NMI: A New Memory Interface to Enable Innovation

David Roberts[1], Amin Farmahini-Farahani[1], Kevin Cheng[1], Nathan Hu[1], David Mayhew[12], Michael Ignatowski[1]

[1]AMD Research, Advanced Micro Devices, Inc. [2]University of San Diego

Introduction – Problems Faced in Memory Systems

Memory subsystem impacts performance, energy, and cost

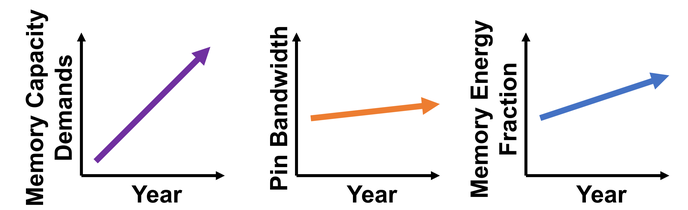

Other challenges: **Interface compatibility** between different memory types, **workload adaptation** to different memory types, and **reliability** due to process technology scaling, multi-level cells, and capacity increase

New Solutions

Processing in Memory (PIM): Processing data near where it resides

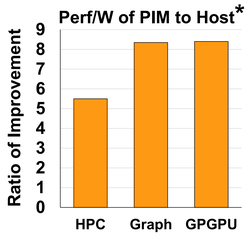

PIM reduces energy and bandwidth requirements by reducing communication and moving compute to the data.

Recent developments in 3D die stacking such as High Bandwidth Memory and the Hybrid Memory Cube are key enablers for PIM.

Standards for Heterogeneous Computation such as the Heterogeneous System Architecture (HSA) allow Host and PIM processors to share memory and work

Abstracted memory interfaces are becoming more important due to the emergence of diverse non-volatile memory technologies

Heterogeneous memory technologies can be used together to reduce cost while providing performance, capacity, and non-volatility

A New Memory Interface To Enable Innovation**

- Existing memory interface protocols present a barrier to overcoming key problems and providing scalable, compatible 'smart' memory components from multiple vendors, from cellphones to supercomputers
- This work is a step in overcoming these key problems

* D. Zhang, et al., "TOP-PIM: throughput-oriented programmable processing in memory," *HPDC*, 2014.
** D. Resnick and M. Ignatowski, "Proposing an abstracted interface and protocol for computer systems," White Paper, Sandia National Lab., 2014.

Novel Features

- Point-to-point networks of NMI Nodes (any combination of ports, switch, memories and processing elements). Master nodes for local control and subnets for massive scalability
- Abstracted, flexible timing interface supporting diverse technologies
- HSA-compatible virtual memory, cache coherency, task dispatch
- Optional feature set profiles to scale cost, area and complexity from embedded to supercomputer systems
- Scalable ECC tunable for application, including memory RAID
- Multiple physical to device mappings for custom interleaving
- Abstracted power modes for global management with mixed devices

NMI Protocol Layers

(1) Physical Layer: Flexible (electrical/optical/other), under development

(2) Link Layer: Packetized, reliable, scalable header overheads, virtual channels for deadlock avoidance, and low latency

(3) Transaction Layer: Classes of optional functionality as follows;

#	Class	Description
0	**Foundation & Computation**	Non-coherent R/W, capability query, logical memory region management, task dispatch
1	**Atomics**	Atomic operations
2	**Virtual Memory**	Address translation, TLB invalidation
3	**Coherence**	Cache coherence
4	**Fixed-Function Units**	Gather/Scatter, reduction, initialization, etc.
5	**Advanced ECC**	Memory RAID support
6	**Persistent Memory**	Fence, flush, freeing address ranges
7-12	**<reserved>**	Reserved for future standard features
13-15	**<vendor-defined>**	Vendor-defined

NMI Power Management

High-level (abstract) power mode and low-level (direct) power mode

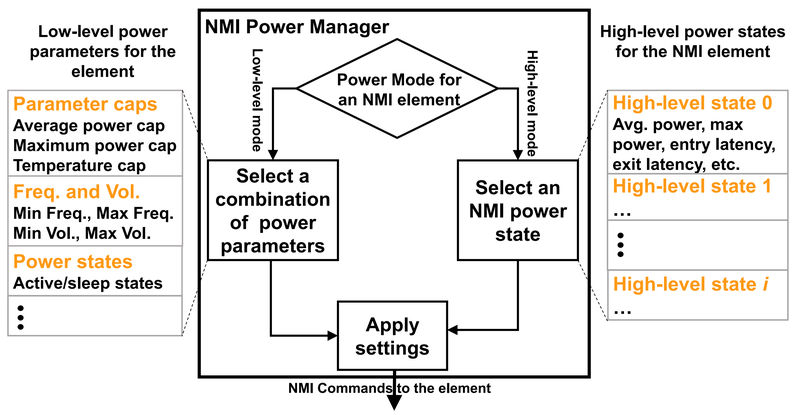

Unmanaged vs Managed NMI Networks

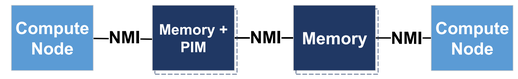

Unmanaged example: Small-scale, low-latency, low-overhead

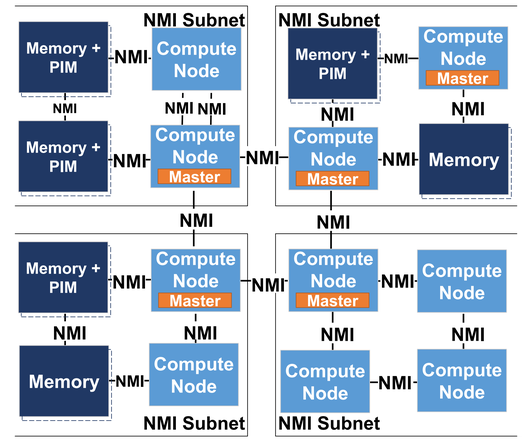

Managed example: Scalable to many nodes, divided to subnets, each managed by a Master node

Heterogeneous System Architecture (HSA)

- Designed with HSA in mind
- Capability & configuration register trees available on each node
- Support for cache-coherent shared virtual memory
- Support for task dispatch via Architected Queueing Language

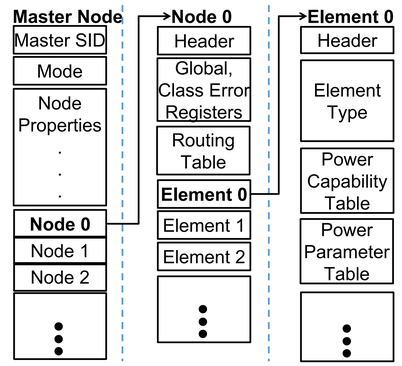

Logical Memory Regions

Logical division of physical memory address space into non-overlapping, contiguous regions

Each memory region has its own:

- Physical-to-device address mapping
- Size and address range
- Multi-level cell configuration

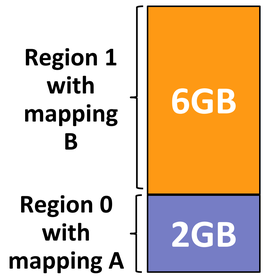

Conclusions

NMI is an abstracted, unified memory interface to support future **scale-out memory capacity, processing-in-memory, I/O devices, emerging non-volatile memories, cache-coherent shared virtual memory**