Multithreaded Algorithms for Approx and Exact Matching in Graphs

M. Halappanavar[1], A. Azad[2], F. Dobrian[3],
J. Feo[1], and A. Pothen[2]

[1] *Pacific Northwest National Laboratory*
[2] *Purdue University*
[3] *Conviva Inc.*

26 January, 2011

First ICCS Workshop

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Motivation: *Irregular applications*

FaceBook - 300 M users

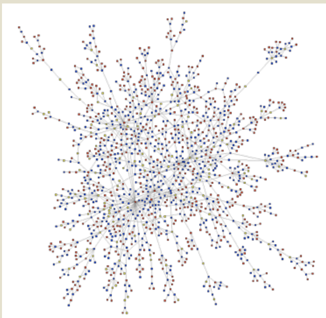

Community Activities

National Security

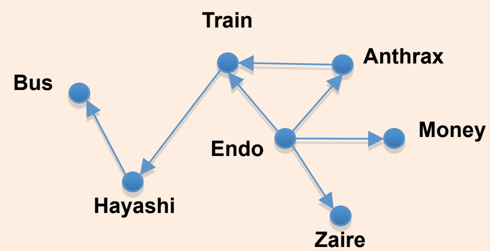

Connect-the-dots

Power Grids

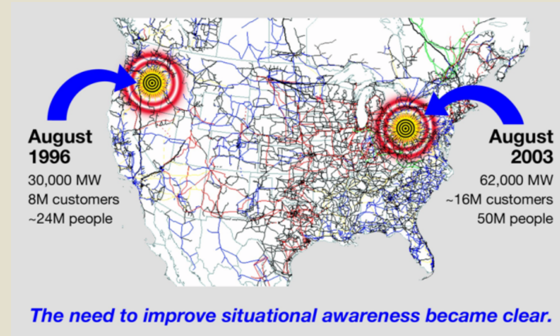

N-x contingency analysis

Blog Analysis

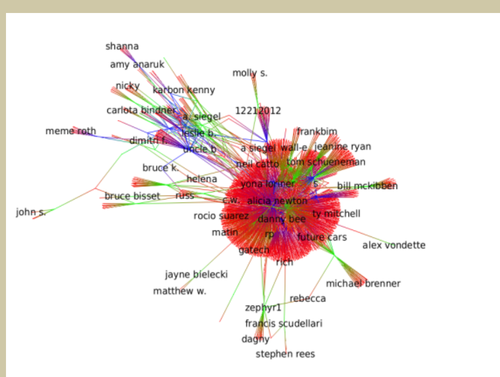

Community thought leaders

Semantic Web

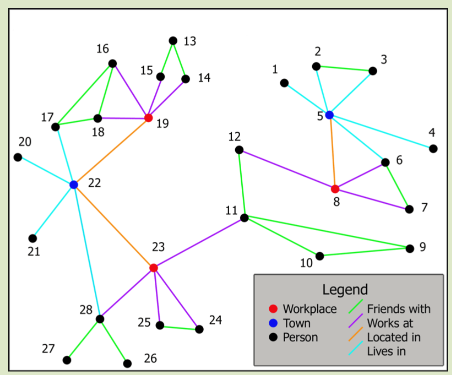

People, Places, & Actions

Security

Anomaly detection

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

Challenges

- Problem size
  - Ton of bytes, not ton of flops
- Little data locality
  - Have only parallelism to tolerate latencies
- Low computation to communication ratio
  - Single word access
  - Threads limited by loads and stores
- Synchronization points are simple elements
  - Node, edge, record
- Work tends to be dynamic and imbalanced
  - Let any processor execute any thread

Key Architectural Features

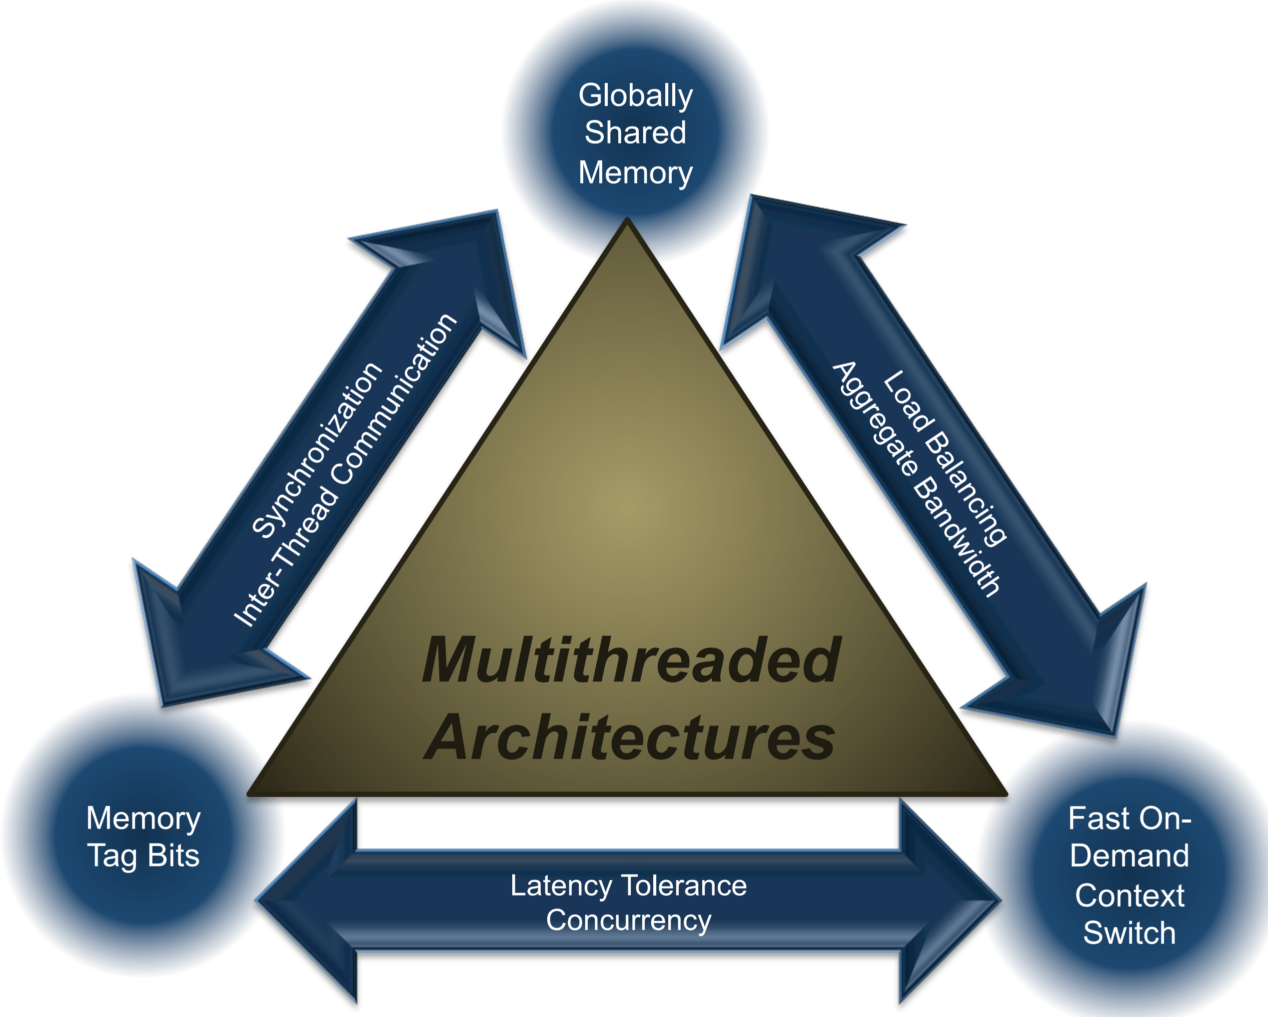

Source: Jace Mogill, PNNL

Overview

XMT

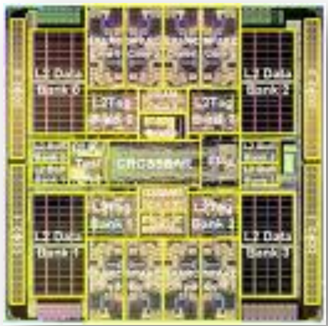

Niagara-2

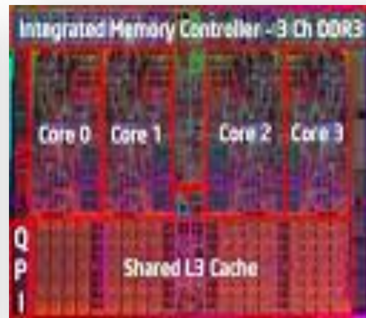

Nehalem

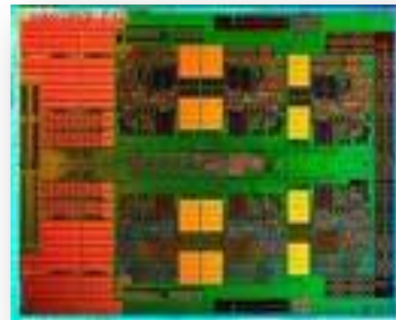

Magny-Cours

Multithreading ⟷ **Caching**

Approx Algorithms:

- Queue-based
- Q + Sorting
- Dataflow

Exact Algorithms

Input:

- RMAT-ER
- RMAT-G
- RMAT-B

Matching

- A matching M is a subset of edges such that no two edges in M are incident on the same vertex
- Maximum matching maximizes some function
  - Number of edges matched (cardinality)
  - Sum or product of edge-weights

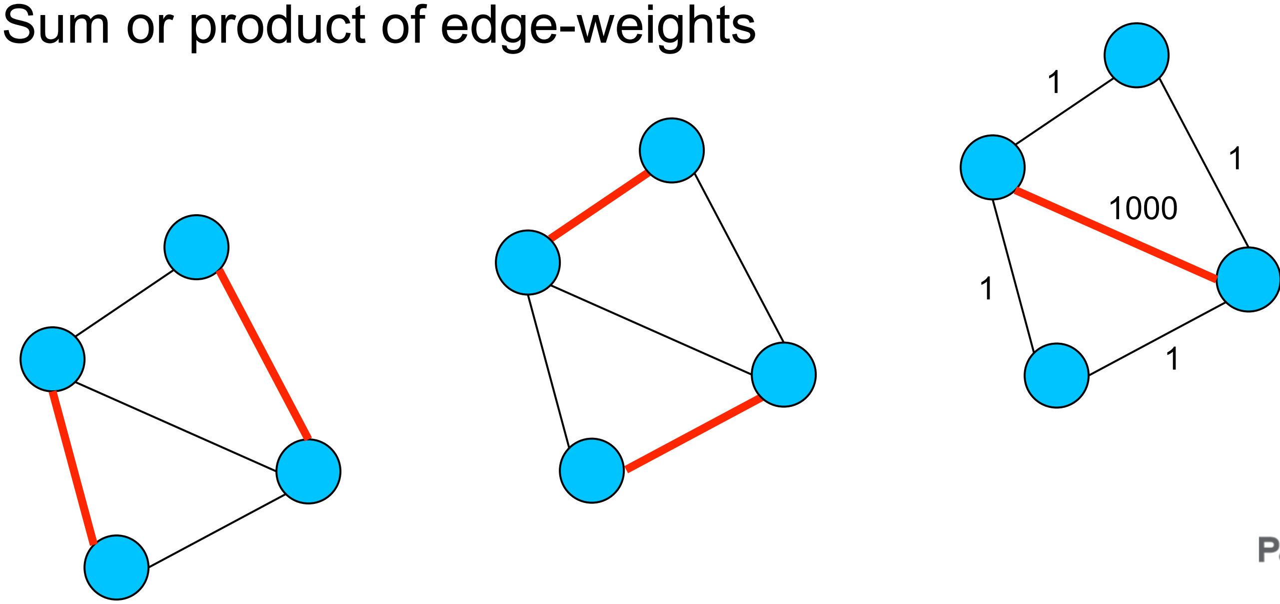

Applications of matching

- Sparse linear solvers
- Block triangular form
- Graph partitioners
- Bioinformatics
- Web technology
- High speed network switching
- …

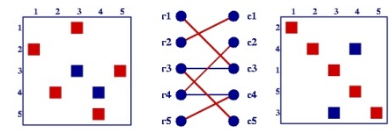

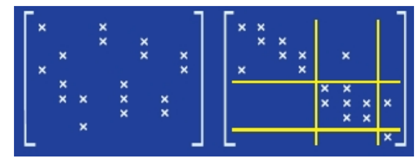

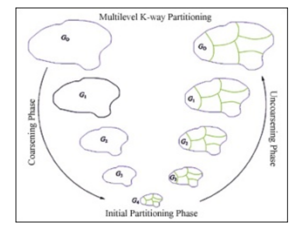

Algorithms

- Exact Algorithms:
  - Polynomial time algorithm first due to Edmonds
  - Maximum matching: Hopcroft-Karp
  - Maximum weighted: Hungarian
- (half) Approx Algorithms:
  - Sorting-based approaches (Global)
  - Search-based approaches (Local)
    - Preis's algorithm and its variants (Hoepman; Manne and Bisseling)

Pointer-based algorithm (Queue-based)

- Identify locally-dominant edges using pointers
- Implement with queues (queue matched vertices)
- Variant: sorted edge-sets

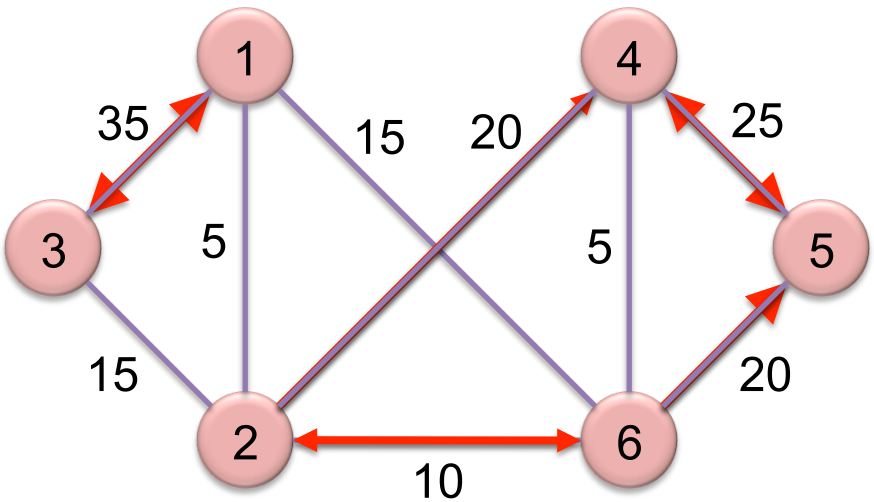

Pointer-based algorithm (Dataflow)

- Queue headers can hotspot
- Dataflow approach

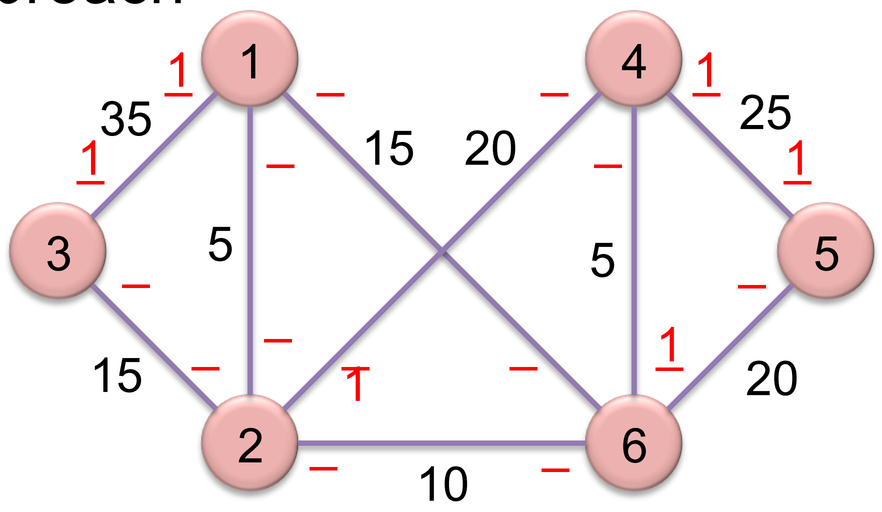

- Each node sets signal on its side of heaviest edge to 1
- Reads companion signal

Dataflow (cont.)

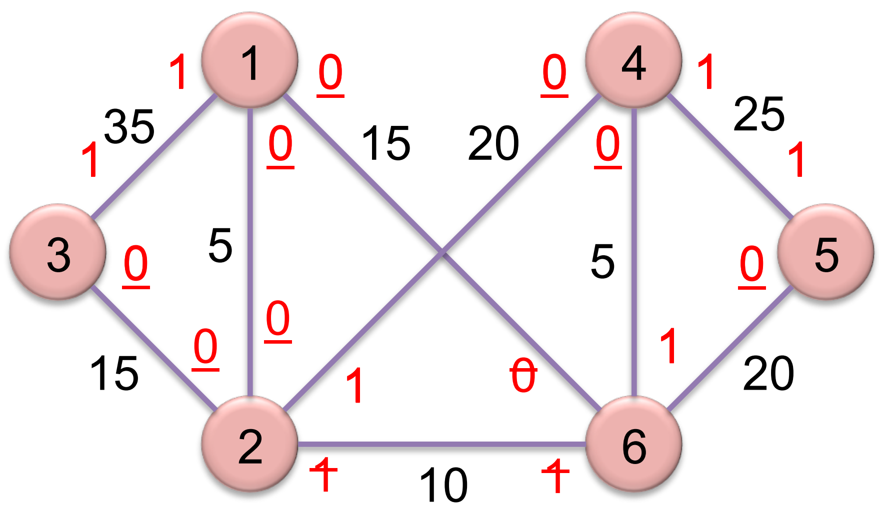

- If companion signal is 1, then set signal of other edges to 0 and stop
- else set signal on next heaviest edge to 1

Overview

XMT

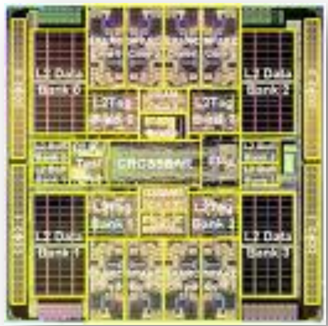

Niagara-2

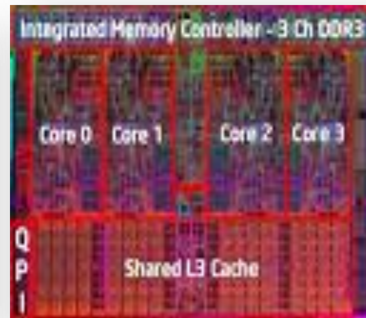

Nehalem

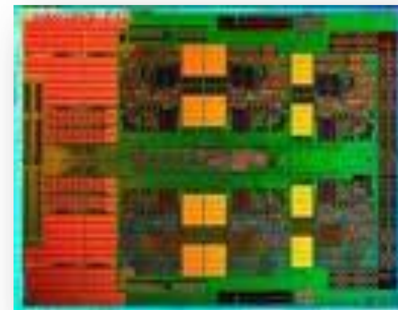

Magny-Cours

Multithreading ⟷ **Caching**

Approx Algorithms:

- Queue-based
- Q + Sorting
- Dataflow

Exact Algorithms

Input:

- RMAT-ER
- RMAT-G
- RMAT-B

Cray XMT: A block view

- Threadstorm Processor:
  - 500 MHz
  - 128 thread-streams
  - VLIW
- 8 GB/proc
- 3D Torus Interconnect

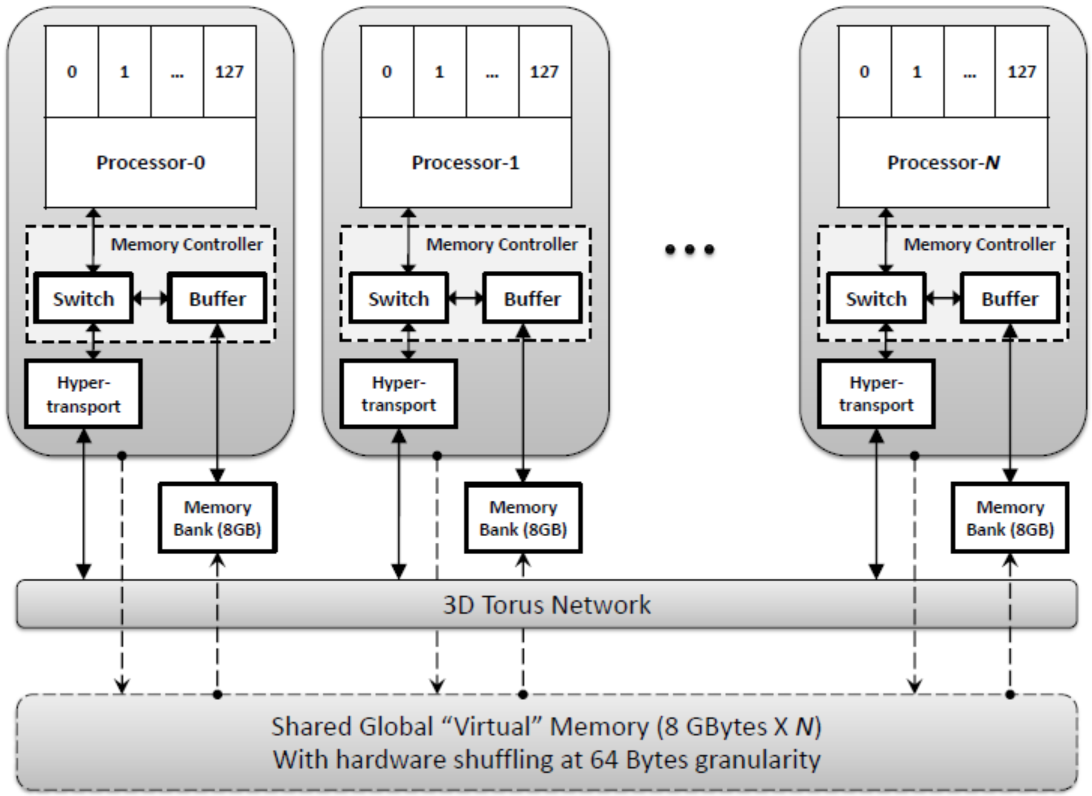

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Cray XMT: Memory

Physically distributed, globally addressable

- 8 GB/proc
- Total = 1TB (128p)
- Byte addressable
- H/w hashing
- 64Byte granularity
- Worst-case latency is 1000 cycles
- Sustainable 60 Megawords/s per processor

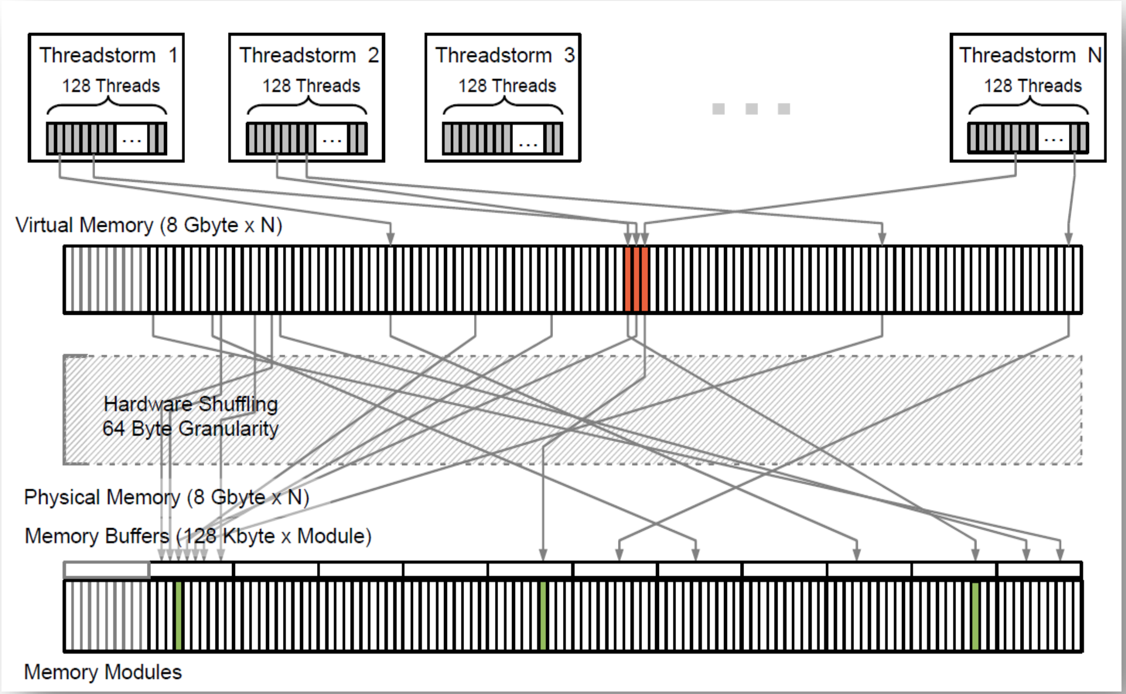

Datasets: Synthetic data with R-MAT

- R-MAT: Recursive MATrix method

- Experiments
  - RMAT-ER (0.25, 0.25, 0.25, 0.25)
  - RMAT-G (0.45, 0.15, 0.15, 0.25)
  - RMAT-B (0.55, 0.15, 0.15, 0.15)

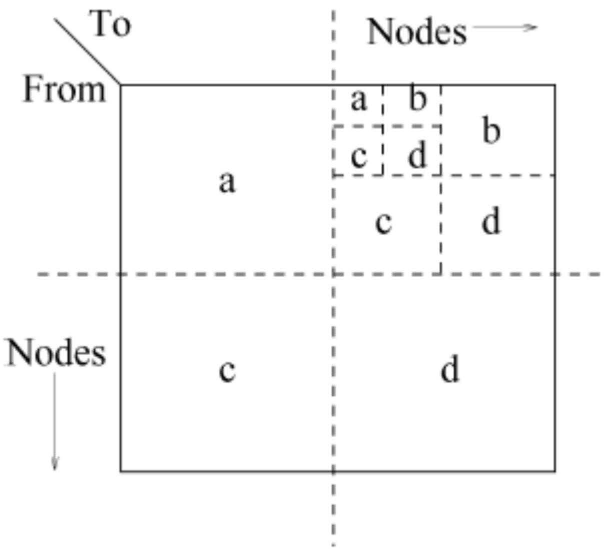

Chakrabarti, D. and Faloutsos, C. 2006. Graph mining: Laws, generators, and algorithms. *ACM Comput. Surv.* 38, 1 (Jun. 2006), 2.

Datasets for experiments

Graph	No. Vertices	No. Edges	Avg. Degree	Max Degree	Variance	Avg. Clustering Coeff.
RMAT-ER	16,777,216	134,217,654	16	42	16.01	$10.0e-7$
RMAT-G	16,777,216	134,181,095	16	1,278	415.72	$12.0e-6$
RMAT-B	16,777,216	133,658,229	16	38,143	8,085.64	$34.3e-5$

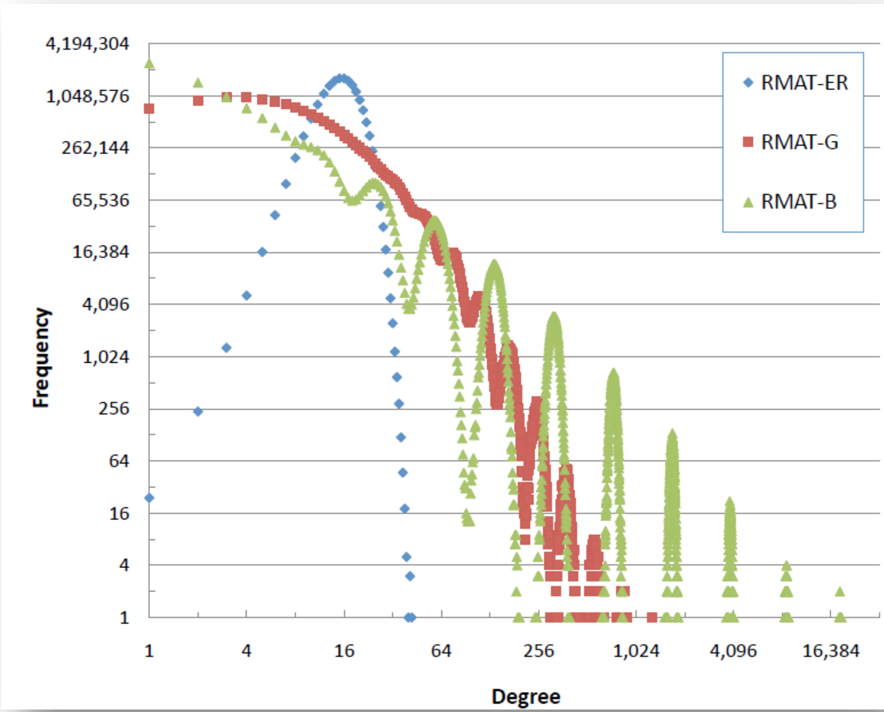

Degree distribution

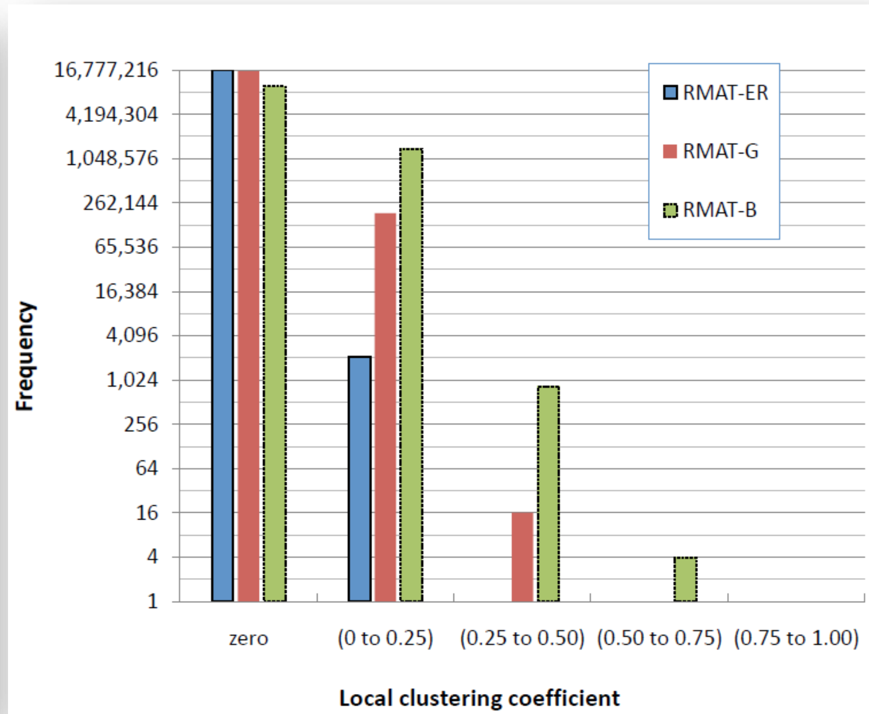

Clustering coefficient

Experimental Results

- ½-approx algorithm
- Magny-cours, Nehalem, Niagara-2, XMT
- RMAT-B

Matching: Cardinality

Graph	Init. (% of final card)	Final (% of |V|)
RMAT-ER	53.14%	94.12
RMAT-G	46.33%	81.70%
RMAT-B	36.06%	44.24%

Matching: Queue status

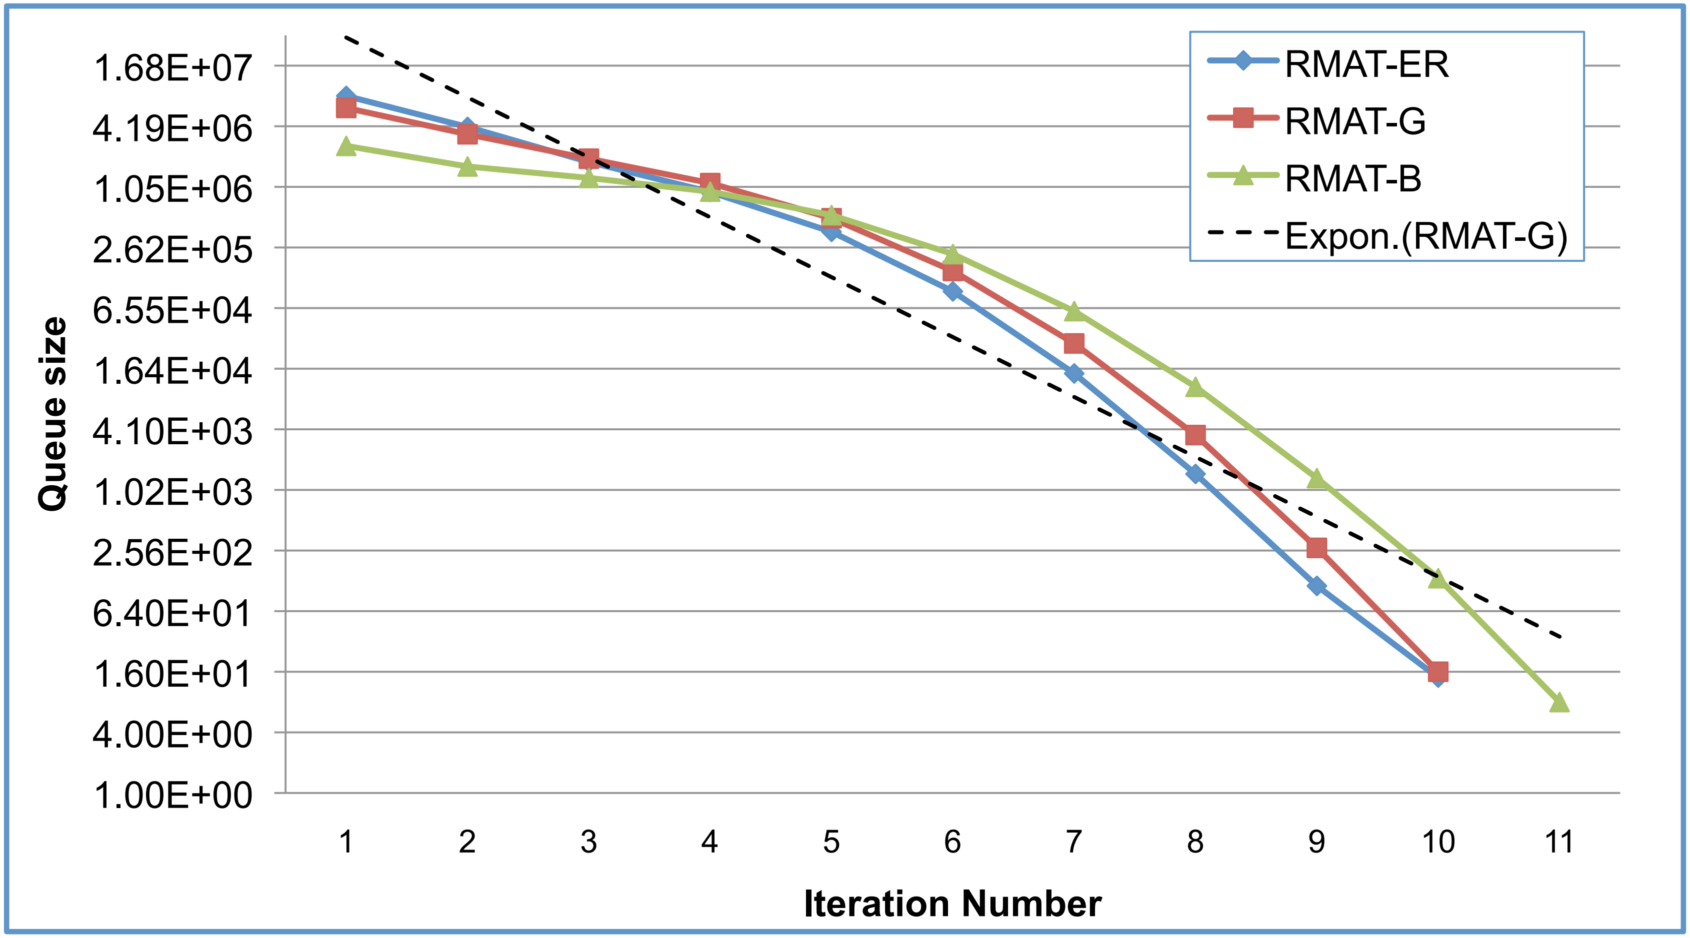

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Strong scaling: Nehalem & Magny-Cours

RMAT-B

Nehalem

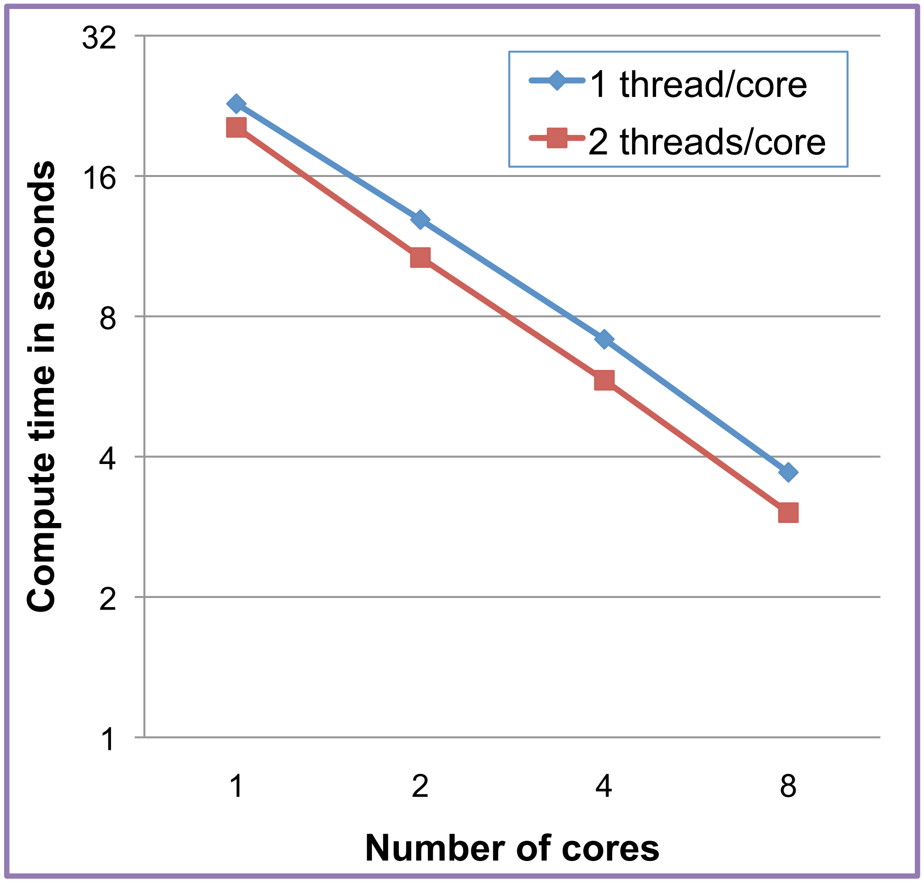

Magny-Cours

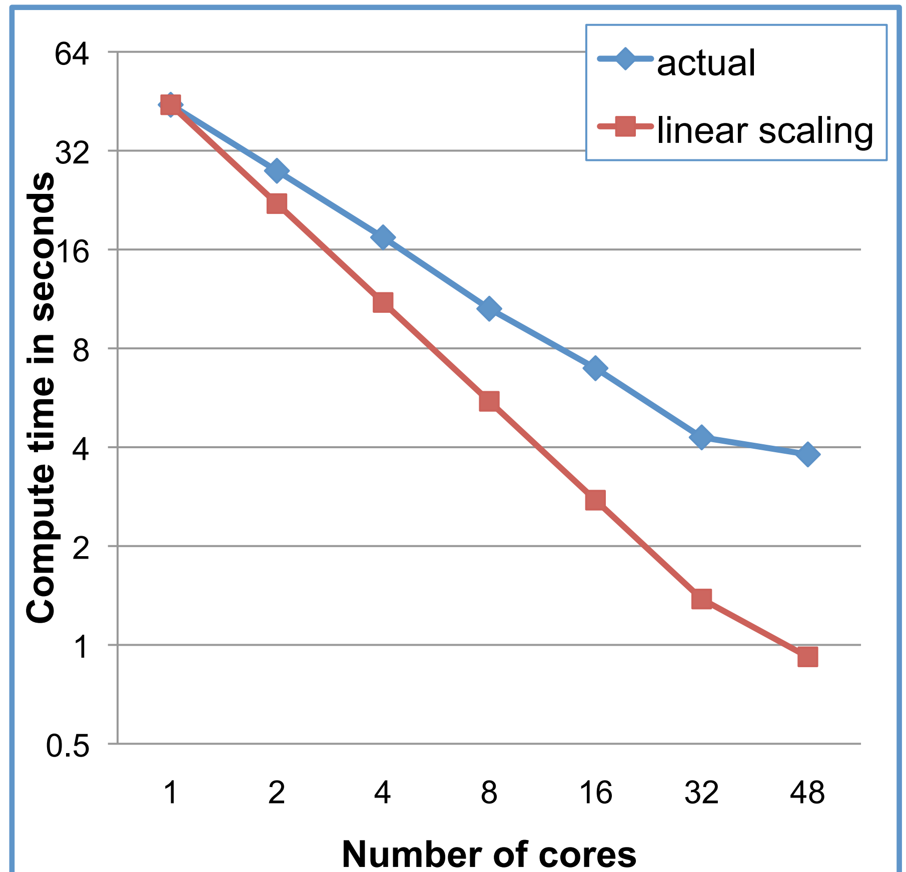

Algorithm: Queue-based

Pacific Northwest NATIONAL LABORATORY

Strong scaling: Nehalem & Niagara-2

RMAT-B

Nehalem

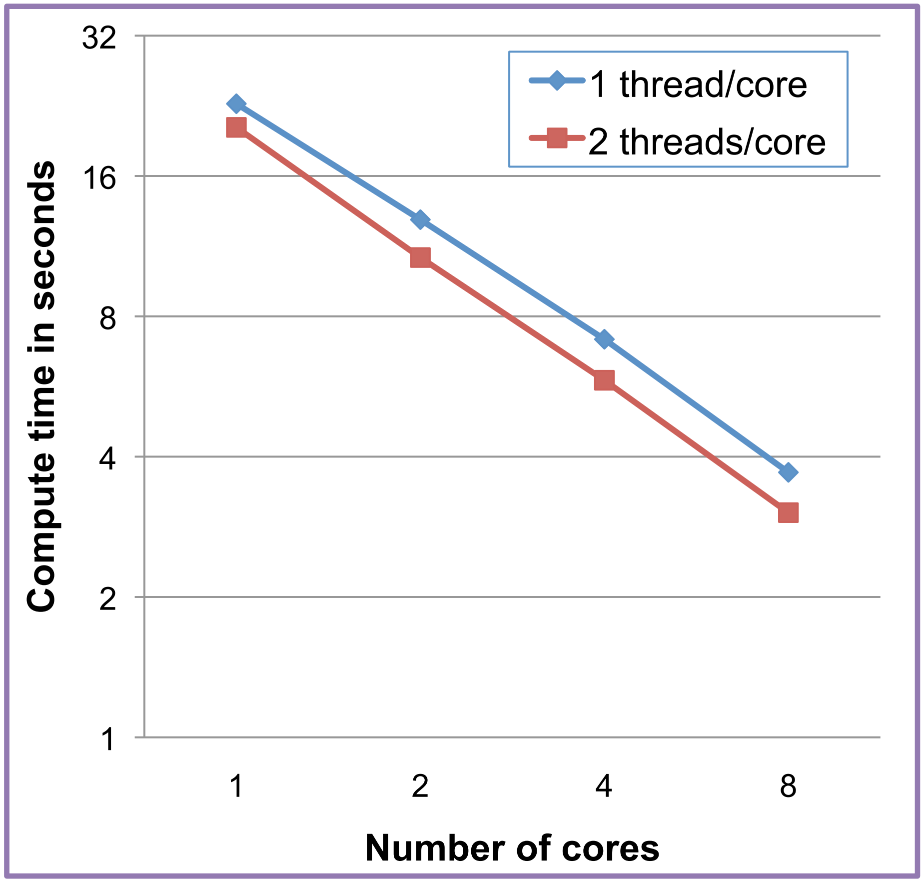

Niagara-2

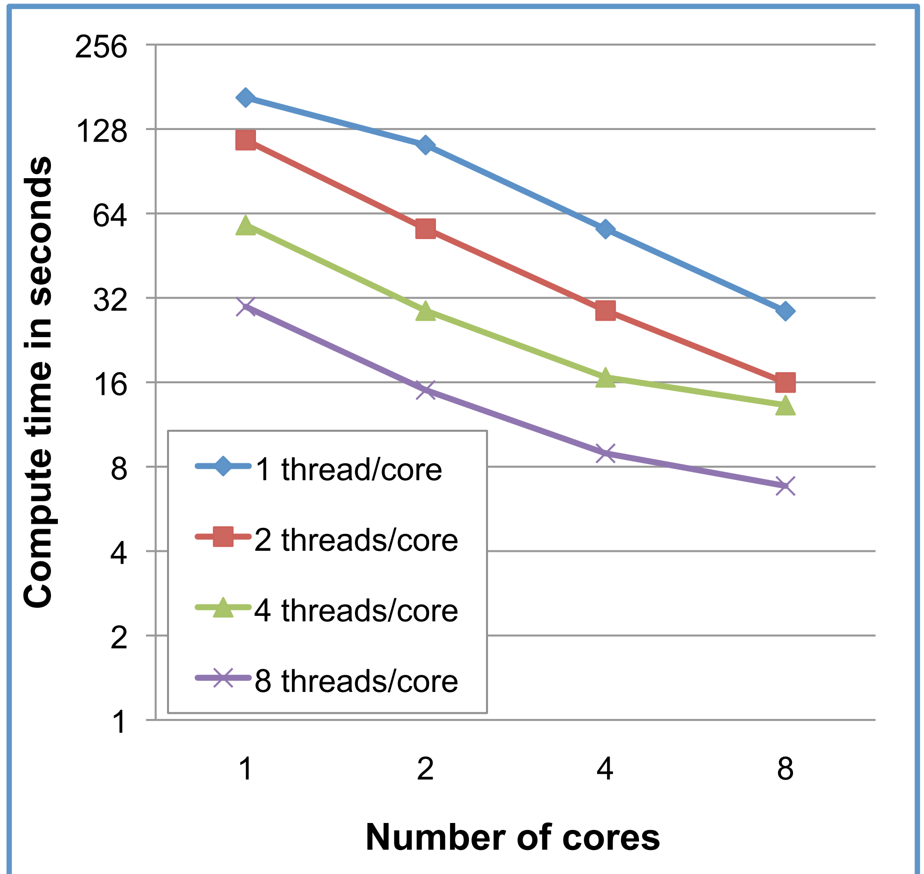

Algorithm: Queue-based

Strong scaling: Nehalem & XMT

RMAT-B

Nehalem

XMT

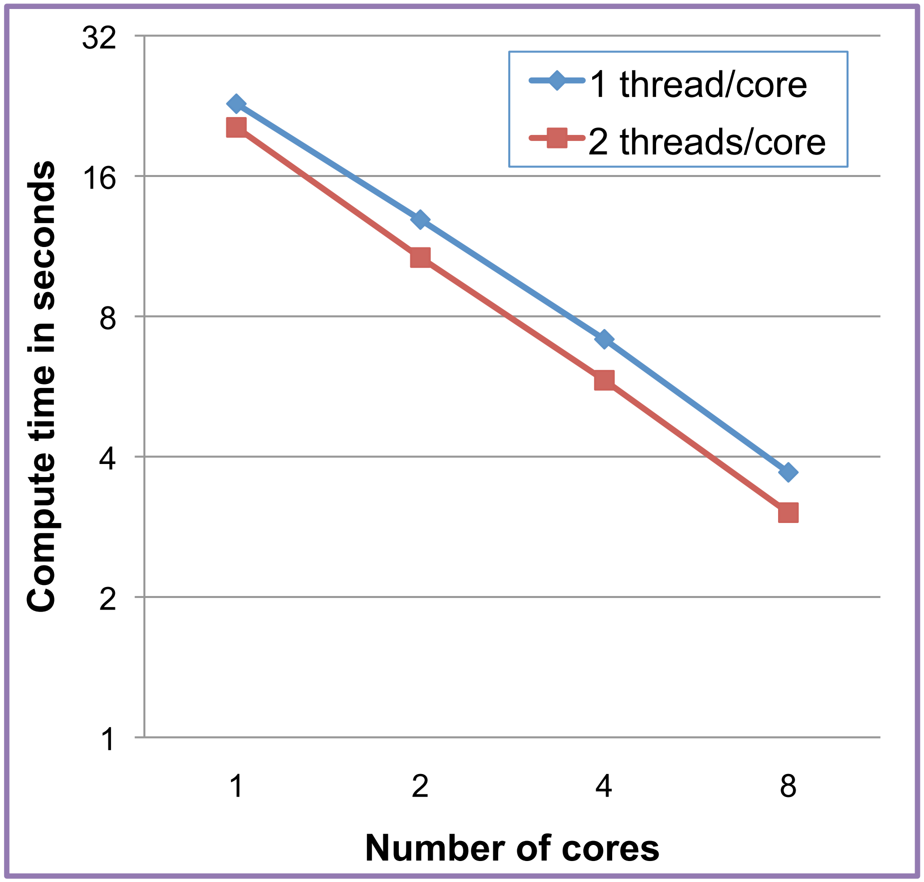

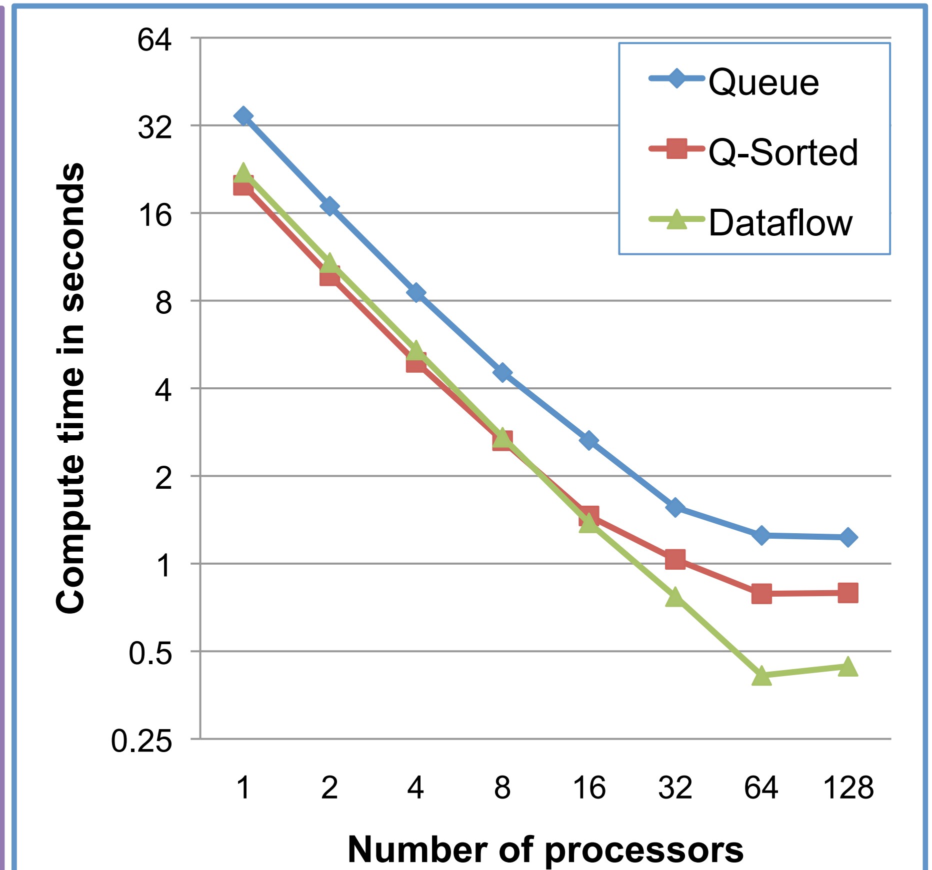

Algorithm: Queue; Q-Sorted; Dataflow

Strong scaling: XMT

RMAT-ER

RMAT-B

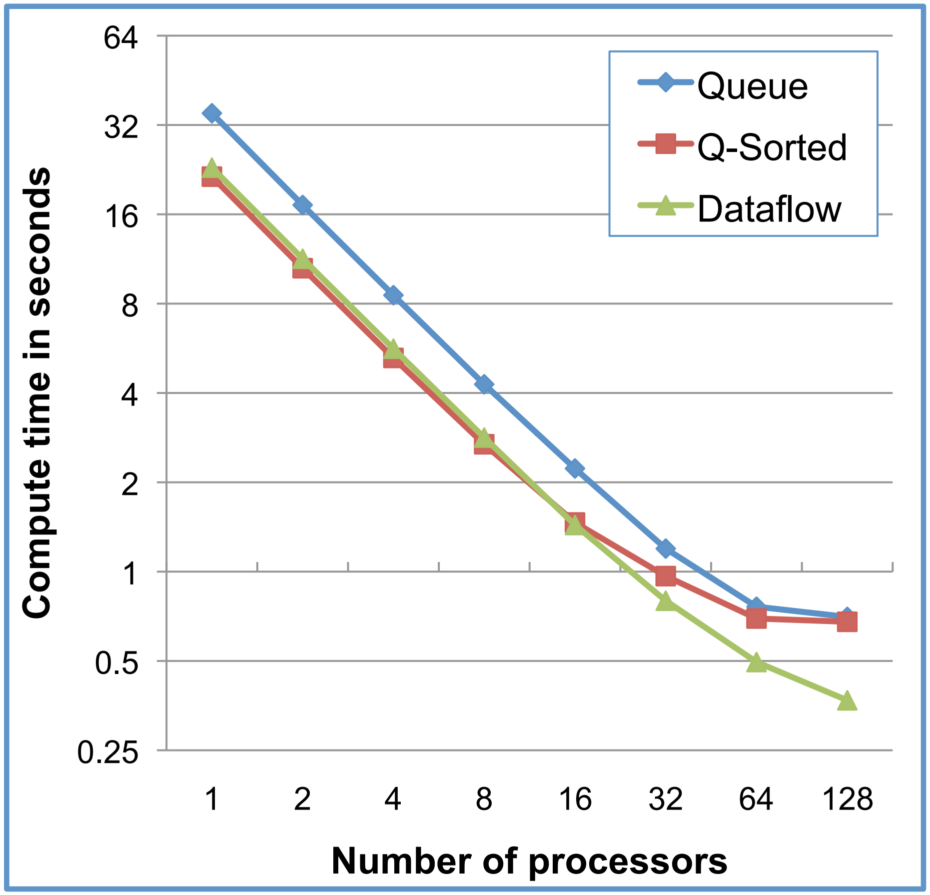

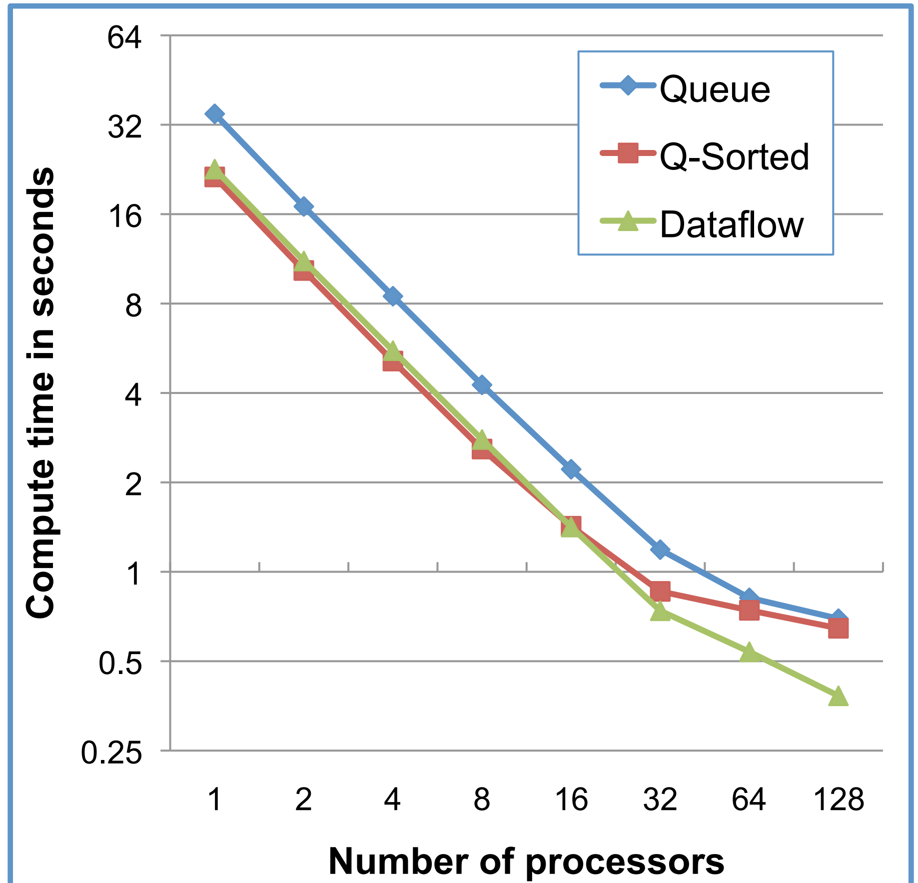

Algorithm: Queue; Q-Sorted; Dataflow

Exact matching

- Augmentation-based approach
- Single-path v/s Multiple-path
- Hopcroft-Karp algorithm:
  - Breadth-first + Depth-first
  - Dynamic: amount and type of parallelism
  - Nested loop structure
- Our approach:
  - Different locking policies (first-visited, last-, random)
  - Disjoint forest (merge BFS+DFS)
- Future: Use futures :-)

Summary & conclusion: The trinity

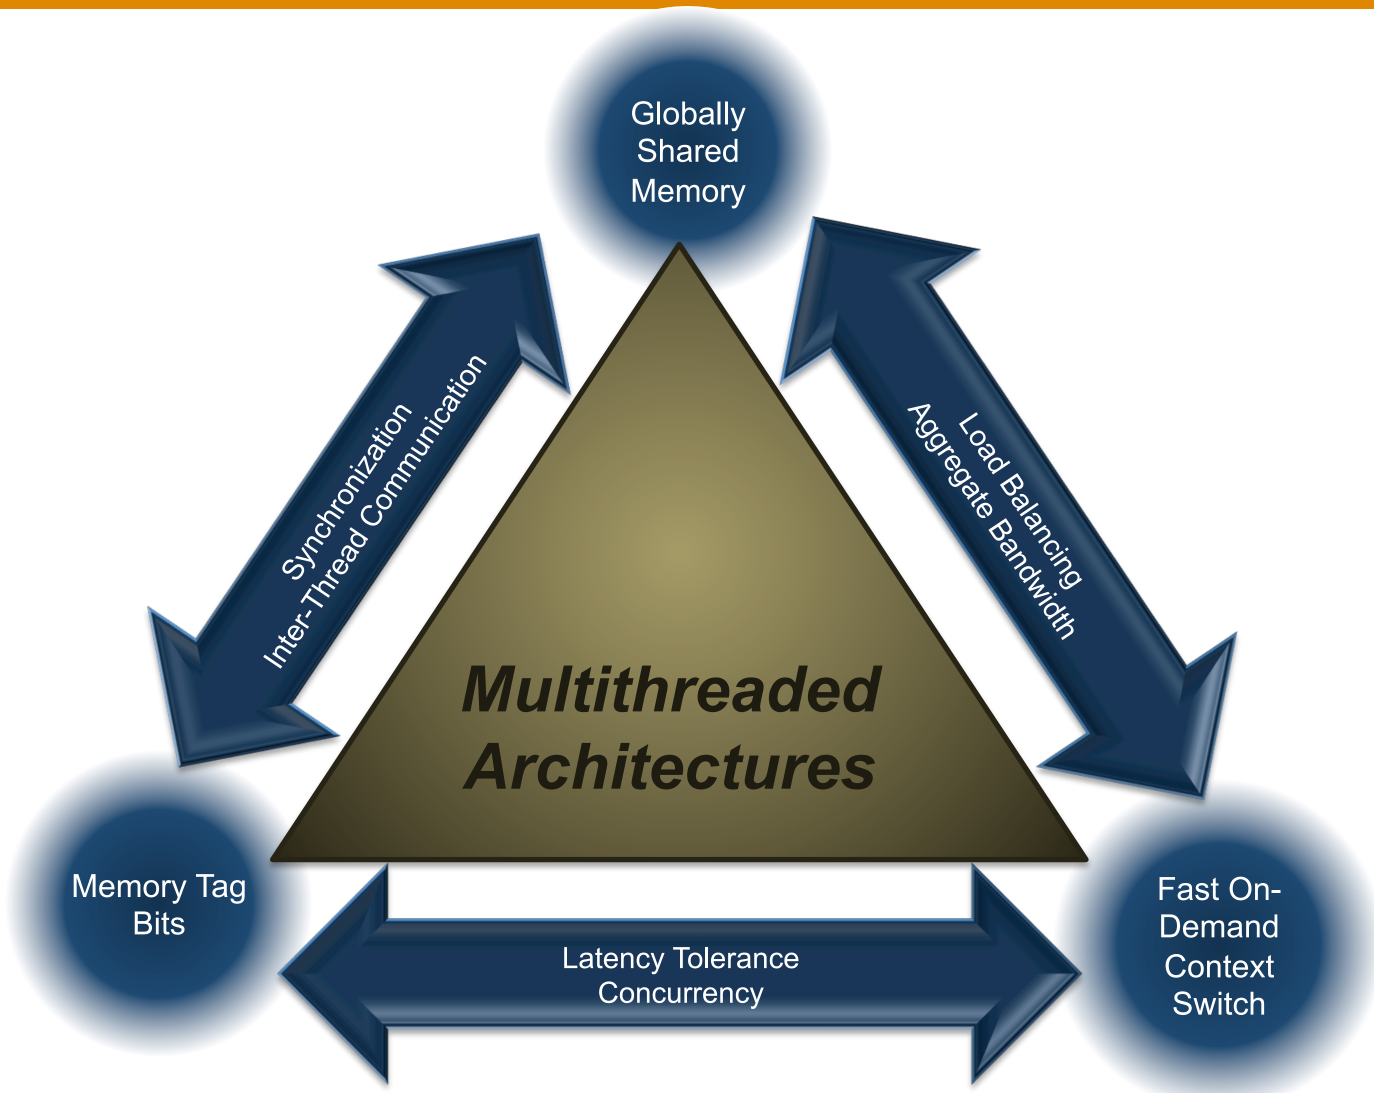

http://cass-mt.pnl.gov/

Thank You!

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*