PARLab Parallel Boot Camp

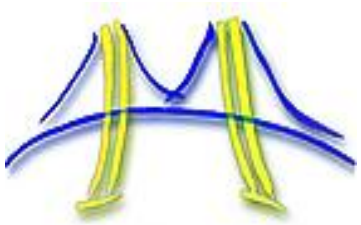

Short Course on Parallel Computing
August 19-21, 2013
parlab.eecs.berkeley.edu/2013bootcamp

Jim Demmel
EECS and Mathematics
University of California, Berkeley

Outline

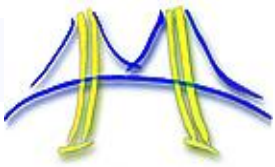

- Motivation and Goals
- Background
  - ParLab and ASPIRE, research centers in Parallel Computing
  - The Designated Emphasis (DE) in Computational Science and Engineering (CSE)
  - CSE at Lawrence Berkeley National Lab
  - XSEDE – Provide access to NSF's cyberinfrastructure, and educate users
- Schedule and Instructors
- Logistics
- The Audience

Motivation (1/2)

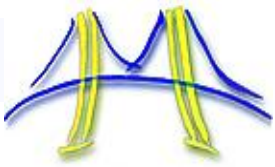

- Parallel Computing is becoming ubiquitous
  - Only way forward for computing industry (unless you don't care if your programs never run faster than in 2008)
  - Unfortunately, parallel programming is (still) harder than sequential programming
  - Better (easier) programming tools under development, but we still need to train people to "think parallel"
- So welcome!

Motivation (2/2)

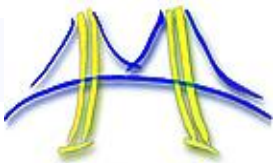

- Recent events at UCB will provide support for many new activities to develop and use parallel computing
  - ParLab established – parlab.eecs.berkeley.edu
    - » Research center about "Multicore Revolution"
    - » Many related, large grants too (AmpLAB, …)
  - ASPIRE established – follow-on to ParLab
  - Designated Emphasis in Computational Science and Engineering (CSE) established – cse.berkeley.edu
    - » Graduate program with 117 faculty from 22 depts.
  - XSEDE – NSF follow-on to Teragrid
    - » Broadcasting Berkeley's parallel computing courses
  - CRT Building @ LBL – new building under construction, to house supercomputers and research teams

Short Course Goals

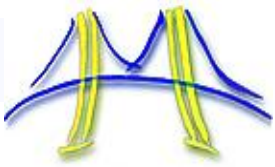

- Teach the basics about parallelism
  - How to program, including hands-on lab
- Tools you can use now (simple and advanced)
- Tools we hope to build, and ongoing research

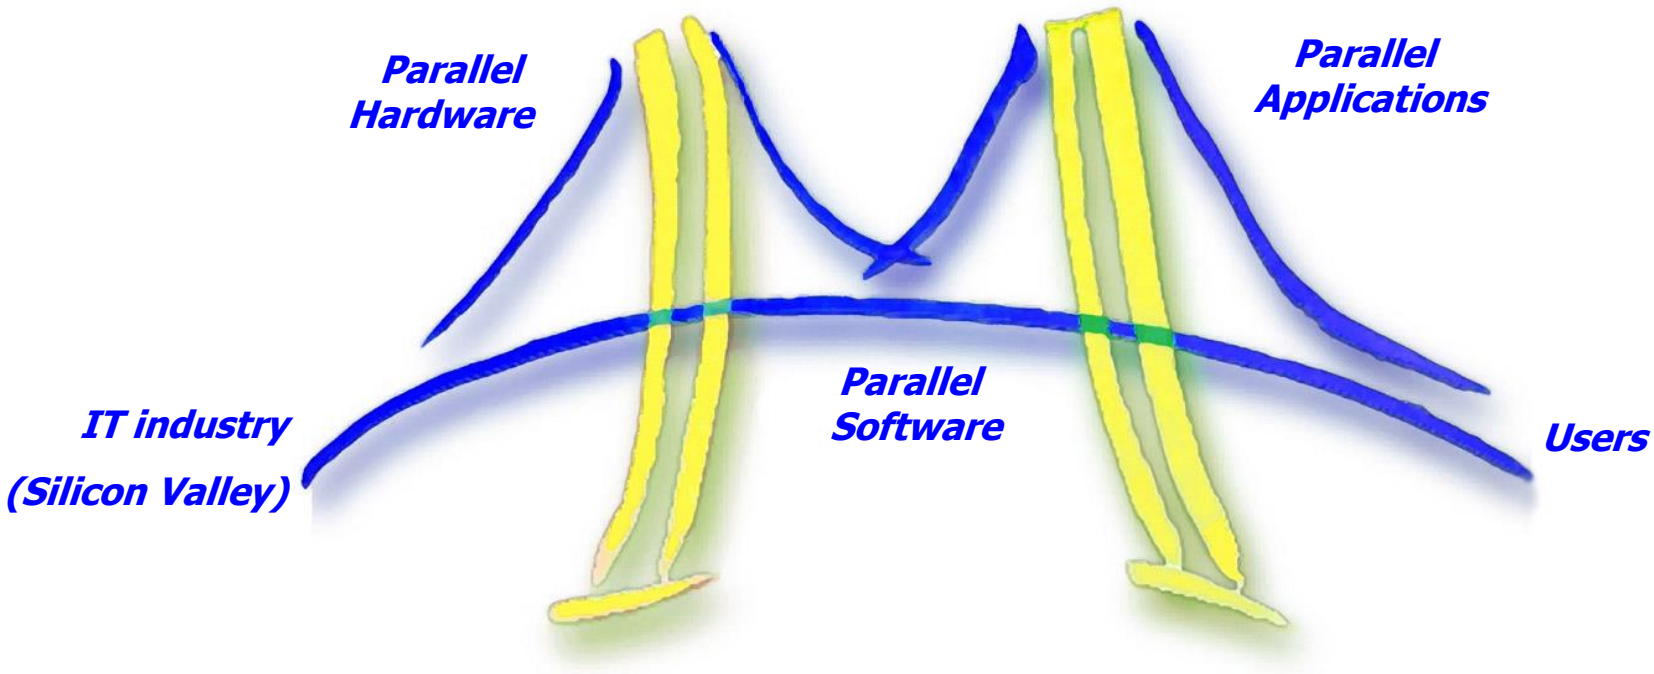

Berkeley ParLab Project

Krste Asanovic, Ras Bodik, Jim Demmel, Tony Keaveny, Kurt Keutzer, John Kubiatowicz, Edward Lee, Nelson Morgan, Dave Patterson, Koushik Sen, John Wawrzynek, David Wessel, and Kathy Yelick

7 Dwarfs of High Performance Computing (HPC)

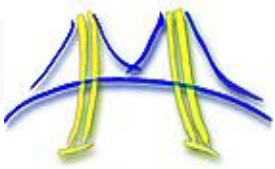

Structured Grid
Dense Matrix
Sparse Matrix
Spectral (FFT)
Particle Methods
Unstructured Grid
Monte Carlo

7 Dwarfs – Are they enough?

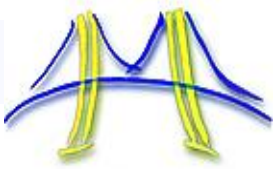

	Embed	SPEC	DB	Games	ML	CAD	HPC
Structured Grid							
Dense Matrix							
Sparse Matrix							
Spectral (FFT)							
Particle Methods							
Unstructured Grid							
Monte Carlo							

13 Motifs (nee "Dwarf") of Parallel Computing

Popularity: (Red Hot / Blue Cool)

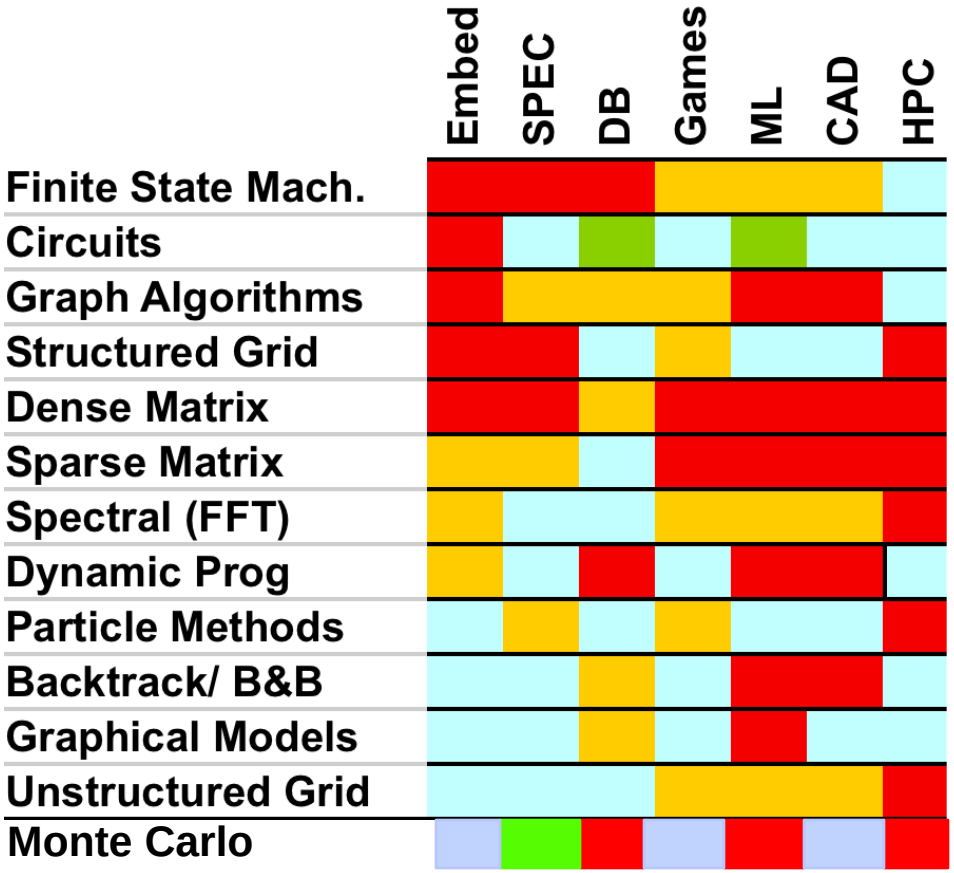

Motifs in ParLab Applications
(Red Hot / Blue Cool)

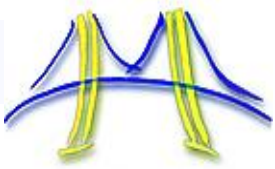

		Embed	SPEC	DB	Games	ML	CAD	HPC	Health	Image	Speech	Music	Browser
1	Finite State Mach.												
2	Circuits												
3	Graph Algorithms												
4	Structured Grid												
5	Dense Matrix												
6	Sparse Matrix												
7	Spectral (FFT)												
8	Dynamic Prog												
9	Particle Methods												
10	Backtrack/ B&B												
11	Graphical Models												
12	Unstructured Grid												

- What happened to Monte Carlo?

Our Pattern Language 2.0

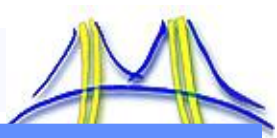

Applications

Choose your high level structure – what is the structure of my application?

Guided expansion

- Pipe-and-filter
- Agent and Repository
- Process Control
- Event based, implicit invocation

Choose you high level architecture? Guided decomposition

Task Decomposition ↔ Data Decomposition

Group Tasks Order groups data sharing data access

- Model-view controller
- Iteration
- Map reduce
- Layered systems
- Arbitrary Static Task Graph

Identify the key computational patterns – what are my key computations? Guided instantiation

- Graph Algorithms
- Dynamic Programming
- Dense Linear Algebra
- Sparse Linear Algebra
- Unstructured Grids
- Structured Grids
- Graphical models
- Finite state machines
- Backtrack Branch and Bound
- N-Body methods
- Circuits
- Spectral Methods

Productivity Layer

Refine the structure – what concurrent approach do I use? Guided re-organization

- Event Based
- Divide and Conquer
- Data Parallelism
- Geometric Decomposition
- Pipeline
- Discrete Event
- Task Parallelism
- Graph algorithms
- Digital Circuits

Efficiency Layer

Utilize Supporting Structures – how do I implement my concurrency? Guided mapping

- Fork/Join
- CSP
- Distributed Array
- Shared Data
- Shared Queue
- Shared Hash Table
- SPMD
- Master/worker
- Loop Parallelism
- BSP

Implementation methods – what are the building blocks of parallel programming? Guided implementation

- Thread Creation/destruction
- Process Creation/destruction
- Message passing
- Collective communication
- Speculation
- Transactional memory
- Barriers
- Mutex
- Semaphores

Algorithms and Specializers for Provably Optimal Implementations with Resiliency and Efficiency

Krste Asanovic, Jonathan Bachrach, Jim Demmel,
Armando Fox, Kurt Keutzer, Borivoje Nikolic,
David Patterson, John Wawrzynek

http://aspire.eecs.berkeley.edu

Designated Emphasis (DE) in Computational Science and Engineering (CSE)

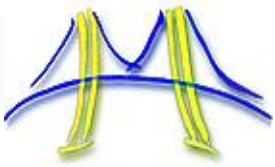

- Goals
- Participants (117 faculty from 22 departments – so far)
- How the DE works
- (New) courses
- Details at citris-uc.org/decse

Designated Emphasis (DE) in CSE

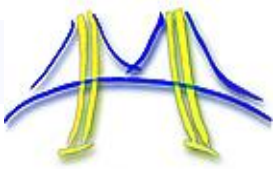

- "Graduate minor"
- Motivation
  - Too little data – too slow (climate), too expensive (building chips) or too dangerous (crash testing) to get experimental data, so need to simulate
  - Too much data – from DNA sequencers, telescopes, WWW, simulations… so need large scale statistical analysis and machine learning
  - Widespread need to train PhD students (eng. to astro to social science to …)
  - Opportunities for collaboration, across campus and at LBNL
- Graduate students participate by
  - Getting accepted into existing department/PhD program
  - Taking CSE course requirements
  - Qualifying examination with CSE component
  - Thesis with CSE component
  - Receive "PhD in X with a DE in CSE"
  - Details at citris-uc.org/decse

Participating Departments (1/2)

(# faculty by "primary affiliation", # courses)

- Astronomy (7,3)
- Bioengineering (3,1)
- Biostatistics (2,0)
- Chemical Engineering (6,0)
- Chemistry (8,1)
- Civil and Environmental Engineering (7,8)
- Earth and Planetary Science (6,3)
- EECS (19,14)
- IEOR (5,5)
- School of Information (1,0)

Participating Departments (2/2)

(# faculty by "primary affiliation", # courses)

- Integrative Biology (1,0)
- Materials Science and Engineering (2,1)
- Mathematics (15, 4)
- Mechanical Engineering (9, 6)
- Neuroscience (7,1)
- Nuclear Engineering (2,1)
- Physics (1,1)
- Political Science (2,0)
- Public Health (2,0)
- Statistics (5, 11)

Example Course – CS267

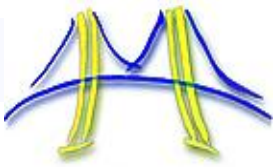

- "Applications of Parallel Computing"
  - Long version of this short course!
  - see www.cs.berkeley.edu/~demmel/cs267_Spr13
- Taught every Spring
  - All lectures on web (slides + video), freely available
  - UC Berkeley, UC Merced, UC Santa Cruz, UC Davis in Spr09
- Provided nationwide by XSEDE starting Spring 13
  - Homework done on NSF supercomputer centers (free)
  - For this bootcamp too!
  - Autograding provided

A few sample CS267 Class Projects (all posters and video on web pages)

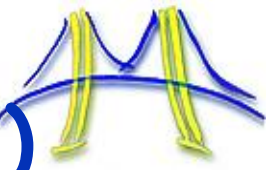

- Content based image recognition
  - "Find me other pictures of the person in this picture"
- Faster molecular dynamics, applied to Alzheimer's Disease
- Better speech recognition through a faster "inference engine"
- Faster algorithms to tolerate errors in new genome sequencers
- Faster simulation of marine zooplankton population
- Sharing cell-phone bandwidth for faster transfers, better gaming experience
- Economic implications of parallel computing

New CSE Courses

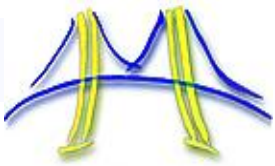

- Python for science – AY250
  - Josh Bloom (Astronomy)
  - 3 day summer short course (Aug 20-22, this week!) + seminar
- SW Eng. for Scientific Computing – CS194/294
  - Phil Colella (EECS,LBL)
  - For non-CS grads and undergrads
- Understanding Molecular Simulation
  - Phil Geissler (Chem) and Berend Smit (ChemE)
  - Matlab based, students from Chem, ChemE, MSE, ME, BioPhys
- Computer Simulations in the Earth Sciences – EPS109
  - Burkhard Militzer (Earth & Planetary Science)
  - Machine learning for understanding simulations/data sets, in Matlab
- Optimization Models in Engineering – EE127
  - Laurent El Ghaoui (EECS)
  - Matlab (CVX) based, models not algorithms

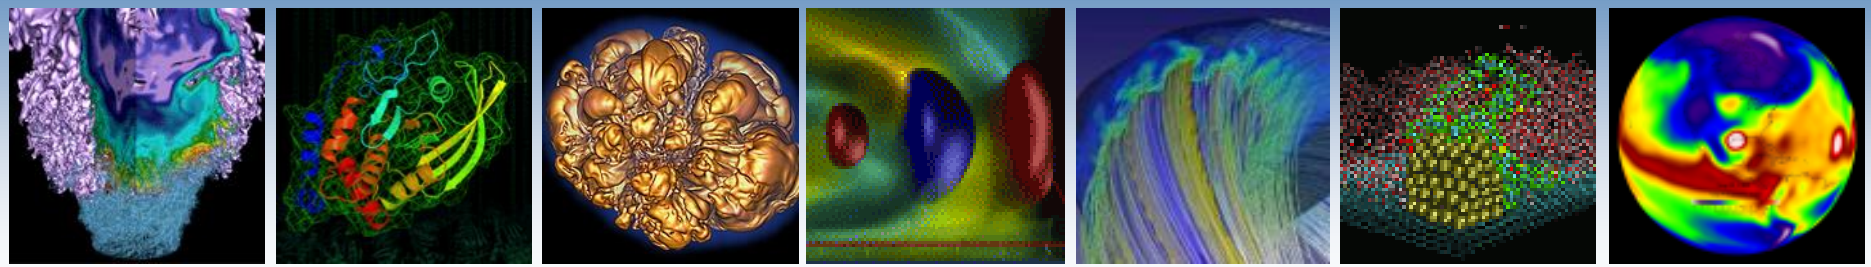

Computing Sciences at Berkeley Lab

Kathy Yelick

Associate Lab Director for Computing Sciences
Lawrence Berkeley National Laboratory

Professor of Electrical Engineering and Computer Sciences
University of California at Berkeley

National Energy Research Scientific Computing Facility

Department of Energy Office of Science (unclassified) Facility

- 4500 users, 600 projects
- 65% from universities
- ***1500 refereed publications per year***
- ***Key to 2 Nobel Prizes (2007, 2011)***

Systems designed for science

- 1.3 PF Hopper system (Cray XE6)
  11th Fastest computer in US, 24th in world
- New system on order

Current NERSC Systems

Large-Scale Computing Systems

Hopper (NERSC-6): Cray XE6

- 6,384 compute nodes, 153,216 cores
- 144 Tflop/s on applications; 1.3 Pflop/s peak

Edison (NERSC-7): Cray Cascade

- Being installed now 2013
- Over 200 Tflop/s on applications, 2 Pflop/s peak

Midrange Systems

140 Tflops total

Carver

- IBM iDataplex cluster
- 9884 cores; 106TF

PDSF (HEP/NP)

- ~1K core cluster

GenePool (JGI)

- ~5K core cluster
- 2.1 PB Isilon File System

NERSC Global Filesystem (NGF)

Uses IBM's GPFS

- 8.5 PB capacity
- 15GB/s of bandwidth

HPSS Archival Storage

- 240 PB capacity
- 5 Tape libraries
- 200 TB disk cache

Analytics & Testbeds

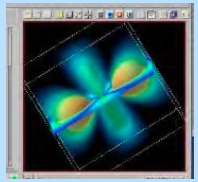

Euclid (512 GB shared memory)

Dirac 48 Fermi GPU nodes

Magellan Hadoop

Computational Research Division

Computational Science

Combustion *Climate* *Energy & Environment* *Cosmology & Astrophysics* *Nanoscience* *Genomics*

Applied Mathematics

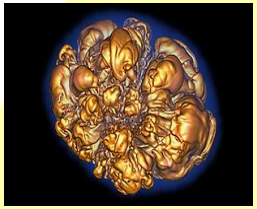

Mathematical Models

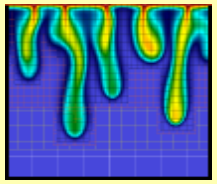

Adaptive Mesh Refinement

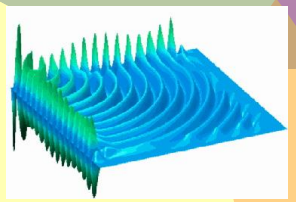

Libraries and Frameworks

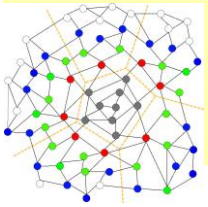

Linear Algebra

Interface Methods

Computer Science

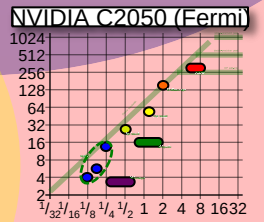

Performance & Autotuning

Cloud, grid & distributed computing

Visualization and Data Management

HPC architecture, OS, and compilers

Computational Science at Berkeley Lab

Large-Scale Simulations and Large Numbers of Simulations

Large-Scale discovery of Events

- Petascale simulations produce data too large for manual analysis
- Data analysis using new algorithms (FastBit, Machine Learning) to discover events

Materials Project (joint with MIT)

- Tens of thousands of simulations screen materials
- Goal: cut in half the 18 years from design to manufacturing
- Advance machine learning and data systems

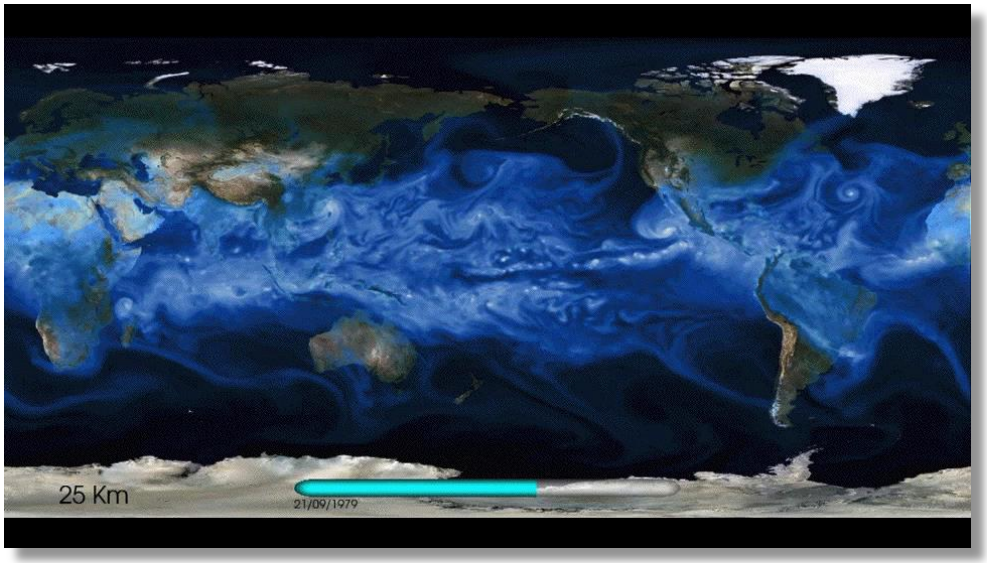

Climate Code is a large community effort.
This set of data analysis and simulations by Prabhat and Michael Wehner, LBNL

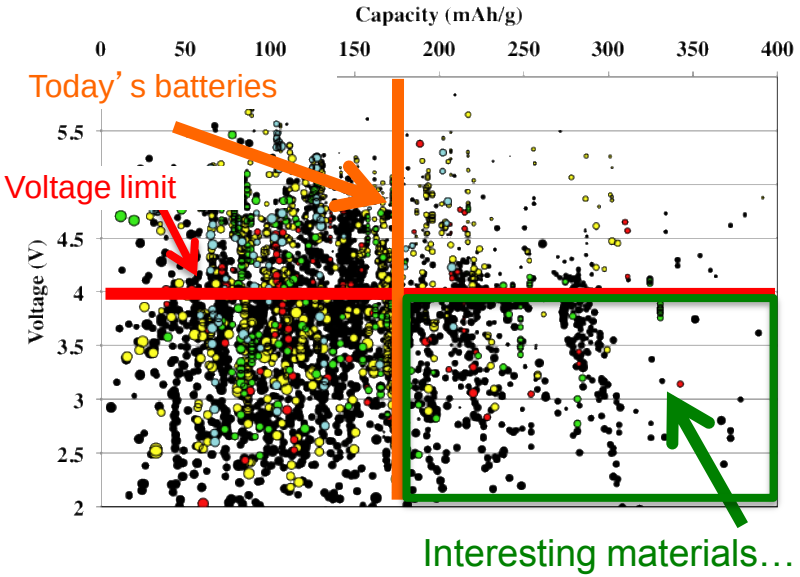

Materials Project PIs: Gerd Ceder, MIT and Kristin Persson, LBNL

From Teragrid to XSEDE

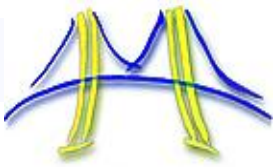

- Teragrid
  - Easy access to NSF cyberinfrastructure
  - Supercomputers, storage, visualization, networks
  - Education, training and support of users

- XSEDE: Extreme Science and Engineering Discovery Environment
  - Next-Generation Teragrid – www.xsede.org
  - Started July 2011, 17 institutions, $121M over 5 years
  - More educational activities
    - Broadcast selected courses
    - 4 courses (so far) from Berkeley:
      - This bootcamp, CS267, ACTS Workshop, Keutzer's CS194
    - Free accounts on NSF supercomputers to do homework

Schedule and Instructors (1/3)

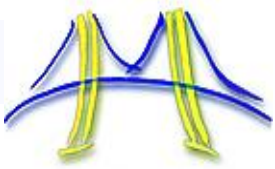

- Monday, Aug 19
  - 9-9:30 am – Introduction and Welcome
    - » Jim Demmel (UCB)
  - 9:30-12pm – Introduction to Parallel Architectures and Pthreads
    - » John Kubiatowicz (UCB)
  - 12-1:15pm – Lunch (see web page for suggested venues)
  - 1:15-2:15pm – Shared Memory Programming with OpenMP - Basics
    - » Tim Mattson (Intel)
  - 2:15-3:00pm – More about OpenMP – New Features
    - » Tim Mattson (Intel)
  - 3:00-3:30pm – Break
  - 3:30-4:30pm – Performance Tuning Random Slowdowns in Recurring Functionalities
    - » Gary Carleton (Intel)
  - 4:30-5:00pm – Break/Transition to Hands-on Lab Rooms
  - 5:00-6:00pm – Hands-on Lab (Parallel Sessions)
    - » Introduction to NERSC Tools (*273, 275, 277 Soda*)
    - » Introduction to OpenMP (*273, 275, 277 Soda*)
  - 6:00-7:00pm –Reception in Soda Hall, 5th Floor!

Schedule and Instructors (2/3)

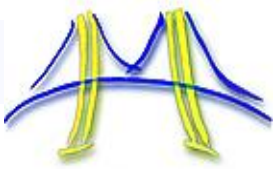

- Tuesday, Aug 20
  - 8:45-9:45am – Distributed Memory Programming in MPI
    - » Tim Mattson (Intel)
  - 9:45-10:45am – Sources of Parallelism and Locality in Simulation
    - » Jim Demmel (UCB)
  - 10:45-11:15am – Break
  - 11:15-12:15am – Architecting parallel software with design patterns
    - » Kurt Keutzer (UCB)
  - 12:15-1:30pm – Lunch
  - 1:30-2:30pm – GPU, CUDA, OpenCL Programming
    - » Bryan Catanzaro (NVIDIA Research)
  - 2:30-3:00pm - Break / Transition to Hands-on Lab Rooms in Soda Hall
  - 3-6pm – Hands-on Lab (Parallel Sessions)
    - » NERSC Tools (*273, 275, 277 Soda*)
    - » OpenMP (*273, 275, 277 Soda*)

Schedule and Instructors (3/3)

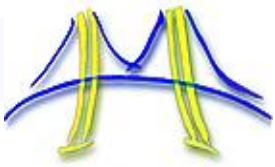

- Wednesday, Aug 21
  - 8:45-9:45am – Partitioned Global Address Space Programming in UPC
    - » Kathy Yelick (UCB)
  - 9:45 – 10:15 – Break
  - 10:15-12:15am –Computational Patterns and Autotuning
    - » Jim Demmel (UCB)
  - 12:15-1:30pm – Lunch
  - 1:30–2:30pm – Performance Debugging: Methods and Tools
    - » David Skinner (LBL)
  - 2:30-3:30pm – Cloud Computing using MapReduce, Hadoop, Spark
    - » Matei Zaharia (UCB)
  - 3:30-4:00pm - Break
  - 4:00-5:00pm – Building Parallel Applications – Browsers, Vision and Music,
    - » Matt Torok, Michael Anderson, David Wessel (UCB)

Logistics (1/2)

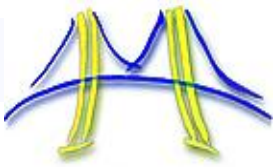

- parlab.eecs.berkeley.edu/2013bootcamplogistics
- Coffee
  - Available outside CITRIS Auditorium, not allowed in lecture hall!
- Live webcast of lectures
  - mms://media.citris.berkeley.edu/parlab2013
- Questions and answers
  - See links on above web page to post questions
  - For on-site and off-site students
- Lecture Materials
  - Slides and archived video will be posted on bootcamp website

Logistics – 3 Kinds of Labs (2/2)

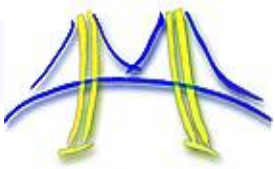

- parlab.eecs.berkeley.edu/2013bootcamplogistics
- Labs for **on-site** students using NERSC tools
  - Bring your own laptop
  - Different kinds of labs: OpenMP, Pthreads, MPI, CUDA
  - We supply wireless access, accounts at NERSC
    - » TAs: Razvan Carbunescu, Michael Anderson, Erin Carson, Nick Knight, David Sheffield
  - Autograding available
- Labs for **on-site** students using OpenMP
  - Bring your own laptop
  - We supply wireless access, VM to download and install
    - » Instructor: Tim Mattson
- Labs for **off-site** students using XSEDE
  - See parlab.eecs.berkeley.edu/2013bootcamplogistics for Q&A
  - Different kinds of labs: OpenMP, Pthreads, MPI

The Audience – you (1/2)

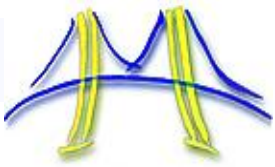

- There are ~~188~~ 209 registrants
  - ~~120~~ 131 on-site, ~~68~~ 78 off-site registrants
- Who are you?
  - 22 software developers or engineers
  - 8 faculty
  - 148 students (undergrad, grad, postdoc)
  - 31 other: architect / director / sysadmin / consultant / …

The Audience – you (2/2)

- Where are you from?
  - 5 companies
  - 23 universities
  - 17 states, 9 countries
  - Many countries over time (2009-2011 data)

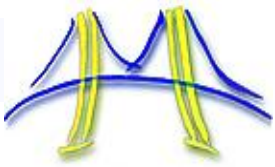

LET'S GET STARTED!