DFS: A Simple to Write Yet Difficult to Execute Benchmark

Richard Murphy, Jonathan Berry, William McLendon, Bruce Hendrickson
Sandia National Laboratories

Douglas Gregor and Andrew Lumsdaine
Indiana University

{rcmurph,jberry,wcmclen,bahendr}@sandia.gov
{dgregor,lums}@cs.indiana.edu

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

The von Neumann Bottleneck and the Demon Moore

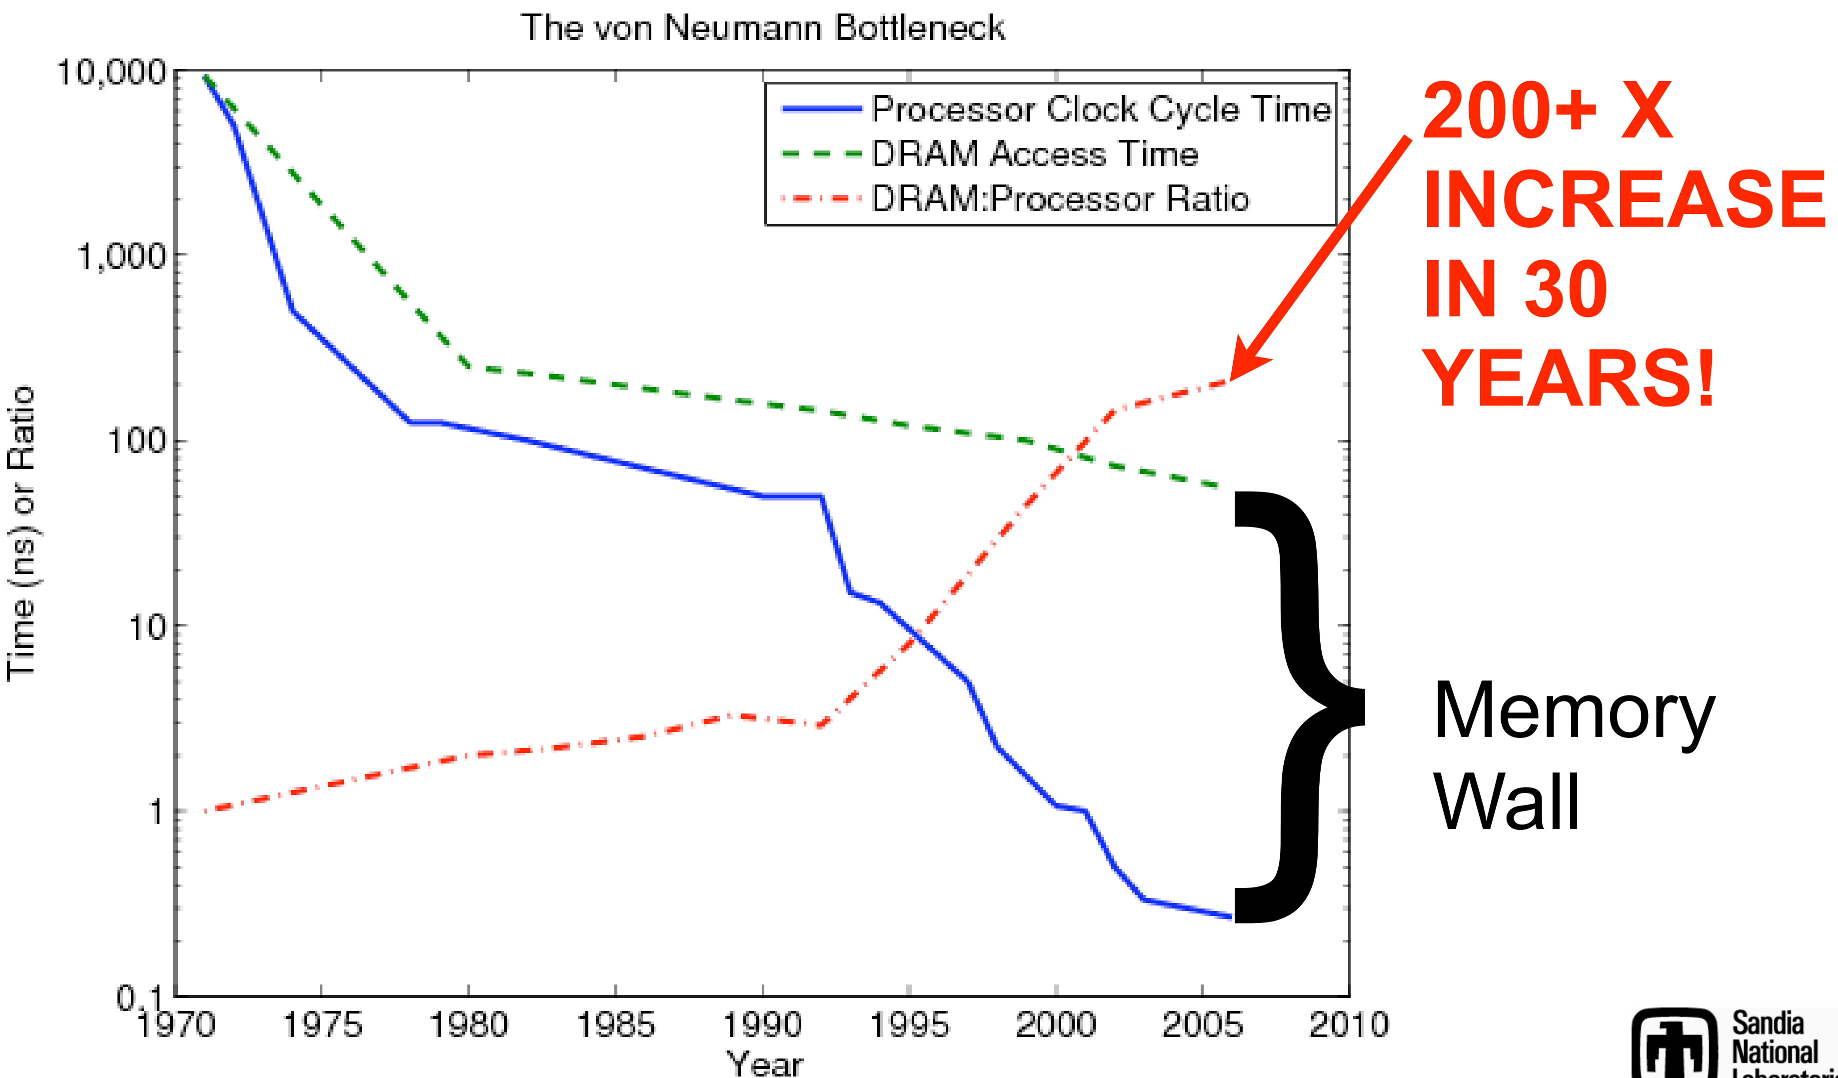

Sandia National Laboratories

The Computer Architect's View of Memory

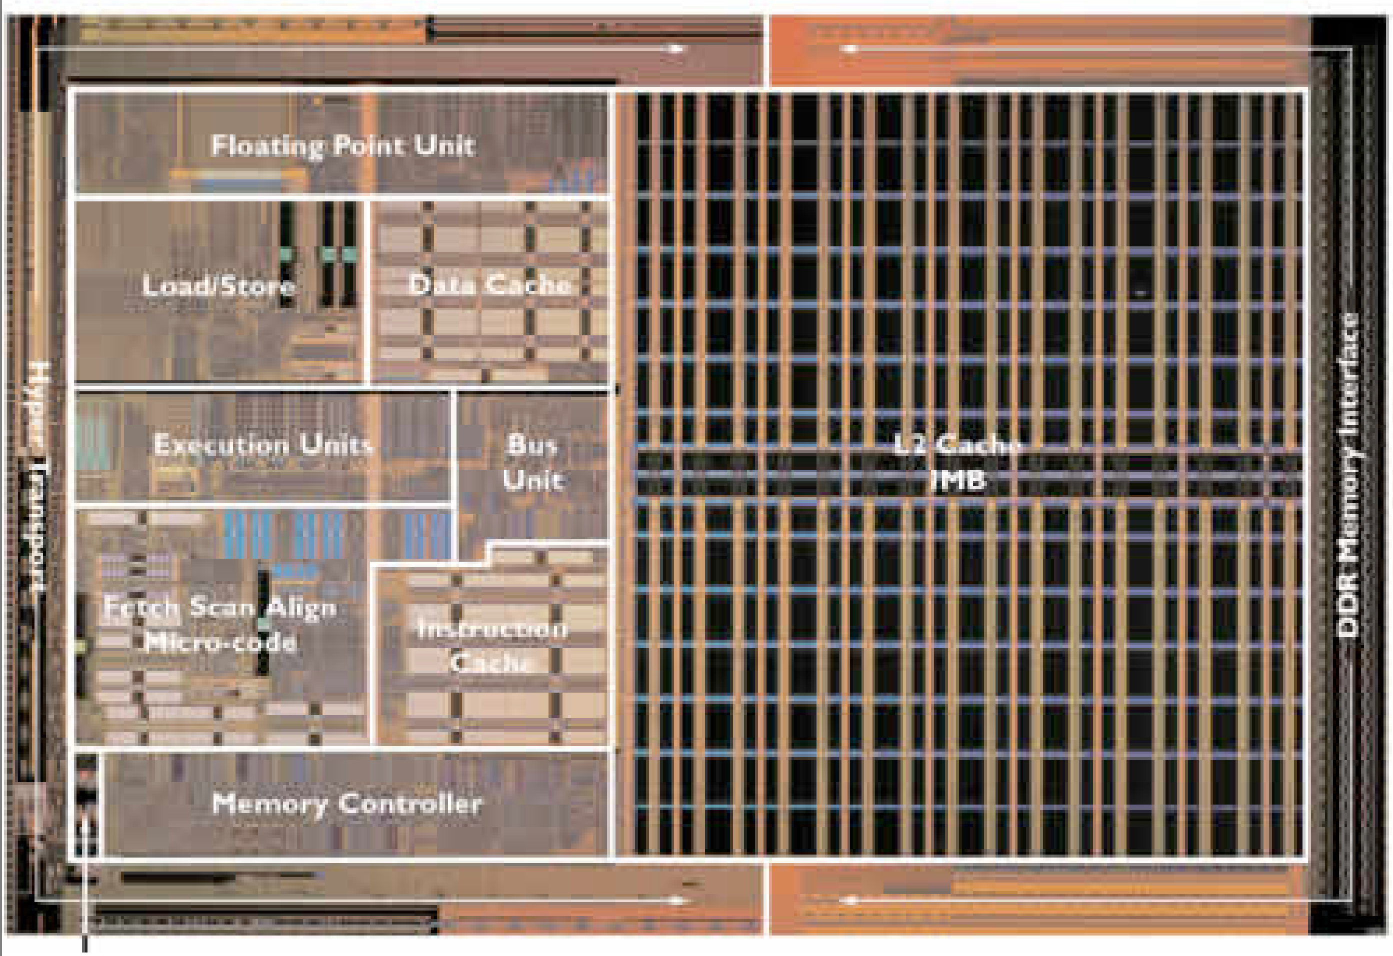

Sandia National Laboratories

The Computer Architect's View of Memory

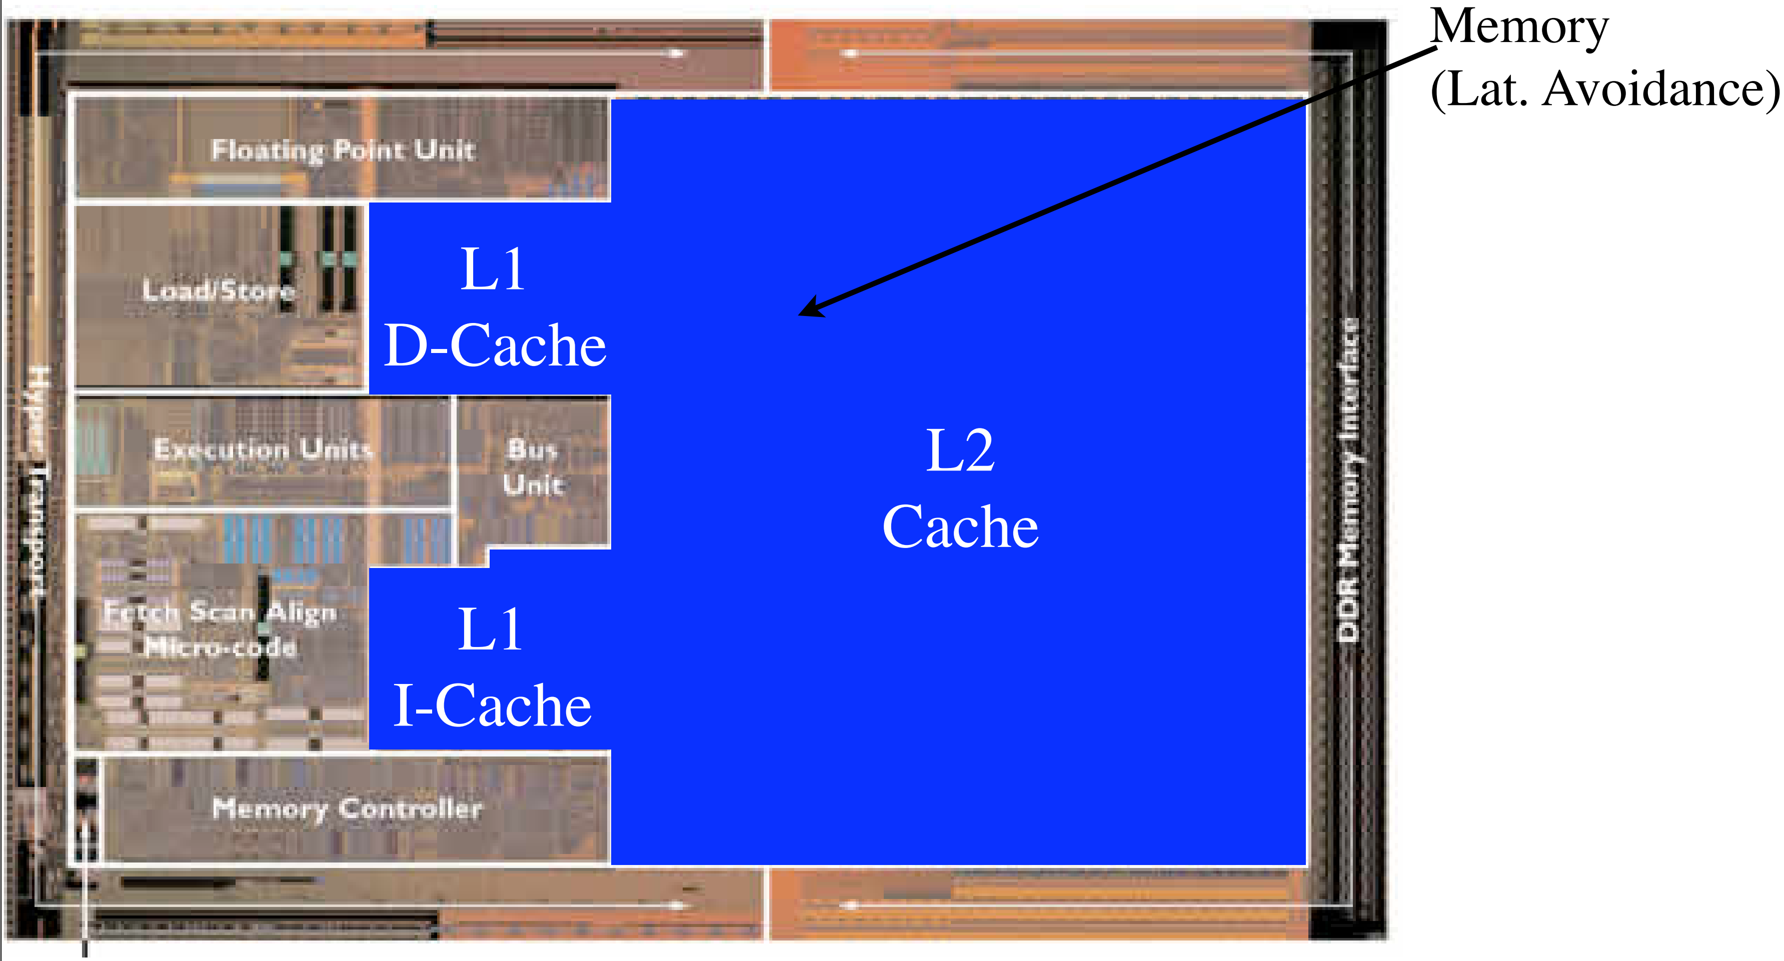

The Computer Architect's View of Memory

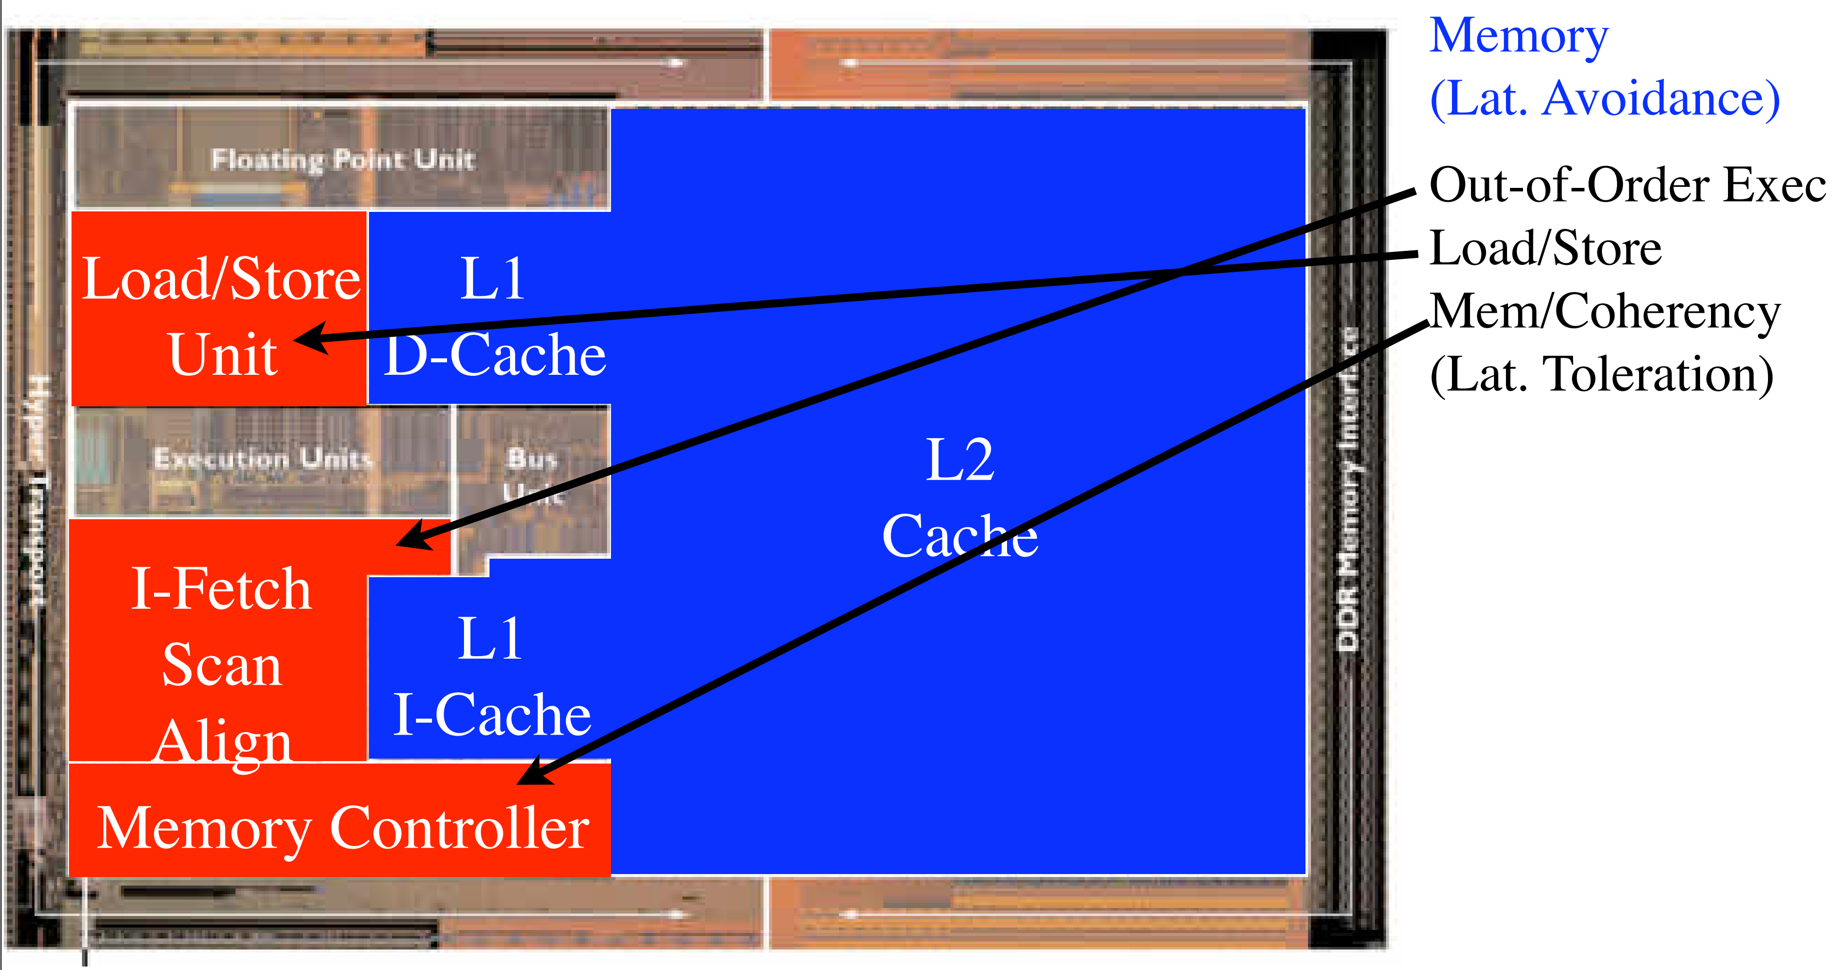

The Computer Architect's View of Memory

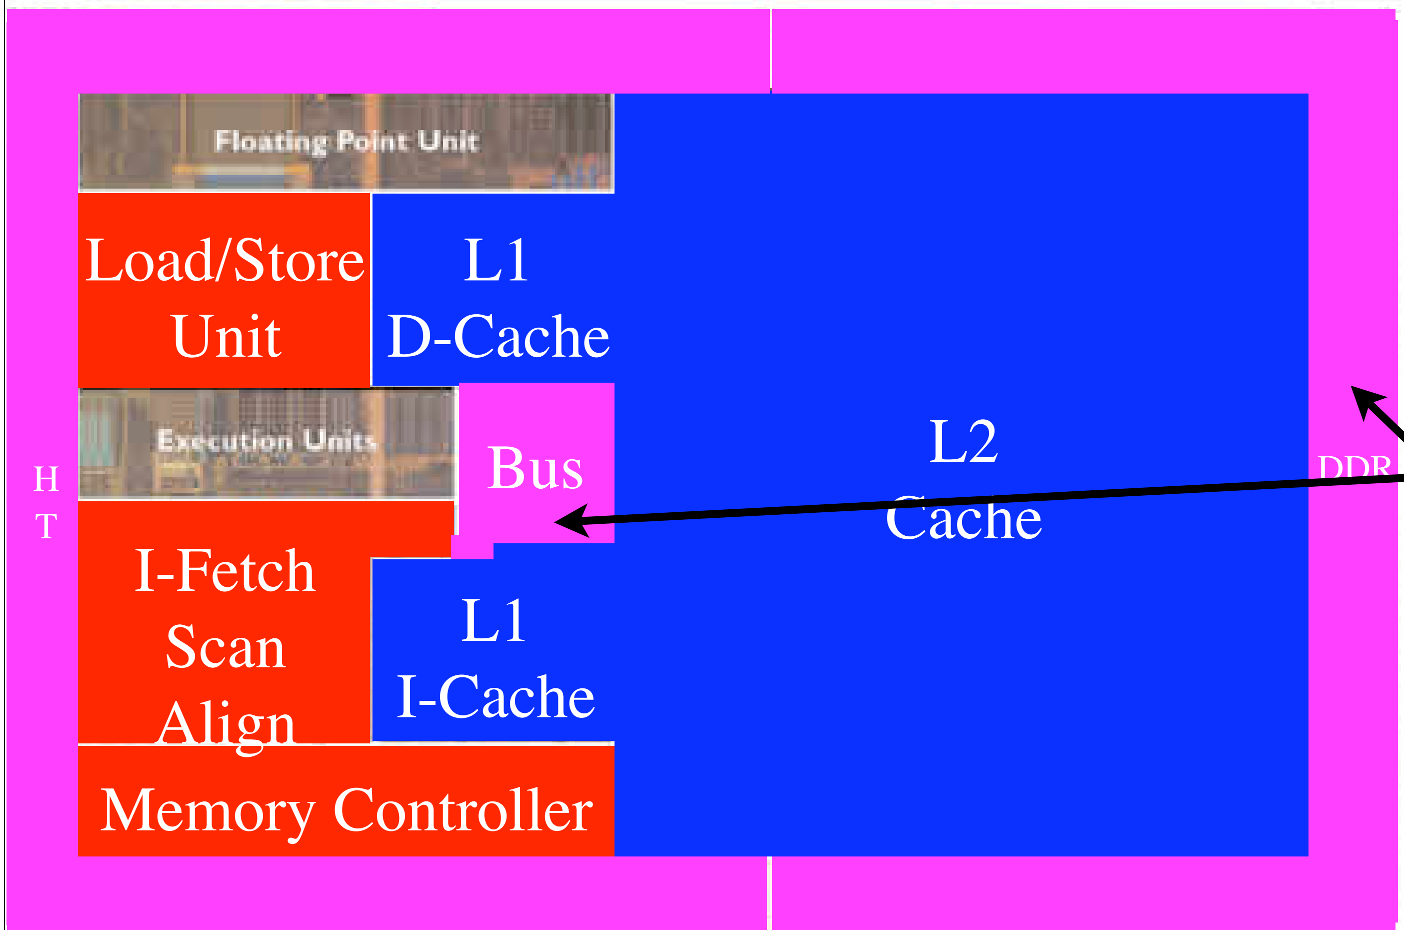

Memory (Lat. Avoidance)

Out-of-Order Exec Load/Store Mem/Coherency (Lat. Toleration)

Memory and I/O Interfaces

The Computer Architect's View of Memory

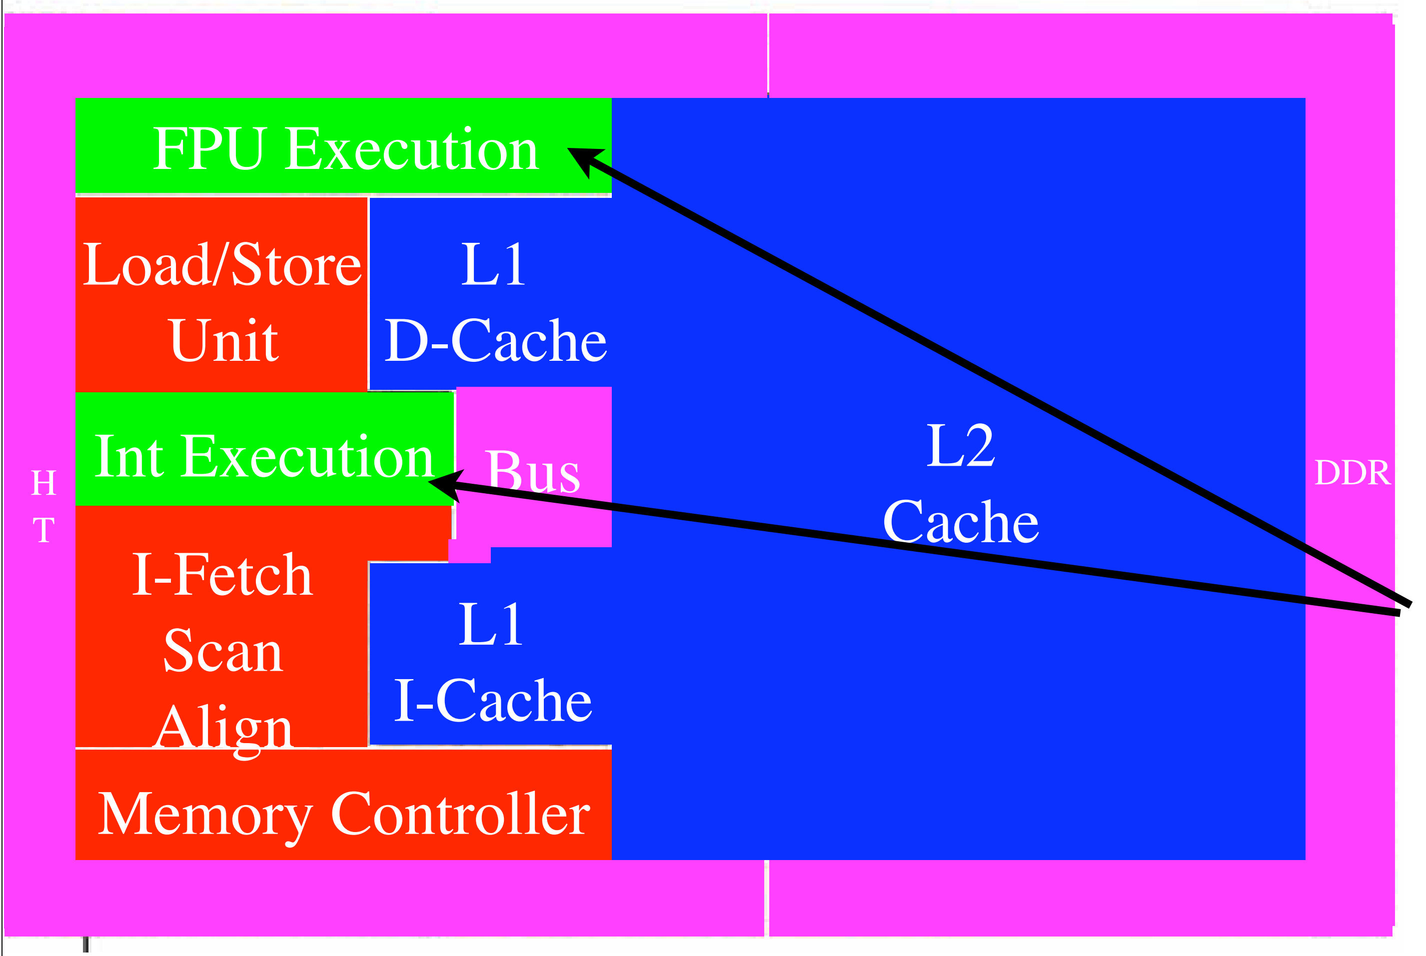

Memory (Lat. Avoidance)

Out-of-Order Exec
Load/Store
Mem/Coherency
(Lat. Toleration)

Memory and I/O Interfaces

COMPUTER

Key Memory-Centric Benchmarks

- Giga Updates Per Second (GUPS)
  - HPCS Discrete Math Benchmark
  - Memory[Random()] ^= Random()
  - Pure Latency Effect Measure
  - Little or no "Real" computation

- Stream (McCalpin's Bandwidth Benchmark)
  - Strided access
    - COPY: a(i) = b(i)
    - SCALE: a(i) = q*b(i)
    - SUM: a(i) = b(i) + c(i)
    - TRIAD: a(i) = b(i) + q*c(i)
  - Prefetcher and compiler options *hugely* impact results!
  - Not presented in this work (too bandwidth-centric)

The Effect on REAL Applications: Latency vs. Bandwidth

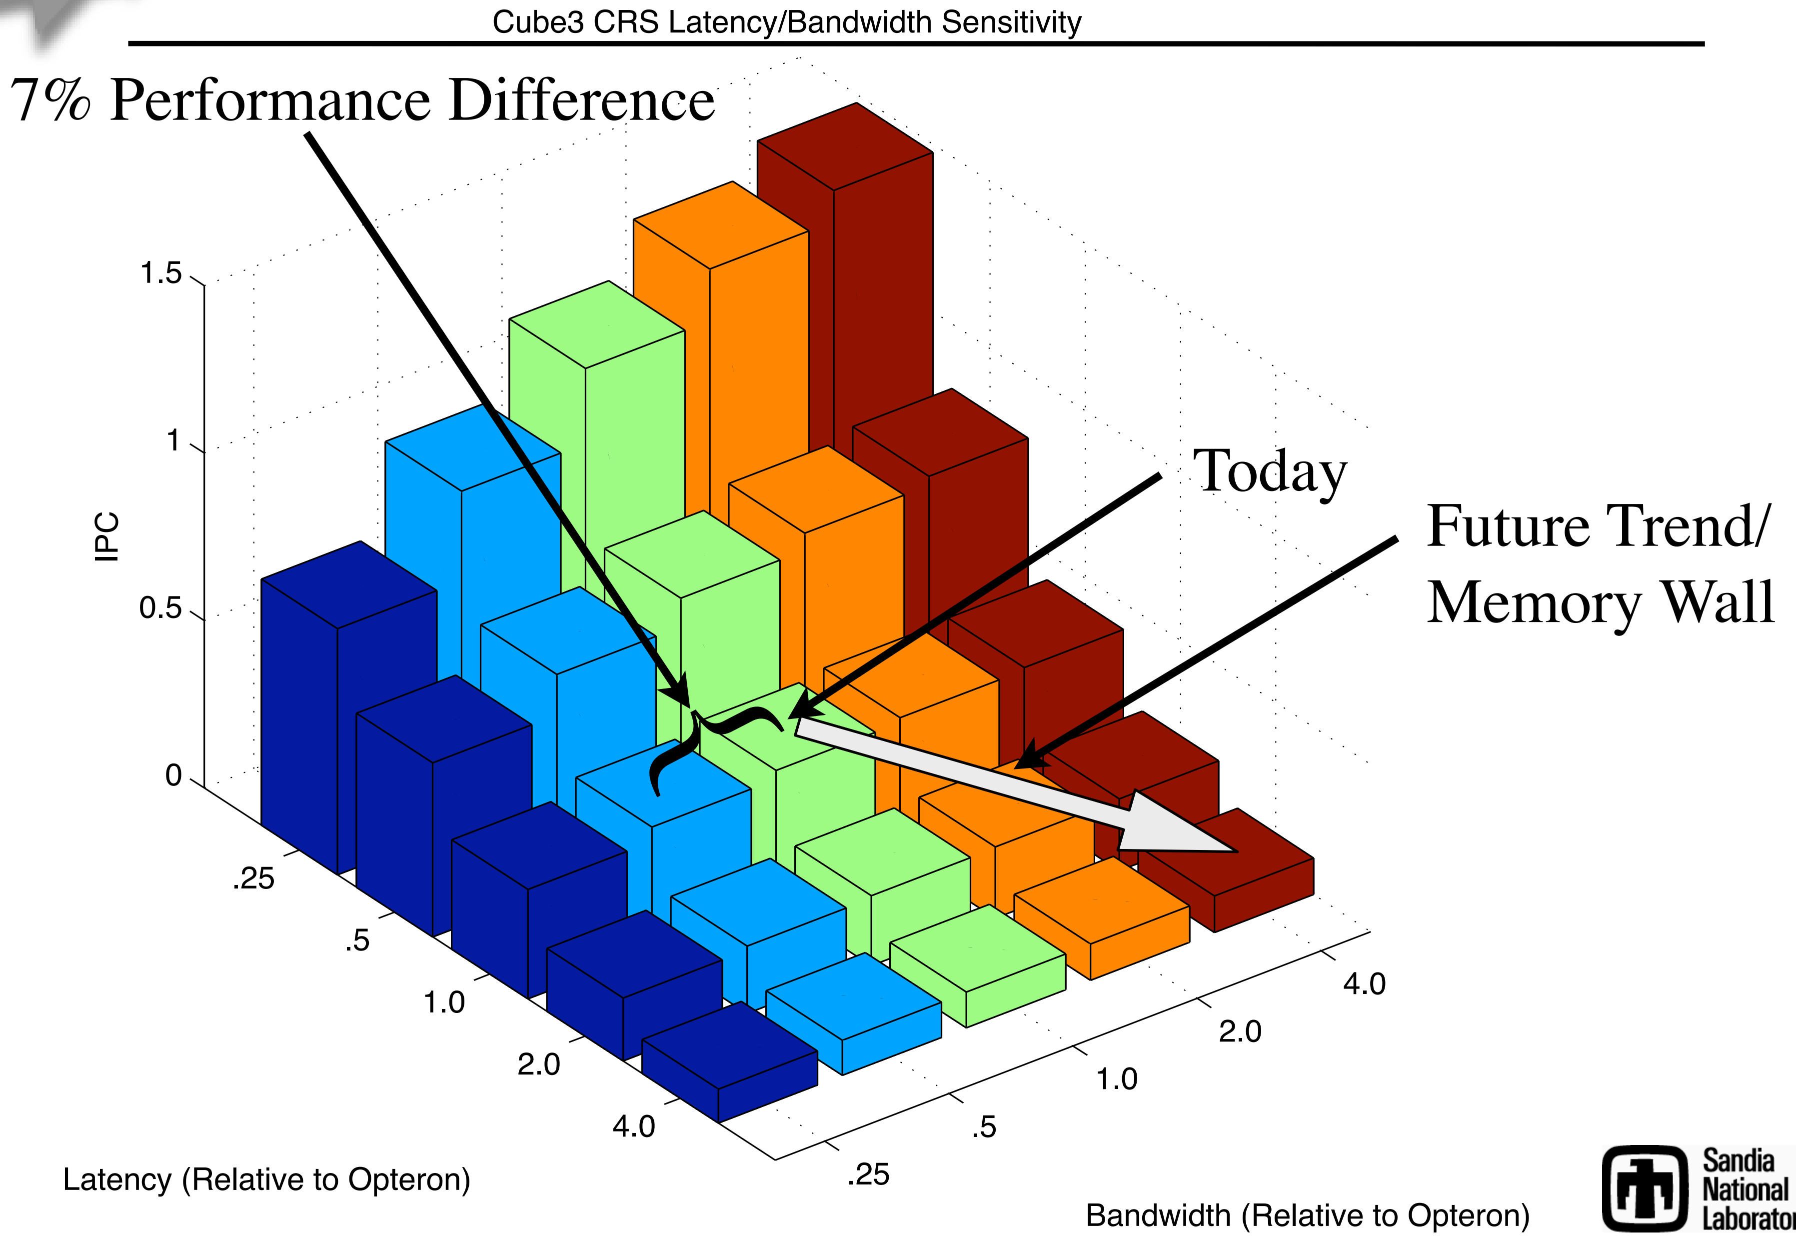

Sandia National Laboratories

Latency

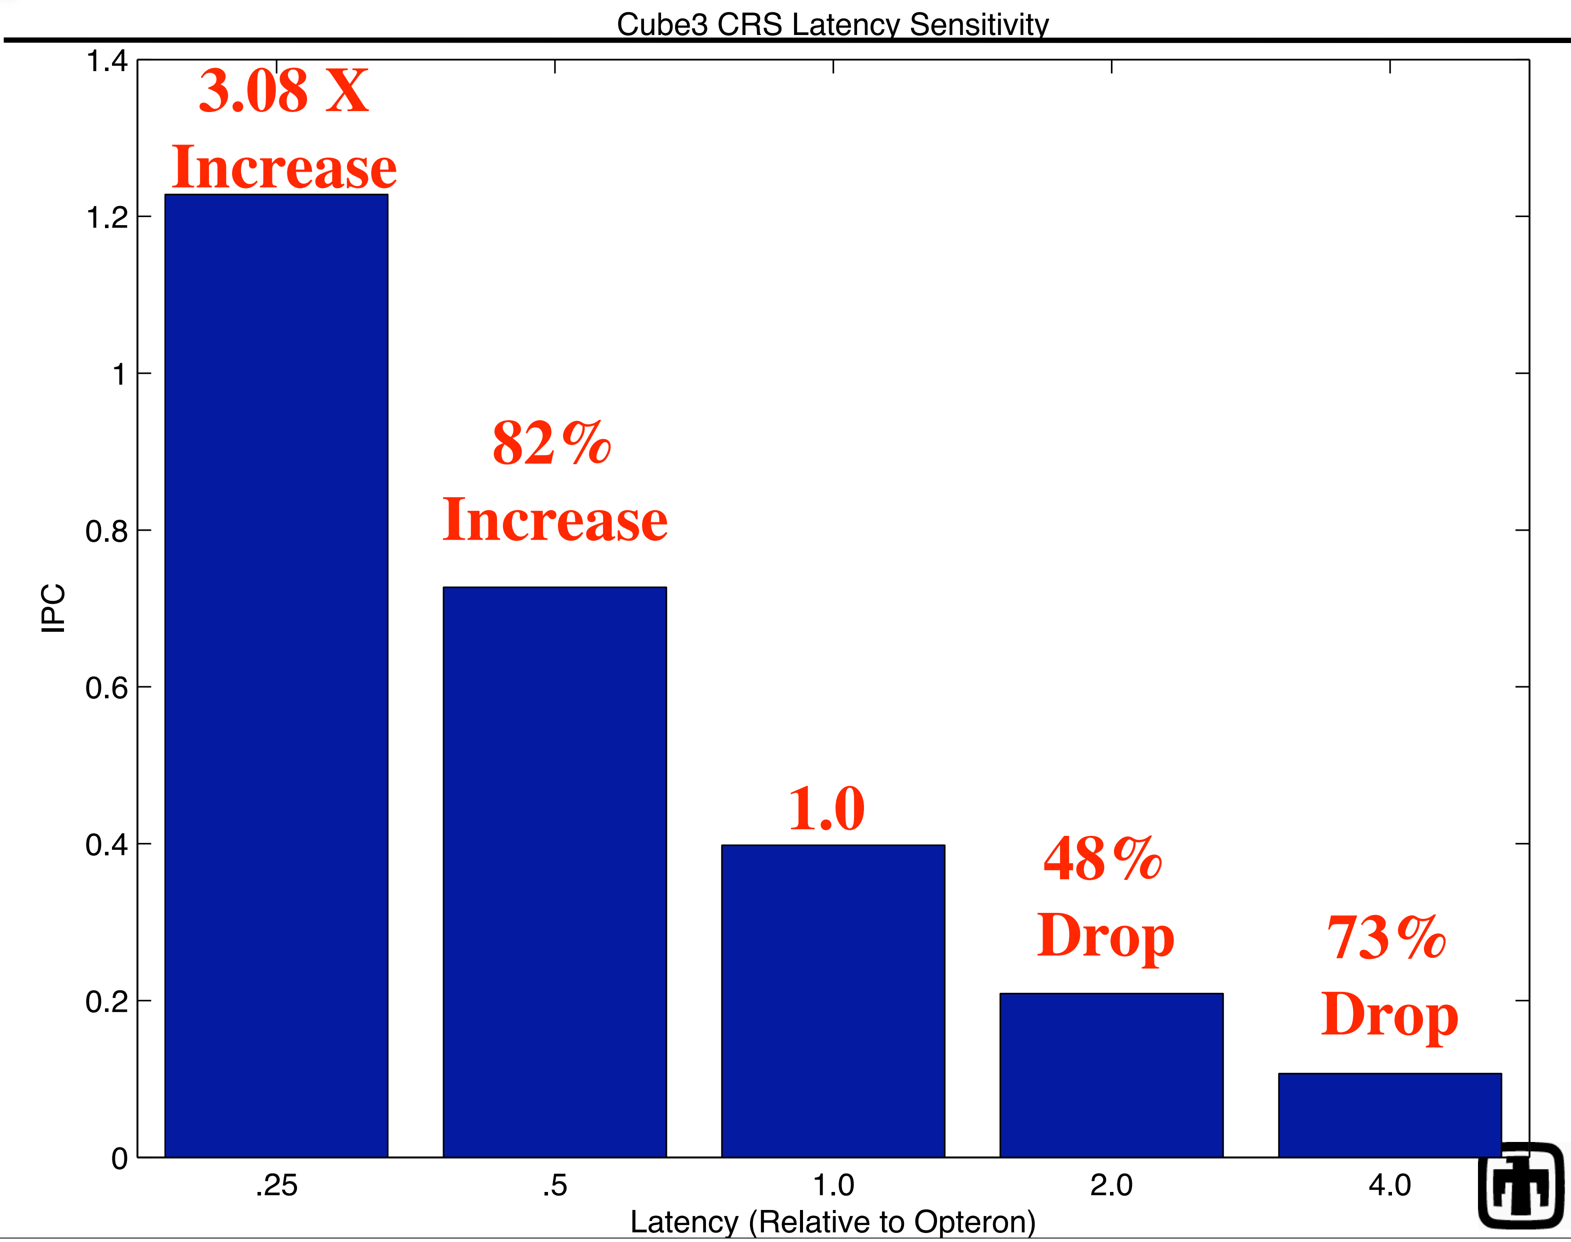
Cube3 CRS Latency Sensitivity
IPC
1.4
1.2
1
0.8
0.6
0.4
0.2
0
3.08 X
Increase
82%
Increase
1.0
48%
Drop
73%
Drop
.25
.5
1.0
2.0
4.0
Latency (Relative to Opteron)

Other Benchmark Suites

- Benchmarks
  - SPEC CPU 2000
    - SPEC-FP and SPEC-Int
    - Primary suite studied by computer architects
  - LINPACK
    - Used to generate the Top 500 Supercomputer List
    - Known to be flawed

Sandia FP Benchmark Suite

- CTH: Shock Physics
  - 2-gas
  - Explosively Formed Projectile (EFP)
  - CuSt AMR
    - 4mm copper ball impacting a steel plate at 4.52 km/s and 90-degrees, uses adaptive mesh refinement
  - LAMMPS: Classical Molecular Dynamics Simulation
    - Lennard Jones Mixture (2048 Atoms)
    - Chain (32,000 Atoms and 31,680 bonds)
  - Cube3: Drives the Trilinos framework for linear solvers
  - sPPM: 3d gas dynamics

Sandia Integer Suite

- Minimum Cut Graph Partitioning
- Concurrent Search
  - Serial Implementation is Depth First Search
- Shortest Path
- Isomorphism
- BLAST DNA Nucleotide Alignment
- zChaff: Satisfiability

Proposed Benchmark: DFS of a Power-law Graph

- Performs a Depth First Search on a "power law" graph
  - Many low-degree vertices
  - Very few high-degree vertices
  - Small Diameter
  - Short distance between two vertices

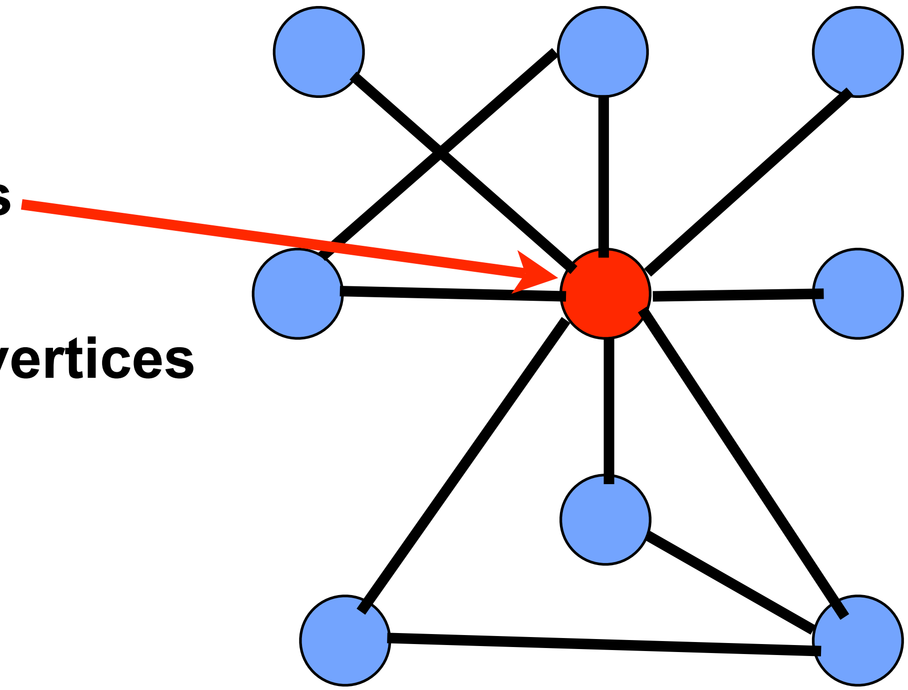

- Relevance
  - Social Networks
  - Biology
  - Data Mining
  - Emerging Integer Applications

DFS Properties

- **Fundamental Building Block for Other Algorithms**
  - Connected Components
  - Subgraph Isomorphism
  - Shortest Path
- ***Parallel/Concurrent Versions available***
  - Multithreaded
    - Can scale to as many or as few threads as required
    - Runs on the Cray MTA
    - Runs on conventional multiprocessors
  - Multicore Architectures
  - Memory Dominated

Measuring Memory Characteristics

- *Temporal Locality*
  - Miss Rate of a Fully Associative 64-kb (L1-sized) cache with machine word (4-byte) block sizes
- *Spatial Locality*
  - the average ratio of bytes consumed in a 64-byte cache line over 1,000 instructions
- *Data Intensiveness*
  - the total number of unique bytes consumed by the instruction stream
- Each measurement was performed on well studied 1 billion instruction traces for PowerPC
  - Murphy, Rodrigues, Kogge, and Underwood. *The Implications of Working Set Analysis on Supercomputing Memory Hierarchy Design*, ICS 2005.

Example

Given the following memory access sequence:
A, B, C, A, A, C

Temporal Working Set

Spatial Working Set

Data Intensiveness

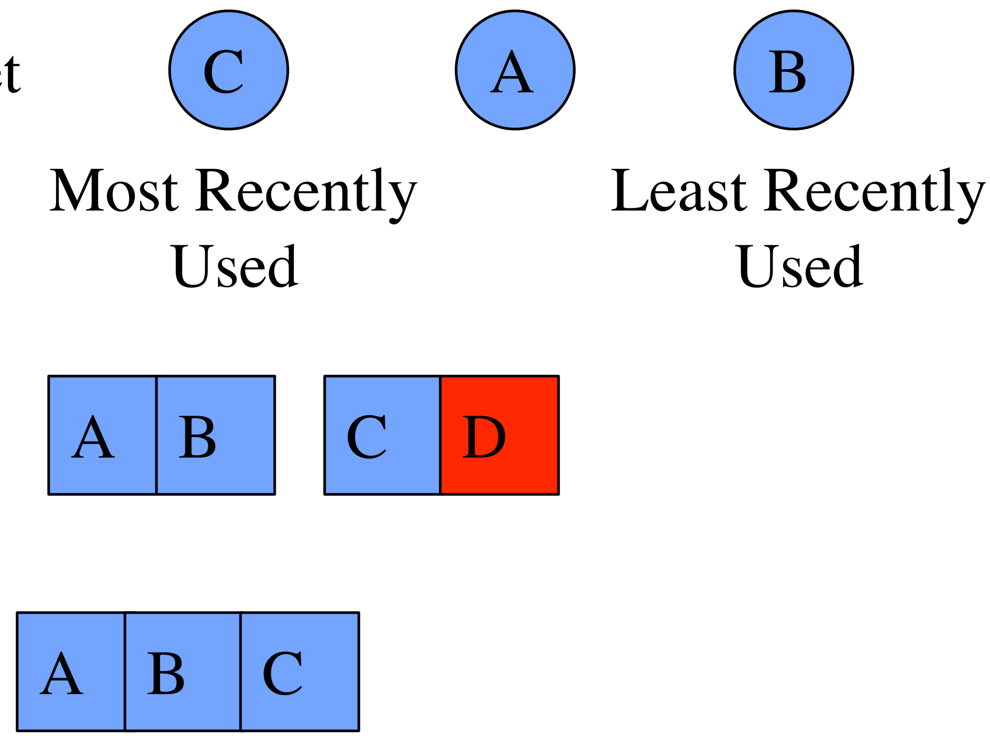

Results

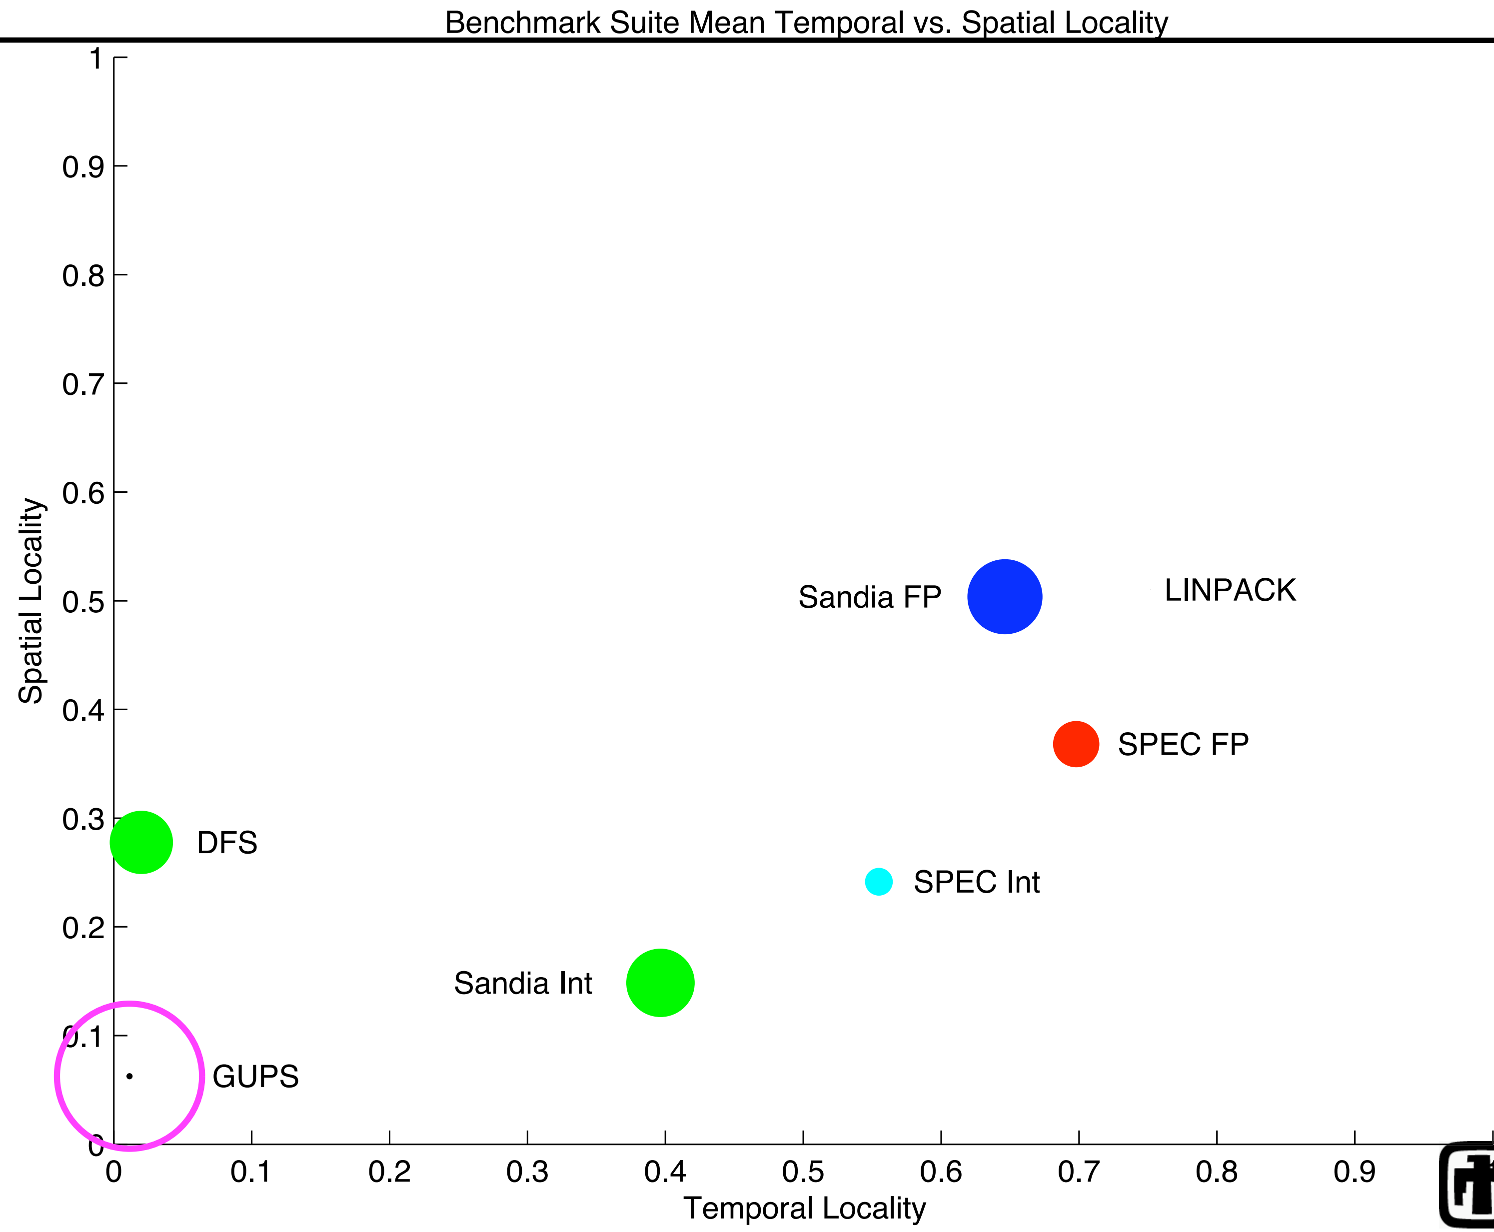
Benchmark Suite Mean Temporal vs. Spatial Locality
Spatial Locality
Temporal Locality
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
Sandia FP
LINPACK
SPEC FP
DFS
SPEC Int
Sandia Int
GUPS

Results Compared to GUPS

4.4x the Temporal Locality
1.75x the Spatial Locality of GUPS
19% the data set size

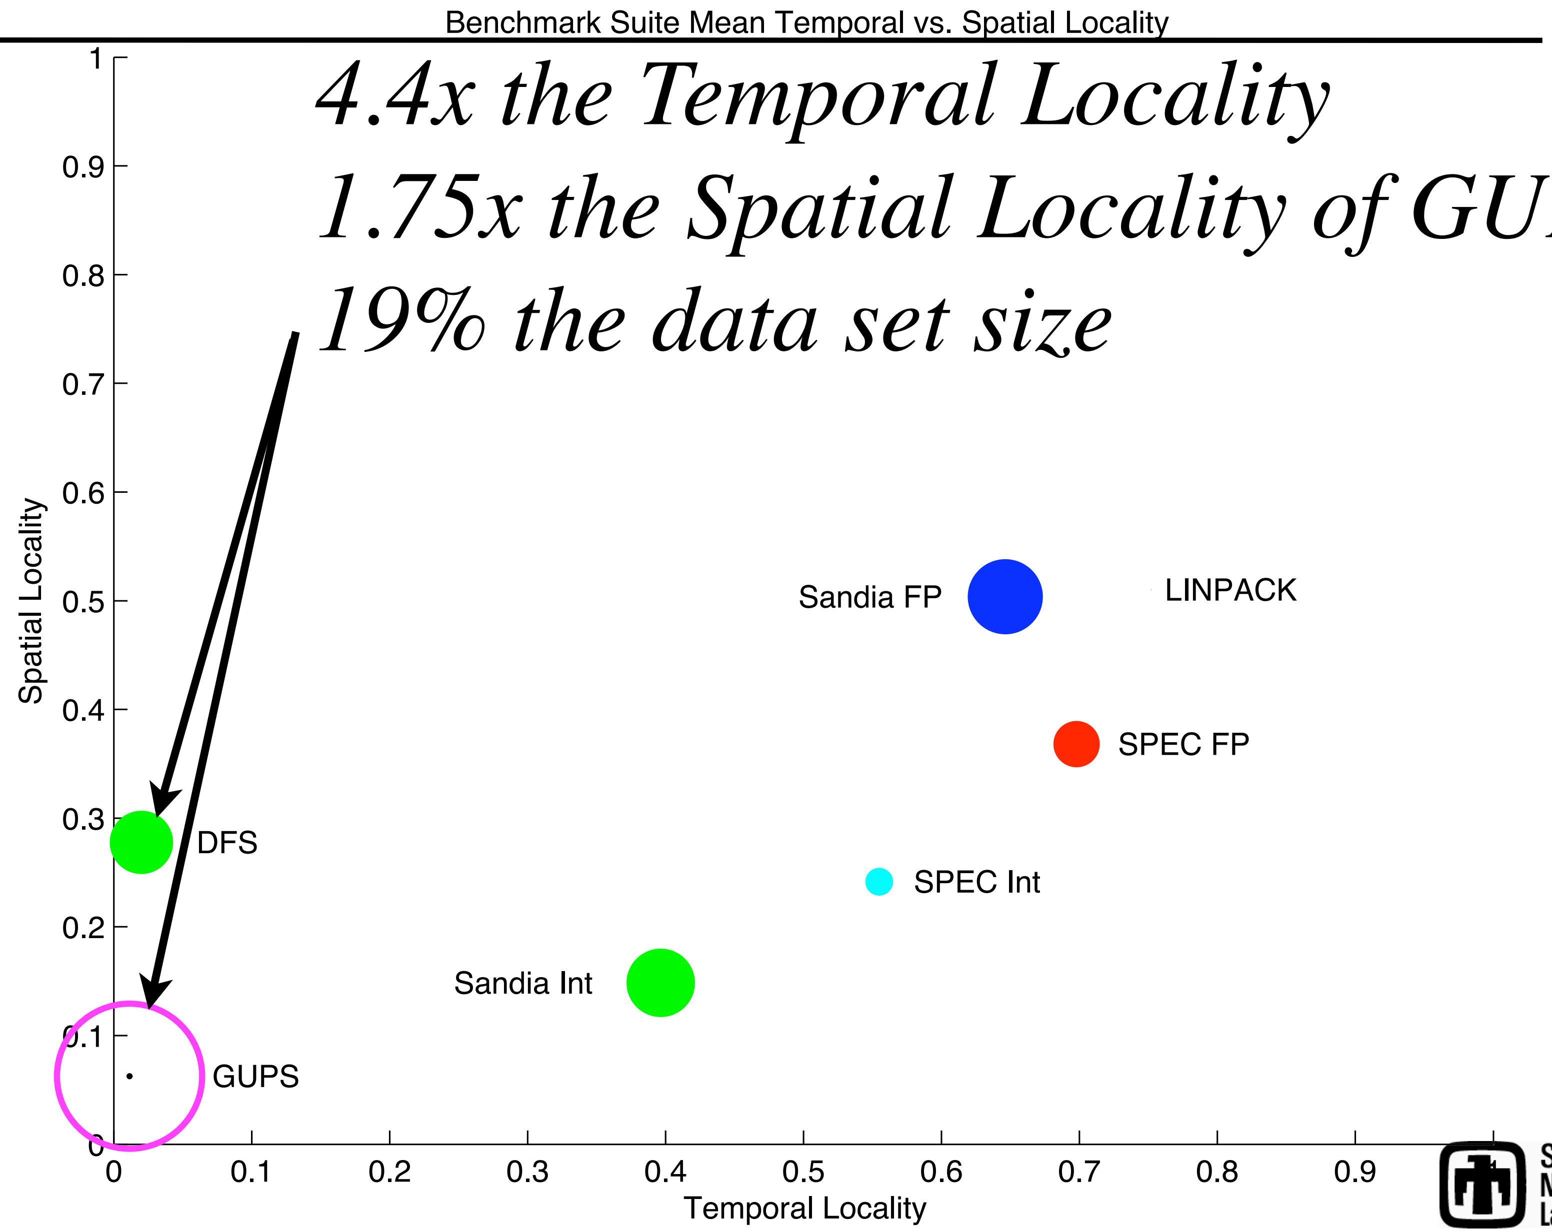

Sandia National Laboratories

Results compared to Integer Applications

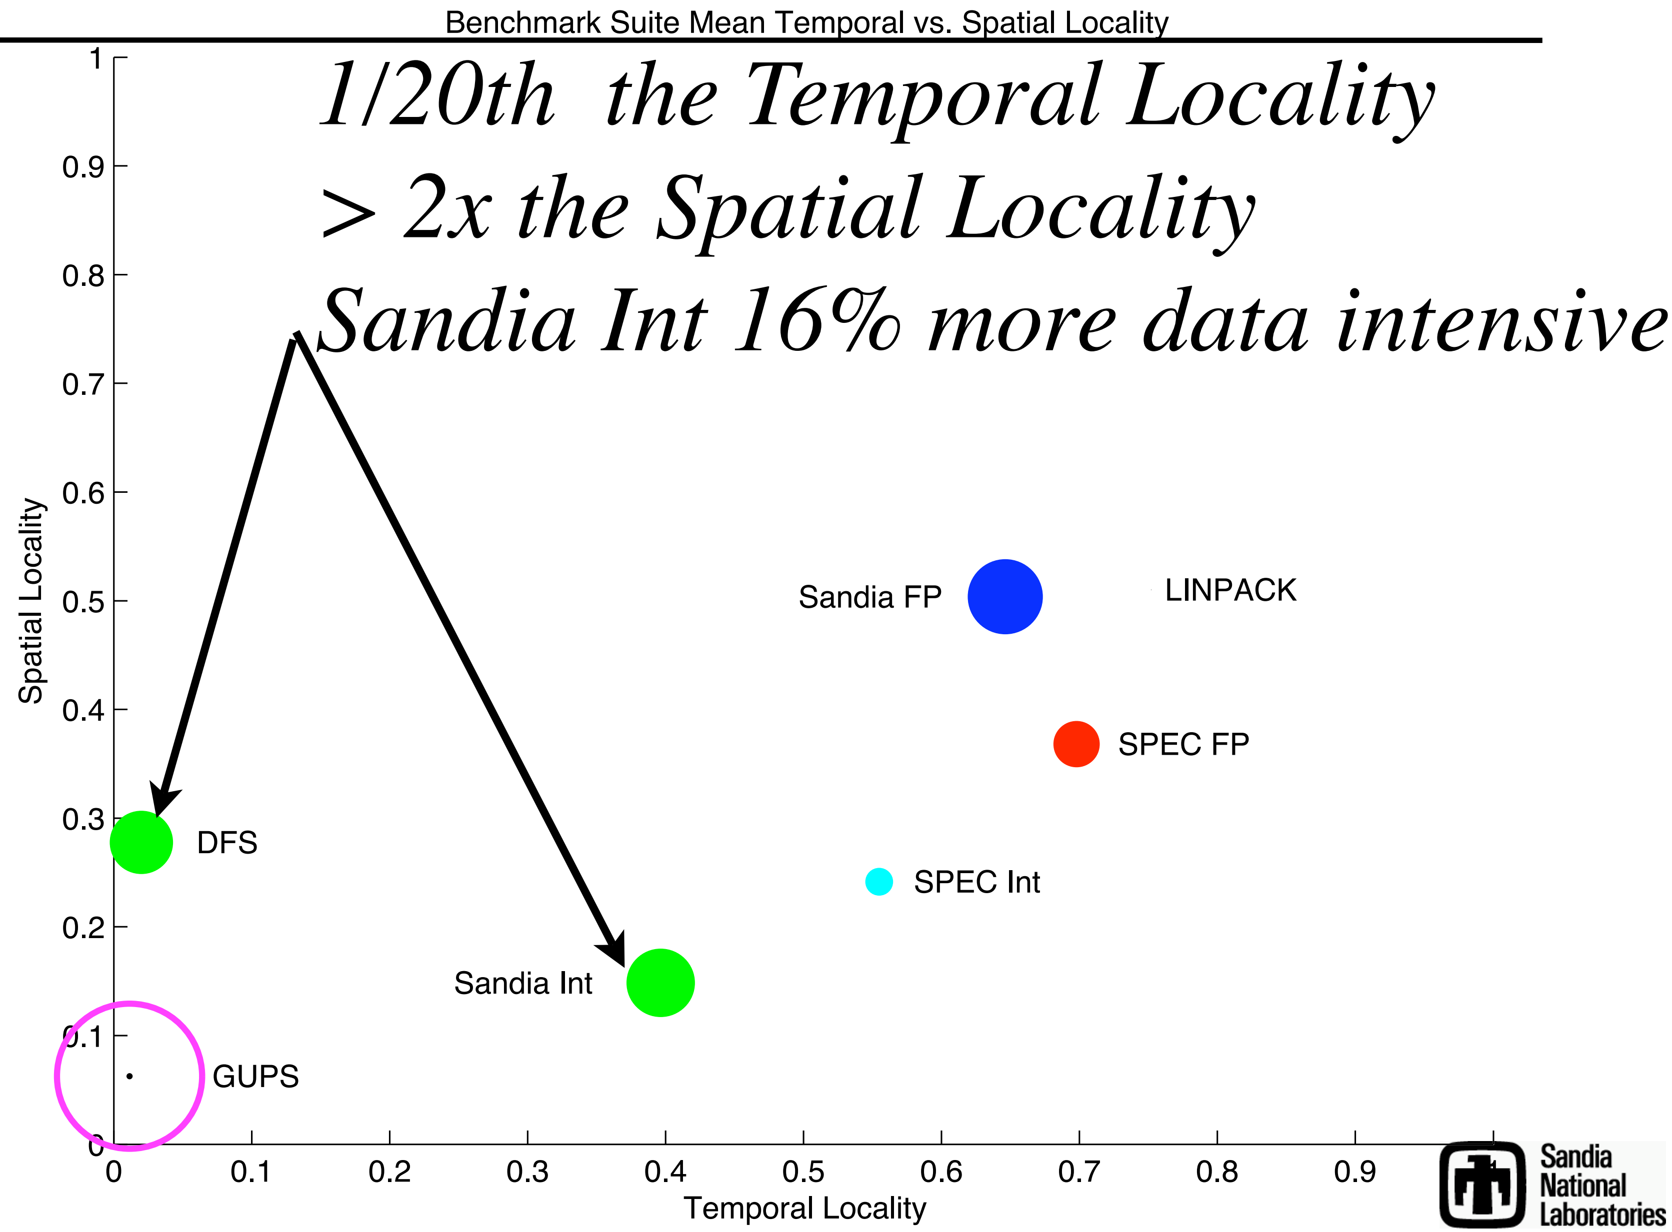

Results compared to LINPACK

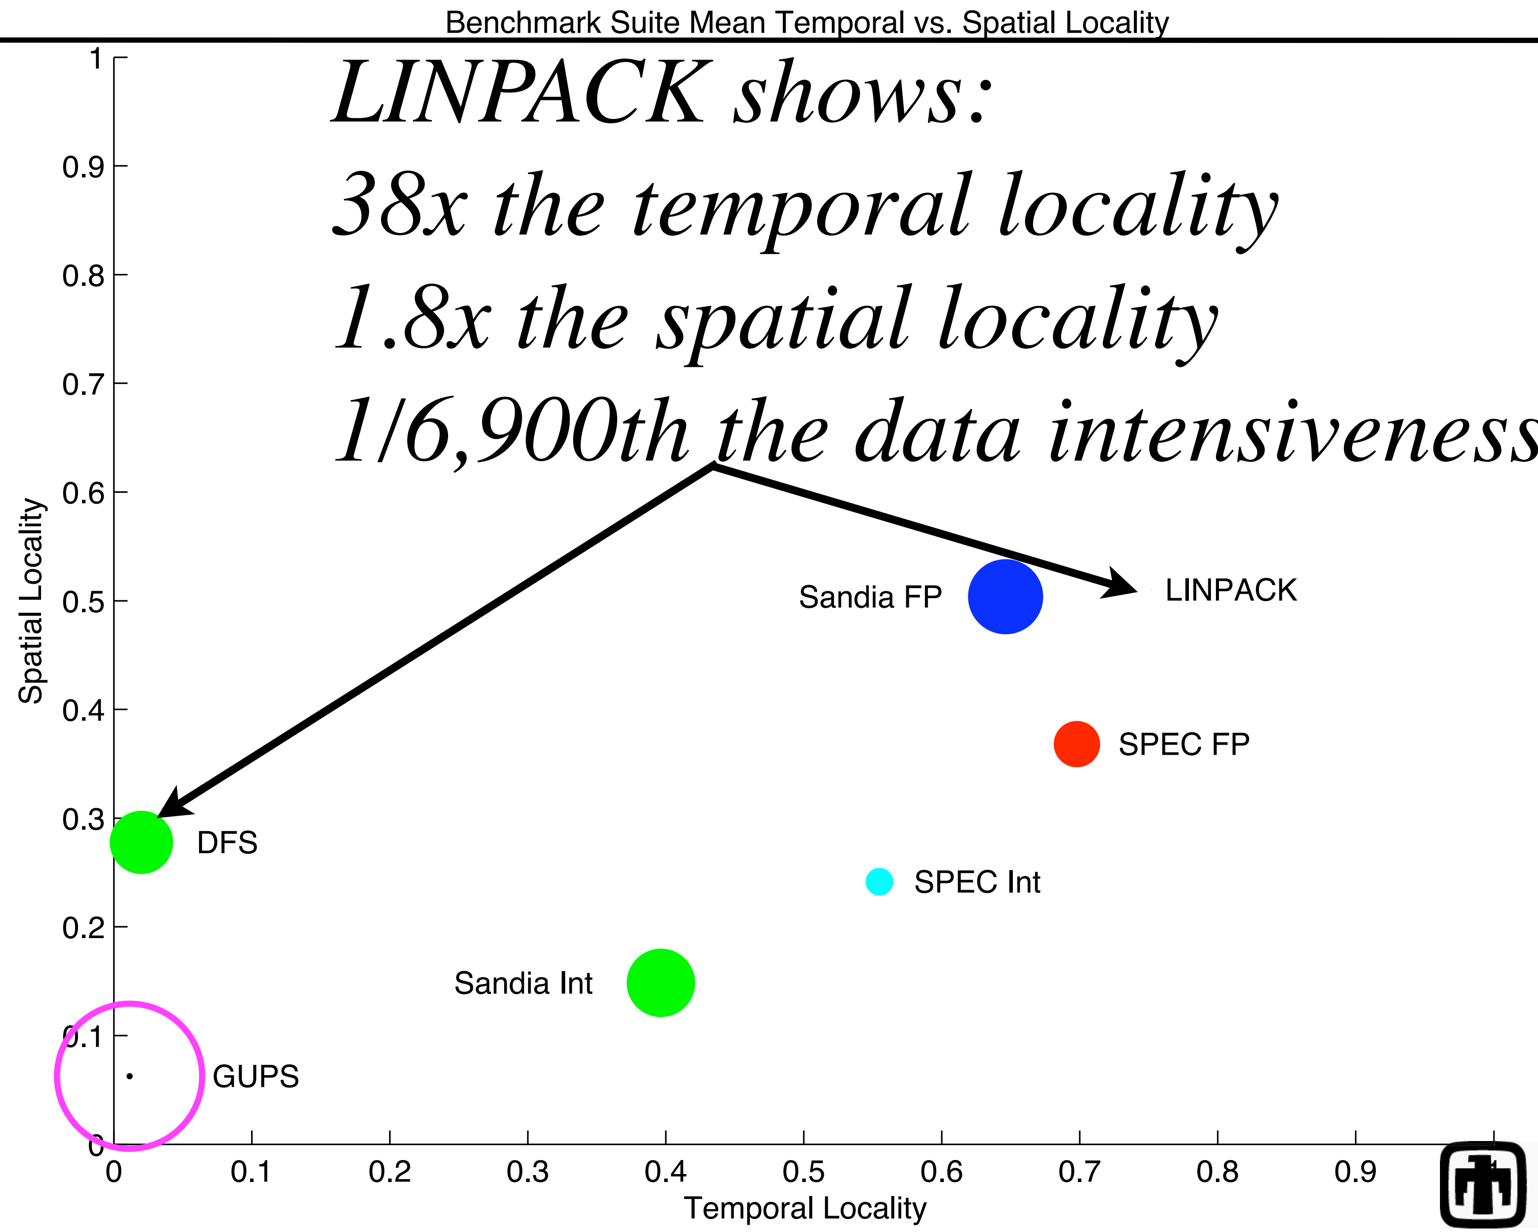

Sandia National Laboratories

Data Intensiveness

- Critical "missing" measurement for differentiation
- Benchmarks are less data intensive
  - 1.86x SPEC-FP
  - 5x SPEC-Int
- Real Applications are more data intensive
  - Sandia FP 41% more data intensive
  - Sandia Int 16% more data intensive

Conclusions

- Memory-centric benchmark that performs a real and relevant computation
- Scalable
  - in memory size
  - in number of threads with the parallel version
- Quantitatively Different from Other Benchmarks
  - Low Temporal and Spatial Locality
  - High Data Intensiveness
- Represents an emerging class of HPC applications
  - Integer rather than floating point based
  - Memory Dominated

THANK YOU