i.e.*, the model cone—represents all valid combinations of HEC values. The Minkowski–Weyl theorem from computational geometry states that every model cone has two equivalent representations: one as the set of points generated by a μ DD, and the other as the set of points bounded by model constraints [44]. We leverage these dual representations by allowing the expert to express their microarchitectural assumptions in the form most natural to them—by encoding their hardware assumptions in a μ DD—while enabling CounterPoint to automatically deduce model constraints as required for user feedback. CounterPoint derives the model constraints using a custom algorithm which calls into an off-the-shelf *convex hull* solver, as described in Section 6.

Generating tight confidence regions of HEC observations. Identifying flaws in the expert’s mental model requires not only a tight model cone but also computing the narrowest possible range of values that can be confidently inferred from the observed HEC values, despite multiplexing noise.

Standard measurement tools (*e.g.*, `perf`) report the mean and standard deviation of the samples for each HEC, which can be used to construct counter confidence regions. Naive methods assume each HEC is independent, resulting in overly loose counter confidence regions (Figure 3d, green box) that reduce the ability to detect violations of model constraints.

Instead, we discover that HEC values are often correlated, a finding we extract from time-series measurements. These correlations mean that the data typically have far fewer degrees of freedom than the number of counters, allowing us to construct much tighter counter confidence regions—even when dozens of counters are measured (Figure 3d, red box). Tighter counter confidence regions uncover more accurate microarchitectural models.

Feasibility testing for guided model exploration. High-quality models demand detail, with μ DDs describing hundreds of unique execution paths. This produces tight model constraints, but also drives rapid growth in complexity. CounterPoint uses linear programming to efficiently determine

model feasibility, and a conic hull algorithm (Section 6) to derive model constraints when infeasible observations occur. Violated model constraints are reported to the expert, who uses this information to formulate refined μ DDs that resolve these discrepancies and represent more accurate models of the hardware. Naturally, the precision which CounterPoint can infer details about the microarchitectural features depends on having a dataset of HEC observations from a rich and diverse set of programs that stress all relevant corners of the microarchitecture. Consequently, to uncover details of the Haswell MMU in our case study, we evaluated about 20 million HEC samples across a diverse range of workloads with dozens of models, each exhibiting a unique combination of bespoke microarchitectural features. As we use CounterPoint to refine our understanding of the hardware, we continue to expand this set of models.

3 From Diagrams to Model Cones

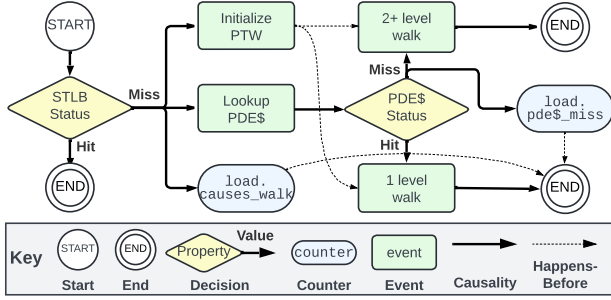
μ path decision diagrams. A μ DD encodes a set of microarchitectural execution paths that individual μ ops may take through part of the microarchitecture (Figure 4a is an example encompassing a subset of address translation hardware). At the core of CounterPoint is the concept of a microarchitectural execution path, or μ path: a happens-before ordered set of hardware events induced by a μ op [52].

Micro-paths are derived from a μ DD by performing a graph search along CAUSALITY edges, with nodes and CAUSALITY edges added to the μ path as they are encountered. When a DECISION node is encountered, there are two possibilities. If the property has been assigned a value (determined by the labels on outgoing edges) earlier in the traversal, then the corresponding outgoing CAUSALITY edge is followed. Otherwise, a concrete property value from the outgoing CAUSALITY edge labels is selected and the corresponding edge is followed. This process continues until all nodes and CAUSALITY edges have been added. HAPPENS-BEFORE edges between node pairs are instantiated in the μ path if there exists a HAPPENS-BEFORE edge between the corresponding μ DD nodes.

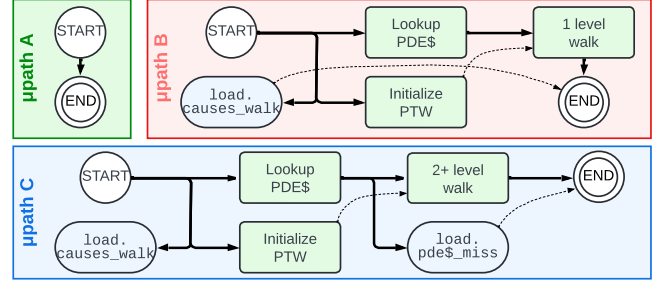
When exhibiting a μ path during its execution, a μ op generates events in a time order that respects both CAUSALITY and HAPPENS-BEFORE edges. Events come in two forms: standard EVENT nodes (green boxes), which represent standard microarchitectural events, and COUNTER nodes (blue pills), which correspond to events directly recorded by HECs.

μ path counter signatures. Each μ path has an associated counter signature—a vector that records how many times each HEC appears within a μ path. This signature captures how a μ op following that μ path increments the HECs.

Counter flow equation. Our first goal is to precisely define when an observed set of HEC values is “feasible” with respect to a μ DD. We do so via the *counter flow equation*, which links HEC values to the number of μ ops traversing microarchitectural execution paths.



(a) Example μ DD that describes how a μ op interacts with the STLB and PDE cache, in the presence of `load.causes_walk` and `load.pde$_miss` HECs. This μ DD encodes three μ paths.



(b) This μ DD describes three unique μ paths, each corresponding to different assignments to microarchitectural properties (e.g., *STLB Status* and *PDE\$ Status*). Edges represent happens-before order.

Figure 4. A μ DD encodes a set of microarchitectural execution paths (μ paths). Each μ path describes a set of events per μ op.

The key insight behind the counter flow equation is that μ ops increment HECs as they traverse a μ path, creating a direct relationship between the microarchitectural flow of μ ops and the resulting HEC values.

Let \mathcal{D} be a micro flow diagram and $\mathcal{P}(\mathcal{D})$ the set of μ paths it encodes. A microarchitectural flow $f(\cdot)$ assigns each μ path $p \in \mathcal{P}(\mathcal{D})$ a non-negative number of μ ops traversing it. Each μ op on μ path p increments the HECs according to the μ path's counter signature $\vec{S}(p)$ —the vector of counter occurrences along p . Thus, the contribution of μ path p is $f(p) \cdot \vec{S}(p)$, and the total HEC value vector \vec{v} is the sum over all μ paths:

$$\vec{v} = \sum_{p \in \mathcal{P}(\mathcal{D})} \vec{S}(p) \cdot f(p) \quad (\text{Counter Flow Equation})$$

This counter flow equation links observed HEC values to the flow of μ ops through the μ DD, and is only valid when $f(p) \geq 0$ for all μ paths, as negative flows of μ ops are impossible. Intuitively, the final counter values are given by the total number of HEC increments across all dynamic μ op instances.

Deriving the model cone. The model cone is the set of all HEC value combinations generated by valid microarchitectural executions (i.e., those with non-negative flows). We determine observation feasibility by testing if it lies within the model cone, a task accomplished using linear programming. This means that feasibility can be determined even without knowing $f(\cdot)$ exactly.

Mathematically, we define a model cone $K_{\mathcal{D}}$ for a μ DD \mathcal{D} as the set of HEC values that are generated by microarchitectural executions with non-negative flow:

$$K_{\mathcal{D}} \triangleq \left\{ \sum_{p \in \mathcal{P}(\mathcal{D})} \vec{S}(p) \cdot f(p) \mid f(p) \geq 0 \right\} \quad (\text{Model Cone})$$

Geometrically, $K_{\mathcal{D}}$ is a convex⁴ polyhedral⁵ cone⁶ defined purely by the μ path counter signatures in the μ DD (Figure

⁴Convex: If $x, y \in K_{\mathcal{D}}$ then $\alpha x + (1 - \alpha)y \in K_{\mathcal{D}}$ for $0 \leq \alpha \leq 1$.

⁵Polyhedral: Defined by a finite number of equalities and inequalities.

⁶Cone: If $x \in K_{\mathcal{D}}$ then $\alpha x \in K_{\mathcal{D}}$ for $\alpha > 0$.

5a). Intuitively, the model cone represents the space of all allowed HEC combinations.


Generalizability. Our decision to design CounterPoint so that it links fine-grained microarchitectural events and interactions—represented as μ paths of μ ops—with HEC updates is intentional. μ op-centered execution paths have been used extensively in prior work to model low-level hardware, including formal verification of memory consistency and its interactions with coherence and virtual memory [51, 66, 68, 71], side-channel security [52, 105], and functional correctness [52, 66]. Because these approaches cover many aspects of CPU pipelines, they suggest that CounterPoint is well positioned to extend to other microarchitectural components.

4 Feasibility Testing with Noise

An observation is feasible if it resides within the model cone; a problem solvable with linear programming. Unfortunately, observations are subject to multiplexing noise which must be accounted for to prevent false violations (e.g., Figure 5b).

Handling noise with counter confidence regions. Counter confidence regions handle noise by treating each observation not as a single value, but instead as a point drawn from a set of values within which the true value is likely to occur, given the presence of noise in the measurement. The likelihood of the region capturing the true value is given by the confidence level, fixed to 99% for our analyses. The size and shape of the counter confidence region depends on parameters of the underlying distribution, which can be inferred from the HEC measurements themselves. CounterPoint computes covariances (in addition to means and variances computable by perf), producing tight counter confidence regions that are more likely to catch violated constraints.

CounterPoint requires HEC vector samples $\{Y_i\}_{i=1}^M$ recorded at regular time intervals (e.g., every 10 seconds) over the course of a program's execution. Such functionality is provided by standard tools (e.g., perf). CounterPoint computes the sample mean \bar{Y} as a HEC vector representative of



(a) The model cone (valid HEC combinations) is determined by μ path counter signatures. (b) Noise introduced by multiplexing can make valid combinations appear infeasible. (c) Confidence region construction oriented along eigenvectors of covariance matrix.

Figure 5. The model cone is determined purely by the μ path counter signatures (Figure a). Testing observations for inclusion in the model cone is complicated by noise, which can cause observations to spuriously appear infeasible (Figure b). CounterPoint handles noise by constructing confidence regions at the 99% confidence level (Figure c). The counter confidence region is an ellipsoid which CounterPoint approximates by its bounding box, enabling a linear programming formulation. The scale and orientation of the confidence region is determined by (i) the confidence level and (ii) correlations in the observed data. λ_k and \vec{e}_k denote the k th eigenvalue/eigenvector of the estimated covariance matrix.

the entire execution. Statistically, the sample mean is drawn from a Gaussian distribution per the Central Limit Theorem. With Gaussian distributions, the confidence region is fully determined by the sample mean and sample mean covariance [107]. We calculate the HEC covariance matrix Σ_Y . We estimate the sample mean covariance with the plugin estimator $\Sigma_{\bar{Y}} = \frac{1}{M}\Sigma_Y$. This defines the confidence region:

$$\{ \vec{v} \mid (\vec{v} - \bar{Y})^T \Sigma_{\bar{Y}} (\vec{v} - \bar{Y}) \leq \chi_{N,\alpha}^2 \} \quad (\text{Confidence Ellipsoid})$$

Intuitively, this means that the confidence region is an ellipsoid in shape, and that the ground truth (e.g., noise-free) counter value is contained within the ellipsoid with $(1 - \alpha)$ -confidence. We adapt it to a linear program in the following section. The confidence region can be made tighter by obtaining more samples (e.g., with longer running programs), providing the program has consistent steady-state behavior.

Determining feasibility with a linear program. Given a model cone and a counter confidence region, we can assess the feasibility of an HEC observation at a specified confidence level. If the counter confidence region intersects the model cone, the observation is deemed feasible. If there is no intersection, the observed HEC values must violate at least one model constraint at that confidence level. For example, Figure 5c shows a counter confidence region which intersects with the model cone, indicating a feasible observation.

To test for feasibility, CounterPoint uses linear programming because of its efficiency, relative simplicity, and availability in mature software libraries [5, 30, 37, 77]. CounterPoint constructs a linear program⁷ by instantiating non-negative variables for the flow and counter values (see Section 3). The flows and counter values are related by the **Counter Flow Equation**, implicitly describing the model cone.

⁷The linear program is provided in our technical report [64].

The counter confidence region, being a quadratic form, cannot be directly encoded. Instead, we approximate it with a bounding hyper-rectangle (Figure 5c). This bounding box is aligned with the principal components of the data, producing the tightest rectangular bound on the confidence region. Empirically, our bounding box approximation detects many surprising constraint violations (Section 7). We leave alternatives like quadratic programming for future work.


5 Guided Model Exploration

Discovering microarchitectural features. The specific model constraints that are violated guide the expert in identifying which microarchitectural features need to be added or modified in the μ DD to make it feasible. When a model constraint of the form $a \cdot x \leq b \cdot x$ is violated, then for all feasible μ DDs, there must exist a μ path whose μ path counter signature $\vec{S}(p)$ satisfies $a \cdot \vec{S}(p) > b \cdot \vec{S}(p)$.

We illustrate this with an example (Figure 6). Figure 6a is a simple μ DD for load μ ops upon a TLB miss. We assume that the load μ op first initializes the page table walker - incrementing `load.causes_walk`- before looking up the PDE cache. In the event of a cache miss, `load.pde$_miss` is incremented. This model implies model constraint C (6b).

CounterPoint identifies that Constraint C is violated by one or more HEC observations⁸. Therefore there are workloads where `load.pde$_miss` exceeds `load.causes_walk`. To explain this apparent contradiction, we must introduce one or more microarchitectural features into the μ DD that allow for this constraint to be broken. This corresponds to


⁸When an observation is deemed infeasible with respect to an μ DD, CounterPoint automatically tests the observation against each feasibility constraint to identify violations. Deriving and testing the feasibility constraints is a non-trivial procedure; we describe our implementation in Section 6.



(a) Initial model.

$$C \triangleq \text{load.pde\$_miss} \leq \text{load.causes_walk}$$

(b) Violated model constraint.



(c) Refined model.

Pde\$ Status	Abort	load.causes_walk	load.pde\$_miss	Satisfies C
Miss	Yes	0	1	No

(d) Properties of $\mu\text{path } p$ including μpath counter signature.

Figure 6. Modifying μDD to remove constraint violations is equivalent to identifying candidate microarchitectural features. (a) Initial μDD of page table walk. (b) Model implies this model constraint, which is violated. (c) μDD is updated by (i) assuming PDE cache is looked up prior to starting walk, and (ii) allowing translation requests to be aborted before starting a walk. (d) Model no longer implies constraint C as $\vec{S}(p)$ does not satisfy constraint.

modifying the μDD such that it contains $\mu\text{path}(s)$ whose μpath counter signatures explicitly violate C .

One way to resolve this is to assume that (i) the PDE cache is accessed before starting a page table walk, and (ii) translation requests can be aborted between the PDE cache lookup and the start of the walk. This allows lookups to access the PDE cache without incrementing `load.causes_walk`. Applying these assumptions produces a new μDD (Figure 6c), with a new $\mu\text{path } p$ whose μpath counter signature $\vec{S}(p)$ explicitly violates constraint C (Figure 6d). Analysis of this μDD confirms that C is no longer implied, resolving the violation. In practice, many constraints often need resolution, requiring an iterative μDD refinement process.

Classifying microarchitectural models. Feasibility testing partitions the set of μDD into subsets of feasible and infeasible μDD s. It is possible for different μDD s representing different microarchitectural assumptions to be feasible. When this happens, experts can identify common structures




Figure 7. Microarchitectural models (boxes) classified by their features and consistency with hardware performance counter data. Red: inconsistent. Green: consistent.




Figure 8. Expert-in-the-loop heuristic search algorithm navigates the model space without needing to explore the full cross-product of microarchitectural features. F . denote microarchitectural features; models (boxes) are identified by (i) their set of features and (ii) their consistency with HECs.

in feasible μDD s to determine likely hardware features despite the ambiguity introduced by multiple feasible μDD s.

Consider Figure 7. There are four models, each identified by the presence or absence of features F_X and F_Y . Consistent models are highlighted in green and inconsistent models in red. There are two consistent models; introducing ambiguity in what model is the best fit. However, *all* consistent models contain Feature F_Y . If we have covered the relevant feature space—by running a wide enough range of programs to ensure that the hardware we are trying to reverse-engineer is adequately exercised—then CounterPoint can reliably conclude that Feature F_Y *must* be present. On the other hand, Feature F_X in isolation is insufficient to explain the performance counter observations, but it is possible that Feature F_X and Feature F_Y are both present. Given a feature space (e.g., $F_X \times F_Y$), we can infer viable feature combinations.

Enumerating models. Feature discovery and model classification can be employed together to infer the presence of microarchitectural features. We propose an expert-in-the-loop algorithm for this purpose.

Our algorithm accepts an initial μDD and a dataset of HEC observations, and returns a set of μDD characterized by their features and feasibility. The algorithm consists of two phases: *discovery* and *elimination*. We advocate starting with a conservative model to ensure a more informative set of constraints that enable discrimination between candidate features, but the expert can start with any model. Features are discovered through the *Discovery* phase. Figure 8 shows an example search graph generated by the algorithm.

Discovery phase. Constraint violations are detected by CounterPoint, and the expert user eliminates the constraints by introducing new microarchitectural features or modifying existing ones. When more than one feature can eliminate a constraint, all features should be added to their model. This process is repeated until a feasible μ DD is obtained.

In Figure 8, the initial μ DD is shown at the bottom left. Feature F_1 is added to produce a new model $\{F_1\}$; features F_2 and F_3 are then added to create the μ DD at the top of the tree $\{F_1, F_2, F_3\}$. This μ DD is a candidate microarchitectural model for the system. At each iteration step, the model cones are verified to ensure that the model cone is expanded.

Elimination phase. The candidate μ DD may contain more features than required for a feasible model. In the elimination phase, we recommend recursively pruning microarchitectural features until infeasible μ DDs are obtained. This is based on our empirical observation—captured through refinement of close to a hundred models—that pruning infeasible models tends to produce infeasible models, so the sub-tree need not be explored further. In Figure 8, features F_1 - F_3 are removed from the top μ DD to create separate μ DDs. The μ DD $\{F_1, F_3\}$ remains feasible, so features F_1 and F_3 are removed separately, resulting in infeasible μ DDs.

6 Implementing CounterPoint

We implement CounterPoint as a Python library with roughly 3K lines of code, integrating with Pandas [83] for convenient data processing. To support broader community adoption, CounterPoint is designed for easy portability using a reproducible Docker [75] environment. We will share our MMU μ DDs to help seed the development of improved MMU models in widely used software simulators [38, 39, 65, 114].

Domain-specific language for μ DDs. We introduce a simple DSL for specifying μ DDs: ACTION and COUNTER nodes are single-line statements, DONE nodes use the done keyword, and DECISION nodes are expressed with C-style switch cases. The DSL does not support functions, loops, or variables beyond μ path properties. Our DSL serves as a reference implementation, avoiding errors that could arise from deriving μ DDs directly from RTL or C/C++ simulator specifications.

Feasibility testing. Given a μ DD and a set of HEC observations, CounterPoint tests each observation for feasibility by constructing and solving a linear program (the complete linear program is provided in our technical report [64]). This entails enumerating every counter and μ path counter signature, implemented by a breadth-first traversal of the μ DD. The back-end LP solver we use is pulp [77]. Constraint violations are identified by testing infeasible observations against the half-space defined by each constraint.

Deducing model constraints. Model constraints are derived from μ path counter signatures as follows. First, μ path counter signatures are normalized by dividing each element

by the greatest common factor, and duplicates are removed. Second, Gaussian elimination identifies equality constraints and eliminates redundant HECs⁹. Third, μ path counter signatures that lie fully within the interior of the model cone are identified using linear programming and removed. Fourth, the *conic hull* is computed by: (i) adding the zero vector to the set of μ path counter signatures; (ii) computing the *convex hull*; (iii) selecting all faces which contain the origin, corresponding to the faces of the cone. The inequality model constraints are given by the planar equations of the resulting faces. We implemented this custom solution because no Python library computes conic hulls, and standard numeric methods (e.g., QR factorization) are ill-conditioned, whilst symbolic operations preserve exact integer values.

7 A Case Study: The Intel Haswell MMU

We demonstrate CounterPoint’s capabilities and evaluate its usability and performance on the Intel Haswell MMU. This case study shows how CounterPoint can uncover the behavior of advanced microarchitectural components, even when they interact deeply with complex systems software. The Haswell MMU is a strong case study target due to its rich set of HECs [53, 104] and its frequent use in prior research on address translation [4, 7, 33, 63, 85, 115]. Haswell also exhibits complex microarchitectural interactions across data and instruction activity [29], as well as under native and virtualized execution [3, 6, 18, 19, 28, 45, 84, 90, 115]. For these reasons, validating and/or refuting assumptions about the Haswell MMU represents a strong test of CounterPoint’s effectiveness. This foundation positions us to extend our study to more modern architectures. For this study, we focus specifically on data-side activity in native execution. While full confirmation of our findings would require proprietary RTL, CounterPoint enables high confidence conclusions possible even without direct access to the RTL.

7.1 Guided Model Exploration.

Our initial model of the Intel Haswell MMU includes features that are well-established through documentation and prior research [20, 33, 47, 88, 89], and are typically integrated in software simulators. We assume a two-level TLB hierarchy and a four-level page table. Building on reverse-engineering studies of Haswell MMU caches [41, 108], we assume the presence of a PDE cache and an additional MMU cache for the page table level immediately preceding the PDE level. Consistent with conventional wisdom, we further assumed that the PDE cache is consulted once during every walk.

We refine the model using a diverse set of HEC observations from workloads that stress the MMU. We measured workloads from the GAPBS [16], SPEC2006 [49], PARSEC [26], and YCSB [31] benchmark suites, sweeping memory footprints from 250 MB to 600 GB using input generators.

⁹For example, consider the following relationship:
load.stlb_hit = load.stlb_hit_4k + load.stlb_hit_2m.

We also collected HEC data for two microbenchmarks: a linear access pattern (parametrized by footprint, stride, and load-store ratio) and a random access pattern (parametrized by footprint and load-store ratio). Through ablation studies, we found that removing these microbenchmarks causes us to miss violations of key model constraints (e.g., Constraint (1) in Table 1) that are essential for reverse-engineering the presence and trigger conditions of the TLB prefetchers described below¹⁰. To stress different MMU behaviors, experiments were repeated with 4 KB, 2 MB, and 1 GB page sizes. Together, these workloads and configuration options yield about 20 million HEC samples—enough observations to thoroughly stress-test our model assumptions and drive higher-quality model refinement.

We evaluate our models at the 99% confidence level. Across dozens of representative μ DDs, we found that correlated counter confidence regions detect over 24% more model constraint violations compared to confidence regions that assume HECs are independent. For some models, exploiting correlations revealed over 75% additional violations compared to baseline. CounterPoint’s confidence regions are effective because HECs are highly correlated: in our dataset we find that over 25% of counter pairs have a Pearson correlation coefficient that exceeds 0.9 (where 1.0 indicates perfect correlation, and 0.0 indicates no correlation).

With CounterPoint’s support for guided model refinement, we explored dozens of μ DDs. Our initial μ DD contained 31 constraints, 8 of which were violated. We refined our initial μ DD over several iterations, details of which we provide in our technical report¹¹ [64]. Across all explored models, there were thousands of μ paths and over a thousand model constraint violations. Our guided refinement surpasses prior ad hoc reverse-engineering efforts [10, 41, 111, 121], enabling us to uncover subtleties (with high confidence) in:

Address translation prefetchers. Several studies have proposed address translation prefetching mechanisms, but little is known about how such prefetchers are actually implemented in real-world processors [25, 57, 74, 97, 109, 110]. Recent work suggests that underdocumented translation prefetching features may be at the core of unexplained performance anomalies in real-world workloads [17].

Using CounterPoint, we uncovered hardware in the Intel Haswell MMU that prefetches page table entries into its L1/L2 TLBs as well as PDE cache. Our analysis revealed three key aspects of the prefetcher’s implementation:

First, we identified prefetch trigger conditions. If a workload is feasible with an μ DD that includes the prefetcher but infeasible in one without it, the workload must trigger

prefetches. This helps us deduce that prefetching logic scans virtual page numbers in the load/store queue and is triggered by sequential accesses predicted to cross a page boundary—contradicting the common assumption that prefetches are triggered exclusively by TLB misses. For increasing virtual addresses, prefetching is triggered after consecutive accesses to cache lines 51 and 52 within a page; for decreasing addresses, the trigger occurs at cache lines 8 and 7. No other cache line pairs were observed to initiate prefetching.

Second, we found that the load/store queue logic responsible for virtual-page prefetching relies on the page table walker to resolve translation prefetch requests. In practice, this means that prefetches trigger the walker to inject additional load instructions into the CPU pipeline—the same way it injects loads for demand page table walks themselves (previously called “ghost” or “stuffed” loads [68, 121]). In some cases, the walker generates hundreds of such additional loads. This overturns the prevailing model in prior work, which assumed prefetches bypass the pipeline and enter the memory hierarchy directly, and therefore model prefetches with unrealistically low latency. It also implies that significantly more prefetches can be injected than previously believed.

Third, we found that prefetch-induced page table walks abort when they encounter a page table entry whose access (reference) bit is unset—unlike regular page table walks, which set this bit. Consequently, the TLB prefetch does not complete. This behavior is logical: allowing a speculatively set access bit for an ineffective TLB prefetch could, in principle, lead to suboptimal paging decisions, and permitting TLB prefetches to set the access bit would also introduce additional microarchitectural complexity. Prefetch-induced page walks can still modify cache state, with potential performance and security implications. Some recent TLB prefetcher proposals allow prefetch-induced page walks to set the access bit and complete [109, 110]. While this behavior is architecturally permitted [54], we have not observed it on Haswell.

Page table walk merging. Despite decades of research on TLBs and MMU caches, little is publicly known about how MMUs schedule page table walks. Using CounterPoint, we discover that MMUs can *merge* multiple outstanding walks to the same virtual page into a single page table walk, which we capture by modeling MSHRs within our MMU μ DD.

Historically, MMU MSHRs have not been modeled in address translation studies because their design involves subtleties beyond those of conventional cache MSHRs [60, 61, 106]. For instance, page sizes—and therefore virtual page numbers—are unknown until after translation [33], making MSHR lookup and allocation non-trivial. Further, distinct page table walks have unique rules for updating access and dirty bits, as well as for determining whether they are allowed to touch physical memory regions marked speculative

¹⁰We ensured that all of our HEC measurements were unaffected by any published HEC errata. For errata that are triggered when SMT is enabled (e.g., HSD29/HSM30 affecting `mem_uops_retired`), we addressed this by disabling SMT in the BIOS.

¹¹Our Haswell MMU models continue to be expanded and refined. Our most recent results are available online in our technical report [64].

versus non-speculative [46]. These complexities make it far from obvious how walk merging can be safely implemented.


Our results show that MMU MSHRs are nonetheless critical for performance. For some workloads, page table walk merging reduces the number of distinct walks by nearly half. This finding underscores the importance of explicitly modeling MMU MSHRs in simulators used to evaluate address translation optimizations [4, 10, 12, 15, 20, 22–24, 58, 121].

Finally, CounterPoint reveals a surprising detail: the PDE cache is queried *before* outstanding walks to the same virtual page are merged. Prior studies have not considered this interaction [12, 20]. One might expect walk merging to reduce PDE cache lookups, easing port pressure, cutting bandwidth, and eliminating queuing delays. Instead, our μ DD suggests that the PDE cache is looked up prior to MSHR allocation, likely to reduce latency via pipelining. Importantly, CounterPoint is able to do this because it enables discovery of not just individual hardware components, but also their relative placement within the pipeline.


Root-level MMU cache. A large body of address translation research proposing hardware optimizations [20, 24, 73, 87, 95, 100, 115] assumes the presence of a root-level MMU cache, yet some recent reverse-engineering studies have found no evidence of its existence [41, 108]. CounterPoint demonstrates its compatibility with all other address translation features identified in this paper for the workloads we analyze, giving architecture researchers confidence in including it in their models. When *walk bypassing* is not modeled, several workloads become feasible only with a root-level MMU cache in the μ DD. These workloads use 1GB pages which would stress a hypothetical PML4E cache (which is explicitly for 1GB pages), suggesting that for these workloads, the “missing” page table walker accesses could be explained by PML4E cache.

Aborted page table walks. Recent research has studied page table walks under speculation in modern out-of-order processors [48, 63, 121]. While prior work shows that x86-64 processors can abort in-flight page table walks in response to machine clears [63, 93, 121], the underlying implementation details remain poorly understood.

Using CounterPoint, we find that aborted walks are entirely consistent with all our HEC measurements and newly discovered features. Additionally, they appear to be triggered more frequently by workloads with high walker utilization. While more detailed study is necessary to better understand how aborted page table walks are implemented, CounterPoint suggests that page table walks can be aborted at any point—even before issuing a single memory access. This implies that aborted walks can still consume MMU and memory hierarchy resources, effectively imposing a hidden performance tax that should be explicitly modeled in simulation infrastructures for address translation [4, 58].



(a) Evaluation time per observation per model.



(b) Constraint deduction time per model.

Figure 9. CounterPoint performance (quantified for 20 representative μ DDs) varies with models and scales with HECs. Blue lines represent individual models; groups of semantically related counters are added along x-axis. (a) Feasibility testing time scales linearly. (b) Constraint deduction time scales exponentially. The red lines represent (a) arithmetic, and (b) geometric means.

Page table walk replays. We observe that page table walks can complete without generating any memory accesses. This suggests that the core may include a mechanism allowing walks to finish without engaging the cache hierarchy. Prior work has shown a complex interplay between the hardware page table walker and microarchitectural structures that maintain memory consistency [113, 121]. We hypothesize that these “missing” accesses occur because walks are replayed or handled by hidden internal address translation caching structures not reflected in the `walk_ref` counter. Understanding these structures more concretely would require implementing new HECs or access to proprietary RTL. An alternative explanation is that the “missing” accesses do occur but are not counted because, unlike regular page-walker accesses, they are non-speculative. If we assume that aborted walks are replayed at micro-op retirement as non-speculative walks—as suggested in prior sources [32, 46, 113]—then the resulting μ DD becomes feasible, but only once features such as TLB prefetching and miss merging are incorporated.

7.2 CounterPoint Performance Characterization

We evaluate CounterPoint on a 24-core Intel Xeon E5-2680 v3 CPU running at 2.5GHz. For this analysis, we evaluate every observation against every constraint to produce a worst-case

runtime, even though in practice only infeasible observations need to be checked. We parallelize the parts of CounterPoint devoted to determining feasibility constraints. On average, CounterPoint evaluates a model in 213 seconds, with the majority of time spent assessing individual model constraints. Determining model feasibility is much faster than checking each constraint, as model cones allow all constraints to be tested simultaneously. This makes it practical to evaluate models with large numbers of counters.

Figure 9a plots the time taken to determine observation feasibility as a function of the counters present in the model. For the full suite of counters, CounterPoint takes around 200 milliseconds to determine if an observation is feasible. Empirically, the time taken scales approximately linearly as counter groups are added. Additionally, observation feasibility testing is embarrassingly parallel, allowing large numbers of observations to be tested simultaneously.

Figure 9b shows how the time required to find the model constraints scales with the counter groups. The logarithmic y-scale demonstrates that empirically the constraint deduction time scales exponentially as counter groups are added. Despite this, CounterPoint only takes between 800 milliseconds and 10 seconds to determine the set of constraints for models with all counters present. Note that explicitly determining the model constraints is only used for providing feedback for model refinement; it is not a prerequisite for determining the feasibility of individual observations.

8 Related Work

Computer architects have recently used HECs to reverse-engineer specific microarchitectural features. HECs are used by uops.info [1] and nanoBench [2] to reverse-engineer microop performance and port assignment, as well as cache replacement policies. Several studies have focused on reverse engineering the MMU [10, 41, 111, 121], while Ragab et al [93] use HECs to characterize the security implications of machine clears. Binoculars [121] use HECs to characterize page table walker contention. With μ DD models, CounterPoint offers a more general-purpose approach.

Multiplexing noise is a well-studied problem [8, 11, 69]. Azimi et al [9] quantify multiplexing noise for a range of workloads. CounterMiner [69] replacing outliers with interpolated values. BayesPerf [11] reduces noise by exploiting *known* statistical relationships between counter values. CounterPoint infers correlations to reduce noise impact.

Interpreting HEC values correctly remains challenging. Vendors provide explicit metrics that convert HEC values to standard metrics (*e.g.*, CPI, MPKI, hit rates, *etc.*), but not all HECs are used. The Counter Inspection Toolkit [34] and related work [13, 14] correlate counters with individual microbenchmarks to define new metrics. Top Down Methodology [118] employs metrics and thresholds to enable application developers to identify performance bottlenecks. Unlike

bespoke approaches, CounterPoint μ DDs capture both HEC semantics and microarchitectural features by construction.

Formal modeling for microarchitectures has recently been used for memory consistency [50, 51, 67, 68, 70, 78], cache coherence [35, 79–81, 91, 92], and security [52, 105]. CheckSuite and related tools [50–52, 67, 68, 70, 78, 105] describe microarchitectural executions using μ spec models featuring μ paths and inter- and intra- μ path dependencies. CounterPoint’s μ DD formalism is compatible with these approaches.

9 Conclusion & Future Work

We presented CounterPoint, a framework that transforms large HEC datasets into accurate, high-quality microarchitectural models. By encoding an expert’s mental model as a μ DD and automatically generating model constraints, CounterPoint eliminates the tedium and errors of manual derivation. At the same time, CounterPoint processes noisy HEC measurements into reliable, high-confidence ranges, bridging intuition with data-driven analysis. CounterPoint accelerates and sharpens the modeling of complex architectures, freeing experts to focus on insight rather than bookkeeping, and making advanced microarchitectural modeling faster and more insightful. By accelerating the productivity of these influential experts, insights extracted by CounterPoint’s have the potential to shape the broader field of computing.

While this first paper demonstrates the promise of CounterPoint, several productive directions remain for future work. For example, CounterPoint could potentially be used to reverse-engineer not only microarchitectural details but also the semantics of undocumented HECs (see our work on reverse-engineering the semantics of the walk_ref counter for page table walk replays). Establishing this, however, would require a detailed study beyond the scope of this work, which we leave for future work. Additionally, our current study evaluates the benefit of CounterPoint on CPUs; exploring the utility of CounterPoint to hardware accelerators would broaden its applicability. Finally, this paper presents a first proof-of-concept study of the key ideas and principles behind CounterPoint. Substantial work remains to extend it into a robust, system-wide modeling framework, including support for multiple cores, multiple sockets, hyperthreading, kernel-level activity, and much more.

Acknowledgments

We thank the anonymous reviewers, Neil Zhao, Michael Wu, and Grace Jia for their feedback on drafts of this paper. We thank Manolis Zampetakis, Lieven Eeckhout, Andrew Milas, and Matt Sinclair for technical discussions. This work was supported in part by National Science Foundation awards 2236855, 2112562, and 2047220. We thank Intel for supporting the IISWC’24 paper that preceded and motivated this work.

References

- [1] Andreas Abel and Jan Reineke. 2019. uops. info: Characterizing latency, throughput, and port usage of instructions on intel microarchitectures. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. 673–686.
- [2] Andreas Abel and Jan Reineke. 2020. nanobench: A low-overhead tool for running microbenchmarks on x86 systems. In *2020 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 34–46.
- [3] Keith Adams and Ole Agesen. 2006. A comparison of software and hardware techniques for x86 virtualization. *ACM Sigplan Notices* 41, 11 (2006), 2–13.
- [4] Mohammad Agbarya, Idan Yaniv, Jayneel Gandhi, and Dan Tsafir. 2020. Predicting execution times with partial simulations in virtual memory research: Why and how. In *Proceedings of the Annual International Symposium on Microarchitecture, MICRO*, Vol. 2020-Octob. IEEE Computer Society, 456–470. doi:10.1109/MICRO50266.2020.00046
- [5] Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. 2018. A rewriting system for convex optimization problems. *Journal of Control and Decision* 5, 1 (2018), 42–60.
- [6] Jeongseob Ahn, Seongwook Jin, and Jaehyuk Huh. 2012. Revisiting hardware-assisted page walks for virtualized systems. *ACM SIGARCH Computer Architecture News* 40, 3 (2012), 476–487.
- [7] Hanna Alam, Tianhao Zhang, Mattan Erez, and Yoav Etsion. 2017. Do-It-Yourself Virtual Memory Translation. In *Proceedings of the 44th Annual International Symposium on Computer Architecture (Toronto, ON, Canada) (ISCA '17)*. Association for Computing Machinery, New York, NY, USA, 457–468. doi:10.1145/3079856.3080209
- [8] Reza Azimi, Michael Stumm, and Robert W. Wisniewski. 2005. Online performance analysis by statistical sampling of microprocessor performance counters. In *Proceedings of the 19th Annual International Conference on Supercomputing (Cambridge, Massachusetts) (ICS '05)*. Association for Computing Machinery, New York, NY, USA, 101–110. doi:10.1145/1088149.1088163
- [9] Reza Azimi, David K Tam, Livio Soares, and Michael Stumm. 2009. Enhancing operating system support for multicore processors by using hardware performance monitoring. *ACM SIGOPS Operating Systems Review* 43, 2 (2009), 56–65.
- [10] Vlastimil Babka and Petr Tuma. 2009. Investigating cache parameters of x86 family processors. In *SPEC Benchmark Workshop*. Springer, 77–96.
- [11] Subho S. Banerjee, Saurabh Jha, Zbigniew Kalbarczyk, and Ravishankar K. Iyer. 2021. BayesPerf: Minimizing performance monitoring errors using Bayesian statistics. *International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS (2021)*, 832–844. arXiv:2102.10837 doi:10.1145/3445814.3446739
- [12] Thomas W Barr, Alan L Cox, and Scott Rixner. 2010. Translation caching: skip, don't walk (the page table). *ACM SIGARCH Computer Architecture News* 38, 3 (2010), 48–59.
- [13] Daniel Barry, Anthony Danalis, and Jack Dongarra. 2024. Automated data analysis for defining performance metrics from raw hardware events. In *2024 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE, 716–725.
- [14] Daniel Barry, Anthony Danalis, and Heike Jagode. 2021. Effortless monitoring of arithmetic intensity with papi's counter analysis toolkit. In *Tools for High Performance Computing 2018/2019: Proceedings of the 12th and of the 13th International Workshop on Parallel Tools for High Performance Computing, Stuttgart, Germany, September 2018, and Dresden, Germany, September 2019*. Springer, 195–218.
- [15] Arkaprava Basu, Jayneel Gandhi, Jichuan Chang, Mark D. Hill, and Michael M. Swift. 2013. Efficient virtual memory for big memory servers. *ACM SIGARCH Computer Architecture News* 41, 3 (2013), 237–248. doi:10.1145/2508148.2485943
- [16] Scott Beamer, Krste Asanović, and David Patterson. 2017. The GAP Benchmark Suite. arXiv:1508.03619 [cs.DC] <https://arxiv.org/abs/1508.03619>
- [17] Andrin Bertschi. 2022. Battling the Prefetcher: Exploring Coffee Lake (Part 1). <https://abertschi.ch/blog/2022/prefetching/#intel-manuals-on-prefetching>. Accessed: 2025-04-29.
- [18] Ravi Bhargava, Benjamin Serebrin, Francesco Spadini, and Srilatha Manne. 2008. Accelerating two-dimensional page walks for virtualized systems. In *Proceedings of the 13th international conference on Architectural support for programming languages and operating systems*. 26–35.
- [19] Nikhil Bhatia. 2009. Performance evaluation of Intel EPT hardware assist. *VMware, Inc (2009)*.
- [20] Abhishek Bhattacharjee. 2013. Large-reach memory management unit caches. *MICRO 2013 - Proceedings of the 46th Annual IEEE/ACM International Symposium on Microarchitecture (2013)*, 383–394. doi:10.1145/2540708.2540741
- [21] Abhishek Bhattacharjee. 2017. Preserving Virtual Memory by Mitigating the Address Translation Wall. *IEEE Micro* 37, 5 (2017), 6–10. doi:10.1109/MM.2017.3711640
- [22] Abhishek Bhattacharjee. 2017. Translation-triggered prefetching. *ACM SIGPLAN Notices* 52, 4 (4 2017), 63–76. doi:10.1145/3037697.3037705
- [23] Abhishek Bhattacharjee. 2019. Appendix L: Advanced Concepts on Address Translation.
- [24] Abhishek Bhattacharjee and Daniel Lustig. 2018. *Architectural and Operating System Support for Virtual Memory*. Morgan and Claypool Publishers.
- [25] Abhishek Bhattacharjee and Margaret Martonosi. 2010. Inter-core cooperative TLB for chip multiprocessors. *ACM Sigplan Notices* 45, 3 (2010), 359–370.
- [26] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. 2008. The PARSEC benchmark suite: Characterization and architectural implications. In *Proceedings of the 17th international conference on Parallel architectures and compilation techniques*. 72–81.
- [27] William Lloyd Bircher and Lizy K. John. 2012. Complete System Power Estimation Using Processor Performance Events. *IEEE Trans. Comput.* 61, 4 (2012), 563–577. doi:10.1109/TC.2011.47
- [28] Xiaotao Chang, Hubertus Franke, Yi Ge, Tao Liu, Kun Wang, Jimi Xenidis, Fei Chen, and Yu Zhang. 2013. Improving virtualization in the presence of software managed translation lookaside buffers. In *Proceedings of the 40th Annual International Symposium on Computer Architecture*. 120–129.
- [29] Dimitrios Chasapis, Georgios Vavoulitis, Daniel A. Jiménez, and Marc Casas. 2025. Instruction-Aware Cooperative TLB and Cache Replacement Policies. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1 (Rotterdam, Netherlands) (ASPLOS '25)*. Association for Computing Machinery, New York, NY, USA, 619–636. doi:10.1145/3669940.3707247
- [30] COIN-OR Project. 2024. *coin-or/Cbc: Release releases/2.10.12*. doi:10.5281/zenodo.13347261
- [31] Brian F Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. 2010. Benchmarking cloud serving systems with YCSB. In *Proceedings of the 1st ACM symposium on Cloud computing*. 143–154.
- [32] Peter Cordes. 2021. What happens after a L2 TLB miss? <https://stackoverflow.com/a/32258855>. Accessed: April 6th, 2025.
- [33] Guilherme Cox and Abhishek Bhattacharjee. 2017. Efficient address translation for architectures with multiple page sizes. *ACM SIGPLAN Notices* 52, 4 (2017), 435–448.
- [34] Anthony Danalis, Heike Jagode, Hanumantharayappa, Sangamesh Ragate, and Jack Dongarra. 2017. Counter inspection toolkit: Making sense out of hardware performance events. In *International Workshop*

- on *Parallel Tools for High Performance Computing*. Springer, 17–37.
- [35] Nirav Dave, Man Cheuk Ng, et al. 2005. Automatic synthesis of cache-coherence protocol processors using bluespec. In *Proceedings. Second ACM and IEEE International Conference on Formal Methods and Models for Co-Design, 2005. MEMOCODE'05*. IEEE, 25–34.
- [36] John Demme, Matthew Maycock, Jared Schmitz, Adrian Tang, Adam Waksman, Simha Sethumadhavan, and Salvatore Stolfo. 2013. On the feasibility of online malware detection with performance counters. *ACM SIGARCH computer architecture news* 41, 3 (2013), 559–570.
- [37] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17, 83 (2016), 1–5.
- [38] Gem5 Discussion. 2024. Riscv Virtual Address Translation Process. <https://github.com/orgs/gem5/discussions/1220> (2024).
- [39] Gem5 Discussion. 2025. Adding realistic SMMU invalidation delays to gem5. <https://github.com/orgs/gem5/discussions/2227> (2025).
- [40] Joel S Emer and Douglas W Clark. 1984. A Characterization of Processor Performance in the VAX-11/780. (1984), 301–310. doi:10.1145/800015.808199
- [41] Philipp Ertmer, Robert Dumitru, and Yuval Yarom. 2025. Reverse-Engineering the Address Translation Caches. (2025), 149–168.
- [42] Stijn Eyerman, Lieven Eeckhout, Tejas Karkhanis, and James E Smith. 2007. A top-down approach to architecting CPI component performance counters. *IEEE micro* 27, 1 (2007), 84–93.
- [43] Vadim Filanovsky and Harshad Sane. 2022. Seeing through hardware counters: a journey to threefold performance increase. *Netflix Blog* (2022).
- [44] Komei Fukuda. 2004. What is the Minkowski–Weyl theorem for convex polyhedra? Polyhedral Computation FAQ. <https://people.inf.ethz.ch/fukudak/polyfaq/node14.html> Accessed via ETH Zurich website (people.inf.ethz.ch/fukudak/polyfaq/).
- [45] Jayneel Gandhi, Mark D Hill, and Michael M Swift. 2016. Agile paging: Exceeding the best of nested and shadow paging. *ACM SIGARCH Computer Architecture News* 44, 3 (2016), 707–718.
- [46] Andy Glew, Glenn Hinton, and Haitam Akkary. 1997. Method and apparatus for performing page table walks in a microprocessor capable of processing speculative instructions. US Patent 5,680,565.
- [47] Krishnan Gosakan, Jaehyun Han, William Kuzsmaul, Ibrahim N. Mubarek, Nirjhar Mukherjee, Karthik Sriram, Guido Tagliavini, Evan West, Michael A. Bender, Abhishek Bhattacharjee, Alex Conway, Martin Farach-Colton, Jayneel Gandhi, Rob Johnson, Sudarshan Kannan, and Donald E. Porter. 2023. Mosaic Pages: Big TLB Reach with Small Pages. In *International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS*, Vol. 3. Association for Computing Machinery, 433–448. doi:10.1145/3582016.3582021
- [48] Siddharth Gupta, Atri Bhattacharyya, Yunho Oh, Abhishek Bhattacharjee, Babak Falsafi, and Mathias Payer. 2021. Rebooting virtual memory with Midgard. *Proceedings - International Symposium on Computer Architecture* 2021-June (2021), 512–525. doi:10.1109/ISCA52012.2021.00047
- [49] John L Henning. 2006. SPEC CPU2006 benchmark descriptions. *ACM SIGARCH Computer Architecture News* 34, 4 (2006), 1–17.
- [50] Naorin Hossain, Caroline Trippel, and Margaret Martonosi. 2020. TransForm: Formally Specifying Transistency Models and Synthesizing Enhanced Litmus Tests. *Proceedings - International Symposium on Computer Architecture* 2020-May (2020), 874–887. doi:10.1109/ISCA45697.2020.00076
- [51] Yao Hsiao, Dominic P. Mulligan, Nikos Nikolieris, Gustavo Petri, and Caroline Trippel. 2021. Synthesizing formal models of hardware from RTL for efficient verification of memory model implementations. *Proceedings of the Annual International Symposium on Microarchitecture, MICRO* (2021), 679–694. doi:10.1145/3466752.3480087
- [52] Yao Hsiao, Nikos Nikolieris, Artem Khyzha, Dominic P Mulligan, Gustavo Petri, Christopher W Fletcher, and Caroline Trippel. 2024. RTL2M μ PATH: Multi- μ PATH synthesis with applications to hardware security verification. In *2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 507–524.
- [53] Intel Corporation. [n. d.]. Intel Performance Monitoring Events (perfmon). <https://github.com/intel/perfmon/tree/main>.
- [54] Intel Corporation. 2025. Intel® 64 and IA-32 Architectures Software Developer’s Manual, Combined Volumes (PDF). <https://cdrdrv2.intel.com/v1/dl/getContent/671200>. Published October 2025..
- [55] International Business Machines Corporation. 2018. *POWER9 Performance Monitor Unit User’s Guide* (version 1.2 ed.). IBM Systems, Somers, NY, USA.
- [56] Canturk Isci and Margaret Martonosi. 2003. Runtime power monitoring in high-end processors: Methodology and empirical data. In *Proceedings. 36th Annual IEEE/ACM International Symposium on Microarchitecture, 2003. MICRO-36*. IEEE, 93–104.
- [57] Gokul B Kandiraju and Anand Sivasubramaniam. 2002. Going the distance for TLB prefetching: An application-driven study. *ACM SIGARCH Computer Architecture News* 30, 2 (2002), 195–206.
- [58] Konstantinos Kanellopoulos, Konstantinos Sgouras, F. Nisa Bostanci, Andreas Kosmas Kakolyris, Berkin Kerim Konar, Rahul Bera, Mohammad Sadrosadati, Rakesh Kumar, Nandita Vijaykumar, and Onur Mutlu. 2025. *Virtuoso: Enabling Fast and Accurate Virtual Memory Research via an Imitation-based Operating System Simulation Methodology*. Association for Computing Machinery, New York, NY, USA. 1400–1421 pages. <https://doi.org/10.1145/3676641.3716027>
- [59] Shinichi Kawaguchi and Toshiaki Yachi. 2016. Adaptive Power Efficiency Control by Computer Power Consumption Prediction Using Performance Counters. *IEEE Transactions on Industry Applications* 52, 1 (2016), 407–413. doi:10.1109/TIA.2015.2466687
- [60] David Kroft. 1998. Lockup-free instruction fetch/prefetch cache organization. In *25 years of the international symposia on Computer architecture (selected papers)*. 195–201.
- [61] Chang Joo Lee, Veynu Narasiman, Onur Mutlu, and Yale N Patt. 2009. Improving memory bank-level parallelism in the presence of prefetching. In *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*. 327–336.
- [62] Kyeong-Jae Lee and Kevin Skadron. 2005. Using Performance Counters for Runtime Temperature Sensing in High-Performance Processors. In *Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'05) - Workshop 11 - Volume 12 (IPDPS '05)*. IEEE Computer Society, USA, 232.1. doi:10.1109/IPDPS.2005.448
- [63] Nick Lindsay and Abhishek Bhattacharjee. 2024. Understanding Address Translation Scaling Behaviours Using Hardware Performance Counters. In *2024 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 236–246.
- [64] Nick Lindsay, Caroline Trippel, Anurag Khandelwal, and Abhishek Bhattacharjee. 2026. CounterPoint: Using Hardware Event Counters to Refute and Refine Microarchitectural Assumptions (Extended Version). arXiv:2601.01265 [cs.AR] <https://arxiv.org/abs/2601.01265>
- [65] Jason Lowe-Power. 2025. PIPT/VIPT caches and TLB lookup latency. *Gem5 Discussion* (2025), <https://github.com/orgs/gem5/discussions/2227>.
- [66] Daniel Lustig, Michael Pellauer, and Margaret Martonosi. 2014. PipeCheck: Specifying and verifying microarchitectural enforcement of memory consistency models. In *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE, 635–646.
- [67] Daniel Lustig, Michael Pellauer, and Margaret Martonosi. 2015. Pipe Check: Specifying and Verifying Microarchitectural Enforcement of Memory Consistency Models. *Proceedings of the Annual International Symposium on Microarchitecture, MICRO* 2015-Janua, January (2015), 635–646. doi:10.1109/MICRO.2014.38
- [68] Daniel Lustig, Geet Sethi, Margaret Martonosi, and Abhishek Bhattacharjee. 2016. COATCheck: Verifying memory ordering at the

- hardware-OS interface. *ACM SIGPLAN Notices* 51, 4 (apr 2016), 233–247. doi:10.1145/2872362.2872399
- [69] Yirong Lv, Bin Sun, Qingyi Luo, Jing Wang, Zhibin Yu, and Xuehai Qian. 2018. Counterminer: Mining big performance data from hardware counters. *Proceedings of the Annual International Symposium on Microarchitecture, MICRO 2018-October* (2018), 613–626. doi:10.1109/MICRO.2018.00056
- [70] Yatin A. Manerkar, Daniel Lustig, Margaret Martonosi, and Aarti Gupta. 2018. PipeProof: Automated memory consistency proofs for microarchitectural specifications. *Proceedings of the Annual International Symposium on Microarchitecture, MICRO 2018-October* (2018), 788–801. doi:10.1109/MICRO.2018.00069
- [71] Yatin A Manerkar, Daniel Lustig, Michael Pellauer, and Margaret Martonosi. 2015. CCICheck: Using μ hb graphs to verify the coherence-consistency interface. In *Proceedings of the 48th International Symposium on Microarchitecture*. 26–37.
- [72] Aninda Manocha, Zi Yan, Esin Tureci, Juan Luis Aragón, David Nelsans, and Margaret Martonosi. 2022. The implications of page size management on graph analytics. In *2022 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 199–214.
- [73] Artemiy Margaritov, Dmitrii Ustiugov, Edouard Bugnion, and Boris Grot. 2019. Prefetched address translation. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*. 1023–1036.
- [74] Artemiy Margaritov, Dmitrii Ustiugov, Amna Shahab, and Boris Grot. 2021. Ptemagnet: Fine-grained physical memory reservation for faster page walks in public clouds. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 11–23.
- [75] Dirk Merkel. 2014. Docker: lightweight linux containers for consistent development and deployment. *Linux journal* 2014, 239 (2014), 2.
- [76] Timothy Merrifield and H Reza Taheri. 2016. Performance implications of extended page tables on virtualized x86 processors. In *Proceedings of the 12th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*. 25–35.
- [77] Stuart Mitchell, Michael OSullivan, and Iain Dunning. 2011. Pulp: a linear programming toolkit for python. *The University of Auckland, Auckland, New Zealand* 65 (2011), 25.
- [78] Chase Norman, Adwait Godbole, and Yatin A. Manerkar. 2023. *PipeSynth: Automated Synthesis of Microarchitectural Axioms for Memory Consistency*. Vol. 3. Association for Computing Machinery. 513–527 pages. doi:10.1145/3582016.3582056
- [79] Nicolai Oswald, Vijay Nagarajan, and Daniel J Sorin. 2018. ProtoGen: Automatically generating directory cache coherence protocols from atomic specifications. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 247–260.
- [80] Nicolai Oswald, Vijay Nagarajan, and Daniel J Sorin. 2020. HieraGen: Automated generation of concurrent, hierarchical cache coherence protocols. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 888–899.
- [81] Nicolai Oswald, Vijay Nagarajan, Daniel J Sorin, Vasilis Gavrielatos, Theo Olausson, and Reece Carr. 2022. Heterogen: Automatic synthesis of heterogeneous cache coherence protocols. In *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 756–771.
- [82] Yueyang Pan, Yash Lala, Musa Unal, Yujie Ren, Seung seob Lee, Yizhou Shan, Abhishek Bhattacharjee, Anurag Khandelwal, and Sanidhya Kashyap. 2024. Scalable Far Memory: Balancing Faults and Evictions. *SOSP* (2024).
- [83] The pandas development team. 2025. *pandas-dev/pandas: Pandas*. doi:10.5281/zenodo.15831829
- [84] Ashish Panwar, Reto Achermann, Arkaprava Basu, Abhishek Bhattacharjee, K Gopinath, and Jayneel Gandhi. 2021. Fast local page-tables for virtualized numa servers with vmitosis. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 194–210.
- [85] Ashish Panwar, Sorav Bansal, and K Gopinath. 2019. Hawkeye: Efficient fine-grained os support for huge pages. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. 347–360.
- [86] Chang Hyun Park, Taekyung Heo, and Jaehyuk Huh. 2016. Efficient synonym filtering and scalable delayed translation for hybrid virtual caching. *ACM SIGARCH Computer Architecture News* 44, 3 (2016), 217–229.
- [87] Chang Hyun Park, Ilias Vougioukas, Andreas Sandberg, and David Black-Schaffer. 2022. Every walk’s a hit: making page walks single-access cache hits. (2022), 128–141. doi:10.1145/3503222.3507718
- [88] Binh Pham, Abhishek Bhattacharjee, Yasuko Eckert, and Gabriel H. Loh. 2014. Increasing TLB reach by exploiting clustering in page translations. *Proceedings - International Symposium on High-Performance Computer Architecture* (2014), 558–567. doi:10.1109/HPCA.2014.6835964
- [89] Binh Pham, Viswanathan Vaidyanathan, Aamer Jaleel, and Abhishek Bhattacharjee. 2012. CoLT: Coalesced Large-Reach TLBs. In *Proceedings of the 2012 45th Annual IEEE/ACM International Symposium on Microarchitecture (Vancouver, B.C., CANADA) (MICRO-45)*. IEEE Computer Society, USA, 258–269. doi:10.1109/MICRO.2012.32
- [90] Binh Pham, Ján Veselý, Gabriel H Loh, and Abhishek Bhattacharjee. 2015. Large pages and lightweight memory management in virtualized environments: Can you have it both ways?. In *Proceedings of the 48th International Symposium on Microarchitecture*. 1–12.
- [91] Fong Pong and Michel Dubois. 1993. The verification of cache coherence protocols. In *Proceedings of the fifth annual ACM symposium on Parallel Algorithms and Architectures*. 11–20.
- [92] Fong Pong and Michel Dubois. 1995. A new approach for the verification of cache coherence protocols. *IEEE Transactions on Parallel and Distributed Systems* 6, 8 (1995), 773–787.
- [93] Hany Ragab, Enrico Barberis, Herbert Bos, and Cristiano Giuffrida. 2021. Rage against the machine clear: A systematic analysis of machine clears and their implications for transient execution attacks. *Proceedings of the 30th USENIX Security Symposium* (2021), 1451–1468.
- [94] Benny Rubin, Saksham Agarwal, Qizhe Cai, and Rachit Agarwal. 2024. Fast & Safe IO Memory Protection. *SOSP* (2024).
- [95] Jee Ho Ryoo, Nagendra Gulur, Shuang Song, and Lizy K. John. 2017. Rethinking TLB designs in virtualized environments: A very large part-of-memory TLB. In *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*. IEEE Computer Society, Los Alamitos, CA, USA, 469–480. doi:10.1145/3079856.3080210
- [96] Daniel Sanchez and Christos Kozyrakis. 2013. ZSim: Fast and accurate microarchitectural simulation of thousand-core systems. *ACM SIGARCH Computer architecture news* 41, 3 (2013), 475–486.
- [97] Ashley Saulsbury, Fredrik Dahlgren, and Per Stenström. 2000. Recency-based TLB preloading. In *Proceedings of the 27th annual international symposium on Computer architecture*. 117–127.
- [98] Balvinder Pal Singh and B Thangaraju. 2020. Power, Performance And Thermal Management Using Hardware Performance Counters. In *2020 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECT)*. 1–7. doi:10.1109/CONECT50063.2020.9198372
- [99] Richard L Sites. 2022. Performance Counters I’d Like to See – Part II. <https://www.sigarch.org/performance-counters-id-like-to-see-part-ii/>. [Accessed 05-01-2025].
- [100] Dimitrios Skarlatos, Apostolos Kokolis, Tianyin Xu, and Josep Torrellas. 2020. Elastic cuckoo page tables: Rethinking virtual memory translation for parallelism. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 1093–1108.
- [101] Lingjia Tang, Jason Mars, Neil Vachharajani, Robert Hundt, and Mary Lou Soffa. 2013. The Impact of Memory Subsystem Resource

- Sharing on Datacenter Applications. *ISCA* (2013).
- [102] Lingjia Tang, Jason Mars, Xiao Zhang, Robert Hagmann, Robert Hundt, and Eric Tune. 2013. Optimizing Google’s warehouse scale computers: The NUMA experience. *HPCA* (2013).
- [103] Andrei Tatar, Daniël Trujillo, Cristiano Giuffrida, and Herbert Bos. 2022. {TLB; DR}: Enhancing {TLB-based} attacks with {TLB} desynchronized reverse engineering. In *31st USENIX Security Symposium (USENIX Security 22)*. 989–1007.
- [104] The Linux Kernel Developers. [n. d.]. pmu-events: Performance Monitoring Unit (PMU) events definitions in Linux perf. <https://github.com/torvalds/linux/tree/master/tools/perf/pmu-events>.
- [105] Caroline Trippel, Daniel Lustig, and Margaret Martonosi. 2018. CheckMate: Automated synthesis of hardware exploits and security litmus tests. *Proceedings of the Annual International Symposium on Microarchitecture, MICRO 2018–October* (2018), 947–960. doi:10.1109/MICRO.2018.00081
- [106] James Tuck, Luis Ceze, and Josep Torrellas. 2006. Scalable cache miss handling for high memory-level parallelism. In *2006 39th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO’06)*. IEEE, 409–422.
- [107] Pennsylvania State University. n.d. Chapter 4. In *Statistics 505*. <https://online.stat.psu.edu/stat505/book/export/html/636>
- [108] Stephan van Schaik, Kaveh Razavi, Ben Gras, Herbert Bos, and Cristiano Giuffrida. 2017. Reverse engineering hardware page table caches using side-channel attacks on the mmu. *Vrije Universiteit Amsterdam, Tech. Rep* (2017).
- [109] Georgios Vavouliotis, Lluç Alvarez, Boris Grot, Daniel Jiménez, and Marc Casas. 2021. Morrigan: A composite instruction tlb prefetcher. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*. 1138–1153.
- [110] Georgios Vavouliotis, Lluç Alvarez, Vasileios Karakostas, Konstantinos Nikas, Nectarios Koziris, Daniel A Jiménez, and Marc Casas. 2021. Exploiting page table locality for agile TLB prefetching. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 85–98.
- [111] Alan Wang, Boru Chen, Yingchen Wang, Christopher Fletcher, Daniel Genkin, David Kohlbrenner, and Riccardo Paccagnella. 2024. Peek-a-Walk: Leaking Secrets via Page Walk Side Channels. In *2025 IEEE Symposium on Security and Privacy (S&P)*. IEEE. [Online]. Available: <https://www.computer.org/csdl/proceedings-article/sp/2025/223600a023/21B7QepK7Fm>.
- [112] Xueyang Wang and Ramesh Karri. 2013. Numchecker: Detecting kernel control-flow modifying rootkits by using hardware performance counters. In *Proceedings of the 50th annual design automation conference*. 1–7.
- [113] Henry Wong. 2015. TLB and Pagewalk Coherence in x86 Processors. <http://blog.stuffedcow.net/2015/08/pagewalk-coherence/>
- [114] Yu Xia, Vishnu Ramadas, Matthew Poremba, and Matthew D. Sinclair. 2025. Narrowing the GAP: Enhancing gem5’s GPU Memory Bandwidth Accuracy. *Gem5 Workshop* (2025).
- [115] Zi Yan, Ján Veselý, Guilherme Cox, and Abhishek Bhattacharjee. 2017. Hardware translation coherence for virtualized systems. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*. 430–443.
- [116] Xi Yang, Stephen Blackburn, and Kathryn McKinley. 2015. Computer Performance Microscopy with SHIM. *ISCA* (2015).
- [117] Yihao Yang, Pengfei Qiu, Chunlu Wang, Yu Jin, Qiang Gao, Xiaoyong Li, Dongsheng Wang, and Gang Qu. 2023. Exploration and Exploitation of Hidden PMU Events. In *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*. 1–9. doi:10.1109/ICCAD57390.2023.10323695
- [118] Ahmad Yasin. 2014. A Top-Down method for performance analysis and counters architecture. *ISPASS 2014 - IEEE International Symposium on Performance Analysis of Systems and Software* (2014), 35–44. doi:10.1109/ISPASS.2014.6844459
- [119] Reza Zamani and Ahmad Afsahi. 2012. A study of hardware performance monitoring counter selection in power modeling of computing systems. In *2012 International Green Computing Conference (IGCC)*. IEEE, 1–10.
- [120] Xiao Zhang, Sandhya Dwarkadas, Girts Folkmanis, and Kai Shen. 2007. Processor hardware counter statistics as a first-class system resource.. In *HotOS*.
- [121] Zirui Neil Zhao, Adam Morrison, Christopher W. Fletcher, and Josep Torrellas. 2022. Binoculars: Contention-Based Side-Channel Attacks Exploiting the Page Walker. *Proceedings of the 31st USENIX Security Symposium, Security 2022* (2022), 699–716.