

# CXLMEMURING: A Hardware Software Co-design Paradigm for Asynchronous and Flexible Parallel CXL Memory Pool Access

Yiwei Yang

## Abstract

CXL has been the emerging technology for expanding memory for both the host CPU and device accelerators with load/store interface. Extending memory coherency to the PCIe root complex makes the codesign more flexible in that you can access the memory with coherency using your near-device computability. Since the capacity demand with tolerable latency and bandwidth is growing, we need to come up with a new hardware-software codesign way to offload the synthesized memory operations to the CXL endpoint, CXL switch or near CXL root complex cores like Intel DSA to fetch data; the CPU or accelerators can calculate other stuff in the backend. On CXL done loading, the data will be put into L1 if capacity fits, and the in-core ROB will be notified by mailbox and resume the calculation on the previous hardware context. Since the distance(timing window) of the load instruction sequence is unknown, a profiling-guided way of codegening and adaptively updating offloaded code will be required for a long-running job. We propose to evaluate CXLMEMURING the modified BOOMv3 with added in-core logic and CXL endpoint access simulation using CHI, and we will add a weaker RISC-V Core near endpoint for code offloading, and the codegening will be based on program analysis with traditional profiling guided way.

## ACM Reference Format:

Yiwei Yang. 2023. CXLMEMURING: A Hardware Software Co-design Paradigm for Asynchronous and Flexible Parallel CXL Memory Pool Access. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

## 1 Introduction

We live in the Great Memory Wall era. Resolving the memory wall is useful for either HPC applications, DLRM, or LLM training, either loading memory from the CPU or fetching and swapping the tensor data from the GPU to the memory pool[7]. How to resolve the latency and bandwidth, given the current circumstances, is a top priority for applications to scale without residual replicas. The traditional solution for hiding the latency of memory like ROB, MSHR, read-ahead cache, stack elimination, or TLB doesn't expand to the CXL.mem memory pool.

The CXL brings in the possibility of co-designing the application yourself with coherency support compared to other private standards like NVLink. Also, the current hack of leveraging the PCIe P2P snapshotting the data[8] and loading


from it still wouldn't be ideal for expanding your swapping of VRAM. The underlying law for that may be the memory roofline [2, 3]; the memory requests from the certain accelerator be classified after a single trial. Where the memory bound is hit in the memory roofline is where to optimize the memory bandwidth bound by prefetching or asynchronous access. The prefetch approaches have been explored in RDMA like AIFM[9], Leap[1] or MIRA[12].

Now that the MLIR way of analyzing the memory side Very Long Instruction Word is well established[6], more and more companies who care about the memory wall greatly will apply VLIW based on MLIR. Asynchronous way of accessing data by itself mapping to certain programming paradigms like coroutine[5] or MPI all to all gather[4, 11]. Concerns are those paradigms are too narrow for the entire codebase of C++ that wants to offload memory to remote.

## 2 Implementation

We divide the stack into software and hardware part, where software is a JIT compiler that does analyses of the offloading window dynamically because the access pattern is not statically computable, and hardware is an async loading engine inside the CPU where on getting the data notifies the CPU to resume the previous context and a coprocessor near endpoint that computes the load instruction sequence and simple memory operations. On calling back to async loading engine, it's a CXL.io request with a mailbox inside the core. The in-core logic can also be implemented in GPU EMC for getting other DSA or CPU data.

**Software Stack.** We propose a binary JIT way of seamlessly executing the binary like Apple Rosetta, and the offloading of the load instructions of the program, as described in MIRA[12] will be adaptively reapplied to the remote. We will apply the forward analysis first to translate all the remotable memory accesses and pointer accesses to CXL byte addressable way of access in MLIR. We name this MemUring because this async way of accessing data is just like IOUring. Then we use the backward analysis to get all the functions with remote pointer passed, and may be rewritten with a native local pointer. All the functions and labels will be labeled as profiling guided points for marking the cost model penalty for the timing window. On running, the JIT can modify the code after labels have been called once to reach a better timing window.



**Hardware Stack.** In the above graph, the red part is where we modify the hardware. We propose the evaluation based on BOOMv3 over FPGA since it's an easy-to-modify core, which can be manifest in the NoC or GPU accelerator with the difference of in-core-logic being put in CPU's cores but in GPU's External Memory Controller(EMC). We will add CHI for simulating access to CXL Switches and other accelerators. We also mark the CXL Switch as a possible offloaded point with a different cost model. We think in the future codesign, the CPU is only a hub for combining the DSA's requests and do OLAP that CPU excels at, and for the hardware design, we only need a Co-processor near the endpoint for calculating the memory requests, the endpoint can be either CXL Flash for memory expander, GPU or CXL Switch.

For in-core logic, we decided to make all the memory returns into L1 so that it's unique to this SMT core and will be instantly consumed as calculated by the software part; we try not to use interrupts but only set the ROB metadata to activate the requests.

### 3 Proposed Evaluation

We think we can use the evaluation to answer these questions.

**Effectiveness of capturing window size.** We want to know how effectively the window of instructions has been offloaded and how much other information can be done first before the memory is loaded.

**Relationship of integrating with ROB, MSHR.** We want to know how we should design this to integrate with ROB and MSHR.

**Additional On-chip Size Comparison.** We want to explore whether this can save chip size or not.

**Guiding the programming model.** We want to know how to make the programming model better because we think in the future, offloading control flow but pin the most memory in their local and only communicate little memory.

### 4 Related Work

**Data Streaming Accelerator.** It's definitely possible to put the things in the backend, apply the MLIR JIT to put all loading operations remotely, and let the remote endpoint to request DSA send back the device memory to LLC. However, the DSA is currently only designed for single root CPU and bulk memory load. The way of communicating from host

to DSA takes driver code and auxiliary data transmission, which is tedious compared with designing the load engine inside the core.

**Asynchronous RDMA/SmartNIC way of accessing data.** Mira[12] proposed the far memory operations offloading paradigm using profiling guided program syntheses with online modification of the offloaded code. Directly putting their implementation to CXL is not working since the granularity of accessing data of CXL is 64Byte while RDMA is 4KB, which is way bigger than a C++ object size that most workloads use. Their approach of RDMA works extremely well in the bulk memory load scenario, but not possible to get good results for pointer chasing and indirect memory reading. A way of rethinking those two scenarios in the world of CXL is our thoughts.

**"In-order-core" Asynchronous Memory Unit.** Compared with putting both the loading core and offloaded code inside the cores like [10]. Their evaluation part only uses the in-order core with object scratchpad inserted in L2 and doesn't talk about the relationship of L2 contention, ROB, and MSHR.

### References

- [1] Hasan Al Maruf and Mosharaf Chowdhury. Effectively prefetching remote memory with leap. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*, pages 843–857, 2020.
- [2] Nan Ding, Pieter Maris, Hai Ah Nam, Taylor Groves, Muaaz Gul Awan, LeAnn Lindsey, Christopher Daley, Oguz Selvitopi, Leonid Oliker, and Nicholas Wright. Evaluating the potential of disaggregated memory systems for hpc applications. *arXiv preprint arXiv:2306.04014*, 2023.
- [3] Yehonatan Fridman, Suprasad Mutalik Desai, Navneet Singh, Thomas Willhalm, and Gal Oren. Cxl memory as persistent memory for disaggregated hpc: A practical approach. *arXiv preprint arXiv:2308.10714*, 2023.
- [4] Dhabaleswar K (DK) Panda Gilad Shainer and Nick Sarkauskas. Accelerating scientific applications in hpc clusters with nvidia dpus using the mvapich2-dpu mpi library. <https://developer.nvidia.com/blog/accelerating-scientific-apps-in-hpc-clusters-with-dpus-using-mvapich2-dpu-mpi/>.
- [5] Yongjun He, Jiacheng Lu, and Tianzheng Wang. Corobase: coroutine-oriented main-memory database engine. *arXiv preprint arXiv:2010.15981*, 2020.
- [6] Paras Jain, Xiangxi Mo, Ajay Jain, Alexey Tumanov, Joseph E Gonzalez, and Ion Stoica. The ooo vliw jit compiler for gpu inference. *arXiv preprint arXiv:1901.10008*, 2019.
- [7] Zhuohan Li, Lianmin Zheng, Yinmin Zhong, Vincent Liu, Ying Sheng, Xin Jin, Yanping Huang, Zhifeng Chen, Hao Zhang, Joseph E Gonzalez, et al. Alpaserve: Statistical multiplexing with model parallelism for deep learning serving. *arXiv preprint arXiv:2302.11665*, 2023.
- [8] Zaid Qureshi, Vikram Sharma Mailthody, Isaac Gelado, Seungwon Min, Amna Masood, Jeongmin Park, Jinjun Xiong, CJ Newburn, Dmitri Vainbrand, I-Hsin Chung, et al. Gpu-initiated on-demand high-throughput storage access in the bam system architecture. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 325–339, 2023.
- [9] Zhenyuan Ruan, Malte Schwarzkopf, Marcos K Aguilera, and Adam Belay. {AIFM}:{High-Performance},{Application-Integrated} far

- memory. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 315–332, 2020.
- [10] Luming Wang, Xu Zhang, Tianyue Lu, and Mingyu Chen. Asynchronous memory access unit for general purpose processors. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 2(2):100061, 2022.
- [11] Shibo Wang, Jinliang Wei, Amit Sabne, Andy Davis, Berkin Ilbeyi, Blake Hechtman, Dehao Chen, Karthik Srinivasa Murthy, Marcello Maggioni, Qiao Zhang, et al. Overlap communication with dependent computation via decomposition in large deep learning models. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1*, pages 93–106, 2022.
- [12] Guo Zhiyuan and Yiyang Zhang He, Zijian. *Mira: Towards a Transparent and Efficient Far Memory System*. PhD thesis, UC San Diego, 2023.