

AXI-Pack: Near-Memory Bus Packing for Bandwidth-Efficient Irregular Workloads

Chi Zhang^{†*}, Paul Scheffler^{†*}, Thomas Benz^{*}, Matteo Perotti^{*}, Luca Benini^{*†}

^{*} Integrated Systems Laboratory, ETH Zurich, Switzerland

[†] Department of Electrical, Electronic, and Information Engineering, University of Bologna, Italy
{chizhang,paulsc,tbenz,mperotti,lbenini}@iis.ee.ethz.ch

Abstract—Data-intensive applications involving irregular memory streams are inefficiently handled by modern processors and memory systems highly optimized for regular, contiguous data. Recent work tackles these inefficiencies in hardware through core-side stream extensions or memory-side prefetchers and accelerators, but fails to provide end-to-end solutions which also achieve high efficiency in on-chip interconnects. We propose AXI-Pack, an extension to ARM’s AXI4 protocol introducing bandwidth-efficient strided and indirect bursts to enable end-to-end irregular streams. AXI-Pack adds irregular stream semantics to memory requests and avoids inefficient narrow-bus transfers by packing multiple narrow data elements onto a wide bus. It retains full compatibility with AXI4 and does not require modifications to non-burst-resaping interconnect IPs. To demonstrate our approach end-to-end, we extend an open-source RISC-V vector processor to leverage AXI-Pack at its memory interface for strided and indexed accesses. On the memory side, we design a banked memory controller efficiently handling AXI-Pack requests. On a system with a 256-bit-wide interconnect running FP32 workloads, AXI-Pack achieves near-ideal peak on-chip bus utilizations of 87% and 39%, speedups of 5.4x and 2.4x, and energy efficiency improvements of 5.3x and 2.1x over a baseline using an AXI4 bus on strided and indirect benchmarks, respectively.

Index Terms—Computer architecture, On-chip interconnects, Memory systems, Irregular workloads

I. INTRODUCTION

Growing performance demands and large, sparse datasets in domains like machine learning [1], graph analytics [2], fluid dynamics [3], and recommender systems [4] push data-driven applications toward increasingly irregular data access patterns. This poses a challenge to general-purpose CPUs [5], [6] and GPUs [4], [7] optimized for highly regular compute. To keep their functional units highly utilized and achieve satisfactory performance and energy efficiency, single-instruction, multiple-data (SIMD) architectures require contiguous data chunks not naturally found in irregular workloads. Memory hierarchies are also tuned to contiguous, high-locality data and struggle with irregular access patterns [5], [6], resulting in long latencies, poor bandwidth utilization, and cache thrashing.

Existing research aims to improve irregular workload performance and tackle these shortcomings through *core-side* or *memory-side* hardware extensions. Core-side extensions often use *stream* abstractions [5], [6], [8]–[10] to describe entire sequences of irregular accesses, freeing processors from address calculation and decoupling memory accesses from exe-

cution. Most works focus on accelerating *strided* and *indirect* streams, which are most common in practice [6]. Mapping these streams to architectural registers [6], [8]–[10] further improves functional unit utilization and enables significant speedups. However, these works largely ignore downstream interconnects and memory systems. While some authors propose high-level cache policies to avoid thrashing [5], [6], they do not address fundamental limitations like the high index-fetching overhead of core-side indirection and the inherent inefficiency of narrow bus accesses in address-based interconnects.

In contrast, memory-side extensions prefetch and accelerate irregular accesses using pattern-aware memory controllers [11], prefetchers [12], [13], and data layout transform (DLT) accelerators [14], [15]. Unlike core-side extensions, these solutions reduce access times and prevent narrow bus accesses by *packing* fetched irregular elements into bus-wide lines, which are then mapped to virtual addresses [11], written to internal scratchpads [12], [13], or written back to memory [14], [15]. However, these solutions have their own drawbacks: they occupy virtual or physical memory and lack the tight architectural integration of core-side extensions, limiting their acceleration potential and complicating programming.

Thus, while existing core- and memory-side extensions tackle inefficiencies in their respective domains, they forego each other’s benefits and do not integrate with established on-chip interconnect protocols, failing to provide an end-to-end solution for bandwidth-efficient irregular streams.

To address these shortcomings, we propose AXI-PACK, an extension to Arm’s widespread Advanced eXtensible Interface 4 (AXI4) on-chip protocol enabling end-to-end, tightly-packed strided and indirect memory streams. AXI-PACK transparently extends AXI4’s existing contiguous bursts, leveraging their decoupled, latency-resilient nature. It remains compatible with all existing AXI4 features and even existing interconnect blocks that do not reshape bursts. It encodes stream semantics (stride or index base and size) directly into burst requests, ensuring performance and flexibility even for short streams. Indirection is efficiently handled at memory endpoints. In principle, AXI-PACK supports non-core requestors (e.g., accelerators) and systems with multiple requestors and endpoints.

To demonstrate AXI-PACK in an end-to-end full-system context, we extend an open-source RISC-V vector processor for efficient strided and indexed accesses and design a banked

[†] Both authors contributed equally to this research.




Figure 1. AXI-PACK AR/AW user extensions and strided read example

memory controller efficiently handling irregular bursts. On a system with a 256-bit-wide AXI-PACK bus running various irregular FP32 workloads, we achieve bus utilizations of up to 87% on strided and 39% on indirect benchmarks, resulting in peak speedups of $5.4\times$ and $2.4\times$ over a baseline with a standard AXI4 bus. We implement our evaluation system in GlobalFoundries' 22nm FD-SOI technology and find that AXI-PACK improves energy efficiency by up to $5.3\times$ and $2.1\times$ in strided and indirect benchmarks while incurring only 6.2% of our vector processor's area for our controller. Finally, we analyze the impact of element and index size as well as bank count on AXI-PACK performance and controller complexity.

To summarize, our contributions are as follows:

- 1) We extend the widespread high-performance on-chip protocol AXI4 to support end-to-end bus-packed strided and indirect streams with full backward compatibility.
- 2) We extend an open-source RISC-V vector processor with AXI-PACK to enable high bus efficiencies and significant speedups on irregular workloads, and demonstrate a simple banked memory controller to serve irregular bursts.
- 3) We evaluate AXI-PACK by benchmarking irregular workloads on our extended vector processor, achieving bus utilizations of up to 87% and 39% and speedups of up to $5.4\times$ and $2.4\times$ for strided and indirect workloads.
- 4) We evaluate our AXI-PACK system and controller in terms of timing, area, and energy efficiency benefits, finding energy efficiency improvements of up to $5.3\times$ and $2.1\times$.

II. ARCHITECTURE

A. AXI-Pack Protocol

AXI-PACK extends Arm's AXI4 [16], a widely-adopted high-performance non-coherent on-chip memory protocol. AXI4 defines five independent channels: AR and AW carry read and write requests, R and W carry read and write data, and B carries the write response. Without extensions, linear, fixed, and wrapping bursts are supported. Each channel provisions a user field of parametric width that allows extending functionality without compromising compatibility with the baseline protocol.

AXI-PACK extends the request channels AR and AW with user signals to support packed irregular bursts as illustrated in Figure 1. The `pack` bit indicates whether our extension is used, while the `indir` bit differentiates between strided and indirect bursts. The remaining bits are shared between both burst types; they indicate either the element stride for strided bursts or the index size and base offset for indirect bursts.

While active, the new irregular burst types alter the semantics of existing channel fields. Most notably, data elements of the

requested size, scattered in memory, are tightly packed onto the R and W data buses to fully utilize them. Additionally, the start of irregular bursts is aligned with the bus instead of the address to simplify feeding data to and from vectorized functional units. Finally, the AR and AW `size` fields, usually only changed for narrow beats, indicate the data element size.

In addition to performance and bus utilization, these semantics aim to maximize the *transparency* and *portability* of AXI-PACK: they ensure that any existing AXI4 intellectual property (IP) blocks handling non-modifiable transactions without splitting, such as the routing blocks provided in [17], are already compatible with AXI-PACK without any modifications. IPs that require burst splitting or reshaping, such as bus width converters, can easily be extended to support AXI-PACK by repacking bus-aligned data elements as for existing burst types.

B. Vector Processor Extension

To demonstrate the benefits and flexibility of AXI-PACK, we extend the open-source *Ara* [18] RISC-V vector processor to leverage it for efficient irregular memory accesses. Ara acts as a co-processor to the CVA6 core [19], which dispatches vector instructions to Ara. Both access memory over AXI4.

As mandated by its instruction set, Ara supports three vector memory access types: *contiguous*, *strided*, and *indexed*. Without extensions, only contiguous accesses can leverage bursts. For strided and indexed accesses, Ara must compute the address and issue individual narrow accesses for each element, leaving the data channels severely underutilized as shown in Figure 1.

Figure 2a shows our extensions to Ara. We modify its vector load-store unit (VLSU) to use AXI-PACK for strided and indexed vector accesses. For strided accesses, we simply translate the existing `vlse` and `vsse` instructions to AXI-PACK requests and exchange the read or written data directly with vector registers or functional lanes for chaining. The existing indexed access instructions `vl(o|u)xei` and `vs(o|u)xei` in the RISC-V vector extension presume that indices are already loaded into vector registers, necessitating the move of indices into the core and precluding efficient memory-side indirection. To remedy this, we extend Ara's decoder and introduce two new *in-memory* indexed access instructions, `vlimxei` and `vsimxei`, which use index arrays in memory for indirection. These are directly translated into indirect AXI-PACK bursts and allow for packed data to be exchanged with Ara's unmodified functional units and registers without data format changes.

C. Banked Memory Controller

To demonstrate the efficient handling of AXI-PACK at banked memory endpoints, we design a proof-of-concept controller translating AXI-PACK requests to sequences of parallel banked memory accesses. Our controller is fully backward-compatible with and efficiently handles regular AXI4 bursts.

Figure 2b shows the controller architecture. The *adapter* translates both regular and irregular bursts to sequences of n parallel *word* accesses, where a word is the same width W as the used memory banks and determines the smallest efficiently-handled element size. For D -bit-wide AXI-PACK data buses, $n = D/W$, since we must read or write D/W words in parallel

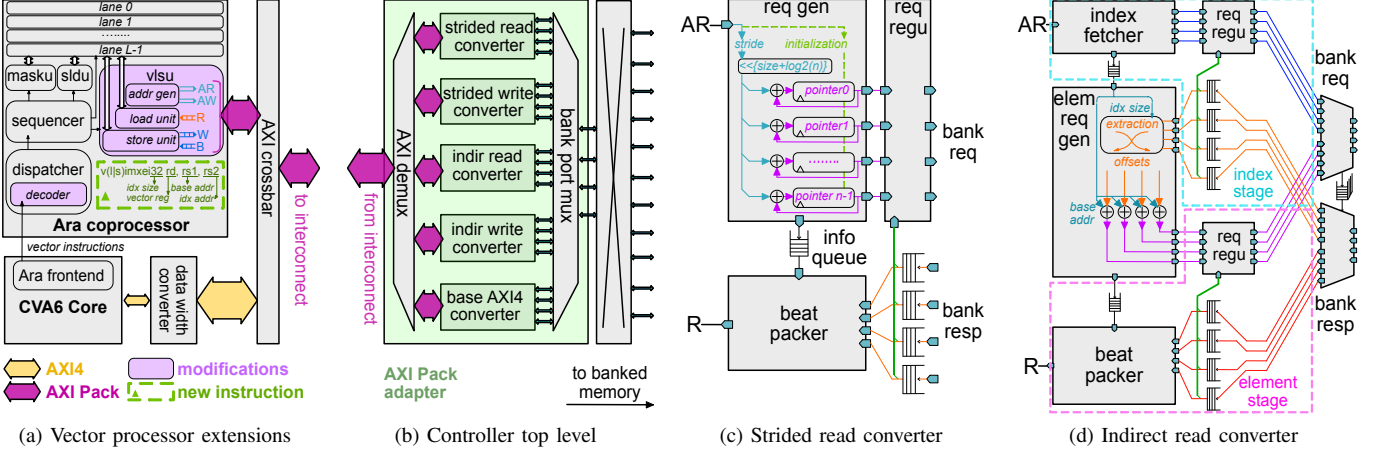


Figure 2. AXI-PACK processor extension, multi-banked controller, and converter architecture

to saturate them. The adapter connects to an $n \times m$ crossbar mapping the n word access ports to m interleaved banks.

Internally, the adapter forwards requests to one of five converters which may concurrently handle bursts. The *base* converter handles regular AXI4 bursts, while the remaining converters are dedicated to strided and indirect read and write operations, respectively. Handling reads and writes individually leverages the inherent concurrency of the R and W channels.

Figure 2c details the strided read converter architecture. For each beat in a burst, the *request generator* issues n parallel word requests fetching the elements to be packed and pushes metadata needed for later packing into an *info* queue. The words read from the banks are stored in decoupling queues and then passed to the *beat packer*, which packs the words as specified by metadata popped from the *info* queue to form the R beats. To prevent word queue overflows, a *request regulator* limits the number of requests in-flight for each word lane.

Figure 2d shows the indirect read converter architecture. It involves two stages sharing the n word request ports through round-robin arbitration: the *index stage* fetches indices from memory and the *element stage* uses these indices to fetch indirect elements and pack them into R beats. The index stage is analogous to the strided read converter, but issues only contiguous word requests. The fetched indices are passed to the *element request generator*, which shifts and adds them to the specified base address to generate word requests for the desired elements. Finally, the requested elements are packed by a beat packer as specified by metadata from the element request generator to form the desired R beats.

The corresponding strided and indirect *write* converters are similar and differ only in the direction of the datapath: a *beat unpacker* splits beats into individual words, which are then used as write data for the generated write requests. The memory responses are combined and forwarded to the B channel.

III. EVALUATION

A. Setup and Workloads

To evaluate AXI-PACK, we consider three RISC-V systems-on-chip (SoCs) using CVA6 with Ara as a vector processor,

AXI4 interconnects, and a banked on-chip SRAM memory:

- **BASE**: unmodified CVA6 and Ara connecting to a regular banked memory over a standard AXI4 bus.
- **PACK**: unmodified CVA6 and AXI-PACK-extended Ara connecting to a banked memory with an AXI-PACK controller over an AXI-PACK-extended bus.
- **IDEAL**: like BASE, but Ara connects directly to an exclusive, idealized memory with one port per lane, serving data with ideal packing, bandwidth, and latency.

IDEAL provides an upper bound for possible AXI-PACK benefits by idealizing interactions between Ara’s VLSU and memory. However, it does not avoid inefficiencies arising from Ara’s internal microarchitecture or CVA6. In all systems, Ara is parameterized to eight vector lanes and 256-bit-wide data buses. The banked memories provide eight 32-bit-wide word ports backed by 17 banks, which we determine in Section III-E to provide a good area-performance tradeoff.

On each system, we evaluate a set of vectorized benchmarks benefiting from efficient strided and indirect memory accesses:

- **ismt**: *in-situ matrix transpose*. We transpose a square matrix in place by swapping and rotating elements above and below the diagonal using strided accesses.
- **gemv**: *general matrix-vector multiply*. We investigate both *row*- and *column*-wise dataflows, with the latter trading vector reductions for strided accesses, and use the fastest approach on each system for fair comparisons.
- **trmv**: *triangular matrix-vector multiply*, a *gemv* with an upper-triangular matrix. Only nonzero elements are streamed, incurring bursts of varying lengths. We again use the fastest dataflow on each system.
- **spmv**: *sparse matrix-vector multiply*, a widespread irregular memory-bound operation using indirect accesses.
- **prank**: *PageRank* [20], which rates each node in a graph based on the edges inbound to it. The graph is represented as a sparse weighted adjacency matrix.
- **sssp**: *single-source shortest path*, which calculates the shortest path from one node to all others in a weighted, directed graph represented as a sparse matrix.

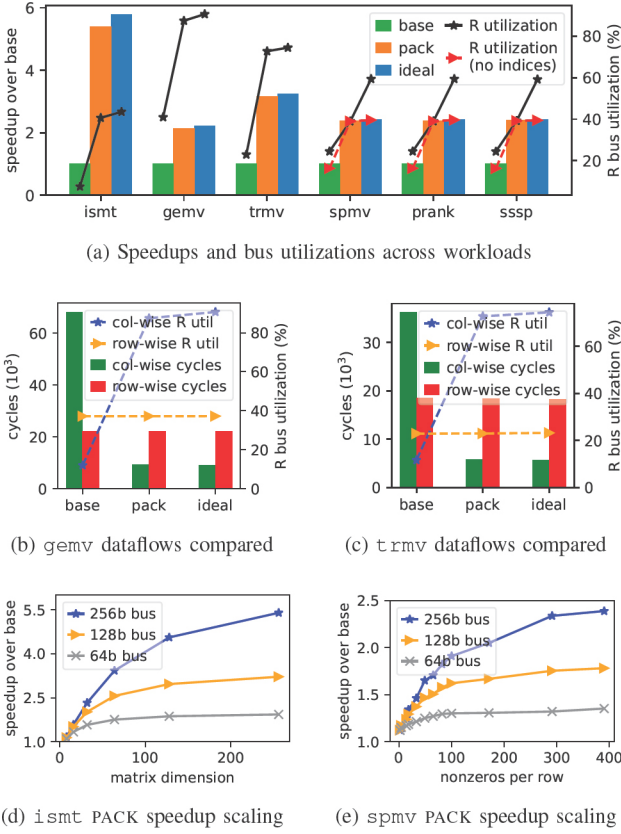


Figure 3. AXI-PACK performance results

We run the first three benchmarks leveraging strided streams on randomly-generated square matrices and the latter three leveraging indirect streams on real-world sparse matrices from the SuiteSparse collection [21] in the widespread compressed sparse rows (CSR) format. Elements are stored as 32-bit floats and indices as 32-bit integers. On indirect workloads, the PACK system uses our extensions to handle indirection in-memory, whereas BASE and IDEAL systems fetch indices into Ara.

B. Performance

We simulate our systems at the register transfer level (RTL) to determine the performance and read bus utilization for each benchmark. We initially assume a fixed matrix size of 256 for strided workloads and the sparse matrix *heart1* (390 average nonzeros per row) for indirect workloads. *gemv* and *trmv* use a column-wise dataflow on PACK and IDEAL and a row-wise dataflow on BASE, which we will show to be optimal for each respective system.

Fig. 3a shows PACK speedups over BASE and read bus utilizations with and without index transfers. AXI-PACK significantly improves bus utilization and performance for all workloads, achieving 97% of the IDEAL performance on average. On strided workloads, we measure peak speedups of $5.4\times$ (*ismt*) and bus utilizations of 87% (*gemv*). We note that read bus utilizations on *ismt* are limited to 50% due to read-write ordering in Ara. On indirect workloads, we achieve speedups of up to $2.4\times$ (*spmv*) and bus utilizations of up to 39% (*sssp*).

PACK handles indirection directly in its AXI-PACK controller, avoiding IDEAL’s waste of up to 20% (*spmv*) of bus time on index traffic and shifting indexed workloads further from the memory-bound toward the compute-bound regime.

Figs. 3b and 3c compare the row- and column-wise dataflow performance for *gemv* and *trmv*. Row-wise flows use only long contiguous accesses, so their performance is identical for BASE and PACK and very close to IDEAL. However, they require costly vector reductions, limiting BASE bus utilizations to 37% and 23%. Column-wise flows avoid reduction by working on multiple results at once, providing higher IDEAL performance and PACK utilizations of 87% and 72%. However for BASE, we stick to a row-wise flow, as the performance impact of strided accesses outweighs that of reductions without our extensions.

We also analyze the impact of input size and bus width on AXI-PACK speedups for representative strided and indirect workloads. Figure 3d shows *ismt* speedups for matrix dimensions of 8 to 256 and bus widths of 64 to 256 bit, corresponding to 2 to 8 Ara lanes and thus a growing number of functional units. As matrix size increases, speedups converge and reach up to 1.9, 3.2, and $5.4\times$; as we widen the bus, the narrow accesses of BASE become less efficient, increasing peak PACK speedups. As matrix size decreases, streams and useful computation phases shorten and become bottlenecked by the overhead of row iteration, decreasing speedups. Figure 3e shows *spmv* speedups for sparse matrices with 2 to 390 average nonzeros per row and the same bus widths as before. Speedups again converge and reach up to 1.4, 1.8, and $2.4\times$. We see similar scaling trends as for *ismt* because in *spmv*, the nonzeros per row determine the computation phase and stream lengths in each row iteration. We note that thanks to our request-bundling approach, using AXI-PACK never results in a slowdown no matter how short streams become.

C. Area and Timing

We synthesize our AXI-PACK adapter with Synopsys *Design Compiler* for GlobalFoundries’ 22nm FD-SOI technology, targeting the SSG corner at -40°C with low- V_t cells, 0.72V supply voltage, and no back-biasing. Unless otherwise specified, we constrain a 1 GHz clock and 100ps IO delays and parameterize the decoupling queues to a depth of four.

Fig. 4a shows the minimum achieved clock period and area for different clock constraints and bus widths of 64, 128, and 256 bit. Our adapter shows good scalability, increasing linearly in area with bus width and incurring 69, 130 and 257kGE at 1GHz. Our full 256-bit controller incurs merely 6.2% of Ara’s area, demonstrating that AXI-PACK handling at banked endpoints is reasonably inexpensive. As we decrease the constrained clock, we see that adapter area scales gracefully past Ara’s 1 GHz clock target and reaches minimum periods of 787, 800, and 839ps with only small increases in area.

Fig. 4b shows a hierarchical area breakdown of the adapter. As expected, the read and write converters are similar in size for both irregular burst types, since they simply reverse each other’s datapaths. While the simpler strided converters are only up to 42% larger than the base AXI4 converter, the indirect converters are nearly double this size due their two stages.

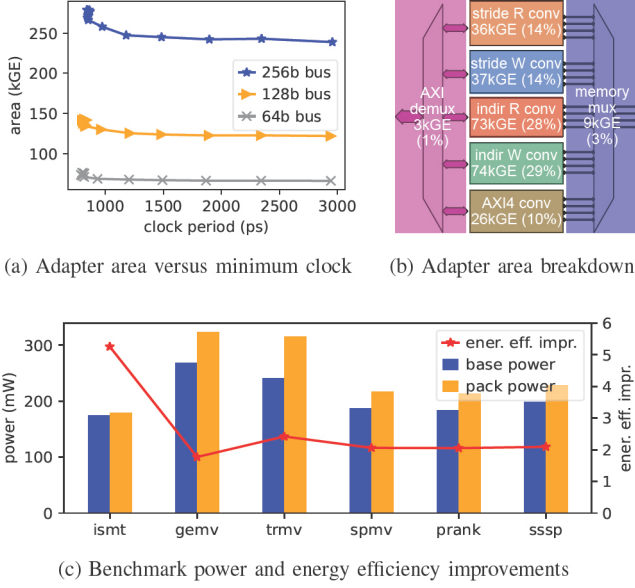


Figure 4. AXI-PACK area, timing, and energy results

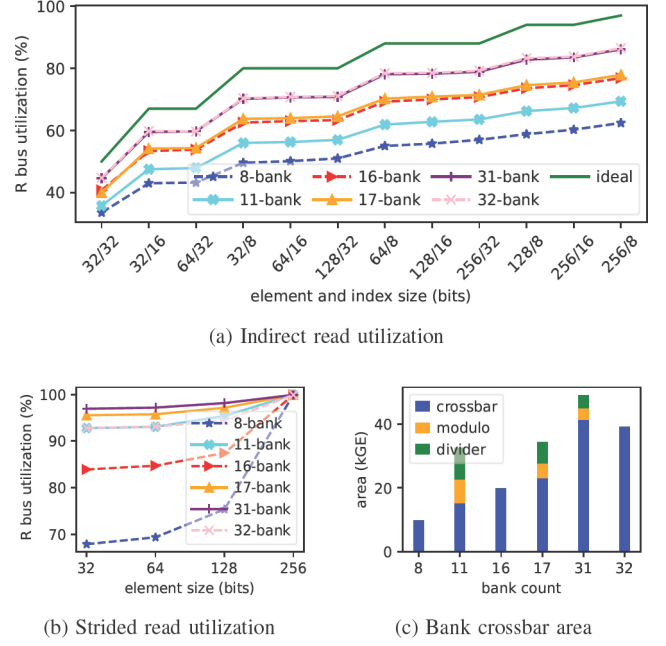


Figure 5. AXI-PACK parameter sensitivity results

D. Energy and Power

We estimate the power consumption of PACK and BASE, excluding the SRAM crossbar and banks, in the TT corner of GlobalFoundries’ 22 nm FD-SOI technology at 1 GHz. We topographically synthesize our system using Synopsys *Design Compiler* and estimate power on the benchmarks from Section III-B using Synopsys *PrimeTime*. Figure 4c shows the average power and energy efficiency improvement of PACK over BASE for each benchmark. Despite small power increases in PACK by at most 31% (*trmv*), all workloads see notable energy efficiency improvements, achieving peaks of $5.3\times$ (*ismt*) and $2.1\times$ (*sssp*) on strided and indirect benchmarks, respectively.

E. Parameter Sensitivity

To gain deeper insight into the scaling of AXI-PACK performance and hardware complexity, we investigate the impact of element and index size as well as bank count on read bus utilization and bank crossbar area. For performance measurements, we connect our controller to an ideal requestor issuing continuous read requests of length 256 and use random indices. Unless otherwise specified, parameters default on their PACK system configuration, but we increase decoupling queue depths to 32 to avoid bottlenecks unrelated to our analysis. We consider power-of-two bank counts from 8 to 32, which result in minimal addressing logic, as well as prime counts in this range, which minimize bank conflicts across different strides. We also consider an ideal memory without bank conflicts.

Indirect accesses: Figure 5a shows the bus utilization achieved on indirect reads for different element-index size pairs and bank counts. For all size pairs, utilization increases monotonically with bank count as fewer bank conflicts occur. Since indirect bursts involve one contiguous and one random but no strided bank access sequences, prime bank counts show

no inherent advantage here. Across all bank counts, utilization improves mainly with the *ratio* r of element size to index size: since we fetch indices as whole bus lines, we must fetch one index line for every r data beats on average, limiting our ideal bus utilization to $r/(r+1)$. For 32-bit elements and index sizes of 32, 16, and 8 bit, this corresponds to ideal utilizations of 50, 67, and 80%. Thus, with larger elements or smaller indices, AXI-PACK indirection bus utilizations can further exceed those shown in Section III-B.

Strided accesses: Figure 5b shows the bus utilization for strided reads for different element sizes and bank counts, averaged across element strides of 0 to 63. As expected, prime bank counts offer significant performance benefits on strided accesses, though more banks further improve performance for both power-of-two and prime bank counts. With increasing element size, conflicts become less likely for all bank counts as there are fewer aligned elements in each bus-wide line.

Bank crossbar area: Figure 5c shows the bank crossbar’s total area for different bank counts, highlighting the overhead prime bank counts incur for modulo and division units to compute bank addresses. Power-of-two-banked crossbars are generally cheaper and prime-banked overheads decrease with increasing bank counts. Since 17 banks provide good area-overhead and area-performance tradeoffs (95% and 81% of ideal performance on strided and indirect reads on average), this is the bank count we chose for our evaluation systems.

IV. RELATED WORK

Existing hardware approaches to efficient strided and indirect streams focus mostly on either end of the memory system.

Core-side extensions decouple accesses from execution and eliminate redundant load-store and bookkeeping instructions. Prodigy [5] prefetches nested indirect streams and proposes

dynamic cache bypass policies. While highly decoupled, it does not simplify program flow, limiting its acceleration. Wang et al. [6] propose strided and indirect streams mapped directly to architectural registers. This eliminates load-store and address iteration instructions, but still incurs dedicated instructions to step streams. Stream semantic registers [8] and their indirection extensions [9] implicitly step streams on access, enabling near-continuous useful instruction issues even on single-issue in-order cores. Domingos et al. [10] extend register-mapped irregular streams to vector processors. Except for cache policies, these extensions ignore interconnect and memory system efficiency. AXI-PACK is largely orthogonal to all of them; it may be used with any bus width, burst length, or mapping mechanism, providing a reusable protocol carrying irregular streams through interconnects and to stream-aware endpoints with high bus efficiency.

Memory-side extensions focus on bus efficiency and access latency. The Impulse memory controller [11] maps irregular streams to virtual pages; it provides inherent, on-the-fly bus packing, but relies on managed virtual addressing. Hussain et al. propose pattern-aware memory controllers [13] and systems [12] prefetching irregular stream descriptors to dedicated scratchpads. This enables fast, packed core accesses, but incurs notable complexity overheads. PLANAR [14] accelerates layout transforms by writing packed, cacheable irregular data to memory ahead of use, and the data rearrangement engine [15] is integrated directly into a hybrid memory cube architecture. While DLT accelerators are highly bandwidth-efficient, they require physical memory buffers and explicit, ahead-of-time invocation to be beneficial. AXI-PACK enables the benefits of all of the above extensions. Bus packing can be done on the fly by our controller or ahead of time by an AXI-PACK-capable direct memory access (DMA) controller. Our lightweight irregular requests provide performance without precluding the use of more complex, memory-mapped stream descriptors. However unlike other proposals, AXI-PACK builds on an established protocol and extends irregular streams throughout interconnects, feeding directly into cores in an end-to-end fashion.

V. CONCLUSION

We present AXI-PACK, an extension to the AXI4 on-chip protocol enabling highly-efficient end-to-end strided and indirect memory streams. To demonstrate AXI-PACK in an end-to-end system, we extend an open-source RISC-V vector processor to use it for strided and indexed accesses and design a banked memory controller serving irregular bursts. AXI-PACK increases bus utilizations up to 87% in strided and 39% in indirect benchmarks, resulting in speedups of up to $5.4\times$ and $2.4\times$. Synthesizing our AXI-PACK controller, we find that it incurs only 6.2% of the area of Ara, but improves energy efficiency by up to $5.3\times$ in strided and $2.1\times$ in indirect workloads.

VI. ACKNOWLEDGMENT

This work has been supported in part by funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101034126 (The EU Pilot) and Specific Grant Agreement No 101036168 (EPI SGA2).

REFERENCES

- [1] T. Hoefer, D. Alistarh, T. Ben-Nun, N. Dryden, and A. Peste, "Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks," *J. Mach. Learn. Res.*, vol. 22, pp. 241:1–241:124, 2021.
- [2] S. W. Min, V. S. Mailthody, Z. Qureshi, J. Xiong, E. Ebrahimi, and W. mei W. Hwu, "Emogi," *Proc. VLDB Endowment*, 2020.
- [3] S. Georgescu and H. Okuda, "Conjugate gradients on multiple gpus," *Int. J. Numerical Methods in Fluids*, vol. 64, 2010.
- [4] H. Li, K. Li, J. Peng, and K. Li, "Cusnmf: A sparse non-negative matrix factorization approach for large-scale collaborative filtering recommender systems on multi-gpu," *2017 IEEE Int. Symp. Parallel and Distributed Process. with Applicat. and 2017 IEEE Int. Conf. Ubiquitous Computing and Commun. (ISPA/IUCC)*, pp. 1144–1151, 2017.
- [5] N. Talati, K. May, A. Behrooz, Y. Yang, K. Kaszyk, C. Vasiladiotis, T. Verma, L. Li, B. Nguyen, J. Sun, J. M. Morton, A. Ahmadi, T. M. Austin, M. F. P. O'Boyle, S. A. Mahlke, T. N. Mudge, and R. G. Dreslinski, "Prodigy: Improving the memory latency of data-indirect irregular workloads using hardware-software co-design," *2021 IEEE Int. Symp. High-Performance Comput. Architecture (HPCA)*, pp. 654–667, 2021.
- [6] Z. Wang and T. Nowatzki, "Stream-based memory access specialization for general purpose processors," *2019 ACM/IEEE 46th Annu. Int. Symp. Comput. Architecture (ISCA)*, pp. 736–749, 2019.
- [7] M. Méndez-Lojo, M. Burtcher, and K. Pingali, "A gpu implementation of inclusion-based points-to analysis," in *PPoPP '12*, 2012.
- [8] F. Schuiki, F. Zaruba, T. Hoefer, and L. Benini, "Stream semantic registers: A lightweight risc-v isa extension achieving full compute utilization in single-issue cores," *IEEE Trans. Comput.*, vol. 70, pp. 212–227, 2021.
- [9] P. Scheffler, F. Zaruba, F. Schuiki, T. Hoefer, and L. Benini, "Indirection stream semantic register architecture for efficient sparse-dense linear algebra," *2021 Design, Automation & Test in Europe Conf. & Exhibition (DATE)*, pp. 1787–1792, 2021.
- [10] J. M. Domingos, N. Neves, N. Roma, and P. Tomás, "Unlimited vector extension with data streaming support," *2021 ACM/IEEE 48th Annu. Int. Symp. Comput. Architecture (ISCA)*, pp. 209–222, 2021.
- [11] J. B. Carter, W. C. Hsieh, L. B. Stoller, M. R. Swanson, L. Zhang, E. Brunvand, A. Davis, C.-C. Kuo, R. Kuramkote, M. A. Parker, L. Schaeckle, and T. Tateyama, "Impulse: building a smarter memory controller," *Proc. Fifth Int. Symp. High-Performance Comput. Architecture*, pp. 70–79, 1999.
- [12] T. Hussain, "A novel hardware support for heterogeneous multi-core memory system," *J. Parallel Distributed Comput.*, vol. 106, pp. 31–49, 2017.
- [13] T. Hussain, O. Palomar, O. S. Unsal, A. Cristal, and E. Ayguadé, "Memory controller for vector processor," *J. Signal Process. Syst.*, vol. 90, pp. 1533–1549, 2018.
- [14] A. Barredo, A. Armejach, J. C. Beard, and M. Moretó, "Planar: a programmable accelerator for near-memory data rearrangement," *Proc. ACM Int. Conf. Supercomputing*, 2021.
- [15] G. S. Lloyd and M. B. Gokhale, "In-memory data rearrangement for irregular, data-intensive computing," *Computer*, vol. 48, pp. 18–25, 2015.
- [16] Arm, "AMBA AXI and ACE Protocol Specification," <https://developer.arm.com/documentation/ih0022/hc>.
- [17] A. Kurth, W. Rönninger, T. Benz, M. A. Cavalcante, F. Schuiki, F. Zaruba, and L. Benini, "An open-source platform for high-performance non-coherent on-chip communication," *IEEE Trans. Comp.*, vol. 71, pp. 1794–1809, 2022.
- [18] M. Perotti, M. Cavalcante, N. Wistoff, R. Andri, L. Cavigelli, and L. Benini, "A "new ara" for vector computing: An open source highly efficient risc-v v 1.0 vector processor design," in *2022 IEEE 33rd International Conference on Application-specific Systems, Architectures and Processors (ASAP)*, 2022, pp. 43–51.
- [19] F. Zaruba and L. Benini, "The cost of application-class processing: Energy and performance analysis of a linux-ready 1.7-ghz 64-bit risc-v core in 22-nm fdsoi technology," *IEEE Trans. Very Large Scale Integration (VLSI) Syst.*, vol. 27, pp. 2629–2640, 2019.
- [20] L. Page, S. Brin, R. Motwani, and T. Winograd, "The pagerank citation ranking : Bringing order to the web," in *WWW 1999*, 1999.
- [21] T. A. Davis and Y. Hu, "The university of florida sparse matrix collection," *ACM Trans. Math. Softw.*, vol. 38, pp. 1:1–1:25, 2011.