

Breaking the Cycle - A Short Overview of Memory-Access Sampling Differences on Modern x86 CPUs

Roland Kühn
roland.kuehn@cs.tu-dortmund.de
TU Dortmund University

Jan Mühlig*
jan.muehlig@tu-dortmund.de
TU Dortmund University

Jens Teubner
jens.teubner@cs.tu-dortmund.de
TU Dortmund University
Lamarr Institute for Machine
Learning and Artificial Intelligence

Abstract

As hardware complexity increases, profiling becomes essential for understanding system behavior. This paper compares different x86 sampling implementations for memory access profiling, revealing their complementary capabilities and limitations. Plus, we demonstrate that current abstractions like the *perf subsystem* inadequately expose platform-specific features.

ACM Reference Format:

Roland Kühn, Jan Mühlig, and Jens Teubner. 2025. Breaking the Cycle - A Short Overview of Memory-Access Sampling Differences on Modern x86 CPUs. In *21st International Workshop on Data Management on New Hardware (DaMoN '25)*, June 22–27, 2025, Berlin, Germany. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3736227.3736241>

1 Introduction

To fully utilize modern hardware, performance-sensitive applications must be designed with hardware-conscious principles (e.g., [5, 7, 13, 16, 20, 25, 27]). However, sophisticated mechanisms such as out-of-order execution and memory prefetching have transformed hardware into a black box—turning hardware-aware optimizations into an uphill battle.

The silver lining lies in Performance Monitoring Units (PMUs)—specialized components embedded within modern CPUs—which allow engineers to examine software execution under a magnifying glass (e.g., [6, 23]). *Sampling*-based profiling techniques, in particular, offer invaluable insights by revealing critical details such as *memory access patterns* throughout execution. But, PMU implementations vary substantially across hardware vendors and CPU generations: Diverse operating modes and consequently different features complicate the comparison of software executions across heterogeneous hardware platforms [26]. This challenge, however, represents two sides of the same coin: These architectural differences can be leveraged advantageously when properly understood.

This paper presents a comparative analysis of two leading PMU-based sampling techniques from the *memory-access sampling* perspective: Intel’s Processor Event-Based Sampling (PEBS) [4] and AMD’s Instruction-Based Sampling (IBS) [10, 11]. We equip engineers with critical knowledge about which platform unveils specific execution details, enabling them to select the right tool for

*The author is now at Huawei Research, Edinburgh, UK.



This work is licensed under a Creative Commons Attribution 4.0 International License. *DaMoN '25, Berlin, Germany*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1940-0/2025/06

<https://doi.org/10.1145/3736227.3736241>




Figure 1: Sequence of a memory access and the hardware components involved¹. The numbers indicate points at which information can be obtained for the memory samples.

each analytical question. Furthermore, we demonstrate how commonly used abstractions—particularly the *perf subsystem* and the command-line interface *perf*—fall short of exposing the full spectrum of capabilities across different hardware platforms, leaving valuable performance insights on the table.

2 Memory Address Sampling on x86 platforms

Nearly all modern x86 architectures are equipped with *memory-access sampling* capabilities that can periodically generate a snapshot of the actual accessed logical and physical memory addresses [23]. However, the collection of samples by the two prominent vendors—AMD and Intel—differs significantly as already shown in [26]. Intel’s PEBS facility explicitly allows to sample load and store instructions (e.g., every *x*-th load instruction will create a new sample) [12]. AMD employs a different strategy with its IBS, where every *x*-th micro-operation (*μOp*), regardless of the type (*arithmetic*, *load/store*, ...), will be tagged and traced throughout the entire processing pipeline [1]. In addition to the accessed addresses, both hardware makers provide additional information in their respective samples that allow deductions about the utilization of critical system resources. To highlight which information can be retrieved in IBS and PEBS, we will follow a memory request through the various execution stages and reveal, at each stage, what information both vendors provide in their samples (Figure 1).

TLB Access (1). If an instruction/*μOp* that accesses memory is executed, the accessed logical address has to be translated into a physical address by consulting the Translation Lookaside Buffers (TLBs), which then returns the page address on a TLB hit, or issues page walk to retrieve the physical page address from the page table (on a miss). While PEBS merely reports the TLB hit/miss status, IBS reports which level was hit and quantifies TLB refill latency [2], i.e., when the L1 TLB was refilled from the L2 TLB or a page walk was issued due to a miss in both TLB levels.

L1d Cache Access (②). After the address translation, the caches will be consulted to find the requested data element¹. If the data element can be found in the L1 data cache (L1d), both sampling implementations report the L1d as the data source. The latency for the cache access, however, will only be reported by PEBS. In contrast, IBS reports the latency when the request missed the L1d.

Line Fill Buffer and Memory Access (③ and ④). If the data element cannot be found in the L1d, the address of the cache line that contains the data element will be written to the Line Fill Buffer (LFB) (or Miss Address Buffer (MAB) on AMD systems²) and will then be serviced by a higher cache level or the memory subsystem. Once the memory request has been processed, both vendors report information about the latency and the data source from which the data element was retrieved (e.g., the last-level cache (LLC) or the main memory). However, some key distinctions exist between the sampling implementations from both vendors. If the cache line is already registered in the LFB, PEBS reports the LFB as the data source, while IBS reports the real source from which the data was finally retrieved. In addition, IBS-samples also contain more information about the memory request itself, like the page size, the number of requested bytes, a flag if an LFB slot was allocated, and the number of actual allocated LFB slots. This information can be crucial, e.g., to identify bottlenecks caused by requests flooding the LFB, since instructions/ μ Ops stall until an LFB slot becomes available [13].

PEBS distinguishes between loads and stores, counting load prefetches as accesses to the L1d. IBS reports software prefetches as such, although it does not report the cache miss latency.

Instruction/ μ Op retirement (⑤ and ⑥). After the memory subsystem has retrieved the requested data element, the instruction/ μ Op will retire. In contrast to PEBS, which then only reports the total latency for the instruction execution, IBS additionally reports the cycles spent between the completion of the μ Op and the point where the μ Op is considered as successfully retired.

Overall, IBS reports four latencies: For refilling the L1 TLB, for fulfilling requests that miss the L1d, from tagging the μ Op until retirement, and separately from completion to retirement. PEBS provides two latencies: data access and instruction retirement. Additionally, further information such as the occupancy of the LFB is also provided by samples from IBS.

Since *sampling* can introduce significant overhead to the operating system, when many samples are created, Intel's PEBS offers the possibility to filter the samples by latency and keep only samples with latency higher than a configurable threshold. AMD introduced this feature also with the latest *Zen 5* micro-architecture [3].

Perf subsystem. The *perf subsystem*—baked into the Linux kernel—allows to interact with PMUs and builds the foundation for the *perf* command-line-interface. The accessible information from the actual version of the *perf subsystem* seems to be leaned on the details provided by PEBS. On Intel systems, it provides access to nearly all sampled information, whereas on AMD systems many details, such as the latency for TLB refills or the number of occupied

LFB entries will not be shown. One way to get this information is to read the raw samples that are provided by the *perf subsystem*. However, these raw samples need to be manually processed to yield the required information, for example, by using libraries like *perf-cpp* [21].

A brief analysis of the IBS driver³ for AMD CPUs in a recent Linux kernel version indicates that its current implementation primarily focuses on populating the existing *perf* data structures. To accommodate e.g., additional latencies from IBS memory samples, modifications to the driver and the `PERF_SAMPLE_WEIGHT_STRUCT` may be considered. This struct is specifically designed to capture different latencies [17] and may be easily adapted to support the additional latencies offered by AMD CPUs.

Although implementing these modifications may require only minor changes to existing data structures, the introduction of new structures into the *perf* subsystem—such as those needed to report outstanding memory requests via allocated MAB slots—could prove more complex, as it would, e.g., necessitate more invasive adaptations in profiling tools like *perf*.

3 Practical Insights

To briefly illustrate the architectural divergence between these sampling mechanisms, we use a B⁺-tree [15, 28]⁴ *lookup* operation as our case study on a AMD Zen 4 system and a machine with Intel's Sapphire Rapids architecture.

Recording Data Access Information. We leverage *memory-access sampling* (e.g., [18, 19, 22, 23]) through the *perf subsystem* via the *perf-cpp* library [21]. Our implementation periodically captures memory addresses and associated metrics—including comprehensive latency characteristics—at fixed sampling intervals: Every 8 000th load instruction on Intel architectures (using the *mem-load* event) and every 8 000th μ Op on AMD platforms (utilizing the *IBS Op* PMU). Notably, extracting granular metrics on the AMD system necessitated explicit configuration of the PMU to record raw values, as critical data points—particularly TLB refill latency—remain obscured behind conventional interfaces.

For our experimental evaluation, we employ lookup operations using the *YCS Benchmark* [9], executing 100 M lookups against a tree populated with 100 M records. Figure 2 visualizes the access latency distributions recorded via AMD's IBS and Intel's PEBS for two critical tree levels: the root node (left) and leaf nodes (right). The plots illustrate the average latency for individual lookups, segmented according to the latency measurement capabilities of each sampling implementation.

Unlike instruction sampling, which allows the correlation of performance data with lines of code and functions, *memory-access sampling* enables the direct mapping of samples including memory addresses to specific tree nodes and their structural components (e.g., headers, keys, and payloads). This distinction is crucial, as all nodes share identical code paths, making instruction-based sampling inadequate for analyzing access characteristics across distinct nodes or tree levels.

¹In virtually-indexed-physically-tagged (VIPT) caches, the address translation and cache access can be parallelized to a certain extent.

²For the sake of simplicity, we use the term LFB in this paper, although we refer to the MAB on AMD systems.

³<https://github.com/torvalds/linux/blob/v6.14/arch/x86/events/amd/ibs.c>

⁴We borrowed the implementation from <https://github.com/wangziqu2016/index-microbench>.




Figure 2: The average access latency during B^+ -tree lookups at the root node (left) and the leaf nodes (right). Since IBS only reports cache-miss latency, no latency is reported for the root node. The μ Op tag-to-completion latency is calculated by subtracting μ Op completion-to-retire latency from the μ Op tag-to-retire latency.

Observation: Memory-Sampling Events. Intel’s PEBS collects samples by triggering on *mem-load* (and also *mem-store*) events [23], whereas AMD’s IBS includes a dedicated PMU that tracks retiring *micro operations* and records their complete memory-access context. Consequently, a nominally identical sampling interval produces architecture-dependent sample counts that scale with the application’s memory intensity. In our benchmark, a sampling interval of 8 000—i.e., every 8 000th load on Intel and 8 000th μ Op on AMD—, yields 8 348 memory samples on the Intel machine, yet the same period results in more than 600 000 samples on the AMD system. Note that we only count samples associated with the B^+ -tree.

To achieve a similar amount of memory samples on the AMD system, it might be necessary to increase the sampling interval by an order of magnitude. An additional caveat is that the IBS μ Op-PMU still records samples not inevitably linked to memory, even when the engineer is only interested in memory access information. Those extra samples are discarded during post-processing but nonetheless add overhead.

Observation: L1d Miss vs. Access Latency. As one would expect, the root node exhibits minimal access latency due to its tendency to reside in the L1d (see left side of Figure 2). However, IBS reports no latency measurements in this scenario, as it exclusively captures cache *miss* events. In contrast, PEBS provides granular insights, reporting 5 CPU cycles for L1d access and an additional 2 cycles for instruction retirement. This subtle difference is particularly noticeable when calculating an average latency over a set of accesses, including both L1d hits and various cache misses. Consequently, for a fair comparison between the two hardware substrates, L1d hits should be excluded from the calculation.

Observation: TLB Latency. Leaf node accesses present a different profile, frequently triggering cache and TLB misses. For the header segment—typically the first segment accessed—IBS delivers detailed timing breakdowns: approximately 200 cycles for TLB refill operations plus 230 cycles for cache miss resolution and ~ 90 cycles for retiring the μ Op. PEBS, however, presents a more consolidated view, reporting 210 cycles for cache operations and roughly 160 cycles for instruction retirement, with the latter inherently incorporating TLB latency.

Observation: Identifying Software Prefetches. Another key divergence between the two sampling facilities appears when demand loads and software prefetches are intermixed. AMD’s IBS tags every sample originating from an explicit software *prefetch** instruction [8]—triggering an asynchronous cache fill—with an appropriate

flag. Intel’s PEBS offers no such feature: all memory reads—demand or prefetch—enter the trace under the generic *load* category. Plus, PEBS reports the data source and cache access latency of software prefetches as if the access occurred in the L1 cache [12]—rendering this information negligible for profiling prefetches.

However, although the *perf subsystem* nominally defines a category for prefetches (along with loads and stores), the PEBS and IBS drivers never raise that bit. Only recording and decoding the raw IBS words allows users to discover data accesses invoked by prefetch instructions as such.

Consequently, AMD systems allow measuring its standalone latency—insights that matter as a prefetch itself may stall when, for instance, the memory address is not found in the TLB or all LFB slots are already occupied [13]. Simultaneously, IBS preserves the data source, revealing exactly where the speculative line originated. Together, the standalone latency and spatial origin form a precise compass for choosing an effective prefetch distance—a task that has often proved cumbersome (e.g., [14, 24]).

Observation: Detailed MAB Information. Beyond the distinction between loads and prefetches, IBS enriches each memory record with two further cues: a bit that signals whether the request hit the MAB and a counter that reveals how many MAB slots were occupied at sample time. PEBS, by contrast, simply reports “LFB” as the data source whenever the line is still in flight and suppresses the line’s ultimate supplier (L2, LLC, or DRAM).

The practical fallout is twofold. First, by correlating the *prefetch* samples with the true lower-level source and the measured latency, IBS lets users observe whether the prefetch arrived in time to hide the miss or whether the demand load overtook it. PEBS likewise tells us that the line was still in the LFB—hence accessed “too early” from the load’s perspective—making it challenging to tune the distance between the prefetch and the actual access quantitatively. Second, the additional MAB-occupancy counter exposes situations in which an otherwise helpful prefetch monopolizes the scarce buffer entries and causes stalls for subsequent misses—buffer pressure that is hardly visible under PEBS and must be inferred from secondary symptoms [13].

4 Conclusion and Outlook

This paper provided a condensed overview of the differences in *memory-access sampling* on recent x86 architectures. We showed that hardware makers provide rich metadata in their memory sampling mechanisms, yet the nature and visibility of this information differs noticeably. For example, AMD and Intel offer distinct sorts

of access latency; AMD exposes prefetch flags and MAB-occupancy metrics. But, not all information is clearly communicated through the *perf subsystem*. These discrepancies complicate cross-platform performance comparison but simultaneously open specialized optimization avenues on both platforms.

This preliminary work serves as a foundation for further research, as we intend to investigate other hardware architectures in greater detail, including the Statistical Profiling Extension (SPE) in recent ARMv8 systems.

Acknowledgments

We thank the anonymous reviewers for their helpful feedback and suggestions. This work has received funding from the DFG Priority Program “Disruptive Memory Technologies” (SPP2377) as part of the project “Memory Diplomat” (grant number 502384507) and has partly been funded by the Federal Ministry of Education and Research of Germany and the state of North-Rhine Westphalia as a part of the Lamarr-Institute for Machine Learning and Artificial Intelligence.

References

- [1] Advanced Micro Devices, Inc. 2024. *AMD64 Technology. AMD64 Architecture Programmer's Manual - Volume 2: System Programming*. Santa Clara, CA, USA.
- [2] Advanced Micro Devices, Inc. 2024. *Processor Programming Reference (PPR) for AMD Family 19h Model 11h, Revision B2 Processors*. Santa Clara, CA, USA.
- [3] Advanced Micro Devices, Inc. 2024. *Processor Programming Reference (PPR) for AMD Family 1Ah Model 44h, Revision B0 Processors*. Santa Clara, CA, USA.
- [4] Soramichi Akiyama and Takahiro Hirofuchi. 2017. Quantitative Evaluation of Intel PEBS Overhead for Online System-Noise Analysis. In *Proceedings of the 7th International Workshop on Runtime and Operating Systems for Supercomputers, ROSS@HPDC*. ACM, 3:1–3:8. doi:10.1145/3095770.3095773
- [5] Cagri Balkesen, Jens Teubner, Gustavo Alonso, and M. Tamer Özsu. 2013. Main-memory hash joins on multi-core CPUs: Tuning to the underlying hardware. In *29th IEEE International Conference on Data Engineering*. IEEE Computer Society, 362–373. doi:10.1109/ICDE.2013.6544839
- [6] Alexander Beischl, Timo Kersten, Maximilian Bandle, Jana Giceva, and Thomas Neumann. 2021. Profiling dataflow systems on multiple abstraction levels. In *EuroSys '21: Sixteenth European Conference on Computer Systems*, Antonio Barbalace, Pramod Bhatotia, Lorenzo Alvisi, and Cristian Cadar (Eds.). ACM, 474–489. doi:10.1145/3447786.3456254
- [7] Peter A. Boncz, Stefan Manegold, and Martin L. Kersten. 1999. Database Architecture Optimized for the New Bottleneck: Memory Access. In *Proceedings of 25th International Conference on Very Large Data Bases*. Morgan Kaufmann, 54–65. <http://www.vldb.org/conf/1999/P5.pdf>
- [8] David Callahan, Ken Kennedy, and Allan Porterfield. 1991. Software Prefetching. In *ASPLOS-IV Proceedings - Forth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM Press, 40–52. doi:10.1145/106972.106979
- [9] 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*. ACM, 143–154. doi:10.1145/1807128.1807152
- [10] Paul J Drongowski. 2007. Instruction-based sampling: A new performance analysis technique for AMD family 10h processors. *Advanced Micro Devices* 1, 3 (2007), 11.
- [11] Paul J. Drongowski, Lei Yu, Frank Swehosky, Suravee Suthikulpanit, and Robert Richter. 2010. Incorporating Instruction-Based Sampling into AMD CodeAnalyst. In *IEEE International Symposium on Performance Analysis of Systems and Software, ISPASS*. IEEE Computer Society, 119–120. doi:10.1109/ISPASS.2010.5452049
- [12] Intel®. 2024. Intel® 64 and IA-32 Architectures Software Developer's Manual. <https://cdrdv2.intel.com/v1/dl/getContent/671200>. Accessed: March 20, 2025.
- [13] Roland Kühn, Jan Mühlig, and Jens Teubner. 2024. How to Be Fast and Not Furious: Looking Under the Hood of CPU Cache Prefetching. In *Proceedings of the 20th International Workshop on Data Management on New Hardware, DaMoN*. ACM, 9:1–9:10. doi:10.1145/3662010.3663451
- [14] Jaekyu Lee, Hyesoon Kim, and Richard W. Vuduc. 2012. When Prefetching Works, When It Doesn't, and Why. *ACM Trans. Archit. Code Optim.* 9, 1 (2012), 2:1–2:29. doi:10.1145/2133382.2133384
- [15] Viktor Leis, Michael Haubenschild, and Thomas Neumann. 2019. Optimistic Lock Coupling: A Scalable and Efficient General-Purpose Synchronization Method. *IEEE Data Eng. Bull.* 42, 1 (2019), 73–84. <http://sites.computer.org/debull/A19mar/p73.pdf>
- [16] Justin J. Levandoski, David B. Lomet, and Sudipta Sengupta. 2013. The Bw-Tree: A B-tree for new hardware platforms. In *29th IEEE International Conference on Data Engineering, ICDE*. IEEE Computer Society, 302–313. doi:10.1109/ICDE.2013.6544834
- [17] Kan Liang and Peter Zijlstra. 2021. *perf/core: Add PERF_SAMPLE_WEIGHT_STRUCT*. <https://git.kernel.org/pub/scm/linux/kernel/git/tip/tip.git/commit/?id=2a6c6b7d7ad346f0679d0963cb19b3f0ea7ef32c>. Online; last accessed May 23, 2025.
- [18] Tongping Liu and Xu Liu. 2016. Cheetah: detecting false sharing efficiently and effectively. In *Proceedings of the 2016 International Symposium on Code Generation and Optimization, CGO*. ACM, 1–11. doi:10.1145/2854038.2854039
- [19] Xu Liu and John M. Mellor-Crummey. 2013. A data-centric profiler for parallel programs. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC'13*. ACM, 28:1–28:12. doi:10.1145/2503210.2503297
- [20] Stefan Manegold, Peter A. Boncz, and Martin L. Kersten. 2002. Optimizing Main-Memory Join on Modern Hardware. *IEEE Trans. Knowl. Data Eng.* 14, 4 (2002), 709–730. doi:10.1109/TKDE.2002.1019210
- [21] Jan Mühlig. 2023. *perf-cpp: Access Performance Counters from C++ Applications*. <https://github.com/jmuehlig/perf-cpp>. Accessed: March 20, 2025.
- [22] Jan Mühlig, Roland Kühn, and Jens Teubner. 2025. Understanding Application Performance on Modern Hardware: Profiling Foundations and Advanced Techniques. In *3rd Workshop on Novel Data Management Ideas on Heterogeneous Hardware Architectures (NoDMC)*. GI.
- [23] Stefan Noll, Jens Teubner, Norman May, and Alexander Böhm. 2020. Analyzing memory accesses with modern processors. In *16th International Workshop on Data Management on New Hardware, DaMoN*, Danica Porobic and Thomas Neumann (Eds.). ACM, 1:1–1:9. doi:10.1145/3399666.3399896
- [24] Georgios Psaropoulos, Thomas Legler, Norman May, and Anastasia Ailamaki. 2017. Interleaving with Coroutines: A Practical Approach for Robust Index Joins. *Proc. VLDB Endow.* 11, 2 (2017), 230–242. doi:10.14778/3149193.3149202
- [25] Michael L. Samuel, Anders Uhl Pedersen, and Philippe Bonnet. 2005. Making CSB+ Tree Processor Conscious. In *Workshop on Data Management on New Hardware, DaMoN*. <http://www-2.cs.cmu.edu/%7Edamon2005/damonpdf/2%20making%20csb+%20trees%20processor%20conscious.pdf>
- [26] Muhammad Aditya Sasongko, Milind Chhabbi, Paul H. J. Kelly, and Didem Unat. 2023. Precise Event Sampling on AMD Versus Intel: Quantitative and Qualitative Comparison. *IEEE Trans. Parallel Distributed Syst.* 34, 5 (2023), 1594–1608. doi:10.1109/TPDS.2023.3257105
- [27] Ambuj Shatdal, Chander Kant, and Jeffrey F. Naughton. 1994. Cache Conscious Algorithms for Relational Query Processing. In *Proceedings of 20th International Conference on Very Large Data Bases*. Morgan Kaufmann, 510–521. <http://www.vldb.org/conf/1994/P510.PDF>
- [28] Ziqi Wang, Andrew Pavlo, Hyeontaek Lim, Viktor Leis, Huanchen Zhang, Michael Kaminsky, and David G. Andersen. 2018. Building a Bw-Tree Takes More Than Just Buzz Words. In *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10–15, 2018*, Gautam Das, Christopher M. Jermaine, and Philip A. Bernstein (Eds.). ACM, 473–488. doi:10.1145/3183713.3196895