

Architectures, Solutions, and Community VIRTUAL EVENT, APRIL 11-12, 2023

Fabric Attached Memory – Hardware and Software Architecture

Clarete Crasta





- Motivation
- Fabric Attached Memory Architecture
- Fabric Attached Memory Software stack
- Results and Use cases
- Summary & future work



Architectures, Solutions, and Community VIRTUAL EVENT, APRIL 11-12, 2023



Motivation

Need quick answers on larger data sizes





Technology Advances

Advances to Memory Technologies





Advances to interconnects – Slingshot, CXL





©Copyright 2023 Hewlett Packard Enterprise Company

Heterogeneous Compute





Architectures, Solutions, and Community VIRTUAL EVENT, APRIL 11-12, 2023



Fabric Attached Memory - Architecture

Fabric-attached Memory

Benefits

- Independent scaling of compute and memory
 - Improved utilization, reduced overprovisioning
 - Decoupling of failure domains
- Reduced depth of I/O stack for large workloads
 - Direct, unmediated access
 - Accessible by all compute resources
- Challenges
 - Coherence domains are limited to individual nodes
 - On RDMA based interconnects
 - Fabric latency is large compared to local memory
 - Software has to be (usually) refactored for performance





Prototype Implementation



interconnected over Slingshot





Architectures, Solutions, and Community VIRTUAL EVENT, APRIL 11-12, 2023



Fabric Attached Memory - Software Stack

Software stack - OpenFAM

OpenFAM

- API and reference implementation to program FAM (fabric attached memory) across scale-up, existing scale-out and emerging FAM architectures.
- API is generic, supports both fabric attached persistent memory and volatile memory
- Supports direct access and RMA based on the environment.





Software stack - Components





Architectures, Solutions, and Community VIRTUAL EVENT, APRIL 11-12, 2023



Results and Use Cases

12 | © SNIA. All Rights Reserved.

FAM access latency and throughput

- Short transfer latency in the order of 3-4 µs with 256 byte transfers
- 16 PEs, 100 operations per PE, 64 MiB transfers for throughput
- Linear scaling close to available bandwidth



1. S. Singhal et. al., OpenFAM: A library for programming disaggregated memory, in OpenSHMEM and related technologies 2021 workshop, http://www.openshmem.org/workshops/openshmem2021/program.html



Motivating Use Cases

FAM Enables

- HPC applications with data sets larger than DRAM of PEs.
- Independent scaling of memory and compute.
- With interoperability, augments existing PGAS libraries -MPI, OpenSHMEM, OpenFAM.
- Dynamic scaling of application workloads.
- Shared distributed heterogeneous configurations with GPUs/accelerators

In applications such as

- Large scale graph analytics applications in
 - Security, Social networks, Advertising, Internet of things for malware detection, Community detection, link prediction etc.
- HPC workflows/pipelines such as those in genomics, AI/ML
 - for transforming data in a workflow with large intermediate data sets
 - fast intermediate storage for staging applications
- FAM as distributed high performance cache through OpenFAM.



Idealized Workflow for HPC and Data Analysis



Summary and Future work

• HPE has built Fabric Attached Memory Architecture

- Defined the Hardware components and configurations
- Built the software stack to access and use FAM
- Performance benchmarks for FAM access latencies and throughput
- Use cases that can benefit from FAM

Future work

- CXL memory, future memory technologies on memory nodes
- More ecosystem enablement OpenSHMEM, Chapel, Arkouda, Data Formats



Acknowledgements

Co-authors:

John Byrne Dave Emberson Harumi Kuno Kristyn Maschhoff Sharad Singhal Ramya Ahobala Faizan Barmawer Gautham Bhat Amitha C Ramesh C Chaurasiya Monisha D Abhishek Dwaraki Darel Emmot Chinmay Ghosh Mashood Abdulla K Venkata Kailash Gumma Jainendra Jain Sridhar Joshi

Cynara Justine Sumant Karla Haardhik M A Ryan Menhusen Lisa Pallotti Shome Porno Sriram Ravishankar Carl Rick Rajak Rishikesh Adya Sharma Prashanth Tamraparni Bheemesh V



Architectures, Solutions, and Community VIRTUAL EVENT, APRIL 11-12, 2023



Please take a moment to rate this session.

Your feedback is important to us.

Architectures, Solutions, and Community VIRTUAL EVENT, APRIL 11-12, 2023



Backup



