SAND2015-10181C

So You Might Want to Work at a National Lab?

Cynthia A. Phillips (Sandia National Laboratories)

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.

Performance

Why would we want to use a parallel algorithm for an application?

- When we have to:
  - Too Slow
  - Too Big
  - Too Inaccurate

- Application evolution
  - More constraints
  - Finer discretization
  - Larger instances

- Parallel computing is a tool for performance
  - We'd really like a faster, bigger serial machine
  - Make sure you beat the serial implementation
  - Better algorithms can be a better way to peformance

Embrace "Embarrassing" Parallelism

- Embarrassing parallelism increases the maximum feasible problem size
- Buys time to do the harder parallelization if necessary
- Using it can present other interesting algorithmic questions

- Some uses
  - Integer programming pseudocost initialization
  - Randomized rounding
  - Independent problems with different parameters for uncertainty quantification
  - Independent simulations to generate data for stochastic programs (sensor placement in water networks)

Parallel computing jobs at Sandia

- Scientific computing
- Data Sciences including streaming
- Operations Research
- Production computing
- Computer architecture and system softwares
  - Light-weight runtime systems, simulators,
  - Work with vendors
  - Neuromorphic computing

I interviewed other Sandia researchers to reflect more of this breadth

Perspective

Sandians solve national-scale problems

Many people who enter workforce with a bachelor's degree will program cell phones, Hadoop, single multicore

- Cost efficient
- Energy efficient

- Undergraduate programs have to serve them too

Scientific Computing

- Physical modeling and simulation (e.g. finite element)
- Sample of applications:
  - Shock physics
  - Molecular dynamics
  - Fluid flow
  - Quantum systems
  - Climate
  - More efficient nuclear energy reactors
- Enable the nuclear test ban
- Science, math, physics very important
  - Many Sandia computational scientists not CS

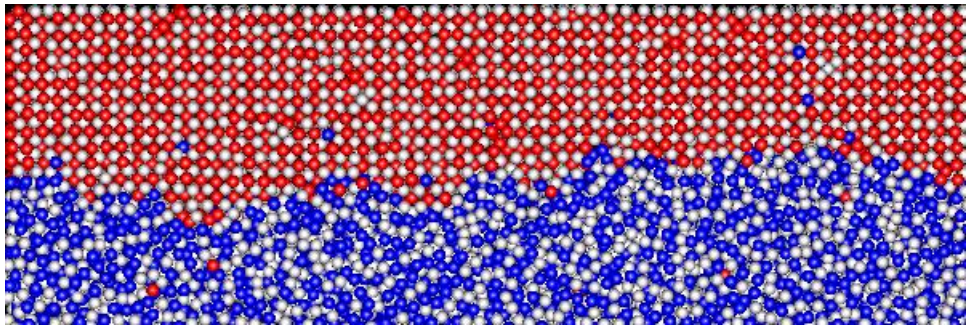

LAAMPS news release 3/3/11

Data Science

- Streaming data (e.g. cybersecurity)
- Streaming images (e.g. non-proliferation)
- Network science
  - Large-scale graph algorithms
- Scientific computing data
  - e.g. visualization

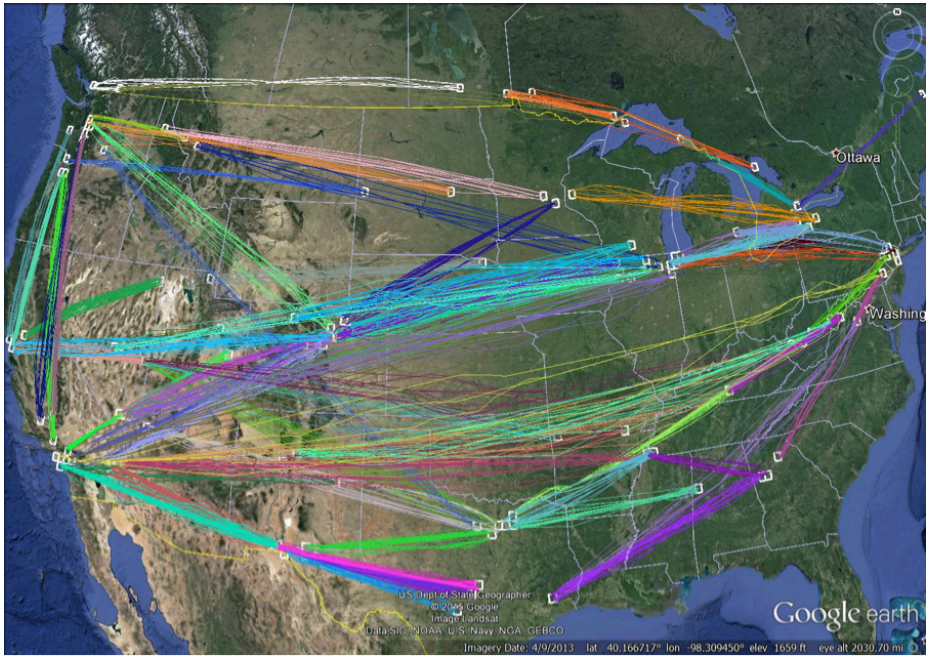

Danny Rintoul, 2015

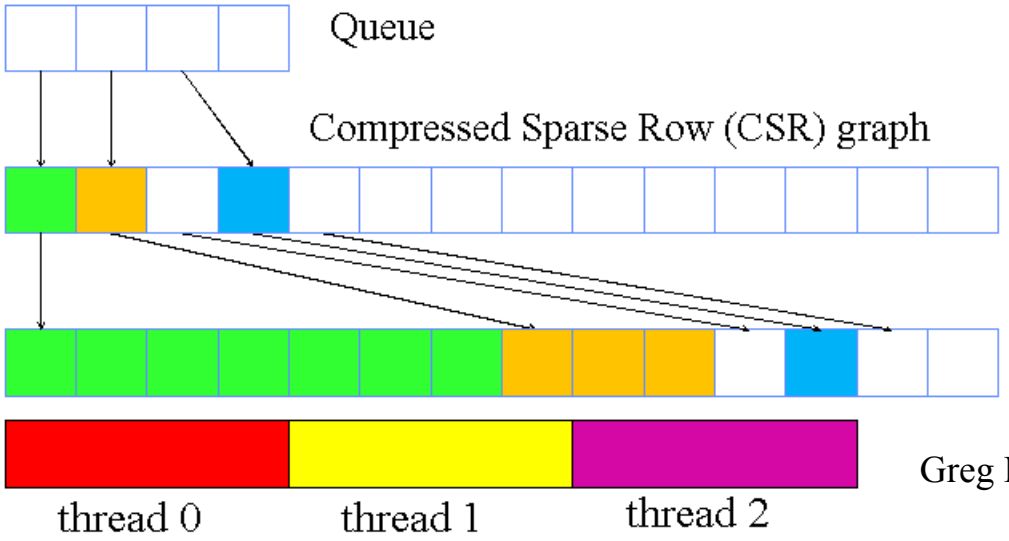

Greg Mackey, 2009

CCR Center for Computing Research

Sandia National Laboratories

Operations Research (Optimization)

- Branch and bound for general combinatorial optimization
- Stochastic programming

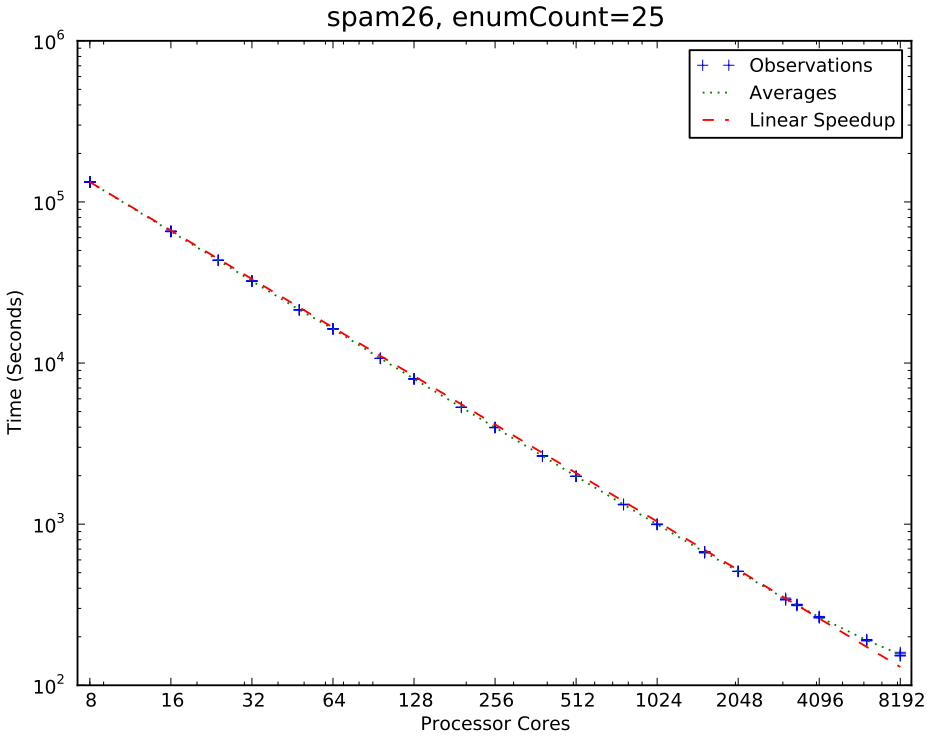

Spam detection, machine learning

Smart grid planning

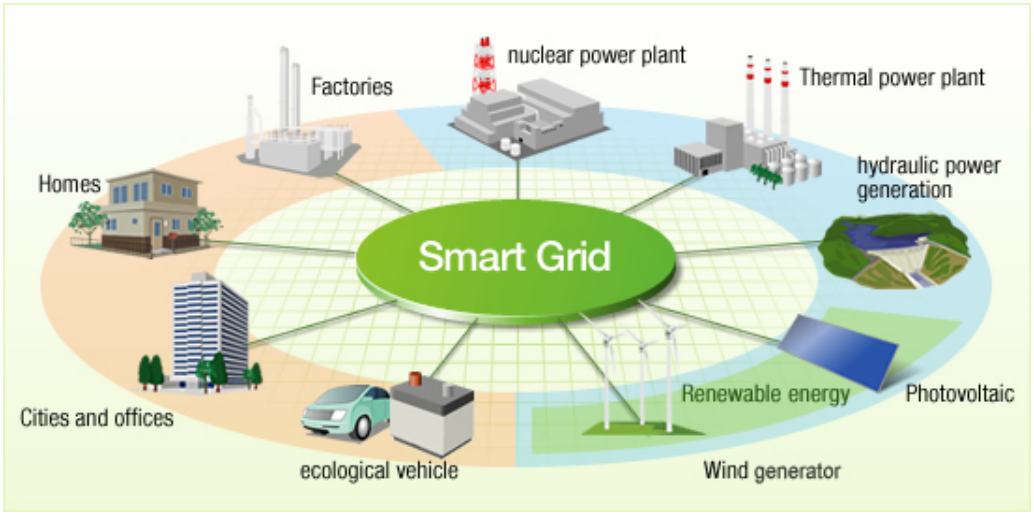

Hitachi.com, 2015

CCR Center for Computing Research

Sandia National Laboratories

Production Computing

- Stand up big clusters and administer
  - Create file systems
- Trying new hardware and software
  - Early prototypes, testbeds
  - No user guide. Become local expert
- User support
  - Teaching
  - Workflow, tool chain
  - code certification, testing, monitoring

- Need a deep understanding of hardware and software layers, and math/cs science for application support. Problems happen at scale that you will not learn about in school

Hardware (sample)

- Scientific computing researchers have access to much larger machines at other labs
  - Jaguar (Oak Ridge), > 200k opterons
  - Titan (Oak Ridge), 18k CPU, 18k GPU
  - Cielo (Los Alamos), 143k cores
  - Sequoia, BlueGene/Q (Lawrence Livermore), ~ 1.6M cores

Jaguar

- A Sandia researcher has run on 1.6M cores.

Hardware (sample)

- Capacity Clusters with thousands to 10,000s of cores
  - Capacity machines are for throughput of moderate jobs
  - Capability machines are for jobs that cannot run anywhere else
- Special data science clusters with more memory and better I/O
- GPUs, etc
- Cloudlike, etc

Sandia was an early
Adopter of the Cray XMT

wwjp.cray.com

Competitors

- Many industries want people who know HPC, big data, or can learn
  - Google, eBay, financial companies, Walmart, AmEx, Amazon
  - HPC on a multicore
  - Microsoft
  - Medical: real-time functional MRI
  - Even things like Shazam have a big search real-time problem

- Many undergraduates are employed early
  - Great for CS department stats
  - Still want to make their on-the-job learning easier

Parallel Thinking

- The details change frequently still
  - Architectures, languages, tools
- Shared memory processing
  - Learn the basics on something as small as a PC
- Distributed memory computing
  - Problem partitioning, message passing, load balancing
- Concurrency
  - Race conditions
  - One advanced topic: lock-free, wait-free algorithms
- Algorithms

Learn things that will stand the test of time, provide basis for career-long learning

Personal Characteristics

- Curious, willing to learn, learns for fun all the time
- Likes to fix things, make them better
- Attention to detail
- Smart, creative, motivated
- Willing to stay for a while
- Nice. Works well with others
- Flexibility in tool choice, etc

A Basic Skill: Memory Hierarchy

- Network is like another layer of the memory hierarchy
- Optimizing cache performance for serial problems directly helpful

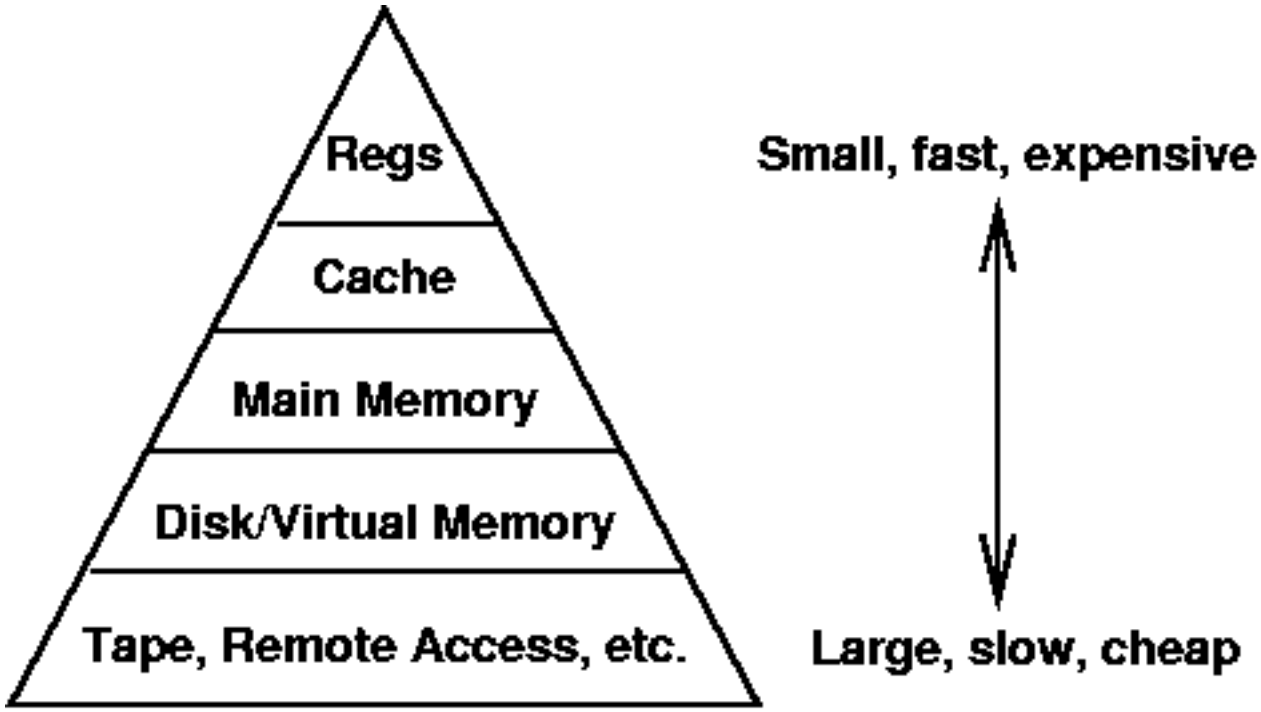

New kind of memory for Trinity

- Next capability machine for NNSA = National Nuclear Security Agency
- Will use Knight's Landing chip:

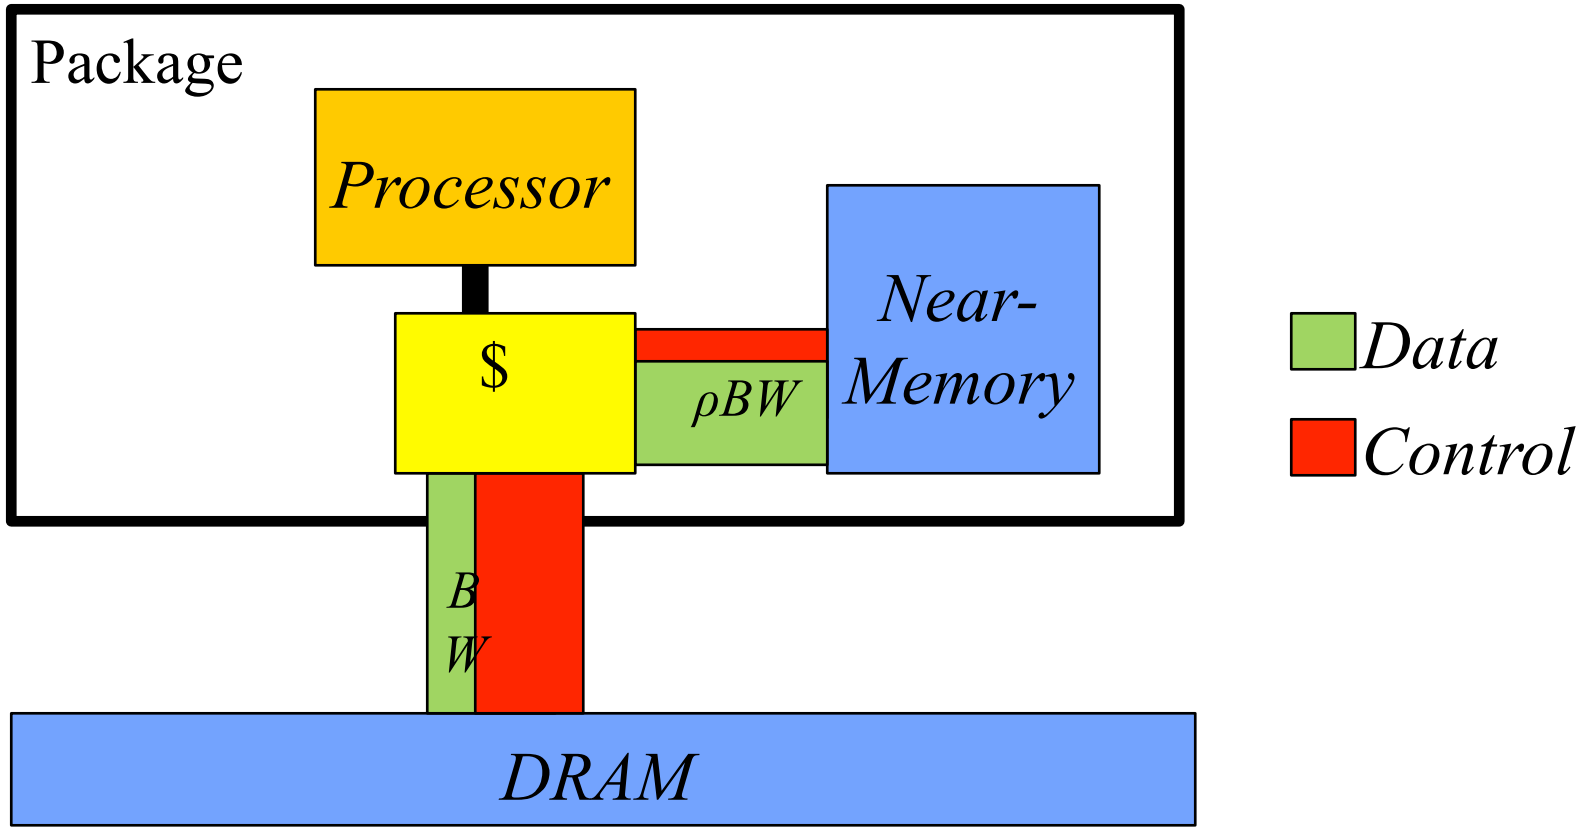

CCR Center for Computing Research

Sandia National Laboratories

New kind of memory for Trinity

- Use as a cache?
- User control: justify the cost of reprogramming

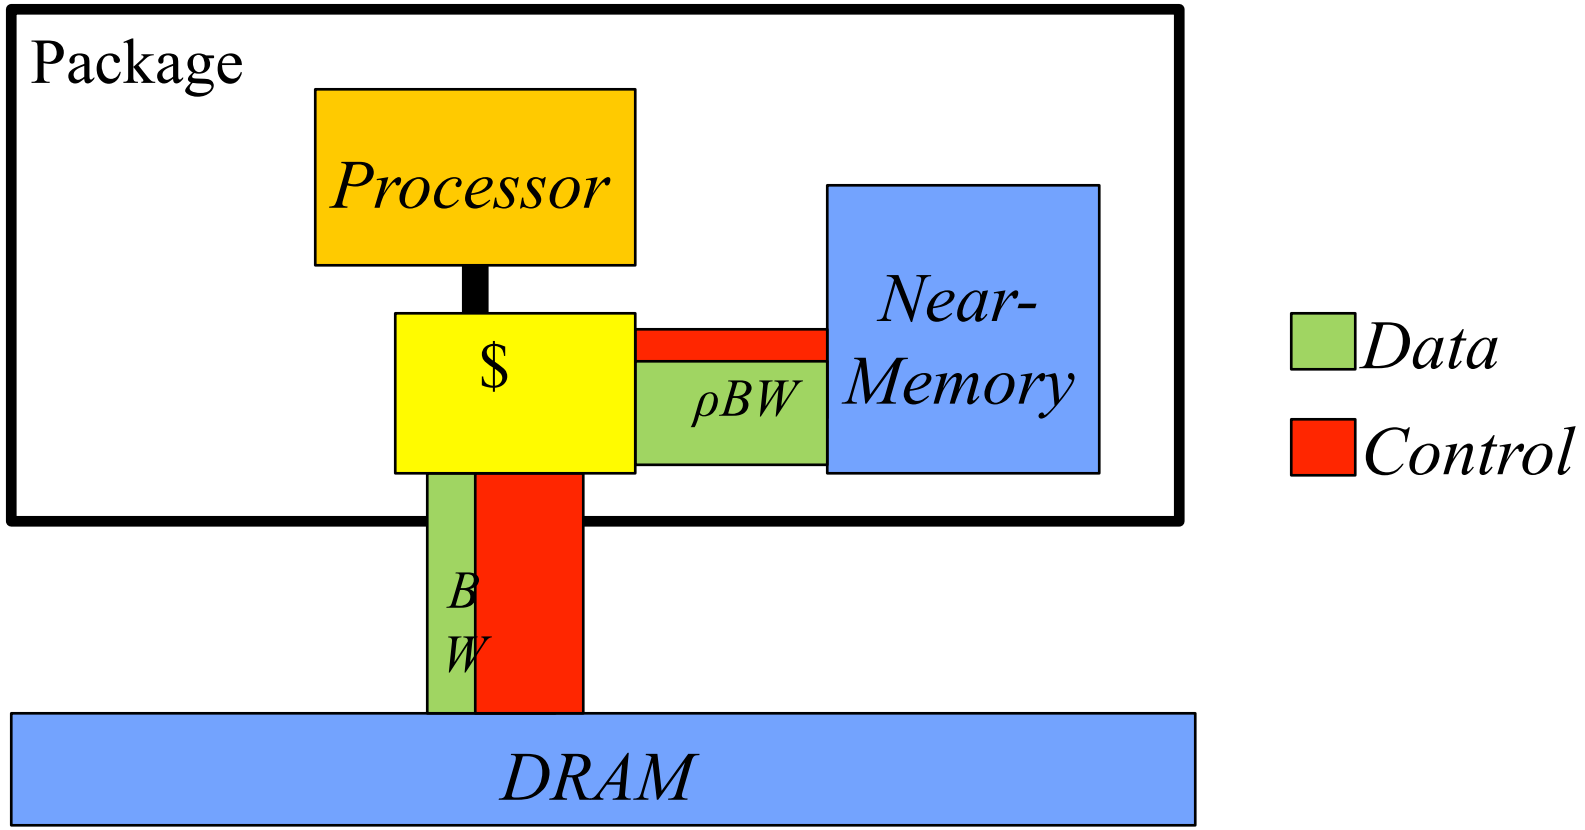

Necessary and Sufficient conditions

- Divide & conquer, memory boundedness analysis (IPDPS 2015, MemSys 2015)

(1) Memory-boundedness

Aggregate DRAM transfer takes more time than aggregate computation

(4) Final check: new algorithm isn't cache chunkable

If the new algorithm works efficiently with cache-sized chunks, scratchpad not needed

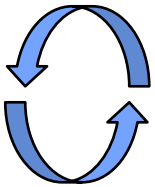

(2) Scratchpad chunkable

The computation can be divided into scratchpad-sized chunks that generate manageable partial results (a new algorithm)

(3) Chunk reuse

Each chunk loaded into near-memory must be reused enough to amortize the cost of the copy

Teams

- Especially for scientific computing problems, it takes a team to make a new code
  - Multiple specialties/disciplines required
    - One person cannot know everything
  - Individuals own different parts of the code/project
  - Shared credit

Software Engineering

- Use repositories, document, unit tests, good coding practice
- The performance expert has to speed up a code without deeply understanding the physics
  - Physics “owner” must set up tests to protect all the behavior he wishes to protect
  - Optimization person can commit anything that passes all tests
- The alternative is a cultural MPI barrier

Issue: Should use use the code written by the best physicist if it is impossible to test/verify?

Frameworks/packages

- Trilinos
  - solution of large-scale, complex multi-physics engineering and scientific problems
  - Many of the pieces one needs
    - Meshes, discretizations, load balancing, linear algebra, etc
- Multi-threaded graph library
  - Developed for the Cray XMT
  - Still used locally

Appropriate performance-optimized frameworks can help weather the changes in hardware.

Frameworks/packages

- Kokkos
  - Performance portability to many architectures
  - C++ templates
  - Abstractions for general performance (for, reduce, scan)
- Developers are part of ISO C++ committee
  - C++ wants to make parallelism “native” by 2020
  - Kokkos wants to meet C++ at that point

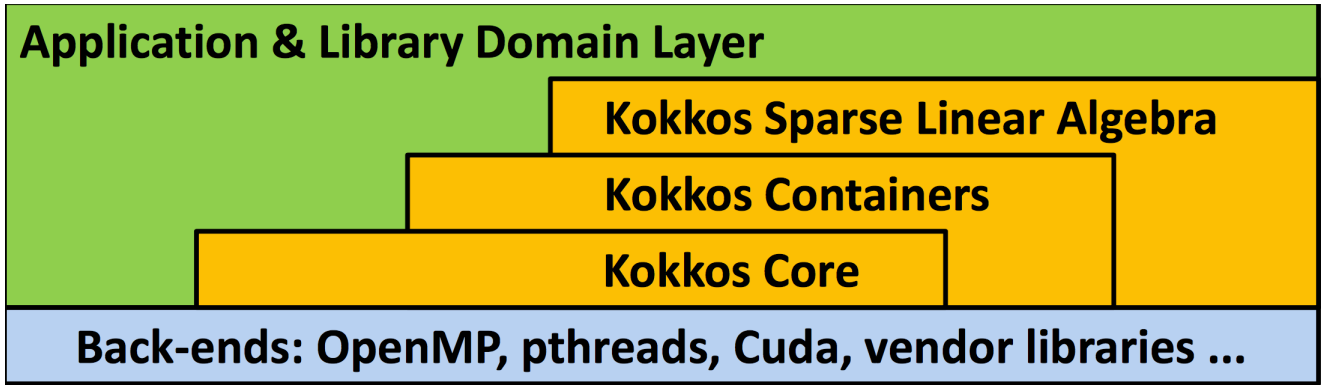

Scientific Computing Skills

- Good numerical math, numerical PDEs, grasp of discretization
- CS distributed memory computing, MPI background
  - Now moving to MPI + X
- Uncertainty quantification, verification (validation harder)
- Exemplar programs: U. Illinois and U. Texas Austin Institute for Computational Engineering and Sciences
- Predictive Science Academic Alliance program (PSAAP)
  - Department of Energy + National Science Foundation
  - Few universities, multi-million $, 4-5 years
  - Multidisciplinary. Each university a different application
- Gaming systems also need physics, visualization, and realism

Computing for Scientists

I think the science is harder

What might a CS department put in a graduate computing class for scientists (chem E, mech E, physics, computational materials)?

- Perspective on algorithms
  - Parallel and Distributed Computation" by Bertsekas and Tsitsiklis. Largely numerical, but high-level, abstract
  - Gives resilience as parallel computing changes
- Basic software engineering
  - Suggested class exercise: use someone matrix-vector multiply in your conjugate-gradient code

Performance Modeling

- Hardware counters, performance tools, vampir, etc
- Strong scalability
- Weak scalability
- Work/span
  - T_p = time to run an instance on p processors
  - T_1 = serial. T_∞ = minimum time (serial part)
  - Work law: $pT_p \geq T_1$ (Except for cache effects)
  - Span law: $T_p \geq T_\infty$
- A basic understanding of general analysis

Streaming

- Perfect answer vs timely answer
  - Ads, bond trading, cybersecurity
- Write-optimized data structures (databases)
- Streaming benchmarking for acceptance testing
- Cyber (monitoring) model: parallel models to approximate

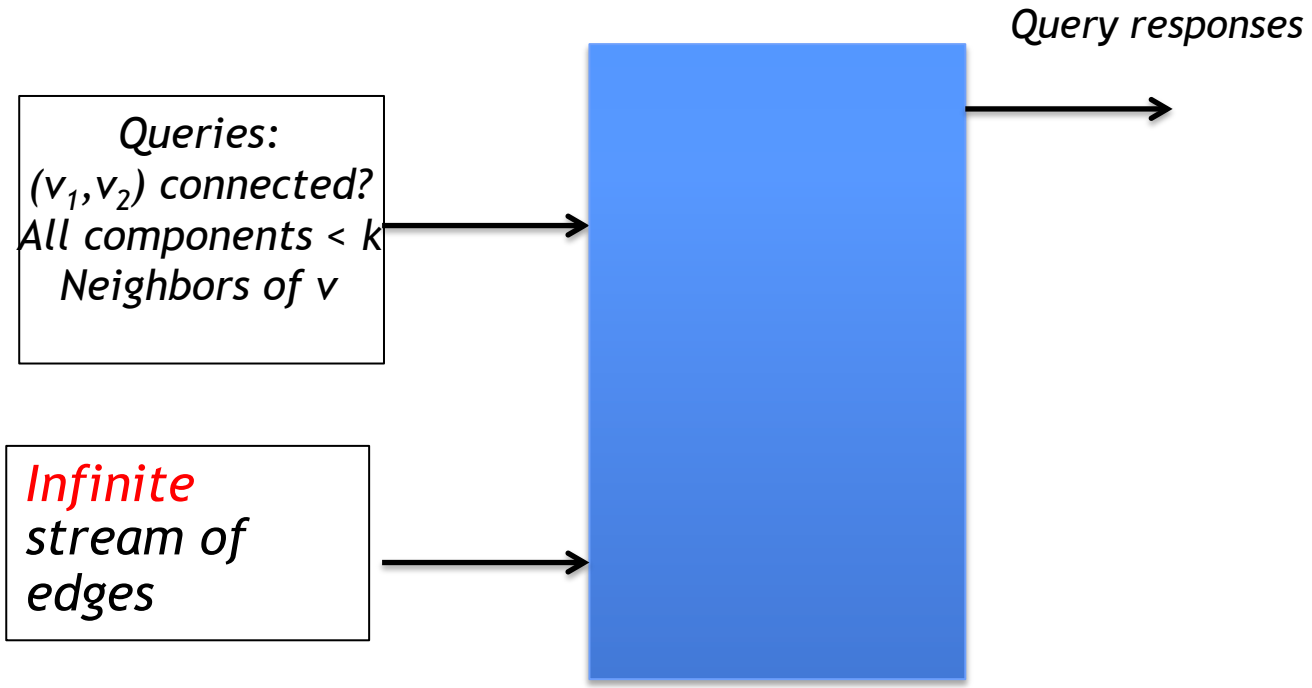

Paid Internships

- 'Tis the season
  - Labs usually handle summer internships November-December
  - Later is possible, but suboptimal
- A good way to write a program for a larger data set