

Vector Runahead for Indirect Memory Accesses

Ajeya Naithani , Ghent University, B-9000, Ghent, Belgium

Sam Ainsworth , University of Edinburgh, EH8 9AB, Edinburgh, U.K.

Timothy M. Jones , University of Cambridge, CB3 0FD, Cambridge, U.K.

Lieven Eeckhout , Ghent University, B-9000, Ghent, Belgium

Vector runahead delivers extremely high memory-level parallelism even for the chains of dependent memory accesses with complex intermediate address computation, which conventional runahead techniques fundamentally cannot handle and, therefore, have ignored. It does this by rearchitecting runahead to use speculative data-level parallelism, rather than work skipping, as its primary form of extracting more memory-level parallelism in runahead mode than a true execution can, which we hope will bring about an entirely new dimension for high-performance processors.

Many modern-day workloads are poorly served by current Out-of-Order (OoO) superscalar cores, since they feature sparse, indirect memory accesses³ characterized by high-latency cache misses that are unpredictable by today's stride prefetchers.⁶ Despite large reorder-buffer (ROB) and issue-queue resources, superscalar cores running these applications have run out of steam, spending the majority of their time stalled since they cannot capture the memory-level parallelism (MLP) necessary to hide today's memory access latencies.

Vector runahead (VR) rearchitects runahead execution to use a new method of generating MLP. Rather than work skipping,⁸ VR extracts MLP as a speculative form of *data-level parallelism*: it groups together independent loads from many different iterations of the same code, allowing them to all follow different sequences of *dependent loads independently*. It further improves throughput by running these newly grouped sequences as vector operations: even when the workload itself is not vectorizable, the prefetching effect from the runahead, which need not be perfectly accurate, is likely to still exhibit data-level parallelism.

On a variety of graph, database, and high-performance computing workloads, VR improves performance by 1.79× compared to a baseline OoO processor with a stride prefetcher. Relative to the state-of-the-art indirect memory prefetcher (IMP)¹² and precise runahead execution (PRE)⁹ VR improves performance by 1.49× on an average. The fundamental reason for this significant performance improvement is illustrated in Figure 1: PRE is unable to accurately prefetch the majority of indirect memory accesses, unlike VR.

EXISTING RUNAHEAD TECHNIQUES

While specialized accelerators are one solution, and programmable forms of prefetching another,¹ the ideal solution would be a pure-microarchitectural technique that could achieve the same benefits without the need for recompilation. Hardware prefetchers can pick up a variety of memory-access patterns, but to achieve the instruction-level visibility necessary to calculate the addresses of complex access patterns in today's workloads,¹ one must operate within the core, instead of within the cache. Runahead execution^{8,9} is the most promising technique to achieve this.

The promise of runahead execution is that the core can continue to perform useful work even while stalled on a long-latency cache miss, by calculating addresses and prefetching data for future memory accesses. By speculatively issuing multiple independent memory accesses, runahead execution significantly increases MLP, ultimately improving overall application performance.

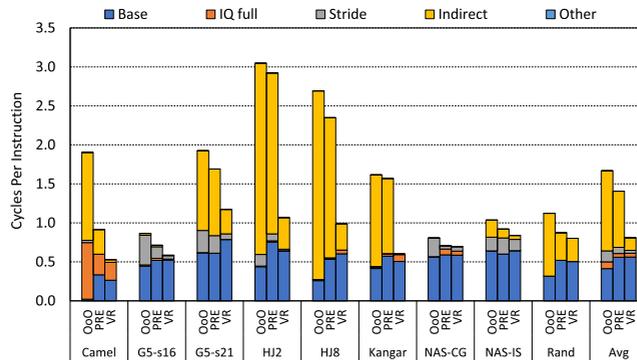


FIGURE 1. CPI stacks for the baseline OoO core, PRE, and VR. The memory component is broken down and attributed to striding loads and indirect dependent-chain loads. The previous state-of-the-art runahead cannot prefetch the majority of indirect memory accesses, unlike VR.

However, conventional runahead comes unstuck by the very mechanism it uses to generate MLP. First, by skipping over loads for which the data source is not yet ready, it is unsuitable for today's complex indirection patterns that consist of chains of dependent load misses. Second, conventional runahead is limited by both the processor's front-end (fetch/decode/rename) width and available back-end resources (issue queue slots and physical registers).⁹ What is needed is a technique that can overcome the limitations of a processor's resources to generate massive amounts of MLP and follow chains of dependent loads to completion, prefetching all data required for many memory accesses in the future. VR is that technique.

VECTOR RUNAHEAD

The key insight behind VR is that many indirect memory accesses occur within loops where each iteration follows approximately the same control-flow path, and that this regularity can be exploited through the parallel execution of multiple iterations simultaneously. The speculative vector execution of multiple future loop iterations is possible and safe, even when the original workload is not vectorizable, since the results will be discarded once VR is terminated and normal execution resumes.

VR addresses the limitations described previously in three ways, as illustrated in Figure 2. First, it deliberately waits for the results of currently unavailable loads, rather than invalidating and skipping them, which enables VR to prefetch entire load chains but causes the technique to quickly exhaust its back-end OoO resources and, thus, stall on waiting for these intermediate results. Second, to fix this, VR vectorizes the runahead instruction stream by reinterpreting scalar instructions as vector operations to generate many

different cache misses at different offsets. This means that despite executing many future iterations of a loop at once, VR only requires the processor resources (both front-end and back-end instruction slots) of a single iteration. In effect, this virtually increases the effective fetch/decode bandwidth during runahead mode by issuing independent operations both in quick succession and merged together into single instructions. Third, it issues multiple rounds of these vectorized instructions through our schemes of vector unrolling and pipelining to speculate even deeper and increase the effective runahead memory bandwidth even further. This has the effect of installing huge numbers of independent loads next to each other in the issue queue and ROB, avoiding the need for OoO structures of unbounded size. Altogether, this means that, while VR must wait for the dependent loads rather than skipping them, it waits on a huge number of them at once, finally allowing the achievement of extreme MLP even on complex workloads.

MICROARCHITECTURE DETAILS

We now describe VR's required changes to the processor pipeline, as illustrated in Figure 3.

Initiating VR

The core enters *runahead mode* when either of the following two conditions is satisfied after a load instruction blocks the head of the ROB: 1) the ROB is filled with instructions or 2) the issue queue is filled to 80% of its full capacity. VR checkpoints the PC and the front-end register allocation table (RAT). This marks the *entry* to runahead mode. After entering runahead mode, the processor continues to fetch, decode, and execute future

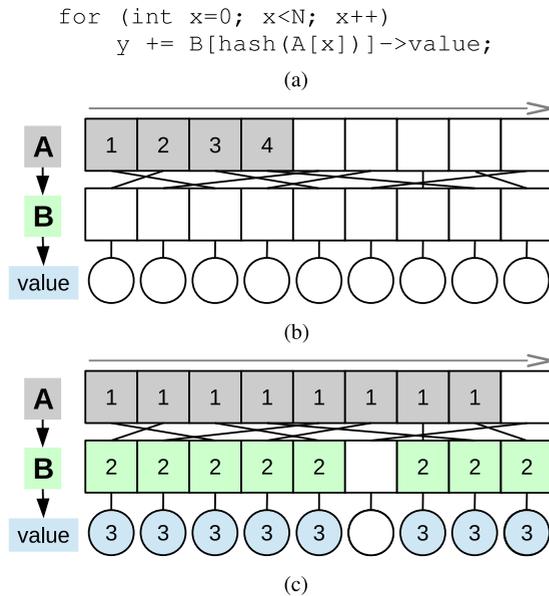


FIGURE 2. VR versus PRE⁹ on an illustrative code example. The loads highlighted in green can only be triggered by stalling on loads highlighted in gray, and those in blue by stalling on gray and green. VR prefetches multiple memory accesses in parallel along the memory dependence chain during runahead mode. (a) Example code, with memory access by array indirection, with intermediate address computation and pointer access. (b) Precise runahead execution (PRE)⁹ is able to prefetch array elements from A. In contrast, the array elements to B cannot be prefetched during runahead mode as they depend on A. Likewise, the data values cannot be prefetched either because they depend on B. Note that the elements in A are accessed serially as indicated. PRE runahead mode is terminated before it can prefetch array elements of B; furthermore, the number of backend resources needed during runahead mode limits the speculation depth. (c) Vector runahead vectorizes memory accesses along the memory dependence chain whilst in runahead mode. Multiple accesses to A happen in parallel, followed by parallel accesses to B, followed by parallel data-value reads. Vector runahead changes runahead mode's termination condition, i.e., instead of returning to normal mode once the blocking load miss returns from main memory, vector runahead continues runahead mode until all loads along the dependent load chain have been issued. This delayed termination condition delivers higher performance by extracting more MLP than an immediate return to normal mode.

instructions. We use a stride detector⁶ to find regular access patterns in the code that can be used as “induction variables” to produce speculative vectorized copies

of code. The detector also keeps track of the last dependent load (known as the “terminator”) on the striding load. Entry to VR mode begins when we decode a striding load. We vectorize the striding load, followed by the sequence of instructions depending on it. We call the dependent instructions between two dynamic instances of a striding load an *indirect chain*.

Detecting Indirect Chains

We use a *taint vector* (TV) to detect the indirect chains depending on a striding load. The TV features an entry for each architectural integer register, and stores two flags: 1) if the previous instruction to write to this register was a vectorized operation (*vectorize bit*), and 2) if the previous instruction to write to this register was invalid (*invalid bit*). The TV is empty at the start of runahead, as it is cleared whenever runahead terminates. Vectorize bits are initially set for the destination architectural register of a discovered striding load. Invalid bits are initially set based on the destinations of unsupported operations, e.g., those that take floating-point operations as input (which are always invalid and so need no TV entry). Both bits are propagated using *vector taint tracking*, a mechanism to propagate vectorization where needed. Instructions with no bits set are issued as conventional scalar runahead operations, and treated as loop invariant with respect to vectorized copies of the instruction sequence in the current VR mode iteration. Instructions with the invalid bit set are discarded, and instructions with only the vectorize bit set are vectorized.

Vectorizing Instructions

A microprogrammed routine vectorizes the indirect chain. For striding loads, the vectorizer generates their vectorized versions by taking the current memory address accessed by the striding load and its stride as inputs. The vectorizer generates one 512-bit vector load instruction and *injects* the vector instruction into the pipeline. Regardless of input bit width, eight scalar operands are fit in this 512-bit vector, such that we can operate on any data size up to 64 bits. We assume that each vector instruction uses 512-bit vector registers (similar to Intel's AVX-512) for its source and destination, and we reuse the microarchitecture's physical vector registers, and the micro-ops implemented by the microarchitecture's vector units. Similarly, we vectorize all arithmetic and load instructions (directly or indirectly) depending on a striding load, and generate their corresponding 512-bit vector versions.

The renamed instructions are dispatched to the processor back-end where they are executed

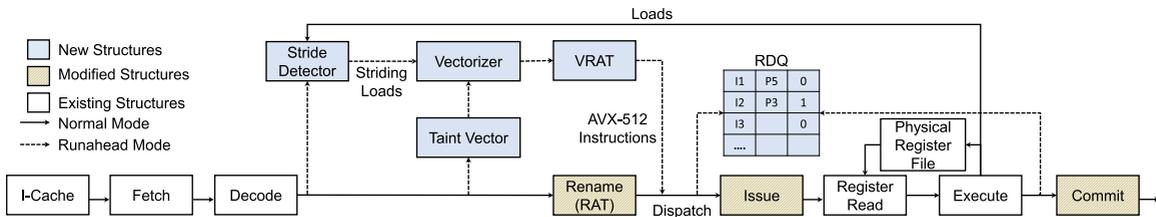


FIGURE 3. Processor pipeline for VR execution.

speculatively. The instructions executed in runahead mode are useful only in generating memory accesses and their state is not maintained in the ROB. Therefore, no ROB entries are allocated in runahead mode. Instead, we use a simpler register deallocation queue (RDQ)⁹ to handle register availability.

Vector Unrolling and Pipelining

To cover more iterations of the indirect chain, we can alternatively generate more than one vector instruction for each scalar instruction in the chain. Depending on the amount of back-end resources available, the generated vector instructions can be dispatched to the processor back-end in two ways. First, through *vector unrolling* [see Figure 4(b)], we can dispatch vector instructions in multiple rounds. For example, we could dispatch $U \times 8$ copies of a loop by issuing the first eight in a single vectorized copy of the instruction stream in round 1, then repeating the process $U - 1$ times, where U is the *unroll depth*. Second, through *vector pipelining* [see Figure 4(c)], we can dispatch all vector instructions for each scalar instruction before dispatching P , the *pipeline width*, vector instructions for the next instruction in the indirect chain. When the amount of back-end resources is limited, vector unrolling is the preferred technique as the processor back-end does not stall due to the lack of available resources to process vector instructions. Vector pipelining, on the other hand, delivers better performance when the back-end has sufficient resources to simultaneously process a large number of vector instructions. A processor microarchitecture can be tuned to dynamically select one of the two techniques for higher performance depending on the availability of back-end resources.

Since we can generate multiple vector instructions for each scalar instruction of the indirect chain, each scalar architectural register first needs to be mapped to multiple vector architectural registers, followed by mapping each vector architectural register to a vector physical register. The complete process of renaming from a scalar architectural register to a vector physical register is accomplished with the help of

the *vector register allocation table (VRAT)*, which maintains P , the *vector pipelining width*, entries per architectural integer register, recording the P destination physical vector registers assigned to the P pipelined copies of the instruction. When we look up these P registers in the VRAT, each of the P copies of the new vectorized instruction uses one of the P entries as its own input. This enables us to distinguish the inputs and outputs of separate pipelined iterations within the vector pipelining arrangement, which, from an instruction fetch point of view, all alias to the same instruction.

THE KEY INSIGHT BEHIND VECTOR RUNAHEAD IS THAT MANY INDIRECT MEMORY ACCESSES OCCUR WITHIN LOOPS WHERE EACH ITERATION FOLLOWS APPROXIMATELY THE SAME CONTROL-FLOW PATH, AND THAT THIS REGULARITY CAN BE EXPLOITED THROUGH PARALLEL EXECUTION OF MULTIPLE ITERATIONS SIMULTANEOUSLY.

Control Flow

All vector lanes follow the same pattern of control flow, apart from when there is a divergence between the lanes in VR mode when they meet a branch instruction. A micro-op converts scalar branches into a predicate mask for the eight vector lanes. Since VR need not cover all code, we use only the results of the first lane to determine the direction of the branch, and mask off any lanes that would have taken a different control-flow path.

Terminating Runahead

VR mode terminates when any of the following four conditions is satisfied: 1) we encounter a dynamic instance of the initial striding load again; 2) we encounter, and issue, the *terminator*: the PC identified by the stride detector as

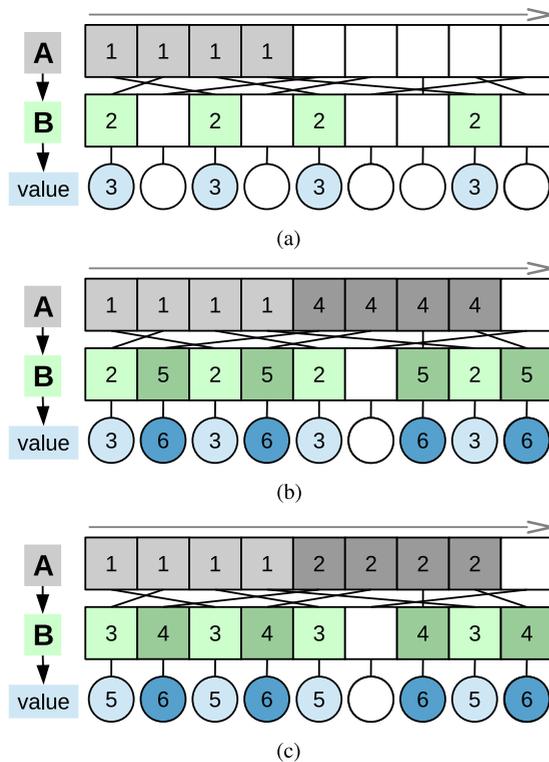


FIGURE 4. VR uses two techniques, vector unrolling, and vector pipelining, to improve the performance by increasing the degree of runahead to allow wider vectors than supported natively by the instruction-set architecture. (a) Basic vector runahead. In this example, MLP is limited to a single vector instruction, so only four outstanding memory accesses can be prefetched at once, and few future memory accesses are covered by the memory-parallel vector runahead, limiting performance gains for future normal execution. (b) Vector unrolling. While the vector runahead operations are still run in sequence, with a maximum MLP of 4, we cover significantly more of the future memory accesses before returning to normal execution, improving the latter's observed performance gain. (c) Vector pipelining. We overlap the independent operations from multiple unrolled iterations. This allows many misses to be handled simultaneously: in this example, 1 and 2 can be executed in parallel, doubling MLP to 8, as can 3 and 4, and 5 and 6.

the last dependent load in the sequence; 3) all vector lanes have been marked as invalid; or 4) we time out (after 200 scalar-equivalent instructions have been executed in VR mode), in the case of traveling down an unexpected code path. When we dispatch multiple rounds of vectorized instructions in vector unrolling, we re-enter VR mode immediately, with the next striding load issuing vector

gathers again. This is repeated until we have issued all the rounds and only then is normal execution resumed. The benefit of vectorizing the entire indirect chain far exceeds the additional duration the core is in runahead mode, as VR yields higher MLP than typical OoO execution.

Upon termination, we restore the front-end RAT to the point of entry into runahead mode, and the TV, VRAT, and RDQ are cleared. The front-end is redirected to fetch from the next instruction after the last dispatched instruction in the ROB.

Hardware Overhead

VR requires only modest changes to the processor pipeline, including the stride detector, TV, and VRAT. The RDQ is already used by PRE.⁹ When put together, the total hardware overhead of VR relative to a baseline OoO core is limited to 1.3 KB, versus 1.24 KB for PRE.

EVALUATION

We compare the following microarchitectural mechanisms, all implemented in Sniper.⁵

- ▶ *Out-of-order*: Baseline OoO core based on Intel's Skylake, with hardware stride prefetcher.
- ▶ *Precise runahead execution*: The state-of-the-art runahead execution technique, as proposed by Naithani *et al.*⁹ We assume an ideal stalling-slice table; therefore, there are no misses in the table.
- ▶ *Indirect memory prefetcher*: The IMP, as proposed by Yu *et al.*¹² IMP is attached to the L1 D-cache, and detects indirect access patterns starting from striding memory accesses.
- ▶ *Vector runahead*: The VR mechanism proposed in this article, assuming an *unroll length* U of 8 and *pipeline depth* P of 8.

We consider a variety of benchmarks featuring complex memory and compute dependencies in their execution stream. These benchmarks are memory-latency bound on today's systems, and are based on high-performance computing (HPC), graph and database workloads evaluated in previous work on programmer- and compiler-managed prefetching mechanisms.^{1,2}

The benchmarks represent a variety of different complex memory-access patterns, with differing indirect chains and compute requirements. We use compiler flag `-ftree-vectorize` (via `-O3`) in all comparisons, but we find that autovectorization does not alter performance because the code is not vectorizable (despite being amenable to VR). We refer to the ISCA 2021 conference paper for details regarding the experimental setup and various sensitivity analyses.

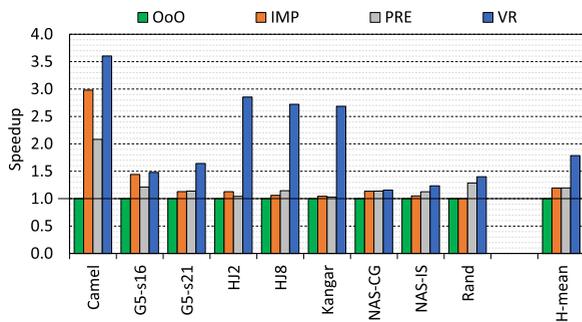


FIGURE 5. Performance of VR execution on a baseline Intel Skylake-style OoO core implemented in Sniper.⁵ VR yields a 1.79 \times and 1.49 \times harmonic mean speedup compared to the baseline OoO core and PRE (and IMP), respectively.

Figure 5 reports speedup for all the evaluated techniques. VR achieves a 1.79 \times harmonic mean speedup across the benchmarks compared to our baseline OoO architecture. The achieved speedup is as high as 3.6 \times (Camel), 2.9 \times (HJ2), 2.7 \times (HJ8), and 2.7 \times (Kangaroo). PRE on the other hand achieves a harmonic mean speedup of 1.20 \times compared to the baseline—in other words, VR achieves a speedup of 1.49 \times relative to PRE. IMP cannot detect complex address-computation patterns and improves speedup by only 1.19 \times relative to the baseline. In short, the significant improvement in performance achieved by VR results from much higher MLP, while fetching in all loads within dependent sequences, and without fetching irrelevant data.

VR achieves higher performance by three main mechanisms. The most important is the software-pipelining effect that reordering of load instructions provides, in that a large number of misses can be serviced simultaneously. This same reordering when implemented with 64 scalar micro-ops instead of eight vector micro-ops is sufficient to gain an average 1.47 \times speedup. The optimization of packing these into fewer vector operations, due to their now-single-instruction–multiple-data (SIMD) layout, increases performance to 1.69 \times by virtue of increasing the effective processor front-end width, and requiring fewer issue-queue slots so that loads can issue earlier. Finally, altering the termination condition, such that VR completes the entire chain of memory accesses before exiting, allows it to cover longer chains of multiple main-memory accesses rather than just the ones it can achieve before the load instruction at the head of the ROB returns, increasing performance to the full 1.79 \times shown in the graph.

Figure 6 shows why VR is able to achieve higher performance. Its pipelined vectors are able to issue many gathers to memory at once, thus hiding the serialization of dependent loads observed by the OoO core and PRE. This also shows us why some workloads

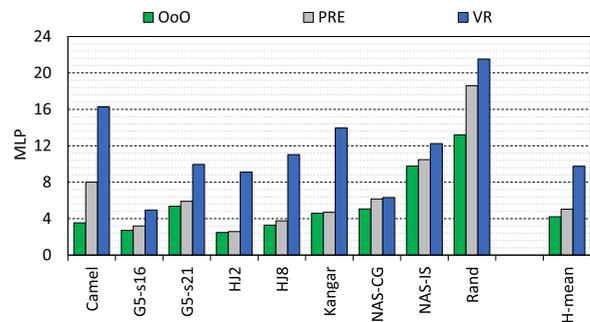


FIGURE 6. MLP measured in terms of MSHR entries utilized per cycle if at least one is allocated. While PRE improves MLP by 1.2 \times , vectorizing indirect chains generates 2.3 \times more MLP than an OoO core.

are sped up more than others. Although our baseline OoO core features a relatively big (224-entry) ROB, which enables it to achieve high MLP on the simplest workloads, we note that VR can extract significantly more MLP. Perhaps unsurprisingly, VR achieves the largest speedups when the OoO core is comparatively weakest: for Camel, HJ2, HJ8, and Kangaroo, there are many instructions (address-computing or otherwise) executing along with the loads, which starve the OoO core of ROB and issue-queue resources,² limiting its memory-reordering ability. By contrast, VR does not rely on the ROB for high MLP, as it can achieve the same effect through its vector gathers.

Some workloads, such as G5-s16 and G5-s21, start from a low baseline and stay relatively low even with VR: complex control flow limits the ability of VR to cover enough of the application’s memory accesses, in effect throttling the vector gathers issued, particularly for the smaller s16 input, which frequently moves between variable-length data-dependent inner and outer loops. Others, such as CG and G5-s16, have small datasets that often hit in the LLC, meaning their L1 data cache misses are serviced quickly with or without VR. Finally, even though the performance of many workloads in VR mode is limited by the number of miss status holding registers (MSHRs), the average MLP is still typically lower than the number of MSHRs available (24 MSHRs at the L1 data cache in our setup): this is because Vector Runahead does not run continuously, and only kicks in when the out-of-order system runs out of resources.

POTENTIAL FOR LONG-TERM IMPACT

VR promises a transformational performance improvement for some of today’s most important and challenging workloads, all in microarchitecture. At a time when other

methods for improving single-thread performance are few and far between, we hope that this work will inspire industry. While the performance improvements are significant, the extra hardware is modest. This reinvention of runahead execution, to be based on speculative data-level SIMD parallelism rather than work-skipping as its primary method for hiding memory latency, could be a fundamental building block for many new techniques both inside and outside the core.

Tomorrow's processors will be able to natively support extreme MLP, even down complex chains. The recent scaling up of other parts of the microarchitecture, such as highly parallel page-table walkers, means that processors will be able to exploit these benefits to the fullest. In turn, we expect processors to adapt their configurations to accommodate forms of extreme MLP as a result: by finally making sparse workloads bandwidth-bound instead of latency-bound, we expect that conventional processors will move to higher latency, higher bandwidth memory.

VR PROMISES A TRANSFORMATIONAL PERFORMANCE IMPROVEMENT FOR SOME OF TODAY'S MOST IMPORTANT AND CHALLENGING WORKLOADS, ALL IN MICROARCHITECTURE.

VR is a qualitative departure from prior solutions. In particular, in contrast to software auto-vectorization, VR does not require the code to be vectorizable to adequately prefetch data into the cache. In contrast to prior runahead techniques, VR presents a solution for achieving MLP down complex dependent memory chains. In contrast to prior pre-execution and helper-thread techniques, VR needs no separate thread, no separate execution units, and neither programmer nor compiler support. Moreover, VR can follow dependent chains, unlike pre-execution and helper threads. In contrast to software prefetching, VR is a pure microarchitecture solution, requiring no changes to the binary or source code, while being able to freely vectorize sequences of instructions that would cause software prefetchers to fault. In contrast to hardware prefetching, VR operates within-core, allowing it to cover arbitrary memory-indirection depths with complex address calculation, as needed in many workloads.⁴ In fact, as we have explored and demonstrated in our ISCA 2021 paper, VR provides significant performance improvements for modern-day workloads with complex indirect memory-access patterns from a wide variety of

application domains including graph analytics, database, and high-performance computing.

Note further that, while VR fundamentally exposes more MLP than OoO execution, it is not fundamentally reliant upon OoO execution. At a time when both OoO execution⁷ and advanced prefetchers¹⁰ have both been exposed for their inadequacies around security, VR proposes a solution for the indirect memory accesses these countermeasures restrict¹¹ that is reliant on neither OoO execution nor out-of-core prefetching. It can preserve secure control flow by being an in-core technique and even despite being speculative itself. We believe that this could finally make such countermeasures,¹¹ and even in-order cores, palatable without severe penalty.

CONCLUSION

VR delivers on what runahead techniques were always designed for, but could never really provide: true latency tolerance for central processing units without OoO resources needing to scale to unbounded dimensions, even for emerging workloads with long and complex chains of dependent memory accesses. We believe that VR provides an opportunity for transformative improvements in single-thread performance, favoring processor designs optimized for MLP rather than being hampered by latency.

REFERENCES

1. S. Ainsworth and T. M. Jones, "An event-triggered programmable prefetcher for irregular workloads," in *Proc. 23rd Int. Conf. Archit. Support Program. Lang. Operating Syst.*, 2018, pp. 578–592.
2. S. Ainsworth and T. M. Jones, "Software prefetching for indirect memory accesses: A microarchitectural perspective," *ACM Trans. Comput. Syst.*, vol. 36, no. 3, pp. 1–34, 2019.
3. K. Asanovic *et al.*, "The landscape of parallel computing research: A view from Berkeley," 2006.
4. G. Ayers, H. Litz, C. Kozyrakis, and P. Ranganathan, "Classifying memory access patterns for prefetching," in *Proc. 25th Int. Conf. Archit. Support Program. Lang. Operating Syst.*, 2020, pp. 513–526.
5. T. E. Carlson, W. Heirman, S. Eyerhan, I. Hur, and L. Eeckhout, "An evaluation of high-level mechanistic core models," *ACM Trans. Archit. Code Optim.*, vol. 11, no. 3, pp. 1–25, 2014.
6. T.-F. Chen and J.-L. Baer, "Reducing memory latency via non-blocking and prefetching caches," in *Proc. 5th Int. Conf. Archit. Support Program. Lang. Operating Syst.*, 1992, pp. 51–61.

7. P. Kocher *et al.*, "Spectre attacks: Exploiting speculative execution," in *Proc. IEEE Symp. Secur. Privacy*, 2019, pp. 1–19.
8. O. Mutlu, J. Stark, C. Wilkerson, and Y. N. Patt, "Runahead execution: An alternative to very large instruction windows for out-of-order processors," in *Proc. 9th Int. Symp. High-Perform. Comput. Archit.*, 2003, pp. 129–140.
9. A. Naithani, J. Feliu, A. Adileh, and L. Eeckhout, "Precise runahead execution," in *Proc. Int. Symp. High-Perform. Comput. Archit.*, 2020, pp. 397–410.
10. J. R. S. Vicarte *et al.*, "Opening pandora's box: A systematic study of new ways microarchitecture can leak private data," in *Proc. ACM/IEEE 48th Annu. Int. Symp. Comput. Archit.*, 2021, pp. 347–360.
11. J. Yu, M. Yan, A. Khyzha, A. Morrison, J. Torrellas, and C. W. Fletcher, "Speculative taint tracking (STT): A comprehensive protection for speculatively accessed data," in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2019, pp. 954–968.
12. X. Yu, C. J. Hughes, N. Satish, and S. Devadas, "IMP: Indirect memory prefetcher," in *Proc. 48th Symp. Microarchit.*, 2015, pp. 178–190.

AJEYA NAITHANI is a postdoctoral researcher with Ghent University, B-9000, Ghent, Belgium. His research interests are in the area of computer architecture with an emphasis on designing novel techniques to improve performance, energy-efficiency, and reliability of modern processors. Naithani received a Ph.D. degree in computer science engineering from Ghent University. He is a Member of IEEE. Contact him at ajeya.naithani@ugent.be.

SAM AINSWORTH is a lecturer in systems and hardware security with the University of Edinburgh, EH8 9AB, Edinburgh, U.K. His research interests include runtime, systems, and hardware security, along with architectural and compiler techniques for data prefetching in software and hardware, and efficient techniques for hardware error detection and correction. He received a Ph.D. degree in computer science from the University of Cambridge, Cambridge, U.K. Contact him at sam.ainsworth@ed.ac.uk.

TIMOTHY M. JONES is a reader in computer architecture and compilation with the University of Cambridge, CB3 0FD, Cambridge, U.K. His research interests span compiler and micro-architectural schemes for performance, reliability and security, especially focused on tackling challenges using different forms of parallelism. Jones received a Ph.D. degree in informatics from the University of Edinburgh, Edinburgh, U.K. Contact him at timothy.jones@cl.cam.ac.uk.

LIEVEN EECKHOUT is a full professor with Ghent University, B-9000, Ghent, Belgium. His research interests include computer architecture performance analysis and modeling, and CPU/GPU microarchitecture and resource management. Eeckhout received a Ph.D. degree in computer science engineering from Ghent University. He is a Fellow of IEEE and ACM. Contact him at lieven.eeckhout@ugent.be.