

How to Be Fast and Not Furious: Looking Under the Hood of CPU Cache Prefetching

Roland Kühn*
roland.kuehn@cs.tu-dortmund.de
TU Dortmund University
Lamarr Institute for Machine
Learning and Artificial Intelligence

Jan Mühlig*
jan.muehlig@tu-dortmund.de
TU Dortmund University
Lamarr Institute for Machine
Learning and Artificial Intelligence

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

ABSTRACT

Software-based prefetching is a powerful method for tolerating access penalties that are encountered by data processing systems: *memory latency*. Although the idea appears straightforward—simply informing the CPU about upcoming data accesses—the intricacies of its implementation remain insufficiently understood. Existing works demonstrate how to rewrite algorithms for prefetching, yet they often overlook the limitations and hardware implications of bringing data into the cache hierarchy. In this paper, we examine software-based prefetching thoroughly by delving into its implementation and identifying pitfalls across various platforms. Furthermore, we provide actionable insights and recommendations for developers seeking to boost their applications through this technique.

ACM Reference Format:

Roland Kühn, Jan Mühlig, and Jens Teubner. 2024. How to Be Fast and Not Furious: Looking Under the Hood of CPU Cache Prefetching. In *20th International Workshop on Data Management on New Hardware (DaMoN '24)*, June 10, 2024, Santiago, AA, Chile. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3662010.3663451>

1 INTRODUCTION

Memory bandwidth and *latency* present one of the highest barriers when aiming for maximum system performance [7]. While many steps have been made to increase the DRAM bandwidth, latency improvements have hit a plateau, leaving a chasm between potential and realized system speed. And new memory technologies, such as high-bandwidth and non-volatile memory, come with heightened access penalties, up to hundreds of nanoseconds [12, 31]. Trying at least to *tolerate* the memory access latency, hardware manufacturers deploy huge caches, invent complex out-of-order execution strategies, and implement sophisticated *hardware prefetching* mechanisms. However, the reliance on hardware alone quickly reaches its limit when the CPU is confronted with non-trivial access patterns. Navigation through tree and hash data structures are prime examples of this limitation.

*Both authors contributed equally to the paper

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
DaMoN '24, June 10, 2024, Santiago, AA, Chile
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0667-7/24/06.
<https://doi.org/10.1145/3662010.3663451>




Figure 1: Results of a micro-benchmark accessing single cache lines in random order, showing the speedup of software-prefetched vs. purely random executions. The data reveals an emerging trend across diverse CPU manufacturers and generations: software prefetching is becoming more advantageous with recent architectures.

Software-based prefetching promises to overcome this situation: Applications can proactively hint to the CPU about upcoming data accesses that are hard to predict by only monitoring the memory access stream. Typically, these hints are communicated via specific instructions embedded within modern ISAs such as x86 and ARMv8. Especially in the context of trees and hash tables, many studies have underscored the significant potential that software prefetching holds (e.g., [9, 11, 16, 18, 24, 28, 30]). However, the discourse often overlooks the incurred costs, specific limitations, and implications across different hardware platforms.

We argue that, in the future, it will be even more important to understand and assess software prefetching in order to optimize data-intensive systems. To substantiate our reasoning, we utilize a micro-benchmark that accesses a series of cache lines in random order on various hardware platforms—with and without software prefetching (more details will follow in Section 3). The resulting *trend* in Figure 1 emphasizes that the potential of software-based prefetching over pure random access has increased dramatically over the past few years. In particular, we could observe that the trend correlates closely with the clock frequencies of DRAM, which have continuously increased. And with an eye on the often-cited *memory wall*, we expect the trend to continue.

As a consequence, this survey conducts an in-depth examination of software-based prefetching. We study how the technique is implemented in various systems and turn our insights into guidance

for developers seeking to optimize their applications using software prefetching. To that end, Section 2 recaps software prefetching from what we found in the literature. In Section 3, we utilize the briefly mentioned micro-benchmark to reveal the precise costs and limitations of executing prefetch instructions and show how prefetchers from hardware and software can work collaboratively. Section 4 demonstrates that our insights are truly valuable for widely-used data structures. Finally, we conclude in Section 5.

2 UNDERSTANDING SOFTWARE PREFETCHING

At first glance, the concept of software-based prefetching appears seemingly straightforward: data accesses, which are challenging or unfeasible for the hardware to predict, can be communicated to the CPU via specific instructions—interpreted as *hints* from the application. Based on these hints, the underlying hardware substrate can move data *asynchronously* into a cache close to the CPU to minimize the access latency—hiding the *actual* latency behind computational work. Data structures that exhibit indirect access patterns provide good examples for instructing the hardware with insights from the software layer, for instance, linked lists, queues, hash tables, and trees.

2.1 Utilize Software Prefetching

However, formulating such predictions is often cumbersome. Consider the traversal through a tree-like data structure: The immediate succession of identifying and accessing a node leaves negligible time for the hardware to load the data, even if the software hints about the soon-to-be-accessed tree node. In order to create a (sufficient) temporal gap between pinpointing the next node and accessing it, the literature explores various methods. The key idea is to break operations into stages and use pipelining, overlaying the prefetch of one item with the compute of others, for example, by processing operations in groups [9, 18]. A more general approach is to utilize asynchronous control flow abstractions, such as coroutines [11, 16, 30, 37] and fine-grained tasks [28]. In addition, the discussion over the strategic implementation of optimization passes to automatically inject prefetch instructions into compiler-generated code has been ongoing for decades (e.g., [6, 15, 22, 27]).

Not only does the redesign of algorithms and data structures play a decisive role in the efficient use of software prefetching, but the instruction must also be issued at the appropriate moment [15]. When software executes the prefetching instruction too early before the actual access, the data may have been already evicted from the cache when needed. Conversely, if prefetched too late, the data might not be transferred to the cache completely, potentially causing the CPU to wait. Accurate prefetch timing necessitates careful understanding of both the execution time of the instructions preceding the actual load and the system’s memory latency [6, 27]. However, the presence of intricate memory systems like NUMA and high-bandwidth memory, as well as the increasing sophistication of CPUs, which now feature out-of-order execution and superscalar architectures, exacerbate the complexity of achieving precise timing. A promising approach to tackle these challenges is the implementation of profile-guided approaches [15].




Figure 2: Lifecycle of various prefetch instructions through the memory subsystem.

2.2 Lifecycle of a Prefetch

Almost all modern server and desktop platforms offer several prefetch instructions to target different levels of the cache hierarchy. Notably, each instruction will initiate the transfer for one cache line. The implementation within various hardware systems has only slight variations, even across different vendors. When the software executes a prefetch, the logical address to be prefetched is translated first into a physical one. This procedure happens *synchronously*, i.e., if the translation is not cached within the Translation Lookaside Buffer (TLB), the CPU will contact a Page Miss Handler (PMH) before continuing the execution of the prefetch instruction. We visualize the *lifecycle* of prefetch requests through the memory subsystem in Figure 2.

Once the physical address is known and the request misses the L1 data cache (L1d), the hardware will trigger the transfer from memory to caches *asynchronously*. To that end, the L1d cache requests the cache line through the Line Fill Buffer (LFB) (or Miss Address Buffer (MAB) on AMD platforms¹). The LFB, as suggested by its name, acts as a buffer-like structure that interfaces between the L1d and L2 cache to communicate requests on a cache-line granularity [8, 33, 34]. For every cache line that misses the L1d, the CPU allocates a slot in the LFB. The L2 cache, on the opposite end, retrieves and delivers the requested lines. This nature makes it straightforward to implement *asynchronous* prefetches: The desired cache line will not cause the instruction to wait until the data is transferred into the cache but only until the miss is communicated to the LFB. Furthermore, the LFB enables the CPU to handle several pending requests simultaneously without blocking a single one, supporting out-of-order execution and performing micro-optimizations like merging multiple requests to the same cache line [34]. When the request also misses the L2 cache, it seeks the data from even higher cache levels or main memory. On Intel platforms, the superqueue [19]—a buffer positioned between the L2 and off-core last-level caches (LLCs)—tracks pending requests.

Upon completing a memory request, the specific software prefetch instruction dictates how close the data is moved to the CPU. In the x86 and ARM ISAs, most prefetch instructions clearly specify the cache level target for the fetched data: Cache lines fetched by `prefetcht1` (and its counterpart `p1d12keep` on ARM) are stored in the L2 cache; `prefetcht0` (`p1d11keep` on ARM) moves data further into the L1d, thereby passing through the LFB. The `prefetcht2` instruction is generally aimed to fetch data into the LLC. However, recent Intel platforms (since Skylake) have implemented the LLC as a non-inclusive victim cache—in that case, data is placed into the L2 instead. Conversely, for the *non-temporal* `prefetchnta` instruction, the x86 ISA does not specify a particular cache target but only the

¹For the sake of simplicity, we only ever refer to the LFB, but also intend the MAB.

objective: reducing cache pollution for data that will be accessed nonrecurring.

AMD’s documentation for earlier processor generations (e.g., Bulldozer [3]) specifies that the prefetch instruction pulls data into the L1d and ensures it is not evicted to the L2 cache unless it originated from there. For later generations (i.e., for Zen), the documentation no longer mentions the exact cache destination but merely states that the L2 is bypassed when evicting non-temporal data. Intel, in contrast, also specifies for modern platforms that non-temporal prefetches move data into the L1d and into the LLC if the LLC is inclusive [13]. Both Intel and AMD note that these non-temporally prefetched cache lines are prioritized for quicker eviction [4, 13].

2.3 Limitations

Modern platforms exhibit two notable limitations in the implementation of software prefetches. First, the execution of a single prefetch instruction stalls until the virtual address is translated into a physical address [13, 30]. Although the TLB might accelerate this process, it is likely that the most prefetched addresses are not present in the TLB as software prefetching is used primarily when the application accesses scattered data objects. Consequently, the latency of the prefetch instruction is dominated by the latency of the PMH performing a page-table walk to translate the address. However, this restriction only applies to the initial prefetch operation within a memory page if several prefetches are performed consecutively to load a block larger than a single cache line. While hardware also implements prefetchers for address translations [35], the software is only empowered to preload data.

Drawing from that observation, it seems advantageous to use prefetching for larger, contiguous data regions by executing multiple instructions while only experiencing a single TLB miss per page. However, this may introduce a secondary constraint: Due to the limited number of LFB entries (typically between 3 and 24), the hardware can accommodate only a small number of outstanding memory requests. If an excessive volume of prefetch requests is sent out rapidly, the LFB becomes filled and cannot accept further requests. Under this circumstance, the CPU will either stall until an LFB slot becomes ready or—according to Intel’s documentation [13]—drop new prefetch requests.

3 SOFTWARE PREFETCHES UNDER THE MAGNIFYING GLASS

To our surprise, these limits have only been studied in depth to a limited extent, although the technique is put to use by several applications in both research (e.g., [9, 11, 16, 24, 25, 30]) and “the real world” (e.g., *Linux* and *TCMalloc*). More than a decade ago, Lee et al. investigated both hardware and software prefetching in the context of widely used applications [20]. In this chapter, we aim to expand upon their foundational research, delving deeper into the nuances of software prefetching. We carefully analyze the costs and implications involved in modern platforms and further examine the synergies between software prefetching and modern hardware prefetchers.




Figure 3: Illustration of the micro-benchmark.

3.1 Micro-Benchmark

To that end, we utilize a micro-benchmark² that aligns with conventional latency benchmarks designed to mimic pointer-heavy access patterns: We employ two arrays, one (*access*) array dictating the order in which the second (*data*) array is accessed. Figure 3 illustrates the structure of the benchmark. The data within this second array is organized as blocks; for different experiments, we vary the size of the blocks from one to several cache lines. For instance, the benchmark underlying the trend in Figure 1 utilizes one cache line per block. Each block’s data is accessed sequentially. To introduce work behind which we can hide memory latency with prefetching, we split each cache line into an array of 32-bit integers and compute hash values (a total of 568 instructions per cache line). All the following experiments are single-threaded.

Hardware Setup. Our evaluation spans a range of different hardware platforms to ensure a broad understanding. However, for illustrative purposes, we will strategically focus on a machine powered by an Intel Xeon Gold 6226 CPU (Cascade Lake architecture). All following plots are performed on this setup. This focus is a result of the unique profiling capabilities enabled by this hardware: Unlike CPUs from other manufacturers or different generations, this system offers specific performance counters and address-sampling features essential for our evaluation.

To further provide a reference for the following measurements: On the mentioned machine, the micro-benchmark’s raw processing time for a single cache line is approximately 193 cycles.

3.2 Quantifying the Costs

So far, the precise costs associated with injecting additional prefetch instructions remain undetermined. Our following examination identifies three principal cost factors during execution: the CPU’s retirement of the instruction, the translation of the address to be prefetched into a physical one, and the allocation of a slot in LFB. Pinpointing the exact source of costs caused by the prefetch instruction remains an intricate challenge. To gain a deeper understanding, we assess the difference in compute cycles between a sequential execution and one that accesses data randomly but uses software prefetching. Thanks to the efficiency of the hardware’s data and TLB prefetchers, the sequential execution encounters almost no stalls, thereby revealing the pure computational cycles of our workload.

²The code is available at <https://dbis.cs.tu-dortmund.de>.




Figure 4: Breakdown of overhead CPU cycles (compared to a sequential execution) associated with a single prefetch for different memory page sizes and instructions.

Execution and Address Translation. For our first experiment, we configure our micro-benchmark to use blocks of only one cache line, accessing all cache lines in random order. By utilizing performance counters, we can categorize the consumed cycles into compute-related and stalled cycles (see more details about performance counters in Appendix A). Furthermore, we can dissect the stalled cycles into two subsets: those occurring due to memory waiting times (using the `CYCLE_ACTIVITY.STALLS_MEM_ANY` counter) and “other” stalls (tracked via `CYCLE_ACTIVITY.STALLS_TOTAL`³). Figure 4 presents the results, showing the cycle *overhead* (beyond the sequential execution) while prefetching single cache lines by software using different memory page sizes. With “common” 4 kB-sized pages, we can observe that a single prefetch induces an additional approx. 60 CPU cycles, most of them related to stalls of unidentifiable sources.

Remarkably, the major amount of these additional cycles diminishes when utilizing Hugepages, simultaneously reducing the number of TLB misses. This observation leads us to deduce that these cycles are caused primarily by stalls while translating the address to prefetch, which is in line with further TLB-focused studies [23]. Only the `prefetcht1` instruction, which directs data into the L2 cache, incurs minimal overhead of up to 9 cycles, even when Hugepages are used. However, the translation (and its related costs) is somewhat inevitable and occurs throughout workloads with random and scattered access patterns; it is only moved forward in time by the execution of the prefetch instruction (and thus associated with the prefetch instead of the actual data access). The reason for stalling lies in the handling of TLB misses during the prefetch execution, which is always processed *synchronously*—compelling the CPU to wait while the PMH takes action. Conversely, the actual data cache misses are processed *asynchronously* with the assistance

³Note that the `STALLS_TOTAL` number already includes memory stalls; thus, it is necessary to subtract the latter.


of the LFB. Note that also memory stalls experience a slight decrease using Hugepages since the PMH can also cause data cache misses.

While the precise categorization of cycles is only possible on a few CPU architectures, we see similar patterns on different CPUs across different hardware manufacturers, like Intel, AMD, and ARM, and architectures (in fact, across all CPUs listed in Figure 1). For instance, on an Intel platform from the Sandy Bridge series, we note a comparable overhead using 4 kB memory pages and only 3 cycles of overhead with the adoption of Hugepages. Conversely, on contemporary AMD platforms (Zen 4), the observed overhead is around 40 cycles for 4 kB pages and ranges from 7 (for `prefetchnta` and `prefetcht0` instructions) to 14 cycles (for `prefetcht1`) when utilizing Hugepages.


Line Fill Buffer. Prefetches—and loads in general—missing the L1d are forwarded via the LFB to the higher memory subsystem. The capacity of these buffers is quite limited: Intel’s architectures provide 10 (e.g., Sandy Bridge and Haswell [13]) to 16 (e.g., Golden Cove [13]) LFB slots, while modern AMD CPUs (like Zen 4) offer up to 24 MAB entries [5]. When the LFB is at full capacity, it cannot accept new prefetch or load requests, causing the CPU to stall until slots become available as previous requests are completed by the L2 cache. Since the software-prefetch instruction allocates a slot *synchronously*, prefetching large data blocks at once (e.g., an entire tree node) can exceed the buffer’s capacity, leading to increased latency when issuing the instruction.

However, precisely breaking down the overhead for LFB-related stalls is not as straightforward as for address translation-related costs since these stalls only occur for requests that flood the buffer. Thus, not all prefetches are impacted equally. To still get an impression, Figure 5 shows the results of a scenario where our micro-benchmark prefetches blocks of 16 cache lines at once. More precisely, Figure 5a shows the average additional cycles per cache line in such an execution over a sequential one (analogously to the comparison for single-line blocks in Figure 4). With 4 kB memory pages, the illustrated overhead includes both TLB miss penalties and LFB-related stalls, ranging around 35 additional cycles per prefetch. However, utilizing 1 GB Hugepages eliminates almost all translation-related stalls, leaving approx. 20 stalled cycles per cache line, which we associate with the waiting times for an available LFB slot. To back this further, we also evaluated the number of requests finding an already full LFB (via the counter `L1D_PEND_MISS.FB_FULL`), which increases from almost zero (when prefetching a single cache line) to 1.2 per cache line when prefetching a block of 16 cache lines.

Again, the stalled cycles demonstrated in Figure 5a are an average over all cache lines of an entire block; but not all prefetches experience identical costs. Delving deeper, Figure 5b dissects the mean *latency* in cycles for each prefetched cache line within a block when utilizing 1 GB Hugepages. We recorded these results using address sampling via the *perf* interface [1], using the `PERF_SAMPLE_WEIGHT` feature to collect latency information. Notably, there is a marked latency escalation starting from the eleventh cache line, where the average latency values leap from 7 to a range of 200 – 340 cycles. This pattern suggests that the underlying Cascade Lake architecture has exactly 10 LFB slots, which is in line with (unofficially



(a) Overhead in cycles



(b) Sampled latency of individual prefetch instructions within 1 GB Hugepages

Figure 5: Results of the micro-benchmark using 16 cache lines per block.

published [2]) numbers for the very similar Skylake architecture. Furthermore, we verified these observations on Intel Sandy Bridge and Haswell systems, which officially have 10 LFB slots, showcasing similar latency patterns as depicted in Figure 5b. Similarly to the previous benchmarks, we can confirm results comparable to Figure 5a on AMD’s Zen 4, which does not enable cycle breakdowns, by utilizing the “common” cycle counter. However, evaluating the latter hardware necessitates prefetching 32 cache lines to detect a similar latency pattern, reflecting the MAB’s larger capacity of 24 slots.

Discussion. Based on our findings, we conclude that the execution of prefetch instructions causes only negligible pressure on the CPU’s instruction bandwidth. Most of the costs associated with a single prefetch instruction are attributable to the (unavoidable) penalties of TLB misses. Mitigating these costs demands hardware-level support, such as adopting larger memory page sizes. Looking ahead, an ideal advancement would be the implementation of asynchronous translation mechanisms, allowing prefetch operations to proceed without stalling. Additionally, the introduction of a specialized instruction or mechanism to prefetch address translations, like it is implemented e.g., in the IAA Accelerator on some CPUs of the Intel Sapphire Rapids Architecture [14], could significantly streamline the process. To avoid expensive LFB-related stalls (see Figure 5), prefetching algorithms necessitate careful engineering (see details in Section 4 below). It is essential to consider the LFB’s capacity, which varies not only across hardware manufacturers but also between different generations—requiring *hardware-conscious* implementations. For instance, AMD’s architectures use up to 24 slots, whereas the latest Intel architectures support up to 16.

3.3 Interplay with Hardware Prefetchers

We saw that software prefetching has its limits when it is necessary to prefetch large data blocks at once. To circumvent this

limitation and associated costs, a collaboration between software and hardware prefetching emerges as a promising solution, forming a symbiotic relationship: Software prefetching takes the initial step by early hinting about unpredictable data accesses; hardware prefetchers can take over, incrementally fetching the remaining data line by line. In this section, we delve into their cooperative dynamics, emphasizing how strategic software prefetching *stimulates* hardware prefetchers beyond the typical dependence on load misses.

Hardware Prefetcher Patterns. In our in-depth analysis, we concentrate on Intel processors, which offer valuable insights into prefetching dynamics by implementing sampling of data addresses and access latency through the Linux *perf* interface [1]. Intel utilizes multiple core-local prefetchers that serve the L1 and L2 caches across both modern and legacy processor generations [13, 32]. The L1d prefetchers monitor load streams within a single cache line, successively requesting the next line. Meanwhile, L2 prefetchers study the patterns of cache line requests traversing through the L2 cache. Here, the *stream prefetcher* initiates prefetches based on a sequence of consecutive cache line requests, while the *adjacent prefetcher* consistently loads a pair of neighboring cache lines at once. This architectural design prompts two critical lines of inquiry: Do software and hardware prefetchers have any interaction? If yes, what are the impacts of different prefetch instructions?

Software Prefetches Trigger the Hardware. Let us revisit the scenario presented in Section 3.2: When software-prefetching is used for randomly accessed data blocks that are 16 cache lines in size, we observed that the smaller-sized LFB becomes overwhelmed. Interestingly, our experiments reveal a compelling interaction wherein software-based prefetching actively triggers hardware prefetchers when injecting prefetch instructions only for the first subset of the needed cache lines.




Figure 6: Prefetching several cache lines by software to explore the hardware/software prefetching interaction.

Figure 6 illustrates the access latency for each cache line within a block when varying the quantity of prefetched cache lines: 8, 10, and 12 lines are prefetched using software (with `prefetchnta` for illustrative purposes), leaving the hardware prefetchers to manage the remainder. To contextualize the results, we also present a scenario without software prefetching (“random”), where we observe the hardware prefetchers quickly adapt to the sequential access pattern within a block. However, the hardware prefetchers are initially too late for every second cache line and cannot hide the entire latency. This could be ascribed to the small workload of our micro-benchmark (around 193 CPU cycles per cache line).

We can recognize a similar pattern when we software-prefetch only the first 8 cache lines. While the software-prefetched cache lines exhibit L1d-related latencies, we observe a noticeable latency spike in the first cache line not prefetched by software. It takes more accesses—up to the 12th cache line—for the hardware to prefetch the remaining data timely. Remarkably, the configuration where 12 cache lines are prefetched by software emerges as the best-performing one. Given that this amount appears sufficient for the hardware prefetchers to operate on time, it strikes the balance between minimizing LFB pressure and reducing access latency in our benchmark. While the results indicate that with doing more computational work, fewer software prefetches are sufficient, the primary conclusion remains: Software prefetches trigger the hardware prefetchers. We notice a comparable trend for AMD’s Zen 4 platform: The throughput reaches a plateau when prefetching only 6 cache lines in software, indicating that the hardware prefetchers are successfully taking over.

Different Instructions Trigger Differently. However, not all software prefetch instructions trigger hardware prefetchers the same way. Our analysis of different instructions reveals that prefetching into the L2 cache (via `prefetcht1`) tends to minimize latency penalties, especially when prefetching fewer than 12 cache lines. Figure 7 illustrates an evaluation where the initial 10 cache lines of a block are prefetched using various instructions, highlighting the differential impacts on latency.




Figure 7: Access-latency when prefetching 10 of 16 cache lines using different prefetch instructions.

Naturally, hardware prefetchers respond to the access patterns they encounter; prefetching directly into the L1d cache (using `prefetchnta` and `prefetcht0`) circumvents requests to the L2 cache—preventing it from “learning” from these access patterns. Consequently, the L2 cache encounters misses only when the first non-software-prefetched cache lines are accessed. This characteristic became particularly evident as we turned off the L1d hardware prefetchers during our study.

Furthermore, Figure 7 demonstrates a rapid convergence of latency for data prefetched into the L2 to those associated with L1d accesses. This occurs as a result of the L1d prefetchers, which detect cache misses and subsequently fetch the relevant cache lines from the L2 to the L1d cache. While address sampling is not applicable for AMD architectures, performance counters (particularly the counter `LS_HW_PF_DC_FILLS.LOCAL_L2`) indicate a similar behavior.

Discussion. The measured results allow us to conclude that hardware and software prefetchers interact with each other, even on different hardware platforms. Leveraging this interaction becomes essential when large amounts of data should be transferred from memory to the CPU cache—given the limitations of LFBs highlighted in Section 3.2. Consider the *range scan* of tree structures as an example: Although accessed in a logical order, leaf nodes are scattered in memory and accessed in a seemingly random fashion from the memory subsystem’s point of view. While software prefetching can signal to the hardware what leaf nodes will be accessed shortly, thus initiating the transfer into the cache, the hardware prefetcher can effectively take over once it recognizes the sequential access pattern. This principle also extends to morsel-driven database engines, such as HyPer [17] and Umbra [29, 36], which process fine-grained packages of tuples.

4 ASSESSING PREFETCHING LIMITS IN A COMMONLY USED DATA STRUCTURE

Based on our findings, it becomes evident that the LFB’s limited capacity presents a new bottleneck when prefetching is the software-side contribution to tolerating memory latency. Given that it is

already profitably implemented for tree-like data structures [16, 24, 28, 30], it is somewhat surprising that these LFB limitations have not been extensively studied in such contexts. When baking software prefetching into tree traversal algorithms, the size of the tree nodes becomes a critical factor that directly influences effectiveness. Historically, the node sizes of index structures have often been aligned with the block size of the underlying storage medium, such as HDDs or SSDs. But even in in-memory databases, where the node size does not have to be aligned with a certain storage medium, many tree-based index structures still use relatively large node sizes. One of the reasons is that trees with large nodes have a comparatively smaller depth, which in turn is associated with less time spent on random accesses, including the address translations during a tree traversal. However, when software prefetching is applied without considering the characteristics of the underlying hardware, these large node sizes can quickly become a challenging problem since fully prefetching a tree node may exceed the capacity of the LFBs within a CPU core, as already highlighted in Section 3.1.

Benchmark Setup. To illustrate the impact of node sizes on software prefetching, we examined a state-of-the-art B^+ -tree as proposed by Leis et al. [21]. We extended the tree to implement prefetching with coroutines similar to the approaches proposed in [16, 30]. The concept of coroutines allows a function to be suspended at a specific point and resumed later, which allows for *interleaved* execution; creating a time frame between the identification and access of a node. We modified the tree so that during a traversal a software prefetch that loads data in all caches (`prefetcht0`) is issued for the succeeding node and the traversal is suspended afterward. If a coroutine suspends, the coroutine scheduler can start the execution of another coroutine.

For our experiments, we executed 50M lookups with different node sizes in the original (hereinafter only referred to as B^+ -tree) and the prefetched, coroutine-based version of the B^+ -tree (B^{coro} -tree) and measured the execution in cycles per lookup. We used the YCSB benchmark [10] to fill the tree with 50 M entries (8 bytes for key and value each). As in Section 3.1, we used a machine with Intel’s Cascade Lake architecture as our primary test platform.

Results. For single-threaded execution, our results showed that a node size of 4 kB performs best for the B^+ -tree, compared to smaller node sizes (1 kB, 256 B), which we also use as a baseline for the B^{coro} -tree. In the B^{coro} -tree, however, we could observe the exact opposite. Here, a node size of 256 bytes performed best and outperformed the best-performing B^+ -tree by a factor of 2.25, while we saw less performance with bigger node sizes. While a node size of 1 kB still experiences a performance boost by factor 1.51, a node size of 4 kB performs worse compared to the baseline (factor 0.8). This behavior can be explained by the amount of cache lines that are needed for every node. While a node with a size of 256 B consists of 4 cache lines, a node size of 4 kB results in 64 cache lines, which exceeds the amount of 10 LFB slots on our primary test platform. As depicted in Figure 8, we could observe a very similar behavior on older Intel architectures, on Intel’s latest Sapphire Rapids architecture, however, a node size of 1 kB has only a small performance loss compared to a node size of 256 bytes. We attribute both observations to the LFB’s capacity since the observed older architectures (Sandybridge and Haswell) also have 10 LFB entries.




Figure 8: Speedup of a single lookup with different node sizes with and without SMT compared to the best unprefetched baseline.

On the Sapphire Rapids platform, the LFB was increased to 16 slots. On recent AMD architectures, we could also see great performance improvements (up to 2.32 \times on Zen 4). In contrast to the surveyed Intel platforms, it is remarkable that the B^{coro} -tree always outperforms the B^+ -tree, even with a node size of 4 kB. We also relate this mainly to the number of fill buffer entries on AMD architectures, which is larger than on Intel platforms (e.g., 24 entries for Zen 4).

Simultaneous Multithreading. Nearly every modern x86 architecture implements simultaneous multithreading (SMT) to better utilize the available resources of the CPU. While some CPU components are duplicated (e.g., some registers), other components, like caches, are shared between the threads. As observed in the previous sections, the LFB represents a critical resource when it comes to using software prefetches beneficially; the use of SMT is expected to bring this scarce resource even more into focus. Therefore, we executed our experiments also with two logical threads that run on the same physical core to measure the impact of sharing this critical resource. On our primary test platform, we could still see a performance improvement (factor 1.8) with a node size of 256 bytes, even if the number of cycles for executing one lookup operation increased, but this accounts also for the B^+ -tree with SMT, that was used as a baseline. Here, we attribute the increased number of cycles primarily to the type of workload we execute, as we do roughly the same kind of work in both threads and, therefore, might use the same shared resources. With a node size of 1 kB, however, we could see a noticeable degradation in performance (factor 0.98) when using SMT on our primary test system. In this case, both threads have to prefetch 16 cache lines each, which massively exceeds the number of the available 10 LFB entries. Similar to the execution with a single thread, we could observe this effect as well on older Intel architectures like Sandy Bridge or Haswell. On recent AMD architectures, we could also see an increase in cycles per lookup when executing two threads on the same physical core, but in contrast to Intel, even prefetching of 4 kB nodes performed better than the baseline. Our findings indicate that modern AMD architectures

are less vulnerable to prefetching larger regions, even when using SMT—in contrast to Intel designs.

The results from our study of the B^{coro}-tree demonstrate that the LFB can indeed become a limitation of software prefetching, also within real-world applications. Especially the use of SMT, which results in sharing this already scarce resource between two (logical) cores, intensifies the competition and potential for stalls when prefetching too large tree nodes. The evaluation shows—once again—the importance of considering the hardware’s characteristics when designing algorithms to achieve optimal performance.

5 CONCLUSION

Software-based prefetching holds enormous potential when seeking to hide memory latency behind valuable CPU cycles. In this survey, we took a closer look under the hood and studied the benefits and limitations of software prefetching on several hardware platforms. And we found that its performance depends dramatically on the developer’s knowledge of the underlying hardware substrate and the ability to adapt to it.

Our evaluations with synthetic and real-world workloads identified address translations and the number of pending memory requests a system can manage as the “next” bottleneck—beyond hiding memory access latency. To leverage the hardware to full performance, developers have to build their software-prefetching algorithms *hardware-conscious*, considering these system parameters. In order to help developers utilize the hardware to its full potential, our top wish towards hardware manufacturers would be for more counters that specifically track the duration of CPU stalls resulting from resource limitations, such as the LFB and PMHs. Additionally, we have experienced challenges with the documentation for available counters, finding it hard to follow and incomplete in some cases.

In the long term, we believe that the hardware substrate has to continue supporting the software to further reduce access penalties, for instance, through enabling prefetching of address translations in software and enhancing the capacity to handle a greater volume of prefetch requests.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and Maximilian Berens for their helpful feedback and suggestions. We also thank Till Miemietz and Adam Lackorzynski from the Barkhausen Institut/TU Dresden and Christian Hakert from the DAES group at TU Dortmund University for providing us access to additional hardware. This work has received funding from the DFG Priority Program “Disruptive Memory Technologies” (SPP2377) as part of the project “Memory Diplomat” (grant number 502384507).

REFERENCES

- [1] [n. d.]. `perf_event_open(2)` - Linux Manual Page. https://man7.org/linux/man-pages/man2/perf_event_open.2.html. Online; last accessed March 20, 2024.
- [2] [n. d.]. Skylake: Intel’s Longest Serving Architecture. <https://chipsandcheese.com/2022/10/14/skylake-intels-longest-serving-architecture/>. Online; last accessed March 22, 2024.
- [3] Advanced Micro Devices. 2014. Software Optimization Guide for AMD Family 15h Processors. https://www.amd.com/content/dam/amd/en/documents/archived-tech-docs/software-optimization-guides/47414_15h_sw_opt_guide.pdf. Online; last accessed March 20, 2024.
- [4] Advanced Micro Devices. 2020. Software Optimization Guide for AMD EPYC™ 7003 Processors. <https://www.amd.com/content/dam/amd/en/documents/processor-tech-docs/software-optimization-guides/56665.zip>. Online; last accessed March 20, 2024.
- [5] Advanced Micro Devices. 2023. Software Optimization Guide for the AMD Zen4 Microarchitecture. <https://www.amd.com/content/dam/amd/en/documents/processor-tech-docs/software-optimization-guides/57647.zip>. Online; last accessed May 05, 2024.
- [6] Sam Ainsworth and Timothy M. Jones. 2017. Software prefetching for indirect memory accesses. In *Proceedings of the 2017 International Symposium on Code Generation and Optimization*. ACM, 305–317. <http://dl.acm.org/citation.cfm?id=3049865>
- [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*. 54–65. <http://www.vldb.org/conf/1999/P5.pdf>
- [8] Pietro Borrello, Andreas Kogler, Martin Schwarzl, Moritz Lipp, Daniel Gruss, and Michael Schwarz. 2022. *ÆPIC Leak*: Architecturally Leaking Uninitialized Data from the Microarchitecture. In *31st USENIX Security Symposium*. USENIX Association, 3917–3934. <https://www.usenix.org/conference/usenixsecurity22/presentation/borrello>
- [9] Shimin Chen, Anastassia Ailamaki, Phillip B. Gibbons, and Todd C. Mowry. 2004. Improving Hash Join Performance through Prefetching. In *Proceedings of the 20th International Conference on Data Engineering, ICDE*. 116–127. <https://doi.org/10.1109/ICDE.2004.1319989>
- [10] 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.
- [11] Yongjun He, iacheng Lu, and Tianzheng Wang. 2020. CoroBase: Coroutine-Oriented Main-Memory Database Engine. *Proc. VLDB Endow.* 14, 3 (2020), 431–444. <https://doi.org/10.5555/3430915.3442440>
- [12] Hongjing Huang, Zeke Wang, Jie Zhang, Zhenhao He, Chao Wu, Jun Xiao, and Gustavo Alonso. 2022. Shuhai: A Tool for Benchmarking High Bandwidth Memory on FPGAs. *IEEE Trans. Computers* 71, 5 (2022), 1133–1144. <https://doi.org/10.1109/TC.2021.3075765>
- [13] Intel. 2024. Intel® 64 and IA-32 Architectures Optimization Reference Manual. <https://cdrdv2.intel.com/v1/dl/getContent/671488>. Online; last accessed March 20, 2024.
- [14] Intel. 2024. Intel® In-Memory Analytics Accelerator (Intel® IAA) Architecture Specification. <https://cdrdv2-public.intel.com/721858/350295-iaa-specification-R03.pdf>. Online; last accessed May 05, 2024.
- [15] Saba Jamil, Tanvir Ahmed Khan, Grant Ayers, Baris Kasikci, and Heiner Litz. 2022. APT-GET: profile-guided *timely* software prefetching. In *Seventeenth European Conference on Computer Systems*. ACM, 747–764. <https://doi.org/10.1145/3492321.3519583>
- [16] Christopher Jonathan, Umar Farooq Minhas, James Hunter, Justin J. Levandoski, and Gor V. Nishanov. 2018. Exploiting Coroutines to Attack the “Killer Nanoseconds”. *Proc. VLDB Endow.* 11, 11 (2018), 1702–1714. <https://doi.org/10.14778/3236187.3236216>
- [17] Alfons Kemper and Thomas Neumann. 2011. HyPer: A hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. In *Proceedings of the 27th International Conference on Data Engineering*. IEEE Computer Society, 195–206. <https://doi.org/10.1109/ICDE.2011.5767867>
- [18] Onur Kocberber, Babak Falsafi, and Boris Grot. 2015. Asynchronous Memory Access Chaining. *Proc. VLDB Endow.* 9, 4 (2015), 252–263.
- [19] Tsvika Kurts, Zelig Wayner, and Tommy Bojan. 2005. Apparatus and method for bus signal termination compensation during detected quiet cycle. US Patent 6,842,035.
- [20] Jaekyu Lee, Hyesoon Kim, and Richard W. Vuduc. 2012. When Prefetching Works, When It Doesn’t, and Why. *ACM Trans. Archit. Code Optim.* 9 (2012), 2:1–2:29. <https://doi.org/10.1145/2133382.2133384>
- [21] 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.
- [22] Chi-Keung Luk and Todd C. Mowry. 1996. Compiler-Based Prefetching for Recursive Data Structures. In *ASPLOS-VII Proceedings - Seventh International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM Press, 222–233. <https://doi.org/10.1145/237090.237190>
- [23] Daniel Lustig, Abhishek Bhattacharjee, and Margaret Martonosi. 2013. TLB Improvements for Chip Multiprocessors: Inter-Core Cooperative Prefetchers and Shared Last-Level TLBs. *ACM Trans. Archit. Code Optim.* 10, 1 (2013), 2:1–2:38. <https://doi.org/10.1145/2445572.2445574>
- [24] Yandong Mao, Eddie Kohler, and Robert Tappan Morris. 2012. Cache craftiness for fast multicore key-value storage. In *Seventh European Conference on Computer Systems*. ACM, 183–196. <https://doi.org/10.1145/2168836.2168855>
- [25] Prashanth Menon, Andrew Pavlo, and Todd C. Mowry. 2017. Relaxed Operator Fusion for In-Memory Databases: Making Compilation, Vectorization, and Prefetching Work Together At Last. *Proc. VLDB Endow.* 11, 1 (2017), 1–13.

- <https://doi.org/10.14778/3151113.3151114>
- [26] Daniel Molka, Robert Schöne, Daniel Hackenberg, and Wolfgang E. Nagel. 2017. Detecting Memory-Boundedness with Hardware Performance Counters. In *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering*. ACM, 27–38. <https://doi.org/10.1145/3030207.3030223>
 - [27] Todd C. Mowry, Monica S. Lam, and Anoop Gupta. 1992. Design and Evaluation of a Compiler Algorithm for Prefetching. In *ASPLOS-V Proceedings - Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM Press, 62–73. <https://doi.org/10.1145/143365.143488>
 - [28] Jan Mühlhig and Jens Teubner. 2021. MxTasks: How to Make Efficient Synchronization and Prefetching Easy. In *Proceedings of the 2021 International Conference on Management of Data, SIGMOD*. ACM, 1331–1344. <https://doi.org/10.1145/3448016.3457268>
 - [29] Thomas Neumann and Michael J. Freitag. 2020. Umbra: A Disk-Based System with In-Memory Performance. In *10th Conference on Innovative Data Systems Research, CIDR*. <http://cidrdb.org/cidr2020/papers/p29-neumann-cidr20.pdf>
 - [30] 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. <https://doi.org/10.14778/3149193.3149202>
 - [31] Georgios Psaropoulos, Ismail Oukid, Thomas Legler, Norman May, and Anastasia Ailamaki. 2019. Bridging the Latency Gap between NVM and DRAM for Latency-bound Operations. In *Proceedings of the 15th International Workshop on Data Management on New Hardware*. ACM, 13:1–13:8. <https://doi.org/10.1145/3329785.3329917>
 - [32] Till Schlüter, Lorenz Hetterich, Leon Trampert, Hamed Nemati, Ahmad Ibrahim, Michael Schwarz, Christian Rossow, and Nils Ole Tippenhauer. 2023. FetchBench: Systematic Identification and Characterization of Proprietary Prefetchers. In *Proceedings of the 2023 ACM SIGSAC*. ACM, 975–989. <https://doi.org/10.1145/3576915.3623124>
 - [33] Michael Schwarz, Moritz Lipp, Daniel Moghimi, Jo Van Bulck, Julian Stecklina, Thomas Prescher, and Daniel Gruss. 2019. ZombieLoad: Cross-Privilege-Boundary Data Sampling. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 753–768. <https://doi.org/10.1145/3319535.3354252>
 - [34] Stephan van Schaik, Alyssa Milburn, Sebastian Österlund, Pietro Frigo, Giorgi Maisuradze, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. 2019. RIDL: Rogue In-Flight Data Load. In *IEEE Symposium on Security and Privacy*. IEEE, 88–105. <https://doi.org/10.1109/SP.2019.00087>
 - [35] 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 *48th ACM/IEEE Annual International Symposium on Computer Architecture, ISCA*. IEEE, 85–98. <https://doi.org/10.1109/ISCA52012.2021.00016>
 - [36] Benjamin Wagner, André Kohn, and Thomas Neumann. 2021. Self-Tuning Query Scheduling for Analytical Workloads. In *Proceedings of the 2021 International Conference on Management of Data, SIGMOD*. ACM, 1879–1891. <https://doi.org/10.1145/3448016.3457260>
 - [37] Xiangyu Zhi, Xiao Yan, Bo Tang, Ziyao Yin, Yanchao Zhu, and Minqi Zhou. 2024. CoroGraph: Bridging Cache Efficiency and Work Efficiency for Graph Algorithm Execution. *Proc. VLDB Endow.* 17, 4 (2024). <https://doi.org/10.14778/3636218.3636240>

A PROFILING SOFTWARE PREFETCHING

In the preceding sections, we discussed that understanding the interplay between hardware characteristics and applications is complex and requires careful profiling. *Performance counters* are valuable tools for gauging hardware behavior, given their widespread integration across consumer platforms. However, the extensive range and diversity of available counters across various manufacturers and generations complicate their use. To maintain the readability of our discussion, we have chosen to include a more comprehensive explanation of the performance counters employed in our study in the appendix. For those seeking a more thorough understanding, the following information includes precise technical details, including the unique names of the counters used in current AMD and Intel microarchitectures.

Execution. Both Intel and AMD offer capabilities to monitor the number of dispatched prefetch instructions targeting various cache

levels (e.g., SW_PREFETCH.NTA on Intel and PREFETCH_INSTRUCTIONS_DISPATCHED.PREFETCHNTA on AMD). AMD, however, takes this a step further by enabling a more fine granular analysis of prefetching activities. Specifically, the Zen 4 micro-architecture enables tracking of the origin of prefetches, for example, those served from the local L2 cache (SOFTWARE_PREFETCH_DATA_CACHE_FILLS.LCL_L2) or from a cache within the local core complex (SOFTWARE_PREFETCH_DATA_CACHE_FILLS.LOCAL_CCX). In total, six different sources can be found in counter-package SOFTWARE_PREFETCH_DATA_CACHE_FILLS.*.

Additionally, AMD facilitates shedding light on the efficacy of prefetch instructions by counting ineffective ones—like those requesting a cache line already allocated in the MAB (INEFFECTIVE_SOFTWARE_PREFETCH.MAB_MCH_CNT) and those fetching data already present in the cache (INEFFECTIVE_SOFTWARE_PREFETCH.DATA_PIPE_SW_PF_DC_HIT). Intel, in turn, provides a counter for tracking *loads* to software-prefetched cache lines that have not yet fully transferred from memory (LOAD_HIT_PRE.SW_PF), which indicates that the prefetch distance may be insufficiently calibrated. While a thorough analysis of prefetch distance tuning is beyond this paper’s scope, prefetch-timing has been explored in [15], including a review of prefetch distance-related performance counters on Intel platforms.

LFB/MAB. In Sections 3 and 4, we discussed how stalls related to LFBs can become a bottleneck when a large number of prefetch requests overwhelm the hardware. Unfortunately, options for precisely quantifying these stalls are limited. To our knowledge, there are no counters that exclusively measure stalled cycles while waiting for available LFB slots. However, on a range of Intel micro-architectures, performance counters can quantify the total number of stalled cycles due to various hardware resources, e.g., PMHs, LFB, and memory using the counter CYCLE_ACTIVITY.STALLS_TOTAL. And on Intel’s Cascade Lake architecture, it is possible to measure the number of stalls directly linked to waiting for memory and caches (CYCLE_ACTIVITY.STALLS_MEM_ANY), not including LFB allocations. In Section 3, we combine both named metrics with counters that track the number of requests facing a fully occupied LFB (L1D_PEND_MISS.FB_FULL) to draw conclusions about the LFB-related stalls. At this point, we would also like to refer to Molka et al. [26], which offer an in-depth analysis of performance counters to identify memory bottlenecks on a range of Intel architectures.

On AMD platforms, the options for determining the overhead related to excessive MAB allocations appear to be more limited. Although it is possible to measure how many MAB slots are allocated (ALLOC_MAB_COUNT) and break it down to the hardware prefetcher (MAB_ALLOCATION_BY_TYPE.HW_PF) and to loads and stores (MAB_ALLOCATION_BY_TYPE.LS), we have not been able to find any counters that measure the extent of MAB flooding. Despite this, the total count of CPU stalls can be determined on AMD’s Zen 4 micro-architecture using the counter CYCLES_NO_RETIRE.NOT_COMPLETE_MISSING_LOAD. This metric aggregates the cycles during which the CPU is forced to halt due to pending loads. While AMD’s documentation remains non-specific about the components involved, our experiments indicate that stalls related to memory, TLB, and MAB are included, aligning closely with Intel’s STALLS_TOTAL counter.

TLBs. When prefetching can effectively hide memory latency, handling TLB misses becomes another critical bottleneck. Intel and AMD provide a range of counters to monitor the frequency of hits and misses at both the first- and second-level TLBs. Some of these are already standardized by the Linux *perf* interface (e.g., `PERF_COUNT_HW_CACHE_DTLB_READ_ACCESS`), others are specific to the underlying platform. For example, Intel uses the counter-package `DTLB_LOAD_MISSES.*` to track the frequency of TLB misses while AMD uses `L1_DTLB_MISS.*`. Additionally, on Intel systems, it is possible to measure the number of cycles during which at least one (`DTLB_LOAD_MISSES.WALK_ACTIVE`) or all (`DTLB_LOAD_MISSES.WALK_PENDING`) PMHs are actively traversing the page table. While this provides insights into the PMHs's activity level, thanks to out-of-order execution, the CPU does not always have to wait actively

for the PMH to finish to proceed with execution. Consequently, these cycles are not necessarily stalled cycles. Despite this, we have not found performance counters quantifying stalled cycles directly linked with TLB misses.

Beyond x86 Architectures. In addition to the x86 architectures considered in this paper, ARM processors are also playing an increasingly important role nowadays. Although ARM processors facilitate application analysis via performance counters, we found it challenging to expose the effects of software prefetching. This difficulty stems from the limited collection of performance counters typically implemented on ARM hardware, especially compared to the latest AMD and Intel micro-architectures. For example, to our knowledge, there are no counters available to track events related to software prefetches and the LFB.