Scalable Data Processing at Network transfer rates with nCorium Compute in Memory Modules

Suresh Devalapalli
nCorium - sureshd@ncorium.com

Brett Neuman
Los Alamos National Laboratory - bneuman@lanl.gov

Arvindh Lalam
nCorium - arvindhl@ncorium.com

***Abstract*—The mismatch between compute and data movement from the network to system memory, system memory to the CPU caches, and system memory to storage devices creates significant bottlenecks in streaming applications. nCorium's technology merges compute capabilities into the memory channel that allows applications to offload compute and I/O intensive tasks to low power reconfigurable compute engines that deliver higher throughputs at lower energy consumption.**

***Index Terms*—processing-in-memory, fpga, object storage**

I. Introduction

nCorium has merged compute capabilities into the memory channel on modules called nCorium Compute in Memory Modules (nCIMMs). nCIMMs replace the traditional DIMMs in the server architectures and offer a path to achieve an order of magnitude better compute performance at the socket level. nCIMMs allow applications to offload compute and I/O intensive tasks to low power reconfigurable compute engines that deliver higher throughput at low energy consumption. Each nCIMM consists of an FPGA for compute and DRAM arrays for memory capacity. It can execute a streamlined datapath that ensures compute is close to data and minimizes the number of times data has to be moved between the CPU caches, Memory and Storage instances. Object Storage Targets (OSTs) with NVMe SSDs are used for data storage and retrieval.

II. Methodology

A data source program was developed that sets up a number of buffers with specific data patterns and transfers the metadata to the device program over standard network sockets. The data is later transferred to the device using Remote Direct Memory Access (RDMA). The testing environment utilizes RDMA over Ethernet (RoCE), FPGAs to perform (10,2) Reed-Solomon encoding for redundancy, and NVMe SSDs for data storage and retrieval. Table 1 shows the four data transfer modes tested:

Mode	Description
RDMA	Transfers remote data stream to FPGA memory
RDMA+ENC	RDMA plus FPGA erasure encoding operations
RMDA+SSD	RDMA to FPGA memory then write to SSDs
RDMA+ENC+SSD	End-to-End solution performing all operations

TABLE I: Data transfer modes.

III. Results

nCorium's server architecture offers linear bandwidth scaling with the number of memory channels. A pair of nCIMMs are associated with a NIC, provide modular in-line processing and a path for storage via NVMe SSDs. We observed 3-5x better performance than Intel CPU based solution. Figure 1 shows the complete data transfer mode bandwidth results.

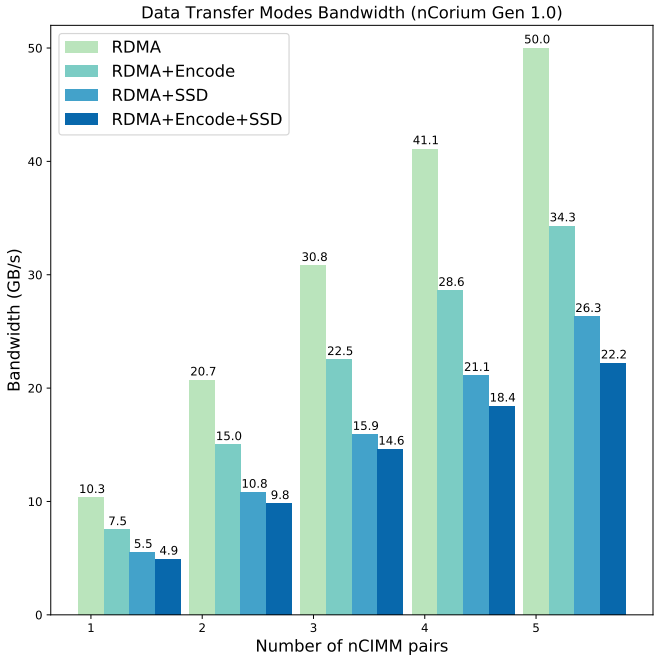

Fig. 1: Bandwidth results for all data transfer modes.

IV. Next Generation Architecture

The next generation of nCorium Accelerated Server will include CPUs with more memory channels, as well as in-line processing for compression and encryption, making them well suited for high throughput storage appliances. Reconfigurable cores can also be used for other applications including AI, ML, and video transcoding.