

MEX: CXL-Based Memory EXpander With Hardware Acceleration

Author1 Seon Young Kim
Electronics and Telecommunications
Research Institute
Republic of Korea
seonyoung8436@etri.re.kr

Author2 Hoo Young Ahn
Electronics and Telecommunications
Research Institute
Republic of Korea
ahnhy@etri.re.kr

Author3 Yoomi Park
Electronics and Telecommunications
Research Institute
Republic of Korea
parkym@etri.re.kr

1. INTRODUCTION

As data-intensive applications represented by artificial intelligence have great attention, the amount of memory required for computing systems is rapidly increasing, especially in HPC (High-Performance Computing) systems. According to the study that analyzed HPL (High-Performance Linpack) performance in an HPC system [1], the memory capacity per CPU core to obtain the theoretical performance of HPL tends to increase in proportion to the number of CPU cores constituting the system. The result of the study suggests that as the number of CPU cores constituting the system increases, the memory capacity required for the system increases more steeply.

Though the required memory capacity is increasing day by day, there is a limit to the memory capacity of a computing node depending on hardware characteristics such as the number of CPU memory channels. To overcome this limitation, various memory expanders are being unveiled by major memory vendors such as Samsung that allow a computing node to utilize additional expanded memory beyond its own local memory [2]. These memory expanders provide expanded memory that can be accessed in a cache-coherent manner through the state-of-the-art interconnects such as CXL [3]. In addition, they also provide hardware acceleration through the built-in accelerator.

In this poster, we introduce our ongoing research to develop our own CXL-based memory expander called MEX (Memory EXpander). MEX not only provides expanded memory through the CXL, but also provides hardware acceleration for the K-NN (K-Nearest Neighbor), a key operation for the similarity search.

2. MEX: CXL-based Memory EXpander

MEX is being developed using a commercial FPGA card with the goal of providing more than 32GB expanded memory and hardware acceleration for the K-NN. Figure 1 shows the brief architecture of MEX. The description of each block constituting MEX is as follows.

2.1 Control Unit

The control unit supports essential functions for MEX such as interfacing with the host or maintaining cache coherency. The control unit is composed of functional blocks operating each function, and these functional blocks are implemented based on IPs provided by the FPGA vendor. The description of each functional block is as follows.

The CXL/PCIe controller handles interfacing with the host over CXL and PCIe. For this, it provides functions such as interface conversion or flow control. The cache coherent block performs the function of maintaining cache coherency between the host and accelerator, and the method of maintaining cache coherency follows the CXL standard [3]. The memory controller manages access to the expanded memory, and it provides several functions such as creating DDR signals or refresh control for this.

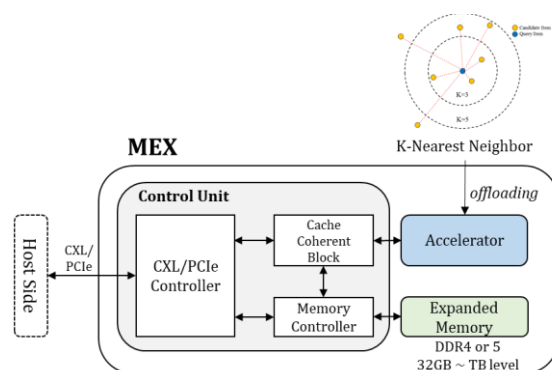


Figure 1. The brief architecture of MEX

2.2 Accelerator

The accelerator leverages the near-memory processing concept to accelerate the K-NN operation offloaded from the host. The K-NN is a key operation of similarity search that is the task of retrieving data items similar to a given query.

We are currently developing the accelerator using Vitis, an HLS (High Level Synthesis) tool. According to our preliminary experiments, the performance of the prototype accelerator on SIFT1M, a large-scale dataset for evaluating K-NN performance, is 3.12x better than CPU-based computing environment.

2.3 Expanded Memory

The expanded memory is additional DRAM memory accessible by the host in a cache-coherent manner via CXL.mem protocol [3]. The expanded memory aims to provide TB-level additional memory for the host, starting with around 32GB of additional memory on the prototype MEX.

ACKNOWLEDGMENTS

This work is supported by the Super Computer Development Leading Program of the National Research Foundation of Korea (NRF) funded by the Korean government (Ministry of Science and ICT(MSIT)) (No.2021M3H6A1017683).

REFERENCES

- [1] D. Zivanovic et al., “Main memory in HPC: Do we need more or could we live with less?,” ACM Transaction on Architecture and Code Optimization, vol. 14, no. 1, 2017, pp. 1-26.
- [2] S. J. Park et al., “Scaling of memory performance and capacity with cxl memory expander,” IEEE Hot Chips 34 Symposium (HCS), pp. 1-27, 2022.
- [3] CXL Consortium, “Compute Express Link Specification Revision 2.0,” 2020.