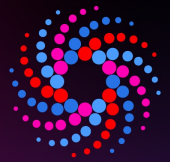

SC22
Dallas, TX | hpc accelerates.

Methodology for Evaluating the Potential of Disaggregated Memory Systems

Nan Ding, Samuel Williams, Hai Ah Nam, Taylor Groves, Muaaz Gul Awan, LeAnn Lindsey, Christopher Daley, Oguz Selvitopi, Leonid Oliker, Nicholas Wright

Lawrence Berkeley National Laboratory

Outline

- Need-to-know about Memory disaggregation
- Disaggregated memory system architecture
- Characterize application performance on a disaggregated memory system
- Case Study

What is memory disaggregation?

Today:

- Compute nodes are the basic unit of today's HPC systems
- Compute and memory resources are tightly coupled in each node
- Users request resources in the unit of a node

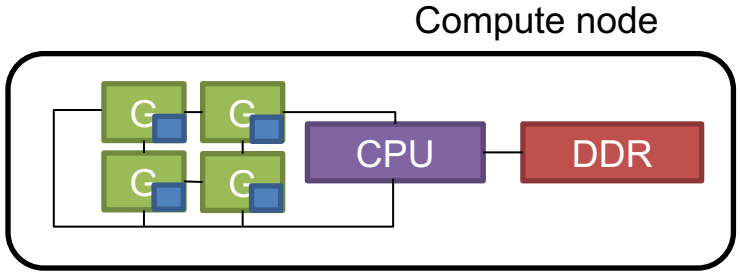

Memory Disaggregation:

- Decouple the compute and memory resources
- Allow for independent allocations of these resources regardless of where a job is placed

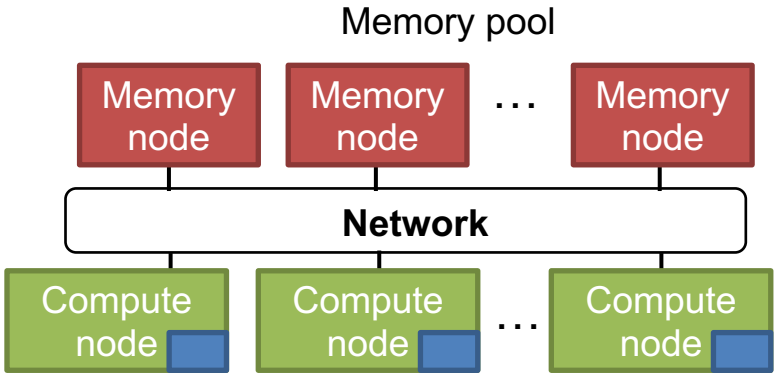

Memory disaggregation addresses memory imbalance and improves memory utilization

Expensive memory is often under utilized:

- At Azure: ~50% of all VMs never touch 50% of their rented memory
- At NERSC: only 15% of the scientific workloads on NERSC's Cori supercomputer use over 75% of the available on-node memory
- At LLNL: 10% of jobs utilize more than 75% of the node memory capacity

Memory disaggregation is practical for public could

- Meet performance requirements and low hardware cost
- CXL-Based Memory Pooling Systems for Cloud Platforms

What impact will these emerging technologies have on HPC?

Memory Disaggregation on HPC: More Memory, Less cost

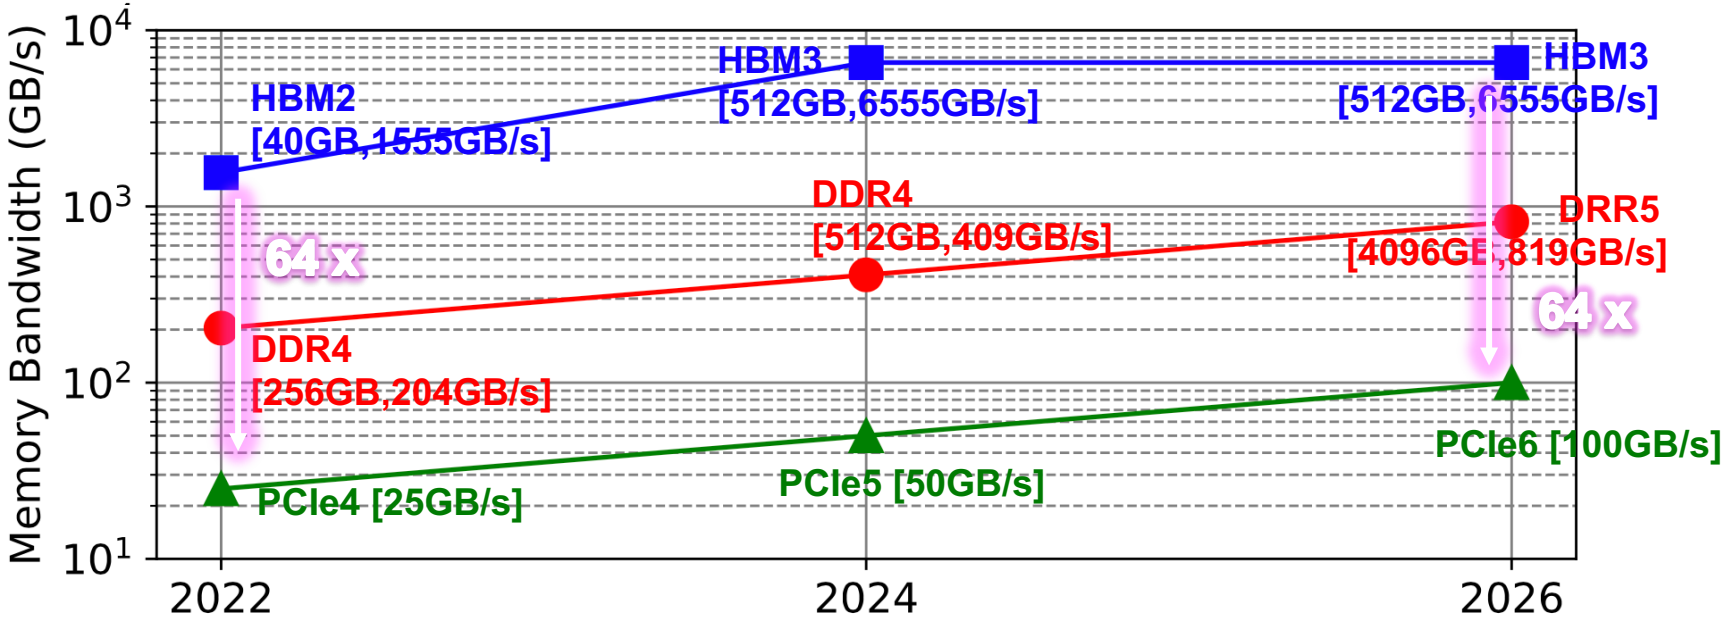

- DRAM w/ 16 DIMMS HBM3 w/ 8 stacks 16-Hi HBM2 w/ 8 stacks 8-Hi

Outline

- Need-to-know about Memory disaggregation
- Disaggregated memory system architecture
- Characterize application performance on a disaggregated memory system
- Case Study

Conceptual disaggregated memory system architecture

- One compute node (C) = one APU + HBM3 (512GB) + one NIC
- One memory node (M) = one DDR5 socket (4TB) + one NIC

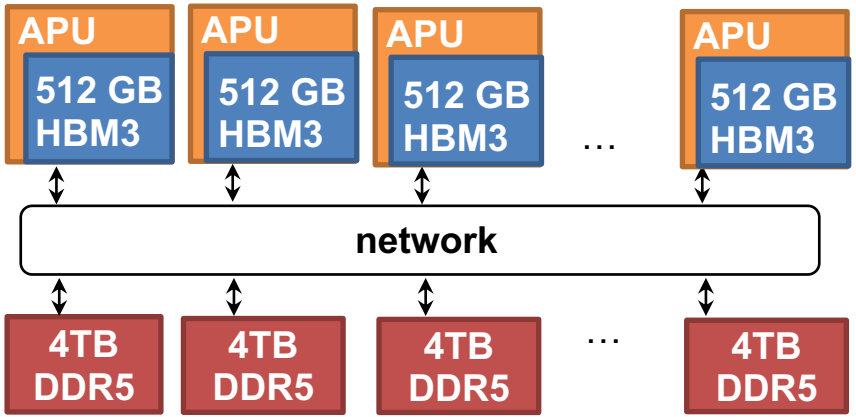

Available Remote Memory Capacity

- C:M = 1: 1

100% bandwidth
100% memory node capacity

- C:M = 2: 1

50% bandwidth
50% memory node capacity

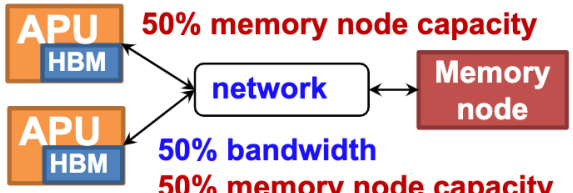

50% bandwidth
50% memory node capacity

- C:M = 1: 2

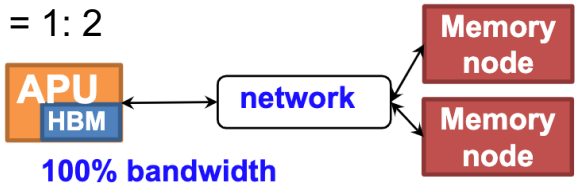

100% bandwidth
200% memory node capacity

Available budget

	TOT MEM: 20PB	40PB	80PB
System architecture $\frac{C}{M}$:	$\frac{10K}{5K}$	$\frac{10K}{10K}$	$\frac{10K}{20K}$
	2TB	4TB	8TB

Less contention

Available Remote Memory Capacity

- Not all jobs need remote memory
- HBM3 provides 512GB
- > 512GB, need remote memory

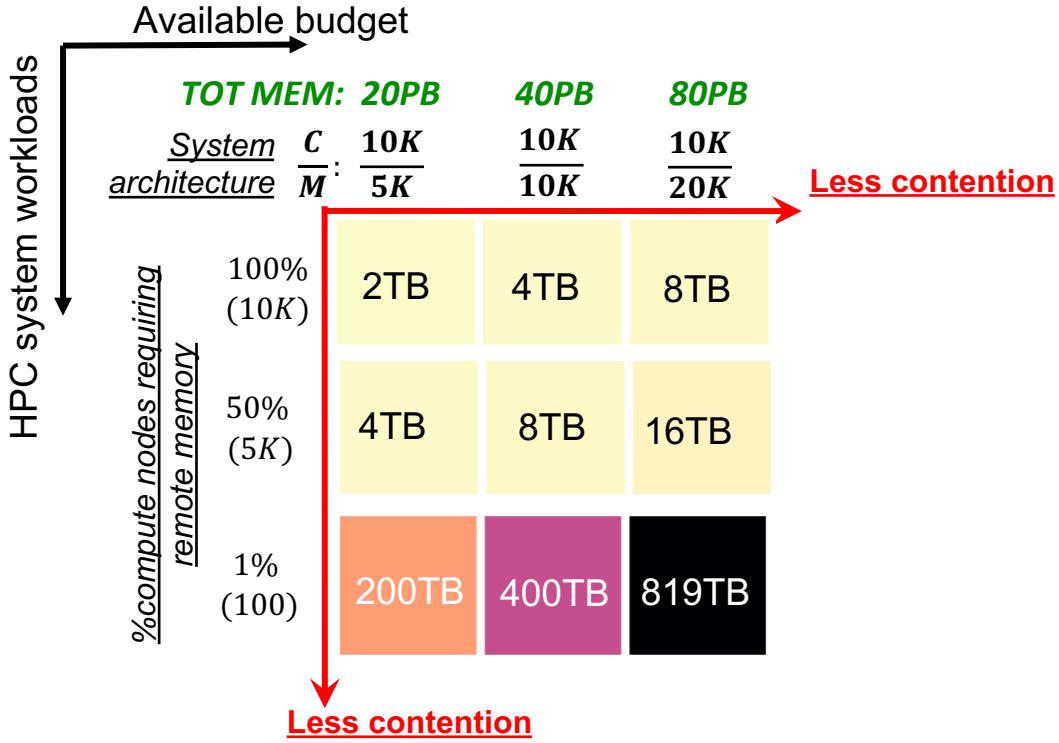

Available Remote Memory Bandwidth

- C:M = 1: 1

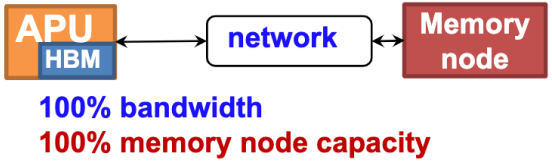

- C:M = 2: 1

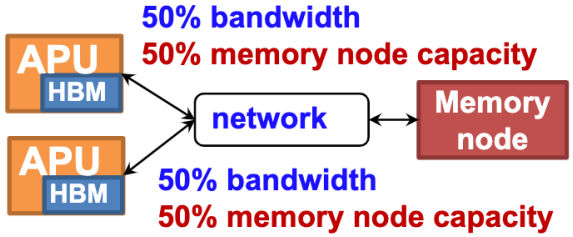

- C:M = 1: 2

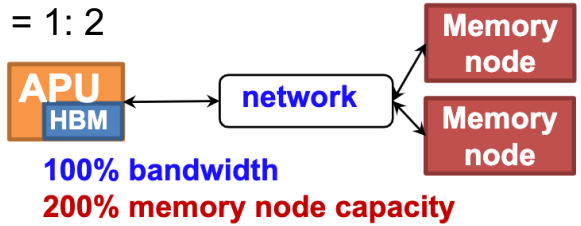

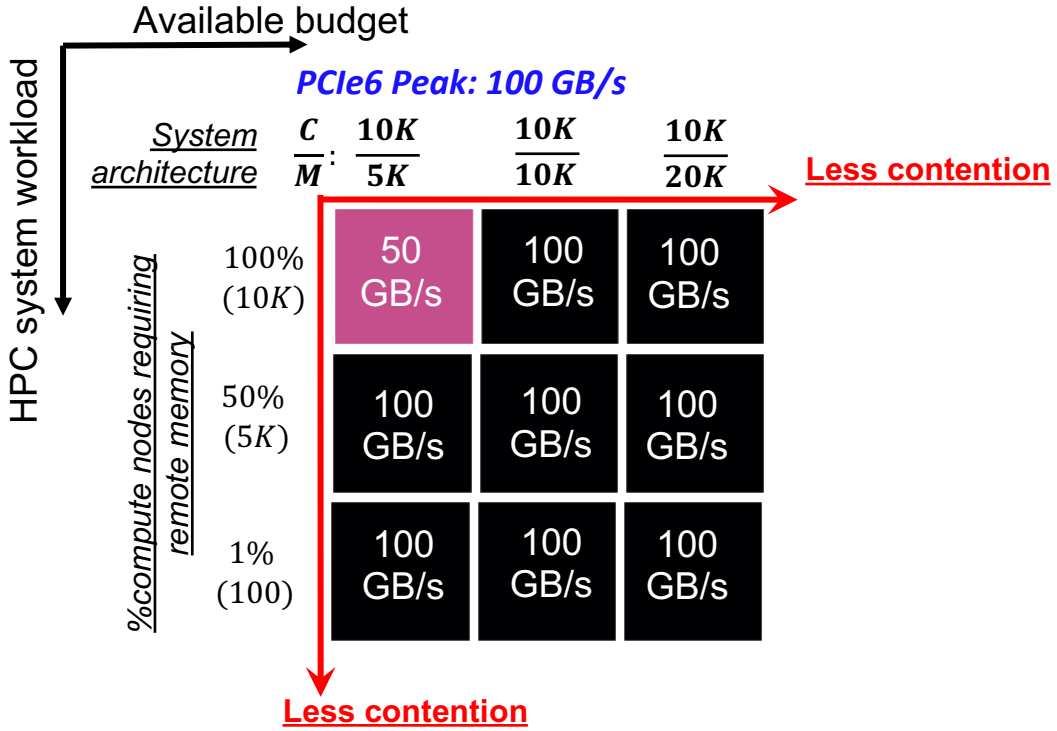

Outline

- Need-to-know about Memory disaggregation
- Disaggregated memory system architecture
  - A structured system design model
  - Budgets, workloads
  - Available remote memory capacity/bandwidth
- Characterize application performance on a disaggregated memory system
- Case Study

Memory Roofline: bound by local memory or remote memory

- FLOP Roofline: Which takes longer?
  - Data movement
  - Compute

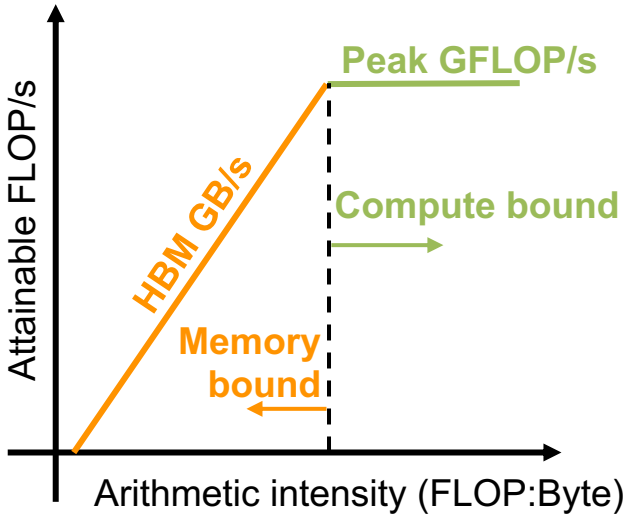

- Memory Roofline: Which takes longer?
  - Local memory (= HBM data movement)
  - Remote memory (= DRAM data movement)

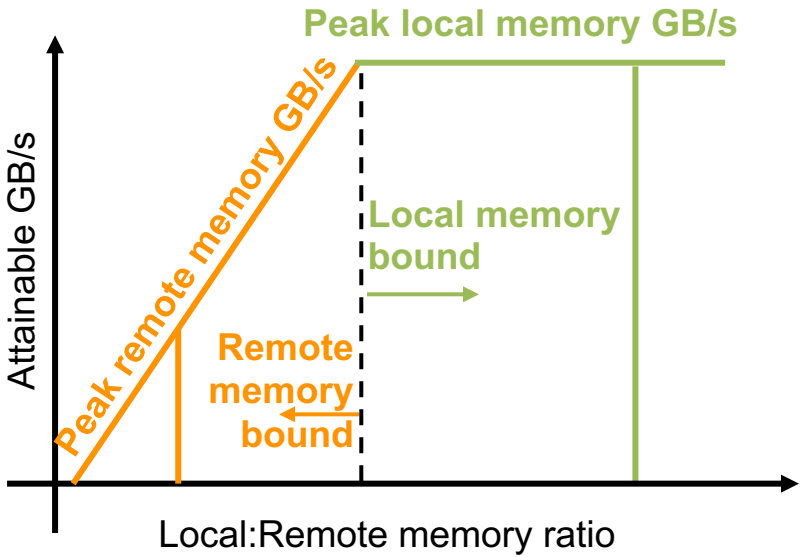

Bounded by Remote Memory Bandwidth? No, if High L:R

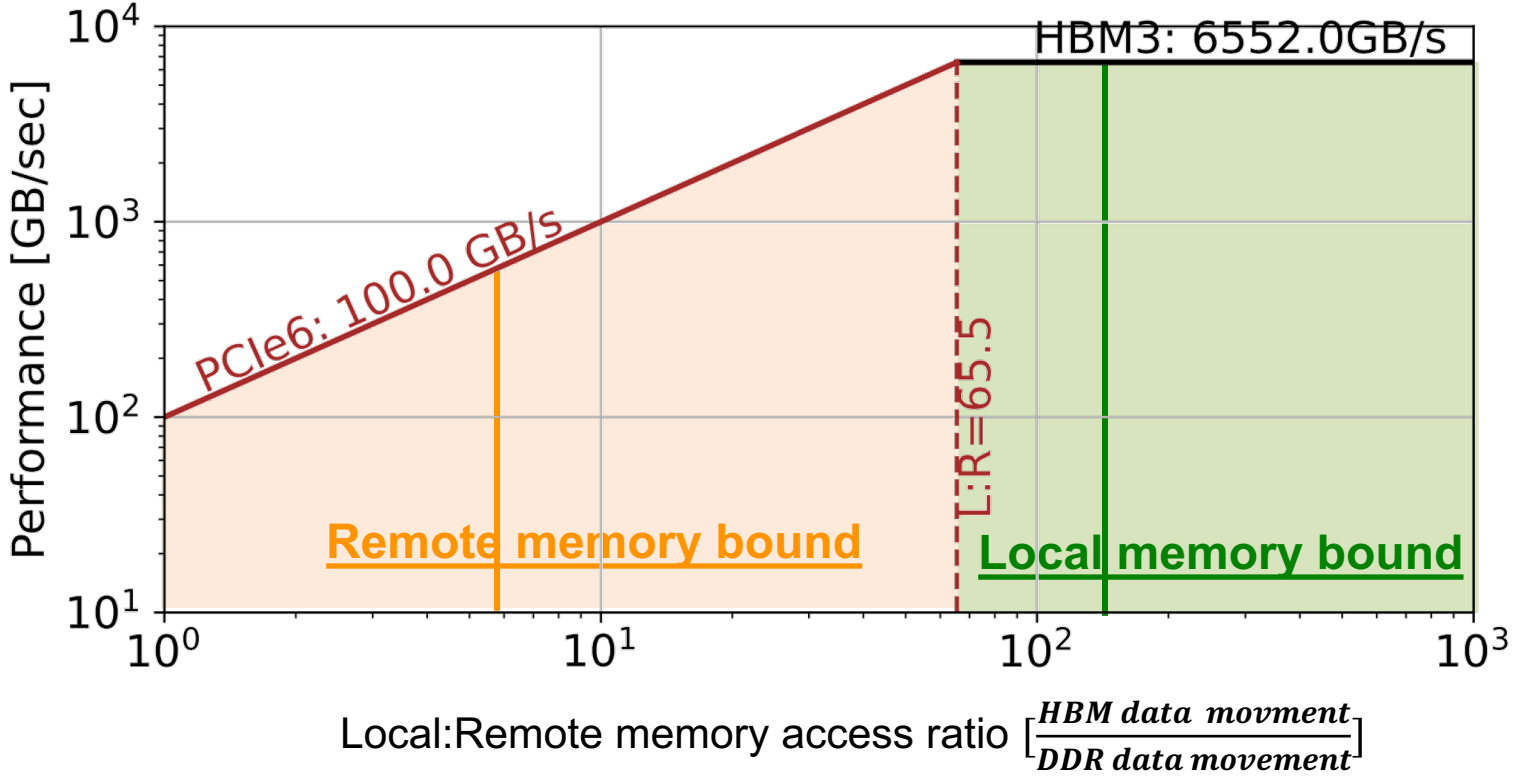

Remote Memory Access Pattern Implication on Different System Configurations [fixed C:M, vary workloads demands]

- C:M=10K:1K

	Remote memory bandwidth	Machine balance
10%	100GB/s	65.5

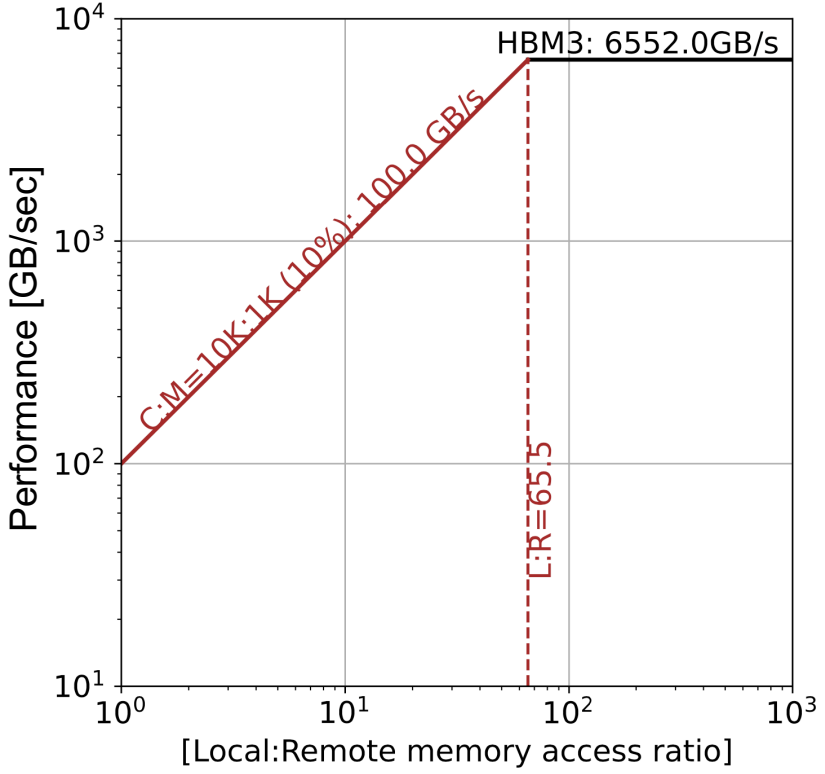

Remote Memory Access Pattern Implication on Different System Configurations [fixed C:M, vary workloads demands]

- C:M=10K:1K

Less contention

	Remote memory bandwidth	Machine balance
10%	100GB/s	65.5
40%	25 GB/s	262

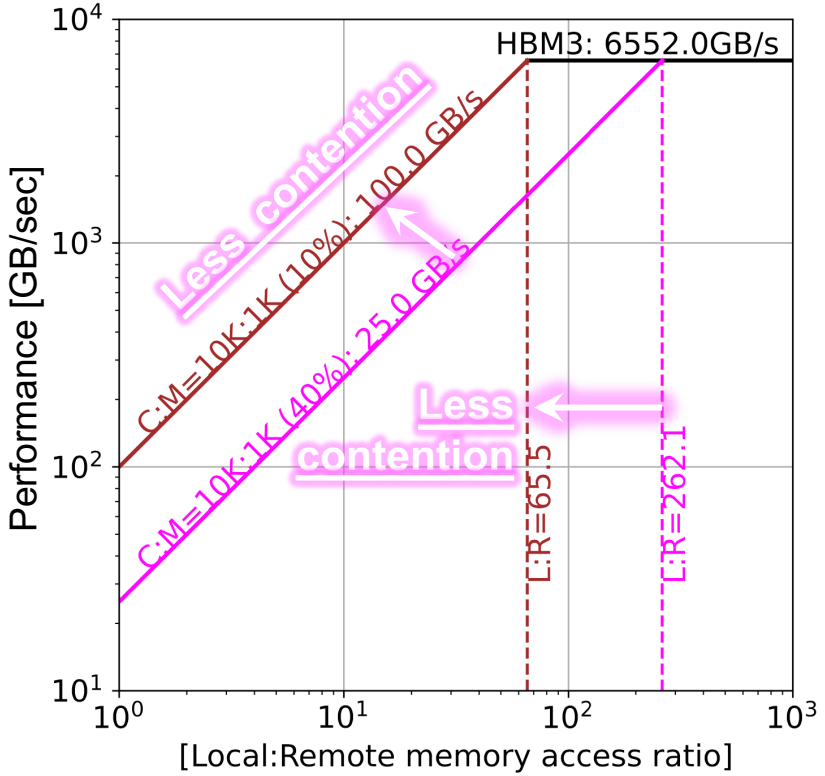

Remote Memory Access Pattern Implication on Different System Configurations [fixed workloads demand, vary C:M]

- %10 compute nodes are requiring remote memory

	Remote memory bandwidth	Machine balance
C:M=10K:1K	100GB/s	65.5

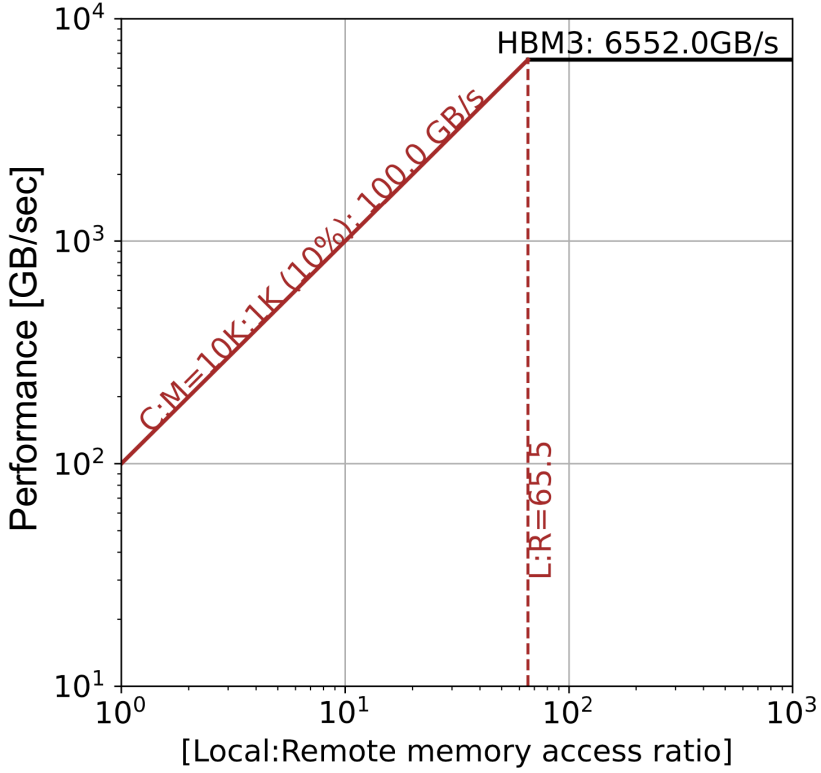

Remote Memory Access Pattern Implication on Different System Configurations [fixed workloads demand, vary C:M]

- %10 compute nodes are requiring remote memory

Less contention

	Remote memory bandwidth	Machine balance
C:M=10K:1K	100GB/s	65.5
C:M=10:0.5K	50GB/s	131

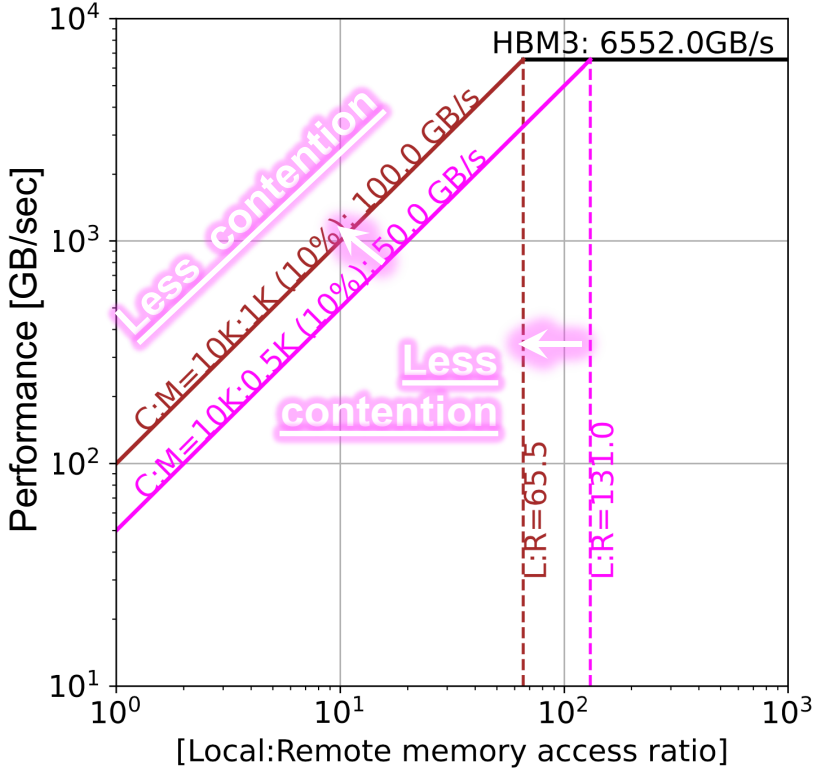

Characterize workloads in a single figure

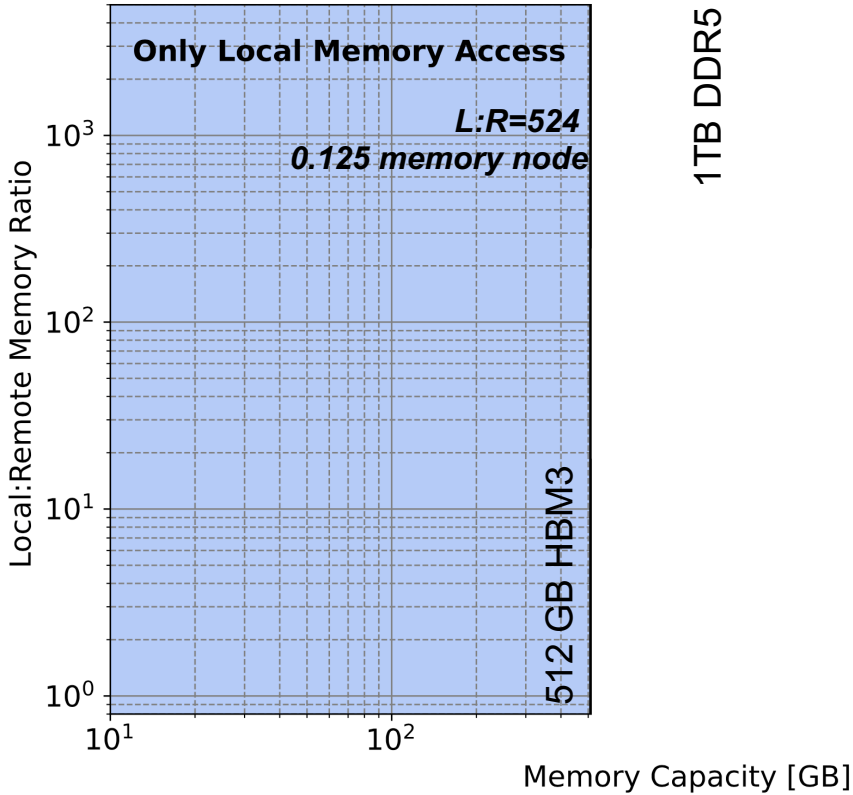

Characterize workloads in a single figure

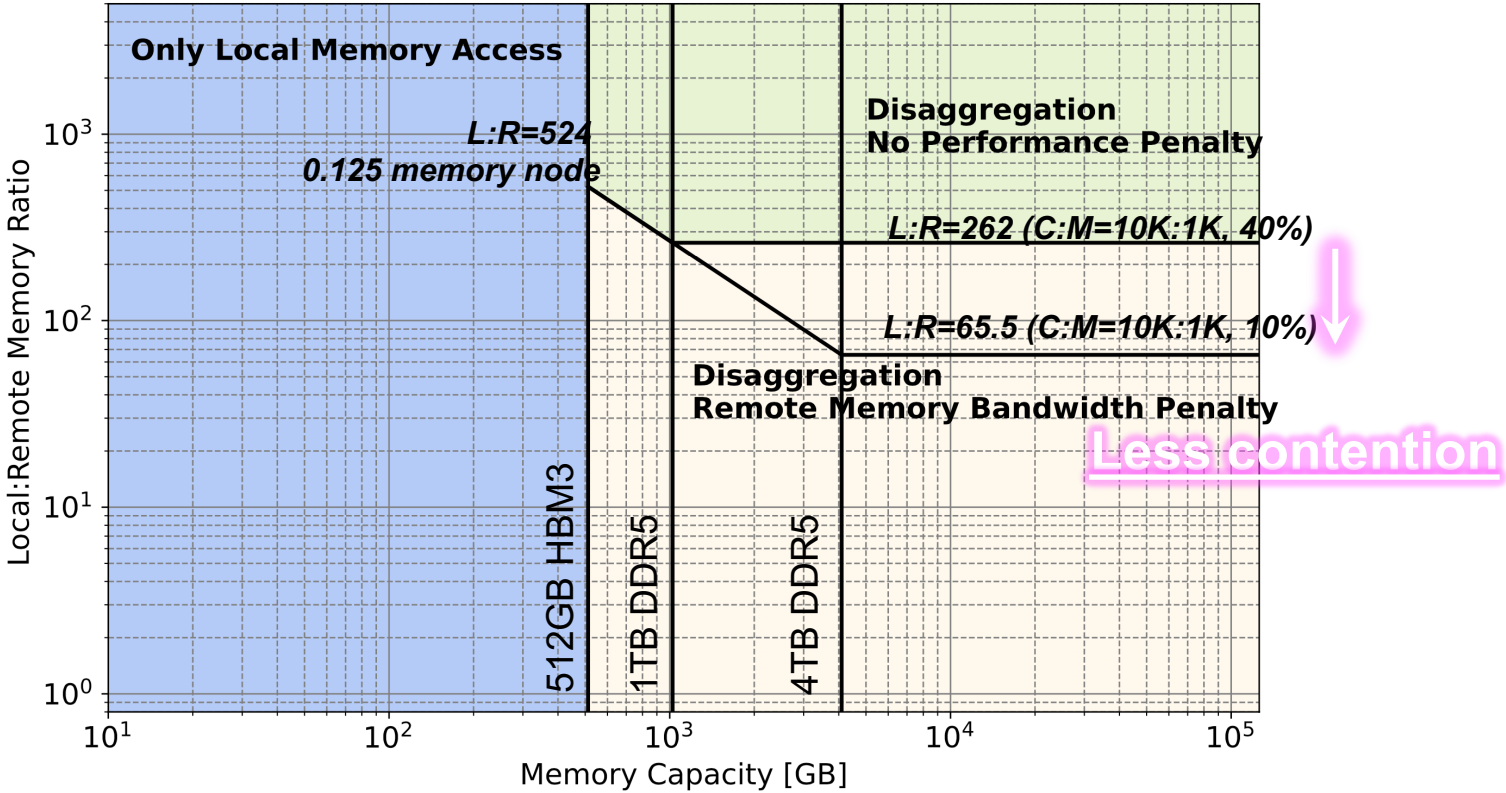

Outline

- Need-to-know about Memory disaggregation
- Disaggregated memory system architecture
- Characterize application performance on a disaggregated memory system
  - Local : Remote memory access ratio
  - Required memory capacity
  - System configurations
- Case Study

Experiment setup

- C:M=10K:1K
- 10% compute node that are requiring remote memory
- Each compute can access 4TB remote DRD5 memory on average
- Each compute can reach peak PCIe6 bandwidth 100GB/s
- L:R=65.5

Case Study: AI workloads

- The DeepCAM climate benchmark is based on the 2018 work of Kurth et al. which was awarded the ACM Gordon Bell Prize
- The L:R memory ratio is characterized by $\frac{\text{FLOP:sample Byte}}{\text{Flop:HBM Byte}}$ from Ibrahim et al.
- Training set = memory capacity (all in memory nodes)

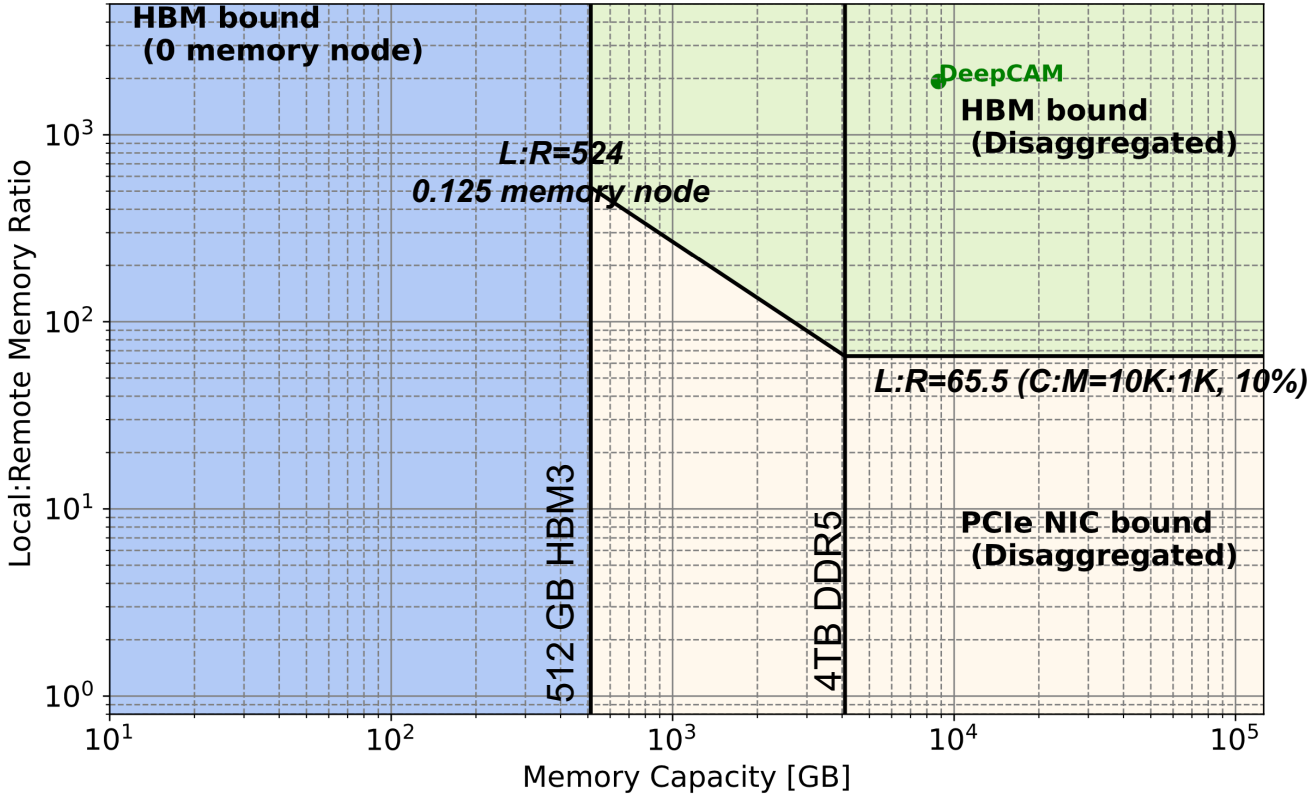

Case Study: Eleven Workloads from Five Computational Scenarios

- 10 out of 11 workloads can leverage disaggregated memory without affecting performance
- STREAM can be a proxy for giant AI=O(1) linear solvers (stencil/sparse) without any multiphysics/AMR

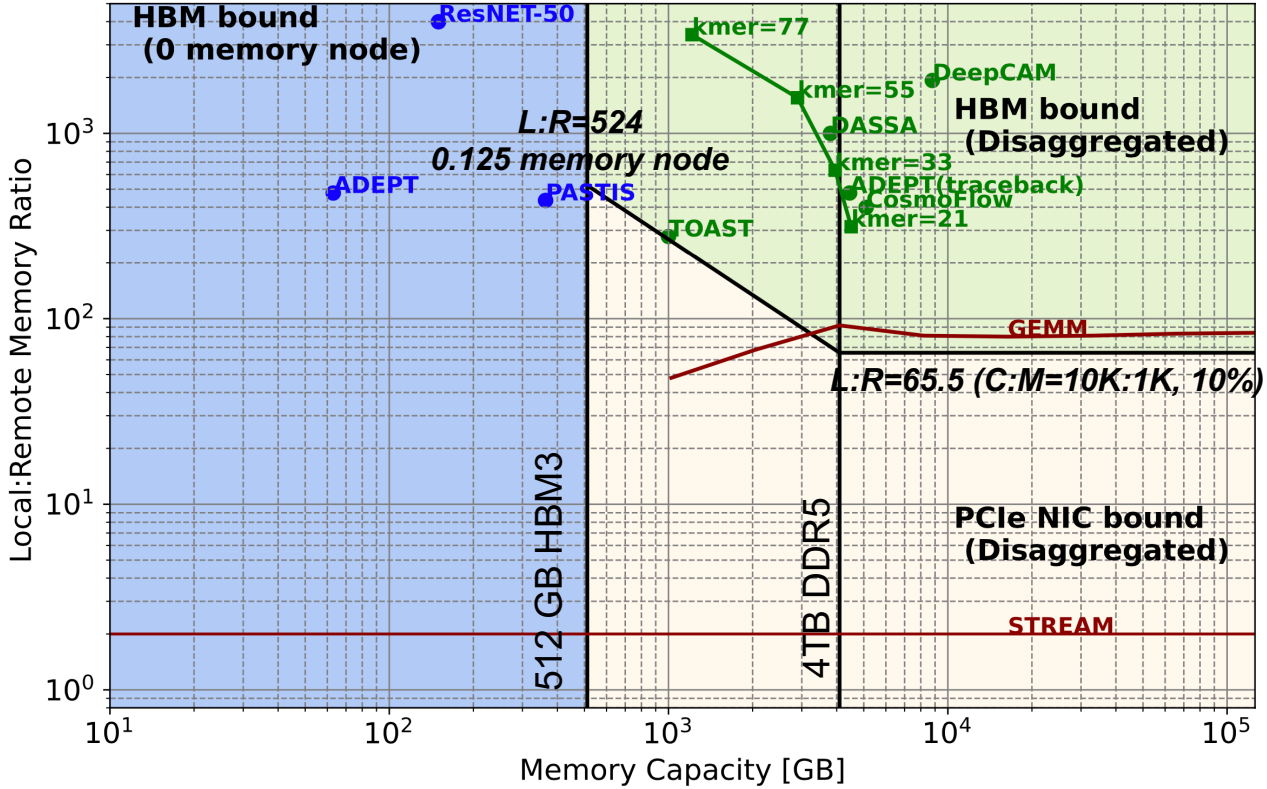

Conclusions

- A practical and intuitive approach to assess how much disaggregation is needed or viable given the technology trend and the impacts to the diverse workload
- Low PCIe bandwidth does not destroy the value of memory disaggregation, combine Local:Remote memory access ratio, memory capacity requirement
- It's promising for HPC applications benefit from disaggregated memory system
- Beneficial to wider groups: HPC architects, scientists