



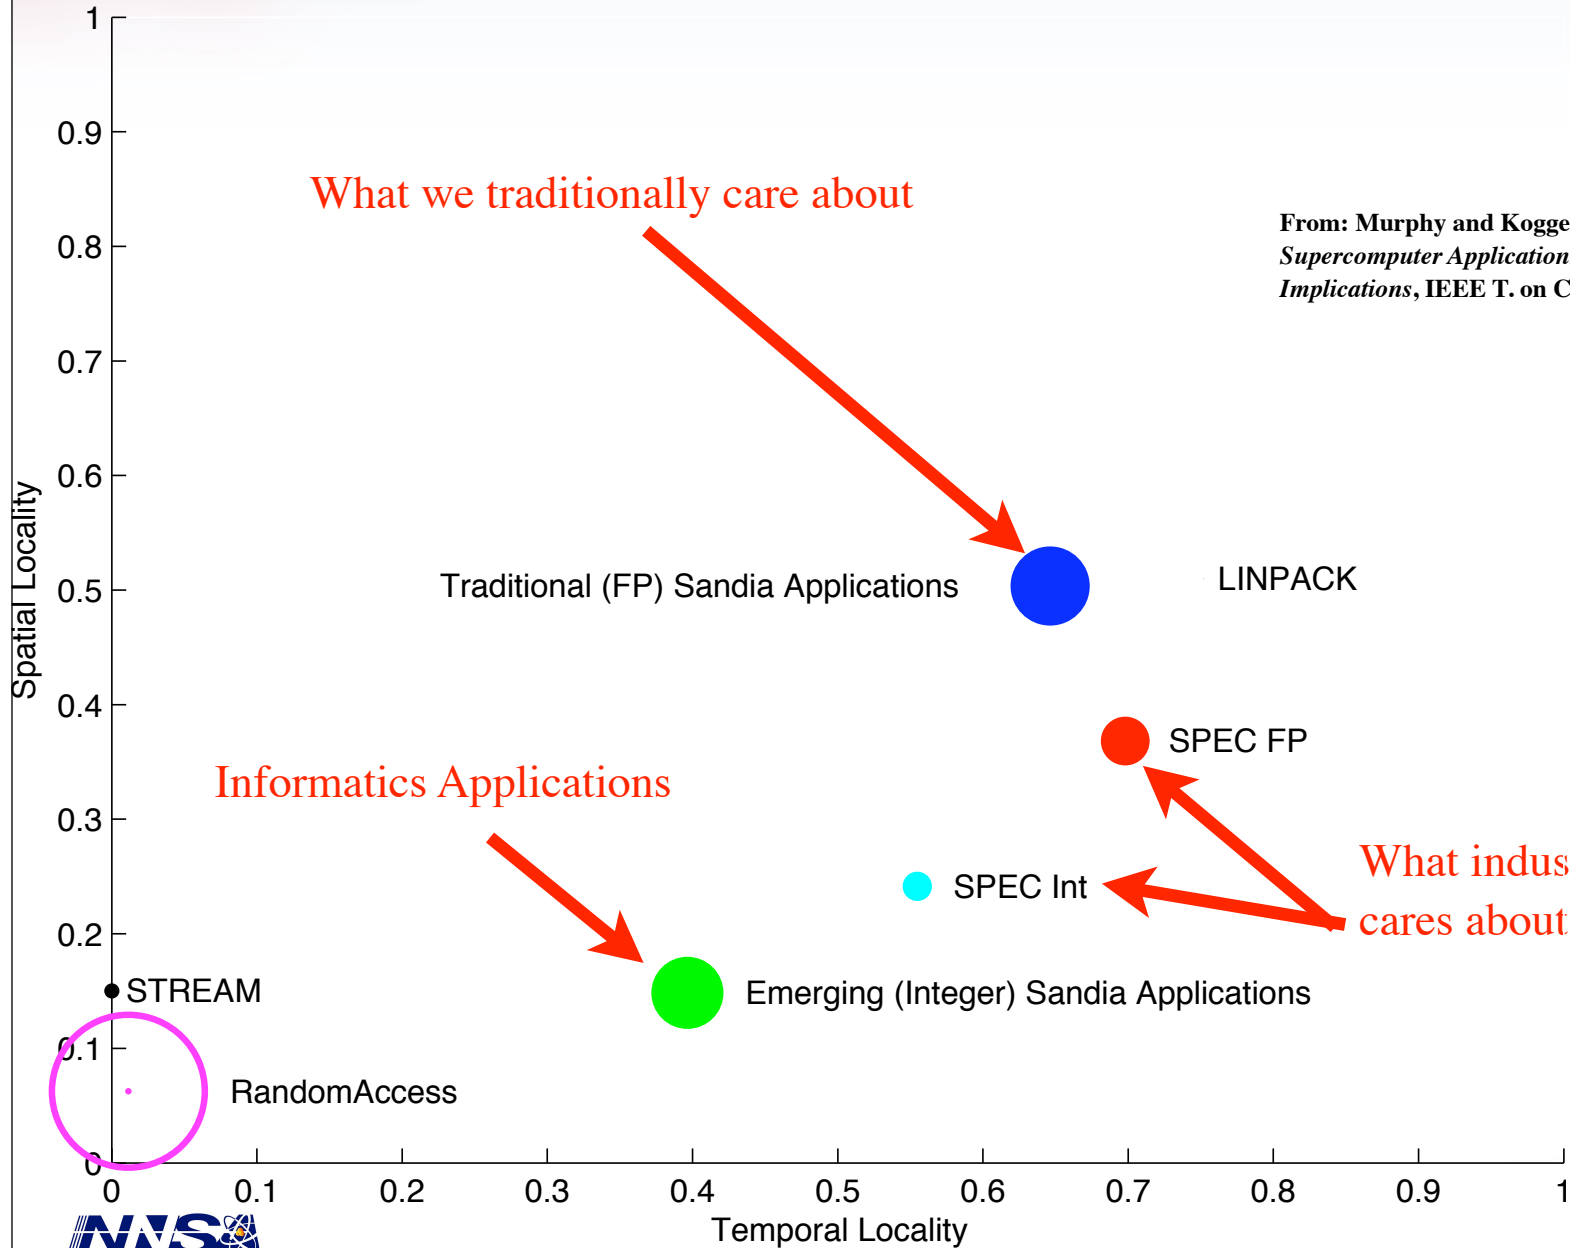
Introducing the Graph 500

**Richard Murphy, Kyle Wheeler, Brian
Barrett, and Jim Ang**

Sandia National Laboratories

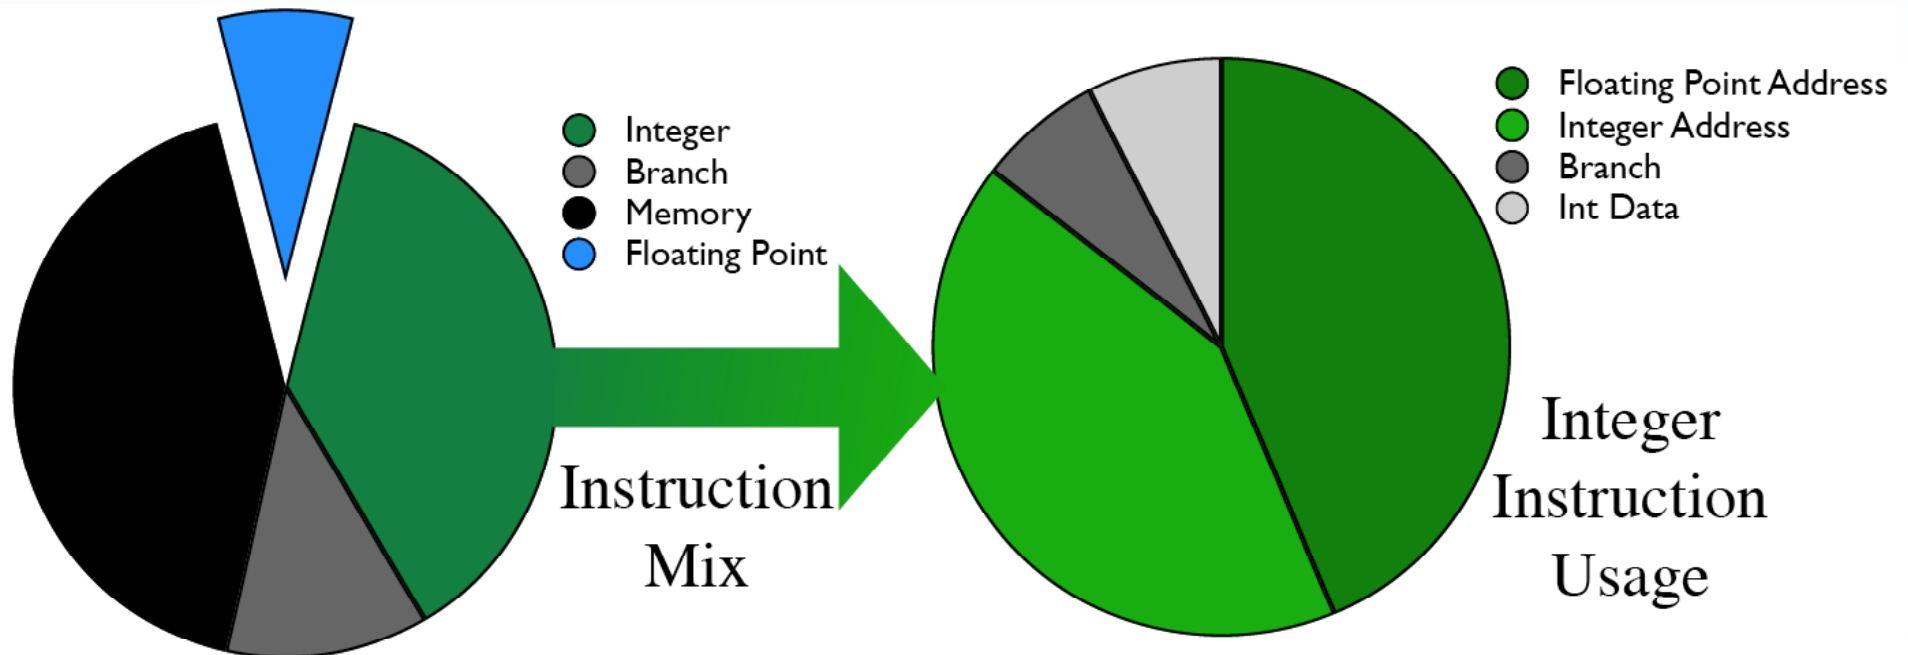
Not All Applications are Floating Point Oriented

Benchmark Suite Mean Temporal vs. Spatial Locality



From: Murphy and Kogge, *On The Memory Access Patterns of Supercomputer Applications: Benchmark Selection and Its Implications*, IEEE T. on Computers, July 2007



Even Floating Point Applications are Memory-Centric



Real Physics Applications Primarily Do SLOW Memory References

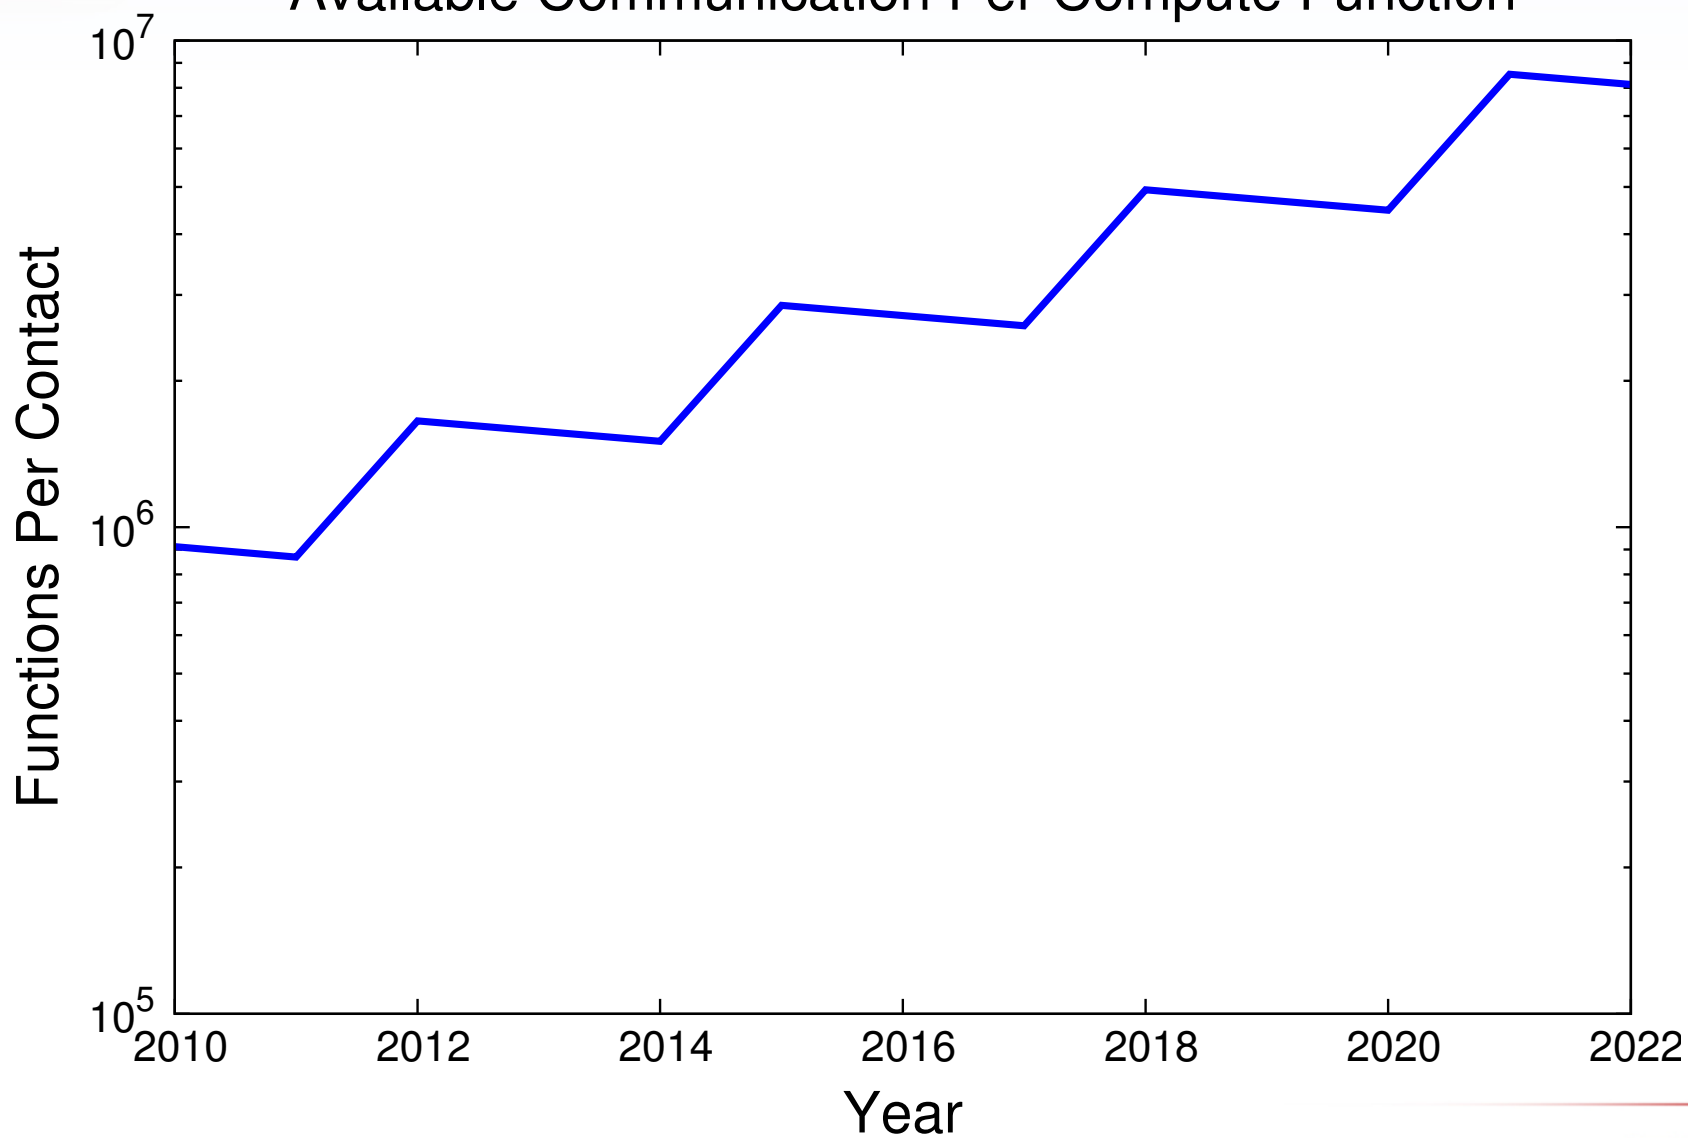
How is memory changing?

$$\text{Throughput} = \frac{\text{Concurrency}}{\text{Latency}}$$



Put Another Way

Available Communication Per Compute Function





What is the Graph 500?

- **New benchmark to complement the Top 500 for large-scale data analysis problems**
- **International Multidisciplinary Steering Committee**
 - Jim Ang, David Bader, Brian Barrett, Jon Berry, Bill Brantley, Almadena Chtchelkanova, John Daly, John Feo, Michael Garland, John Gilbert, Bill Gropp, Bill Harrod, Bruce Hendrickson, Jure Leskovec, Bob Lucas, Andrew Lumsdaine, Mike Merrill, Hans Meuer, David Mizell, Shoaib Mufti, Richard Murphy, Nick Nystrom, Fabrizio Petrini, Wilf Pinfold, Steve Poole, Arun Rodrigues, Rob Schreiber, John Simmons, Marc Snir, Thomas Sterling, Blair Sullivan, T.C. Tuan, Jeff Vetter, Mike Vildibill
- **Three Kernels**
 - Search (Concurrent Search)
 - Optimization (Single Source Shortest Path)
 - Edge Oriented (Maximal Independent Set)
- **Random Algorithms will not be allowed**



What is the Graph 500 (continued)

- **Five “Business Area” Data Sets**
 - **Cybersecurity**
 - **Medical Informatics**
 - **Data Enrichment**
 - **Social Networks**
 - **Symbolic Networks**



Data Sets

- **Cybersecurity**


- 15 Billion Log Entries/Day (for large enterprises)
- Full Data Scan with End-to-End Join Required

- **Medical Informatics**

- 50M patient records, 20-200 records/patient, billions of individuals
- Entity Resolution Important

- **Data Enrichment**

- Easily PB of data
- **Example: Maritime Domain Awareness**
 - Hundreds of Millions of Transponders
 - Tens of Thousands of Cargo Ships
 - Tens of Millions of Pieces of Bulk Cargo
 - May involve additional data (images, etc.)



Data Sets (continued)

- **Social Networks**
 - Example, Facebook
 - Nearly Unbounded Dataset Size
- **Symbolic Networks**
 - Example, the Human Brain
 - 25B Neurons
 - 7,000+ Connections/Neuron



Reference Implementations

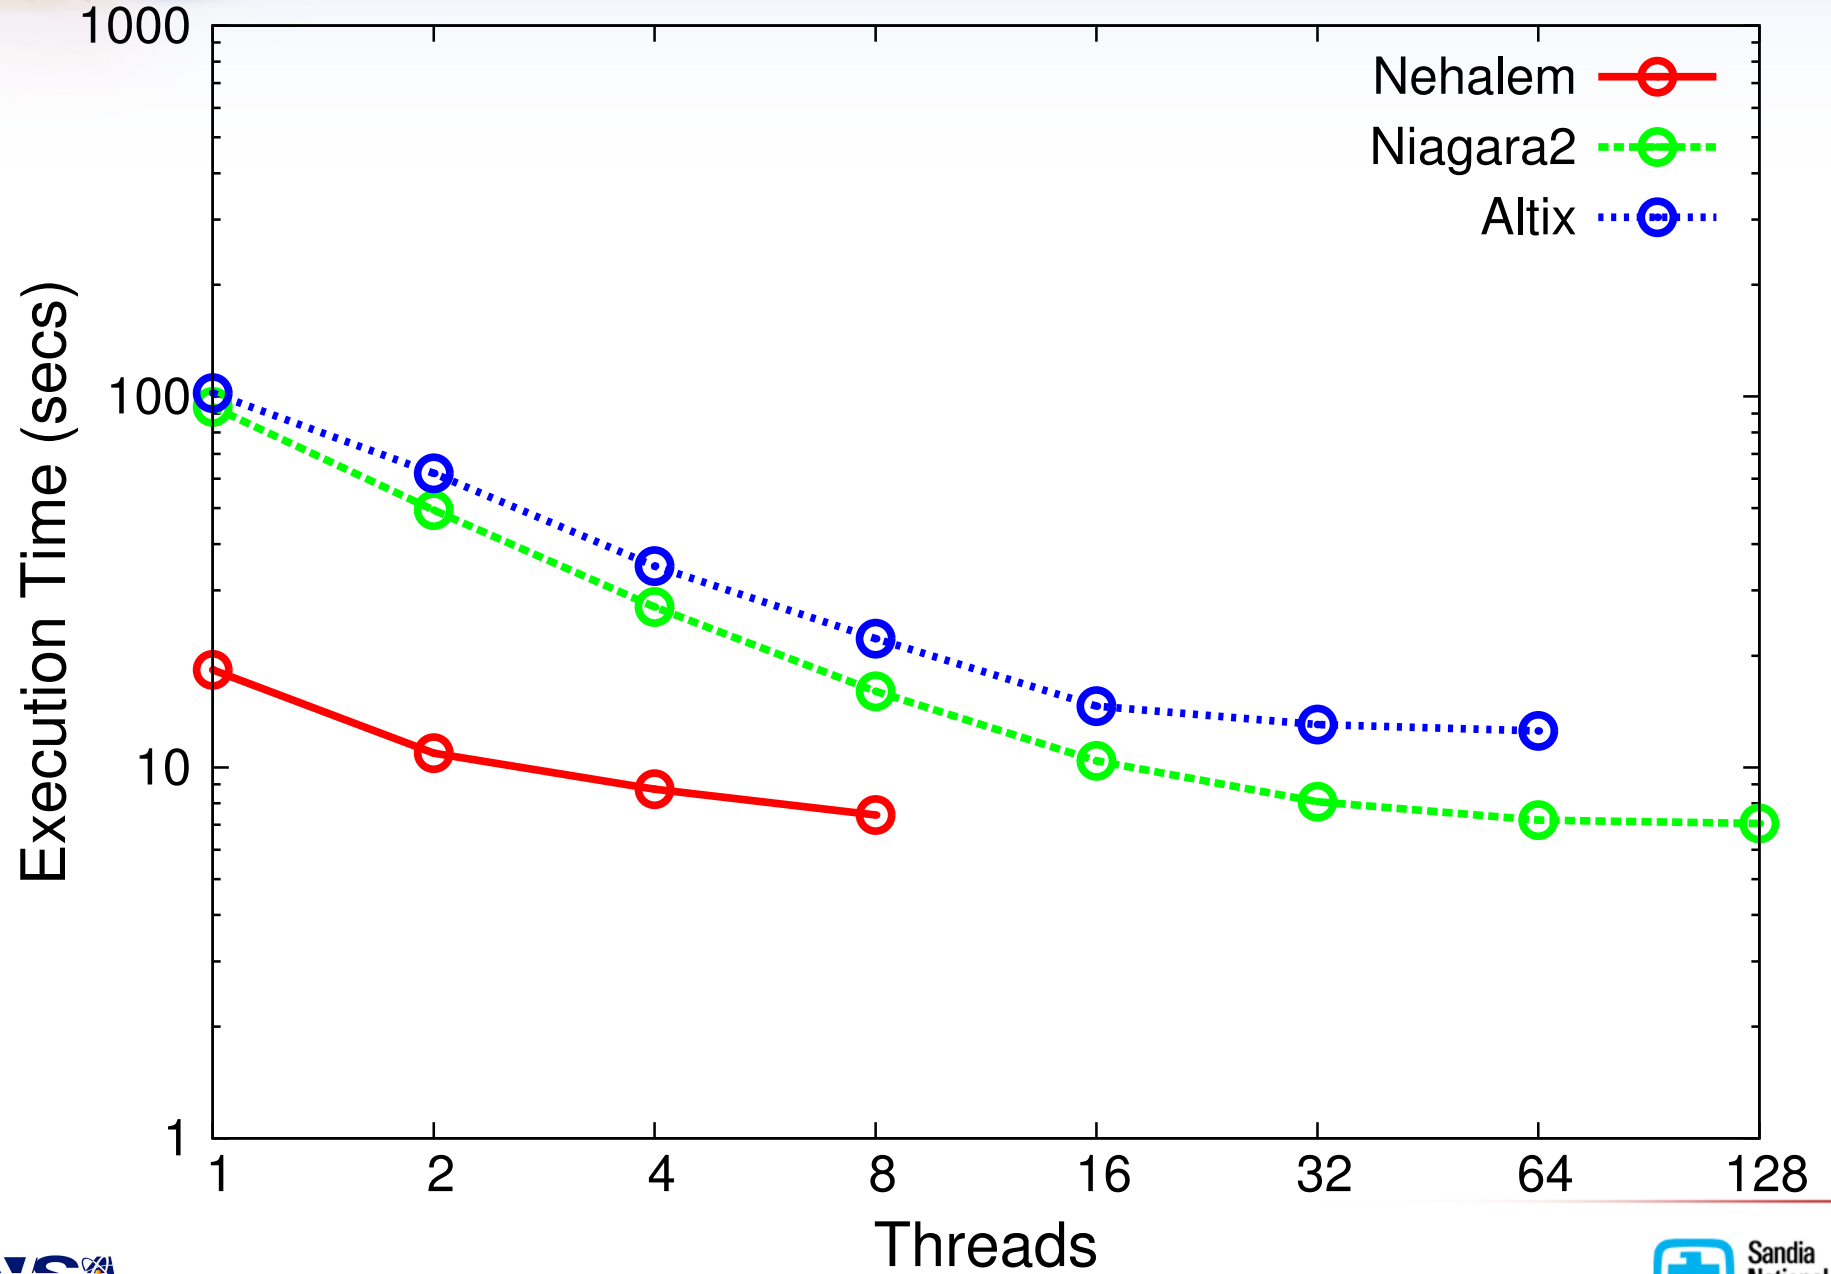
- **Will allow “base” and “peak” results similar to SPEC**
- **Three Reference Implementations:**
 - **Distributed Memory**
 - **Cloud/MapReduce**
 - **Multithreaded/Shared Memory**
- **Industry May implement custom frameworks**
 - **LexisNexis Data Analytic Supercomputer (DAS)**
 - **Custom Software and Programming Language (ECL)**
 - **Commodity Hardware**
 - **Cray XMT may requiring “tuning” of the multithreaded benchmark**



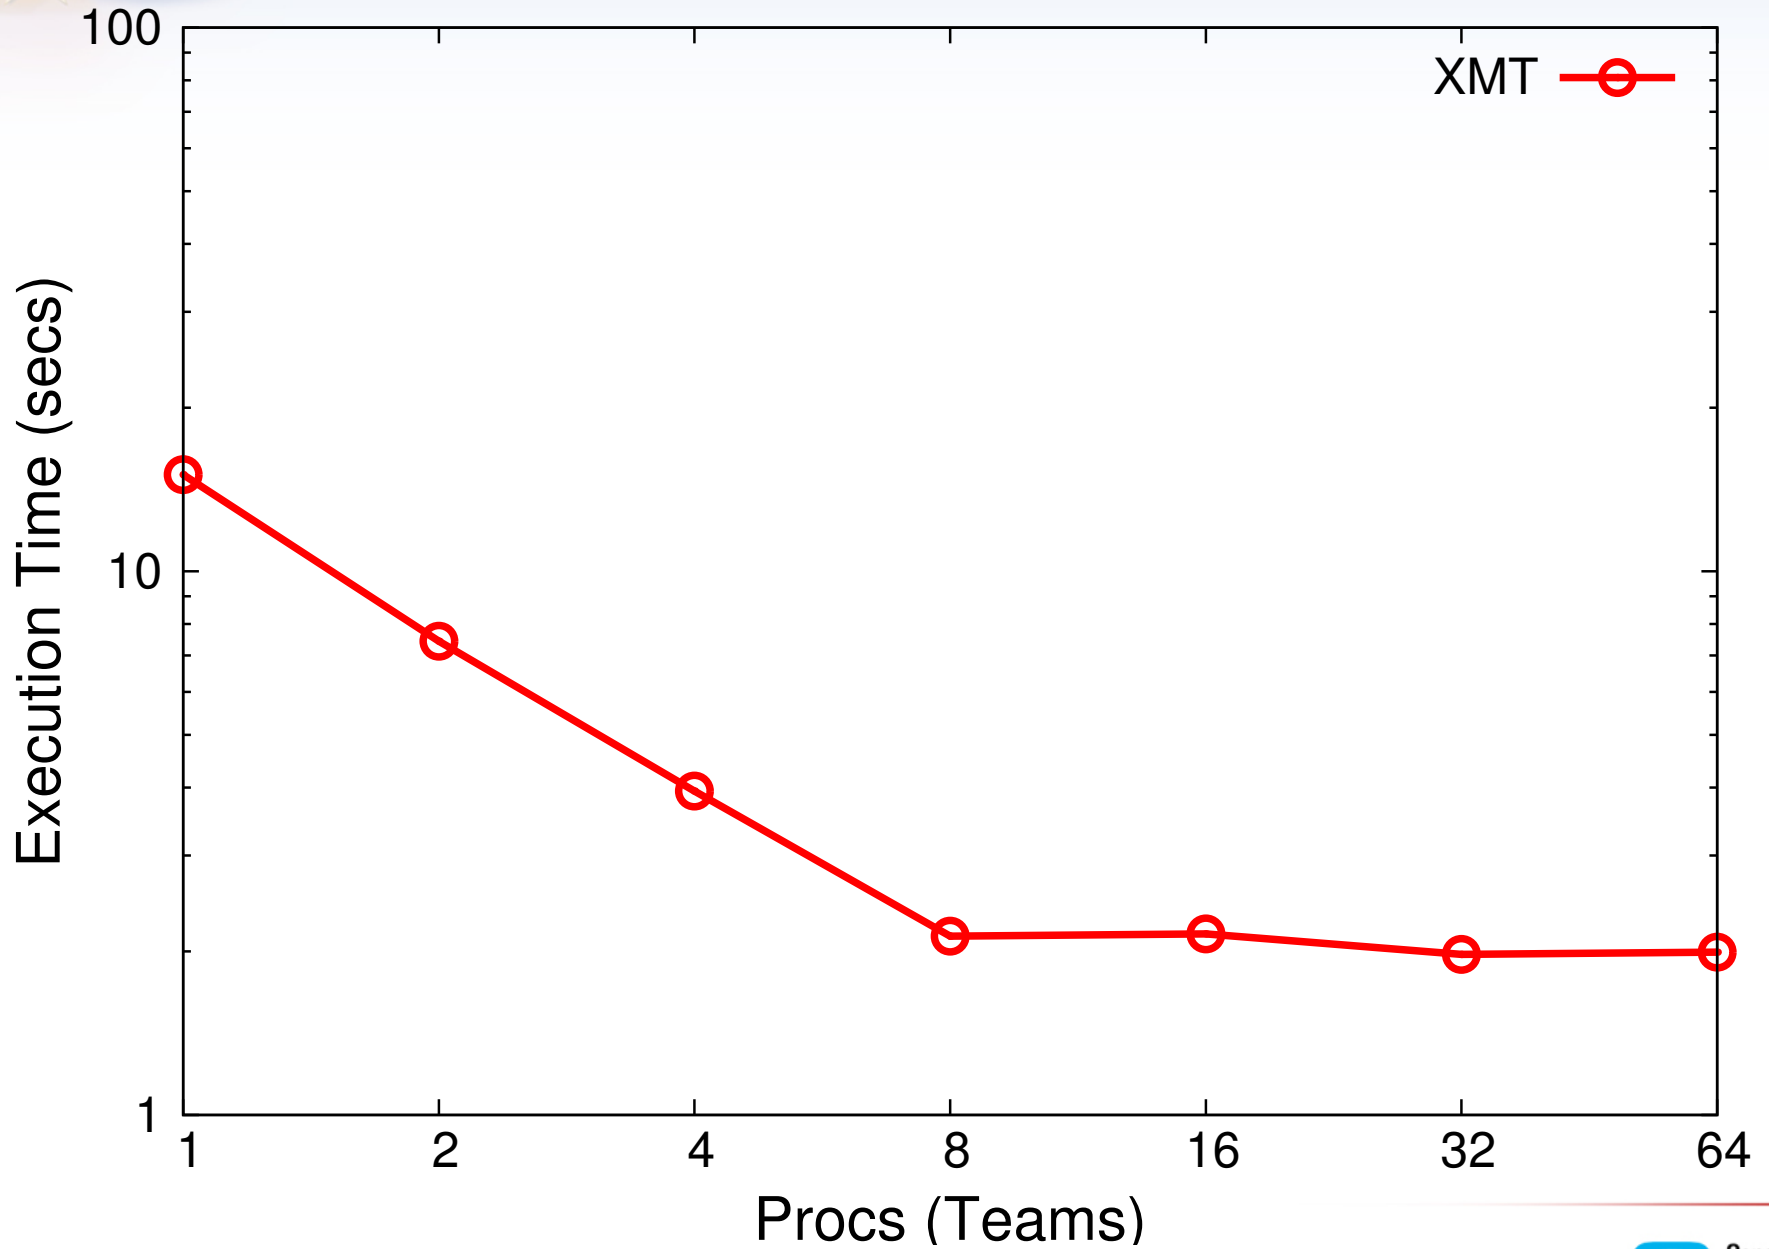
Example Problem

- **Concurrent Search**
- **R-MAT Graph**
 - $a=0.57$, $b=0.19$, $c=0.19$, $d=0.05$
 - Steep Degree Distribution Power Law Graph (max. degree $\sim 200k$)
 - $\sim 2^{25}$ vertices
 - $\sim 2^{28}$ edges

SMP Results



XMT Results






Caution Against Comparing Results

- **The problem is unstructured and responds to increased memory parallelism**
 - XMT has 512 memory controllers to push against any size problem
 - Would have to rewire the machine to compare on a per-controller basis
- **MTGL-based XMT implementation has been significantly performance tuned over many years**
 - Direct apples-to-apples comparison is unfair
 - Performance tuning on the other platforms is in the early stages
- **Graph 500 will have to address precisely these problems**
 - Desire to require “full memory” runs with *a posteriori* normalization of results (into Graph Operations Per Second, GROPS)
 - This is a really hard problem, and we may likely punt



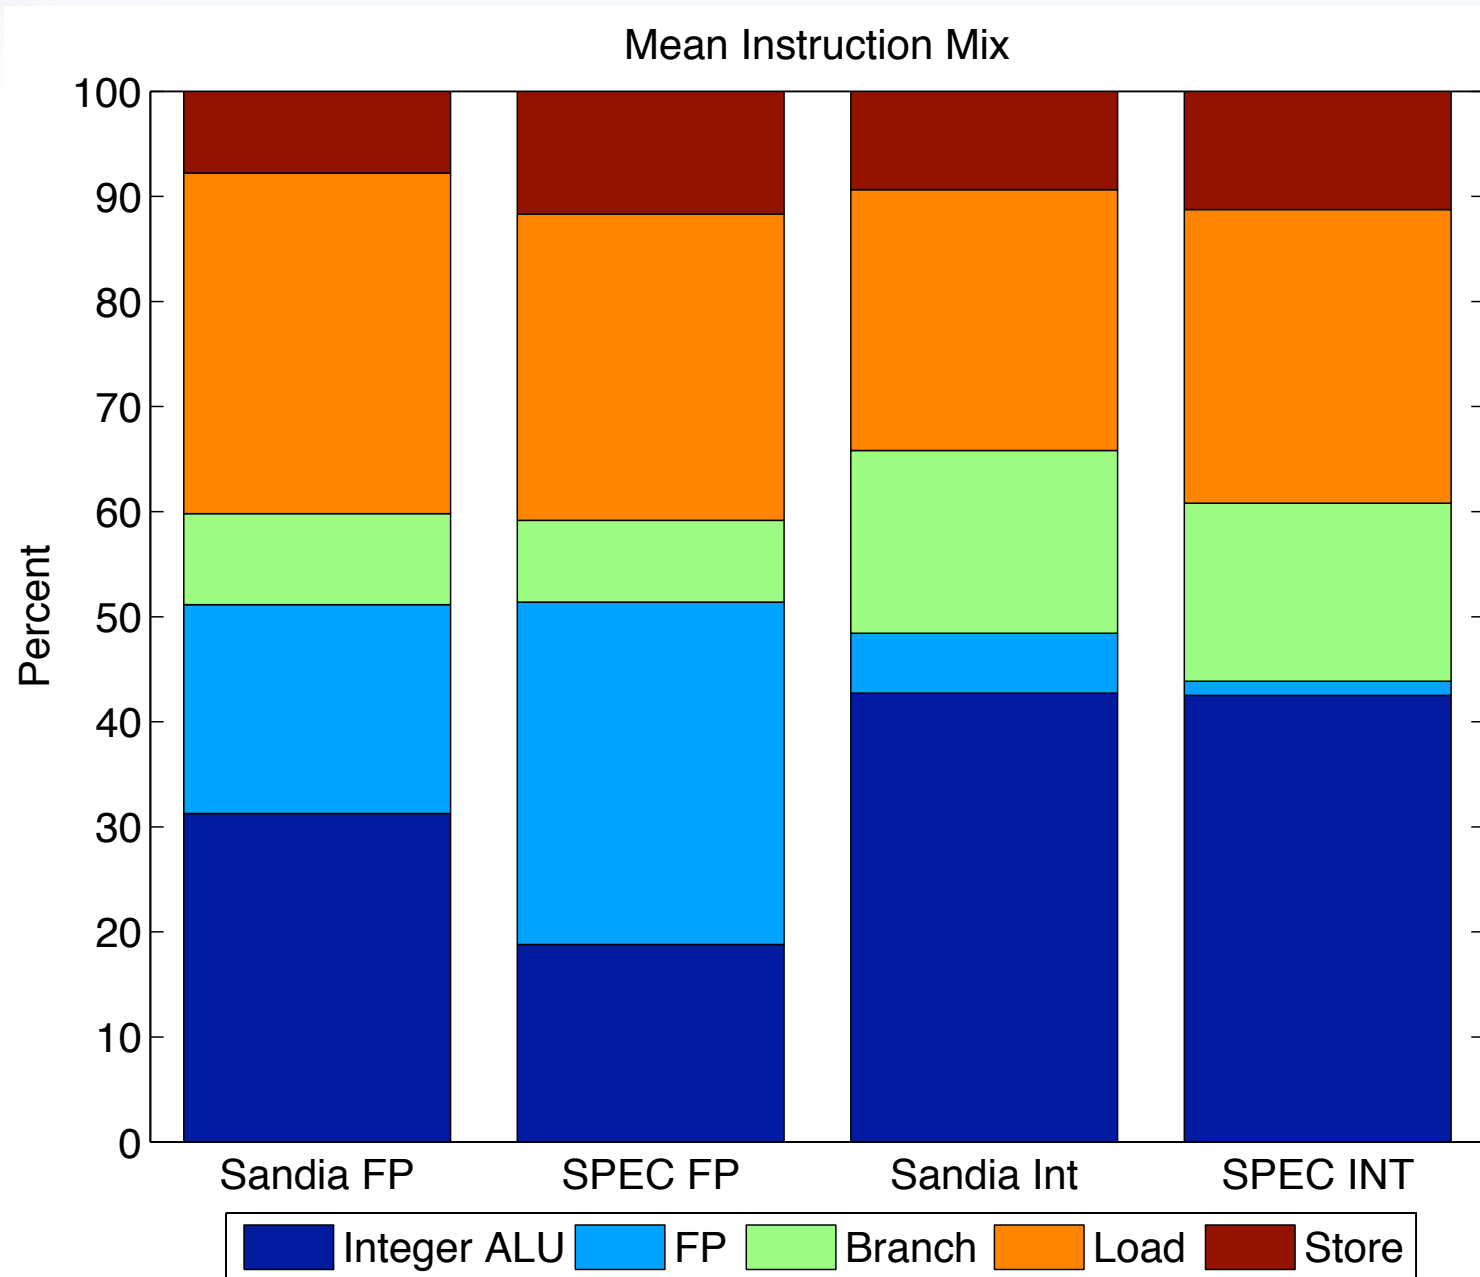
Conclusions

- **Lord Kelvin was Right**
 - “if you cannot measure it, you cannot improve it”
- **Graph 500 is an attempt to measure for an emerging critical problem domain**
- **We hope the five business areas will prove large enough to justify R&D investments**
 - We believe they are already potentially larger than HPC
 - Significant growth possible over the next decade
 - Impact into every day life
- **Roll Out**
 - Open Discussion throughout the summer of 2010 (including ISC BOF)
 - Benchmark Release in the Fall
 - First List at SC10



Thank You!

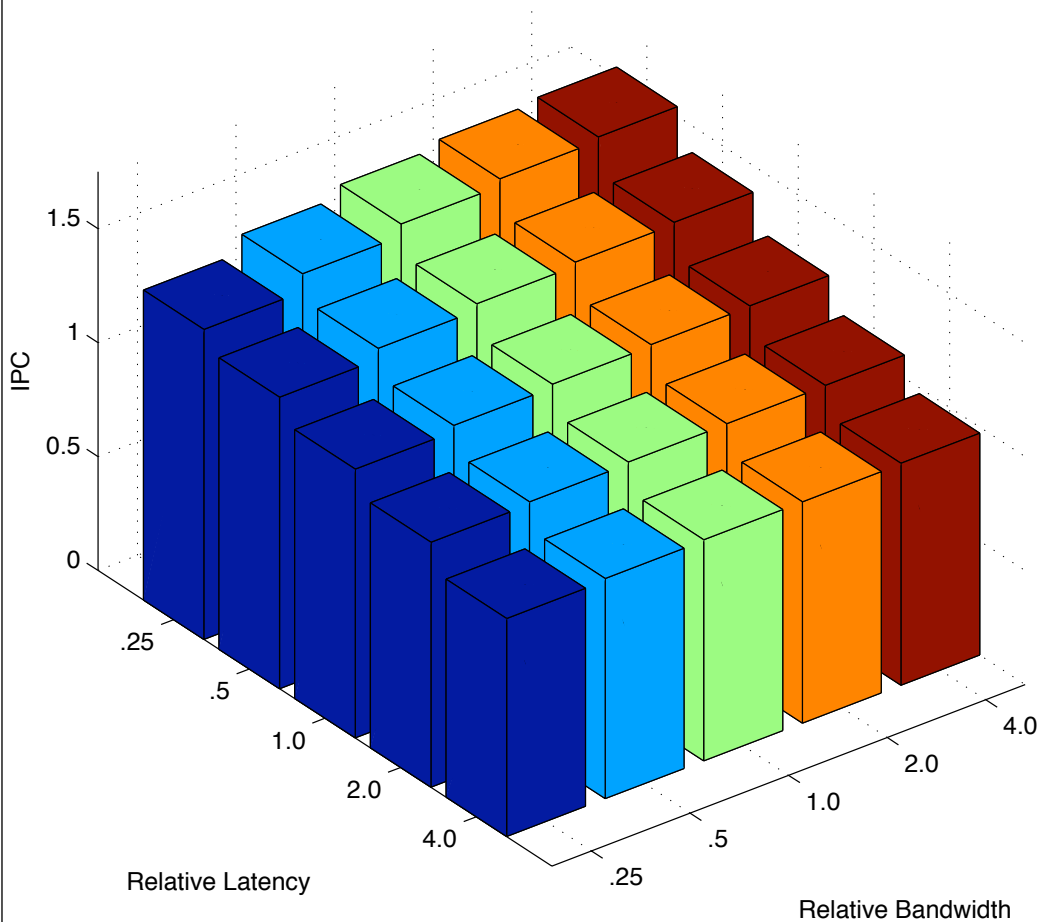
Most Real Applications Do Memory Accesses, Not Floating Point



Latency Dominates Bandwidth (Concurrency Decreases Effective Latency)

Physics

Average Sandia FP Latency and Bandwidth vs. Performance



Informatics

Average Sandia Int Latency and Bandwidth vs. Performance

