CSMD Division All Hands Meeting

Barney Maccabe

March 14, 2011

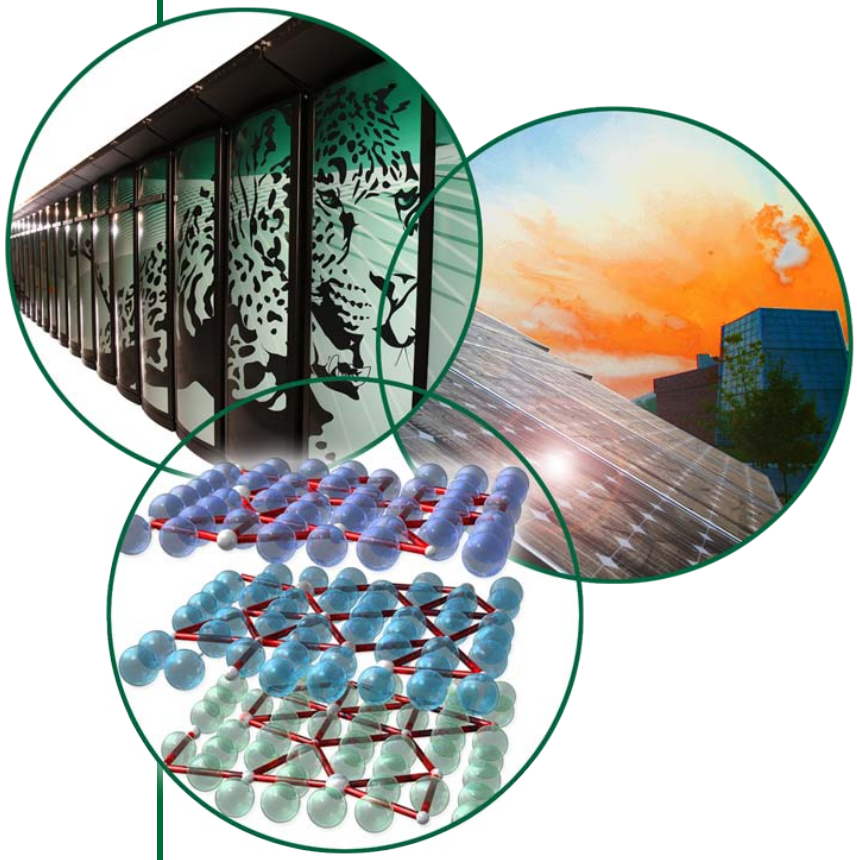

U.S. DEPARTMENT OF ENERGY

OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY

Computer Science and Mathematics Division

Computer Science and Mathematics Division
Barney Maccabe

John Cobb, Teragrid

Thomas Schulthess, Europe

Steve Poole, Chief Scientist Director of Special Programs

Robert Harrison, Director of Computational Science Programs

Apply
Computation for Science

Computational Astrophysics
Anthony Mezzacappa

Computational Chemical and Material Sciences
Bobby Sumpter

Computational Earth Sciences
Danny McKenna

Computational Engineering and Energy Sciences
John Turner

Explore
Computational Approaches

Computational Mathematics
Ed D'Azevado (Interim)

Statistics and Data Science
Barney Maccabe (Acting)

Complex Systems
Jacob Barhen

Build
Computational Systems

Future Technologies
Jeffrey Vetter

Computer Science Research
Al Geist

Application Performance Tools
Rich Graham

Groups

Centers and Institutes

Center for Computational Molecular Biophysics
Jeremy Smith

Experimental Computing Laboratory
Jeffrey Vetter

Center for Engineering Science Advanced Research
Jacob Barhen

Institute for Advanced Architectures and Algorithms
Jeff Nichols

EASI Math/CS Institute
Al Geist

Extreme Scale System Center
Steve Poole

Computing becomes official in 1948

Memo from C. N. Rucker, forming a Central Mathematics and Computing Panel led by A. S. Householder

Dated December 22, 1948

(Found by Tim Gawne)

X-16 H300LM 8751
183 C.E.LAMB

OAK RIDGE NATIONAL LABORATORY
OAK RIDGE, TENNESSEE

OFFICIAL BULLETIN

To: All Exempt Salaried Employees December 22, 1948

Subject: Mathematics and Computing Panel AR-No. 68

A central Mathematics and Computing Panel, with Dr. A. S. Householder in charge, has been established to serve the following functions:

1. Provide mathematical consultative services for the entire Laboratory.
2. Provide general computational services for the entire Laboratory.
3. Provide a central computer pool from which computers may be drawn for individual computing jobs.
4. Carry out the mathematical researches of relevance to the general Laboratory program.

C. N. Rucker
C. N. Rucker
Executive Director

Distribution A

Communications

- Within the division
  - Newsletter
  - Division Town Hall every 3-4 months
  - Revamping of the division web pages
  - Brown Bag Seminars
- External
  - Newsletter
  - FWP summaries (quad charts)
    - These have been great for visits from HQ
    - Lots of use in my talks

(Virtual) Collaborative Space

- For the development of **proposals**, **publications**, and **code**.
- Currently using:
  - SVN, TRAC and Mercurial.
- In the future, the system will also use
  - GIT, Redmine and Drupal.
- To request work space on this system, please send an email to
  - pubs-csm@ornl.gov – to request space for publication development
  - prop-csm@ornl.gov – to request space for proposal development
  - code-csm@ornl.gov – to request space for code development
- Thomas Naughton, Geoffroy Vallee, and Daniel Pack working with ITSD

- This system is not appropriate for PII storage and distribution.

Awards Committee

- Members
  - Phani Nukala, Don Nicholson, Bobby Sumpter, Sreekanth Pannala, David Bernholdt, Nageswara Rao, Bronson Messer, and Jeffrey Vetter
- Activities
  - Catalogued relevant external/internal awards for CSMD
    - Early/Mid/Late career awards
  - Developed guidelines for nomination
  - Identified criteria for mentoring early-career candidates
  - Devised employee incentive programs
    - Ex: Professional membership payments

Upcoming Events

- Computer Science visit, with new program manager
  - Early summer
  - The scope would a review of everything funded by the ASCR Computer Science base, as well as current CS-related SciDAC activity
- ASCR Math PI Meeting in October

Co-Design Centers

Chemistry Exascale Co-design Center (CECC)

The central science theme of the CECC is understanding, controlling, and ultimately designing chemically and electrically active interfaces that are relevant to diverse battery technologies, ultra-capacitors, fuel cells, environmental chemistries, and catalytic processes for sustainable energy conver-sion including biomass conversion.

Team includes members from ORNL, PNNL, ANL, LANL, NREL, Sandia, & SLAC, NC State, Ohio State, UCSD,Florida, UIUC, UT, Wisonsin, Virginia Tech, IBM, Q-Chem

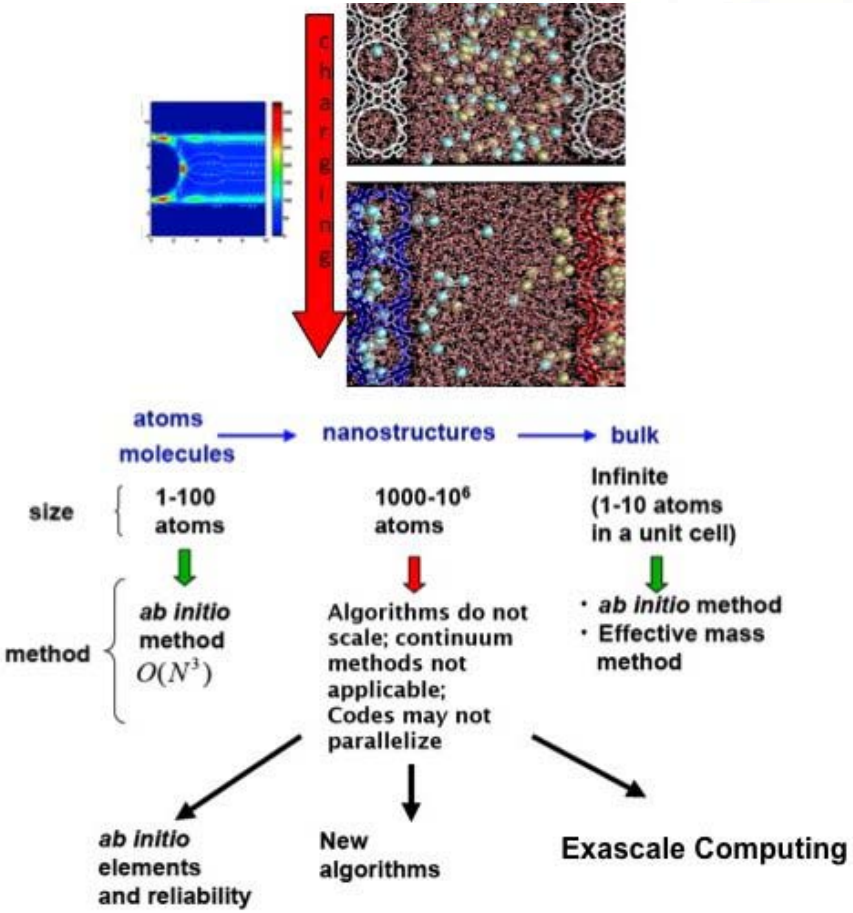

Other Centers

- Rosner (ANL): Center for Exascale Simulation of Advanced Reactors (CESAR)
- Lamb (ANL): FLASH High Energy Density Physics Exascale Codesign Center
- Koniges (LBNL): The CERF Center: Co-design for Exascale Research in Fusion
- Germann (LANL): Exascale Co-Design Center for Materials in Extreme Environments: Engineering-Scale Predictions
- Chen (SNL): Combustion Exascale Co-Design Center
- Jones (LANL) & Hack (ORNL): Exascale Performance Research for Earth System Simulation (EXPRESS)

Exascale Program Awards

Kickoff Meeting was held last month in San Diego

http://exascaleresearch.labworks.org/ascr2011

Advanced Dynamically Adaptive Algorithms for Stochastic Simulations on Extreme Scales

Rick Archibald, ORNL
Dongbin Xiu, Purdue University

FWP #ERKJU80

Objectives

- Extend the existing deterministic adaptive mesh refinement (AMR) algorithms to high dimensional random space to facilitate adaptive stochastic simulations.
- Develop and extend the existing edge-detection methods to stochastic spaces to identity and locate smooth sub-domains.
- Develop advanced, dynamically adaptive, stochastic collocation methods for extreme-scale computing.

Approach

- Minimize cost by *leading* computational simulations and running only the essential computational simulations scenarios.
- Automate *adaptive mesh refinement* (AMR) strategies in stochastic space of arbitrary dimension
- Build *advanced dynamically adaptive stochastic collocation* approaches

Strong (left) and weak (right) scaling data for our block-structured AMR system AMROC.

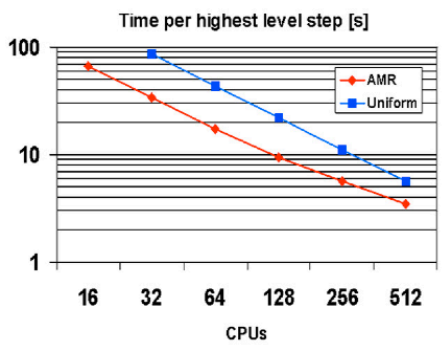

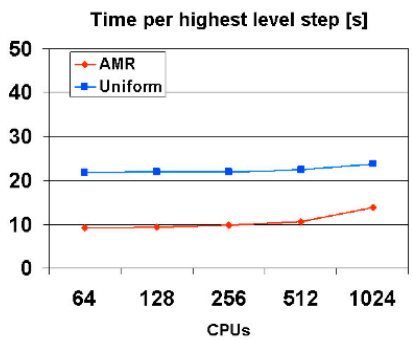

Impact

- This set of advanced numerical algorithms and computational tools for modeling uncertainty in complex stochastic systems will enable application scientists in the high-performance computing (HPC) environment to provide predictive modeling and simulation on the next-generation computer architecture on extremely large scales.
- Essential for predictive modeling and simulation-based decision-making.

Blackcomb: Hardware-Software Co-design for Non-Volatile Memory in Exascale Systems

Jeffrey Vetter, ORNL
Robert Schreiber, HP Labs
Trevor Mudge, University of Michigan
Yuan Xie, Penn State University

FWP #**ERKJU59**

Objectives

- Propose new distributed computer architectures that address the resilience, energy, and performance requirements of future DOE exascale systems:
  - replace mechanical-disk-based data-stores with energy-efficient non-volatile memories;
  - place low power compute cores close to the data store;
  - reduce number of levels in the memory hierarchy.
- Evaluate the impact of the proposed architectures on the performance of critical DOE applications.

Approach

- Identify and evaluate the most promising non-volatile memory (NVM) technologies.
- Explore assembly of NVM technologies into a storage and memory stack.
- Propose an exascale HPC system architecture that builds on our new memory architecture.
- Build the abstractions and interfaces that allow software to exploit NVM to its best advantage.
- Characterize key DOE applications and investigate how they can benefit from these new technologies.

A comparison of various memory technologies

	SRAM	DRAM	NAND Flash	PC-RAM	STT-RAM	R-RAM
Data Retention	N	N	Y	Y	Y	Y
Memory Cell Factor (F^2)	50-120	6-10	2-5	6-12	4-20	<1
Read Time (ns)	1	30	50	20-50	2-20	<50
Write / Erase Time (ns)	1	50	106-10^5	50-120	2-20	<100
Number of Rewrites	10^{16}	10^{16}	10^5	10^{10}	10^{15}	10^{15}
Power Read/Write	Low	Low	High	Low	Low	Low
Power (Other than R/W)	Leakage Current	Refresh Power	None	None	None	None

Impact

- Address energy scalability of future exascale systems; NV memories have zero standby power.
- Increase system reliability; MRAM/PCRAM are resilient to soft errors.
- Develop new programming models that exploit NVMs to improve the fault-tolerance of applications.

In-situ Data Reduction and Analysis for Extreme Scale Science

Scott Klasky, ORNL
Podhorski, Samatova, Wolf, ORNL;
Shosani, Wu, LBNL; Schwan, GT

FWP # ERKJU60

Objectives

- Create a robust I/O staging framework for *in situ* analysis and reduction of extreme-scale application data
- Create and extend programming models for end scientists to utilize *in situ* I/O stream processing
- Create a toolkit of I/O modules such as FastBit indexing that scientists can easily utilize

Approach

- We leverage existing multi-year effort in software tools for scientists
  - ADIOS for I/O abstraction in code, DataTap from GT for staging area, FastBit from LBNL for indexing, Parallel R for analysis
- We will provide run-time parameterization of ADIOS methods through a control layer
- We will introduce a PGAS-like interface for staging area processing to simplify for scientists

Scenario of *in situ* I/O Pipeline analysis and visualization for fusion simulation data

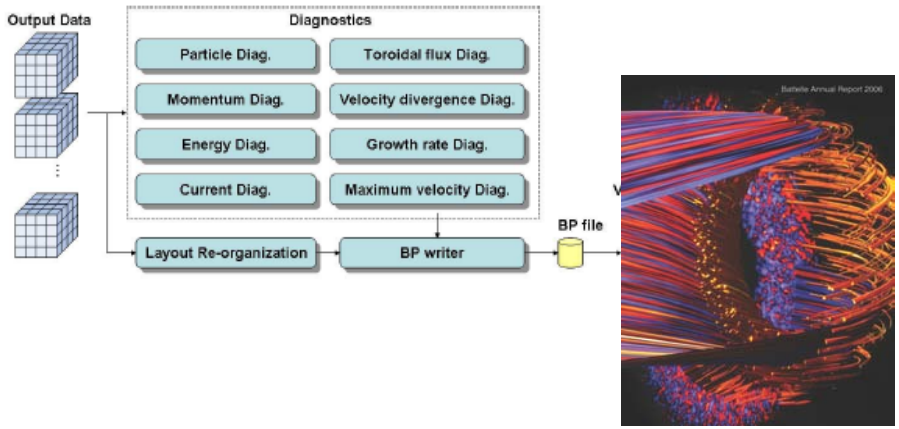

Impact

- ADIOS is already used by many extreme-scale science codes
  - Extensions to ADIOS will be immediately inherited
- Provide future-proofing for scalable I/O needs
- Initial evaluation systems already deployed for codes running at over 100k cores (Fusion, combustion)

Vancouver: A Software Stack for Productive Heterogeneous Exascale Computing

Jeffrey Vetter, ORNL
Wen-Mei Hwu, UIUC
Allen Malony, University of Oregon
Rich Vuduc, Georgia Tech

ERKJU44

Objectives

- Enhance programmer productivity for the exascale
  - Increase code development ROI by enhancing code portability
  - Decrease barriers to entry with new programming models
- Create next-generation tools to understand the performance behavior of an exascale machine

Approach

- Programming tools
  - GAS programming model
  - Analysis, inspection, transformation
- Software libraries: autotuning
- Runtime systems: scheduling
- Performance tools
- Impact on DOE Applications

The proposed Maestro runtime simplifies programming heterogeneous systems by unifying OpenCL task queues into a single high-level queue.

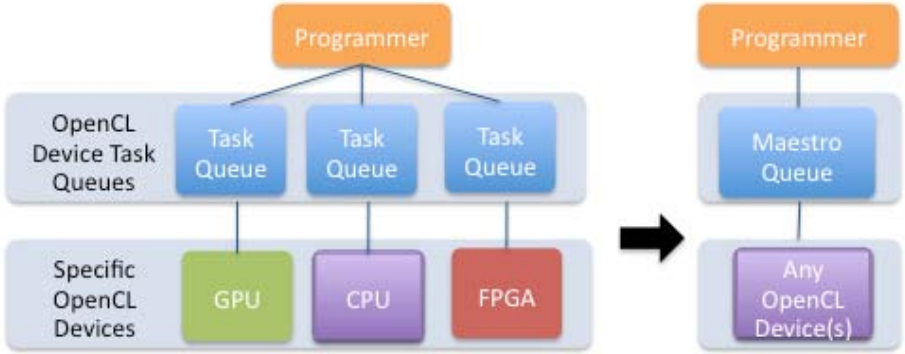

Impact

- Reduced application development time
- Ease of porting applications to heterogeneous systems
- Increased utilization of hardware resources and code portability

COMPOSE-HPC: Software Composition for Extreme Scale Computational Science and Engineering

David E. Bernholdt, ORNL
Tom Epperly, LLNL
Manoj Krishnan, PNNL
Matt Sottile, Galois
Rob Armstrong, SNL

ERKJU68

Objectives

- Develop a flexible, extensible toolkit to help software developers address various kinds of software composition challenges
- Provide examples of how the toolkit can be applied to specific composition problems
- Develop new approaches to facilitate composition of parallelism (threads and processes)

A schematic of the KNOT tool chain

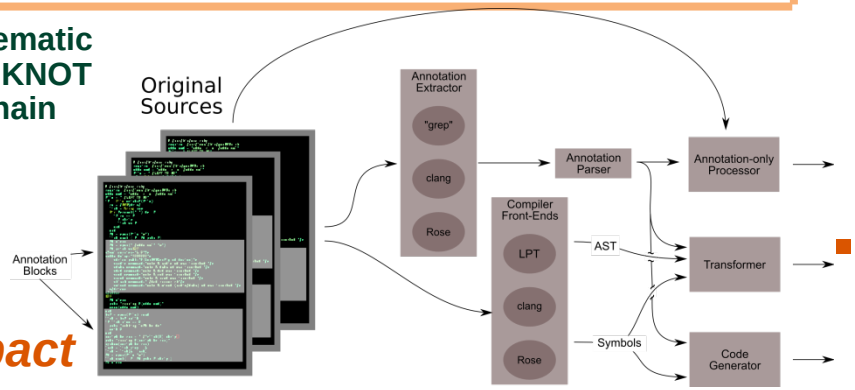

Impact

- Practical tools to help software developers produce higher quality code that works effectively in modern HPC environments
- Bring the capabilities of code transformation and code generation to bear on the challenges of software composition

Approach

- Develop the Knot Nimble Orchestration Toolkit (KNOT), consisting of three components:
  - An annotation parsing facility (PAUL) to interpret guiding annotations embedded as comments in user source code
  - A transformation facility (ROTE) to apply source-to-source transformations to the code, based on annotations and other inputs
  - A code generation facility (BRAID) capable of manipulating compiler-like intermediate representations, transforming, optimizing, and generating source code based on the results
- Develop examples of using KNOT to address different composition problems: language interoperability, contract enforcement, automatic performance instrumentation, data marshalling for GPUs
- Build on this infrastructure to facilitate the composition of parallelism
  - Allow codes to express their preferred threaded execution model, and call modules with different threading models
  - Simplifying the expression and exploitation of MPMD parallelism

Enabling Exascale Hardware and Software Design Through Scalable System Virtualization

Stephen L. Scott, ORNL
Patrick Bridges, UNM
Peter Dinda, NWU
Kevin Pedretti, SNL

ERKJU70

Objectives

- Provide a novel solution for testing at scale
- Ease the transition to production by supporting scaling of legacy system software
- Enable advanced research
  - Architecture research toward exascale
  - New parallel programming models
  - System software research
  - Resilience

Approach

- Extend the Kitten/Palacios prototype
  - Support for modern hardware
  - Port to HPC operating systems
  - Integration of system management tools
- Design & implement new capabilities
  - Integration with micro-architectural simulator
  - Binary translation for the emulation of new hardware
  - Fault injection

Kitten/Palacios Architecture: Virtualization for Exascale

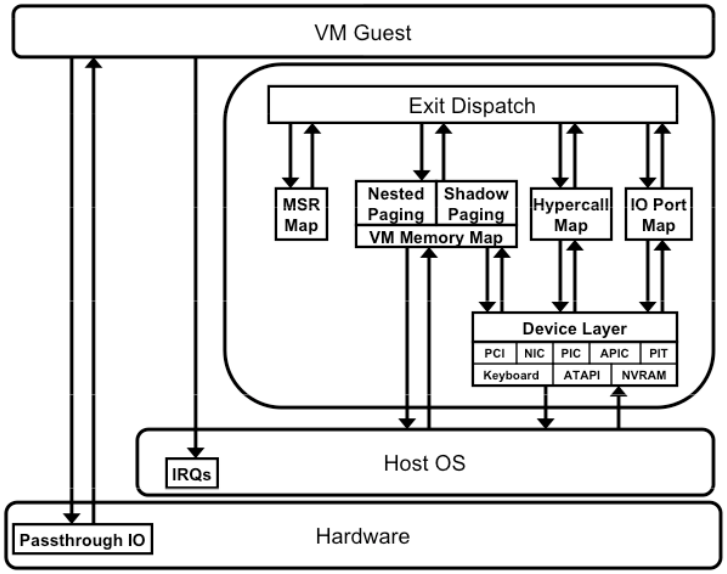

Impact

- Provide a test bed solution for exascale
  - Vertical profiling
  - Fault injection
- Provide a platform for exascale research
  - System architecture research
  - Programming languages research
  - System software research
  - Resilience

LDRD

LDRD Starts for FY11

- Jeremy Meredith
  - Scalable and Efficient Infrastructure for Exascale Analysis and Visualization
- Thomas Maier
  - Towards full first-principles simulations of correlated electron materials
- Oscar Hernandez
  - Hercules: A User Guided Translation Tool to Facilitate Application Porting to New Peta/Exascale Architectures

Upcoming Calls

- SciDAC-3 Institutes
- SciDAC-3 Applications
- SciDAC-3 SAP
- LDRD

Financial

Budget Authority

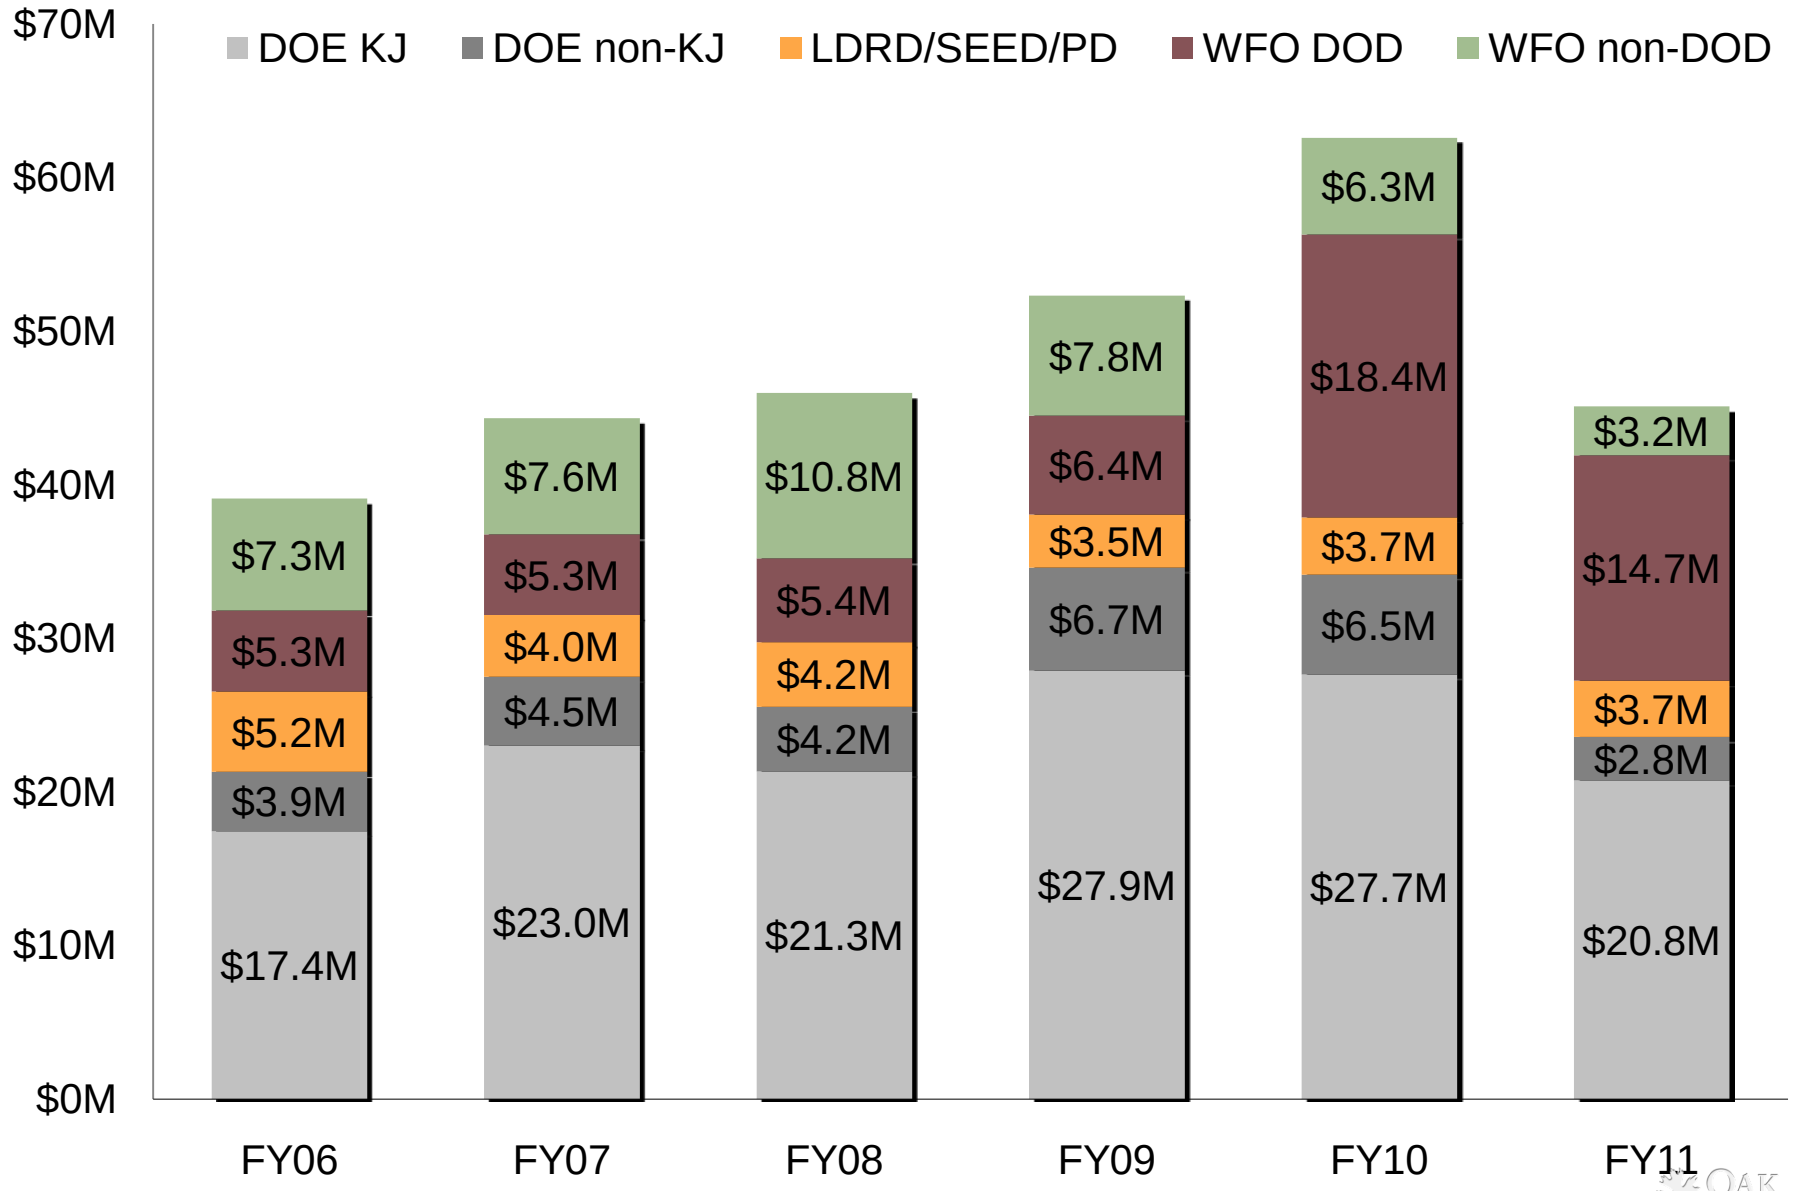

FY09 Snapshot

DOE	**66.1%**
ORNL	**6.6%**
WFO	**27.2%**

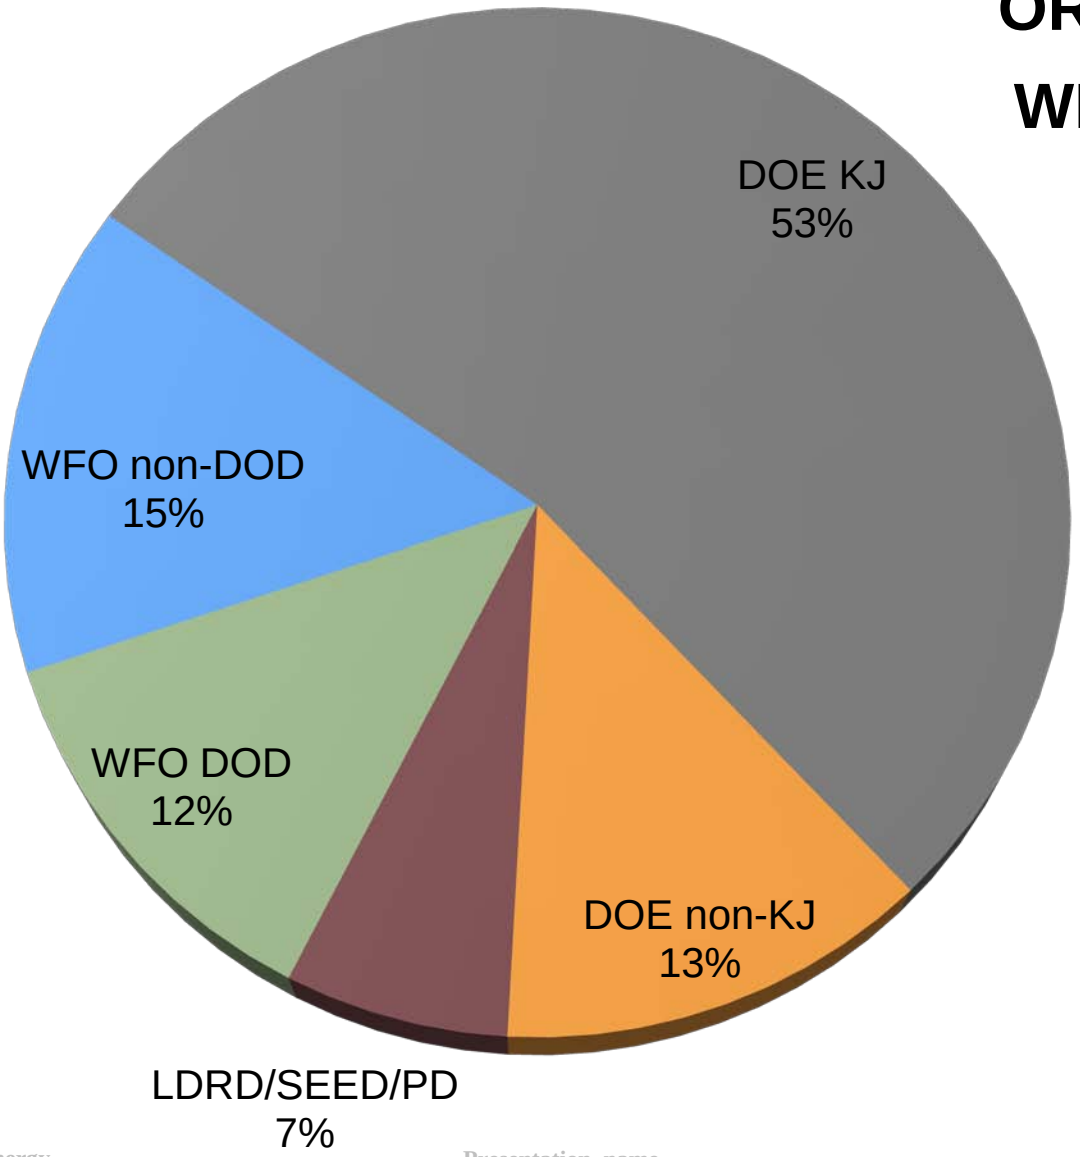

FY10 Snapshot

DOE	**48.9%**
ORNL	**6.7%**
WFO	**44.4%**

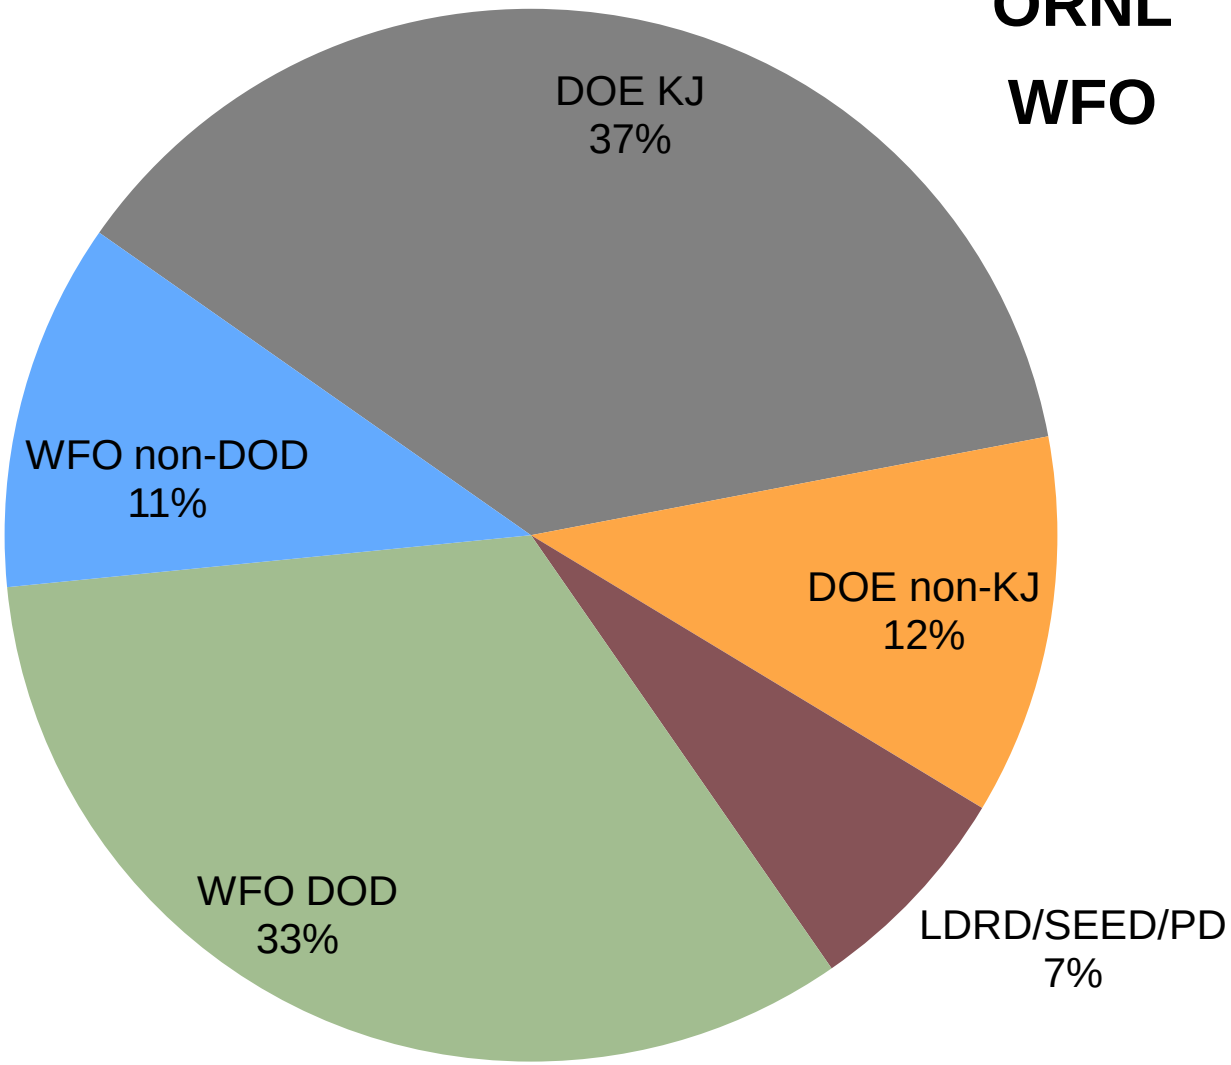

Productive Hours

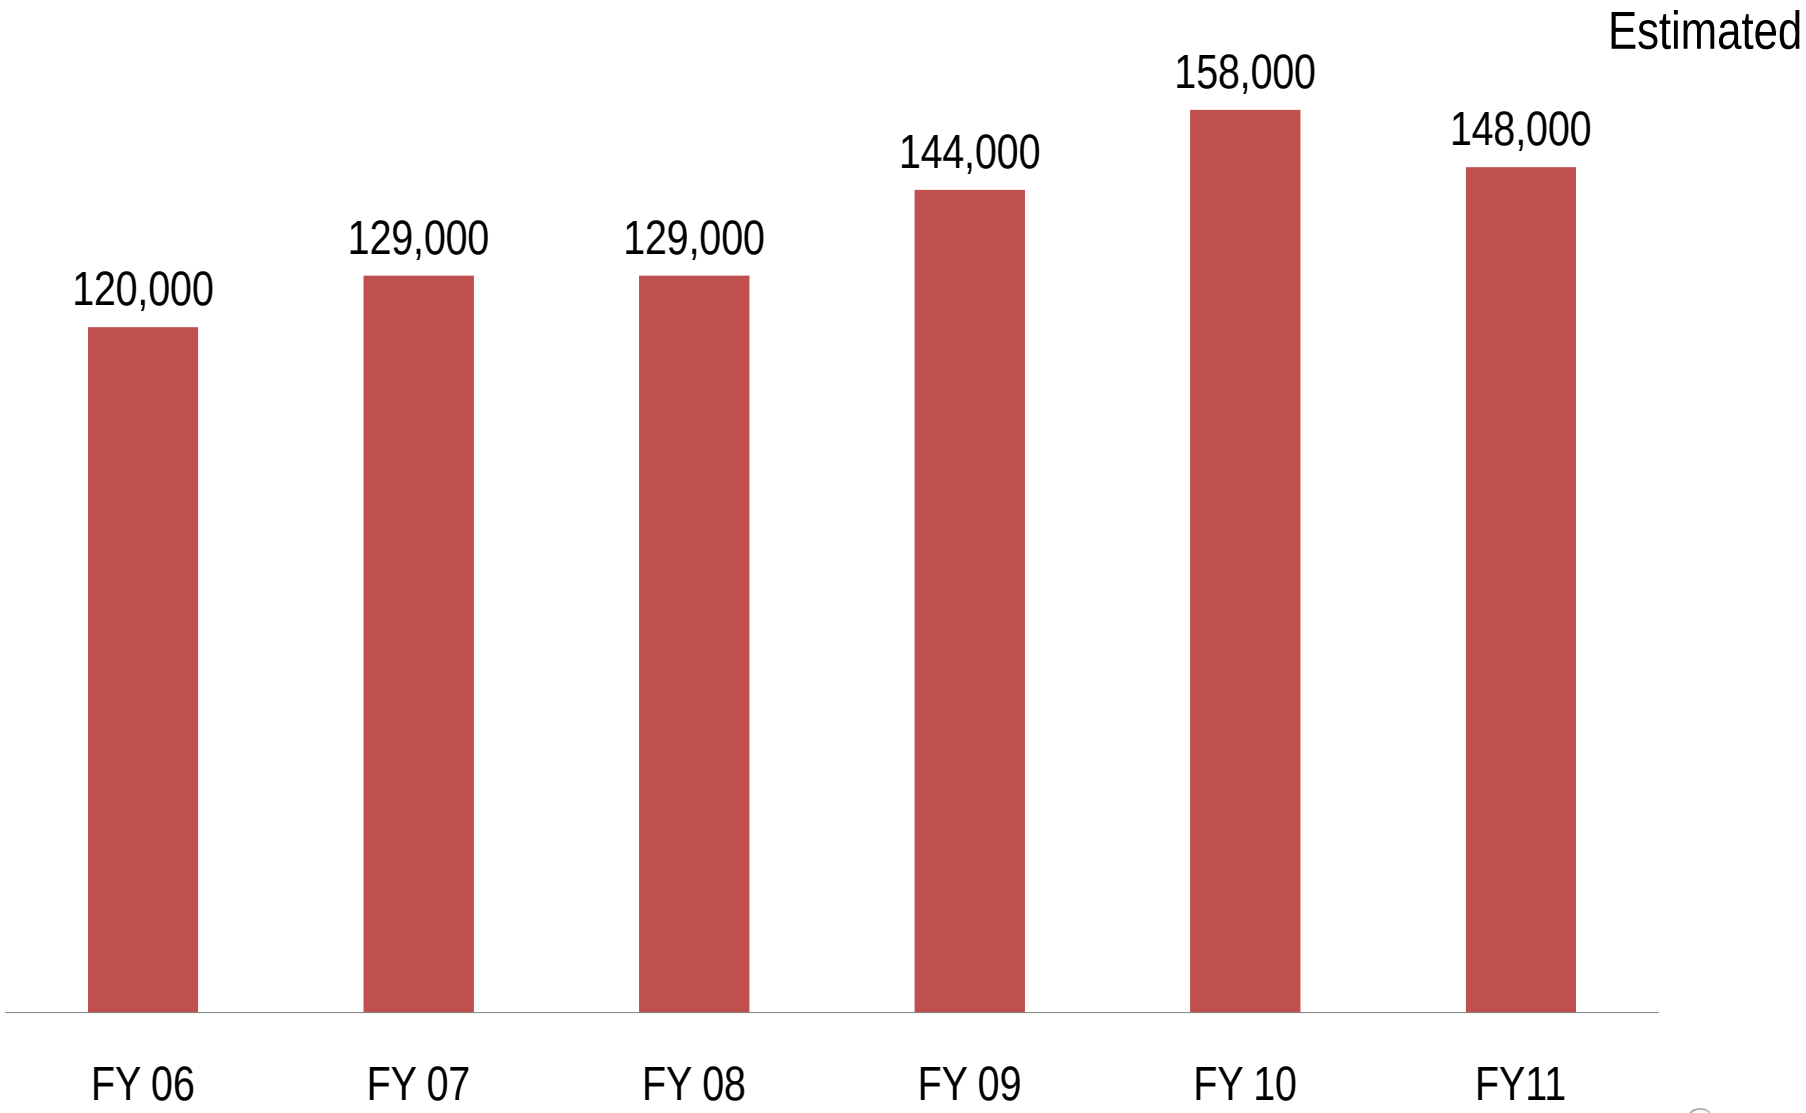

Personnel

Service Anniversaries

2010

Name	Years
Jacob Barhen	25
John Cobb	15
David Bernholdt	10
Wael Elwasif	10
David Erickson	10
Richard Archibald	5
Neena Imam	5
Jeremy Meredith	5
Igor Jouline	5
Cindy Sonewald	5
Olaf Storaasli	5

2011

Name	Years
Ahmed Khamayseh	10

New Hires

Application Performance Tools	
Matthew Baker	11/22/2010
Pavel Shamis	06/01/2010
Christos Kartsaklis	10/04/2010

Complex Systems	
William Ray	10/18/2010
Satyabrata Sen	02/21/2011
Mark Gomien	03/21/2011

Computational Chemistry and Materials Sciences	
Jingsong Huang	08/07/2010

Computational Earth Sciences	
Salil Mahajan	08/16/2010
Daniel McKenna	09/13/2010

Computational Engineering and Energy Sciences	
Mark Berrill	08/23/2010

Computer Science Research	
Guruprasad Kora	06/01/2010
Jay Billings	07/12/2010

Future Technologies	
Collin McCurdy	01/03/2011
Dong Li	02/21/2011

Highlights

Gordon Bell Award, Nov 2010

Simulating Blood Flow with FMM

- Multiphysics particle flow of deformable cells in viscous fluid with non-uniform distribution
- **Jeffrey Vetter** and **Rahul Sampath**

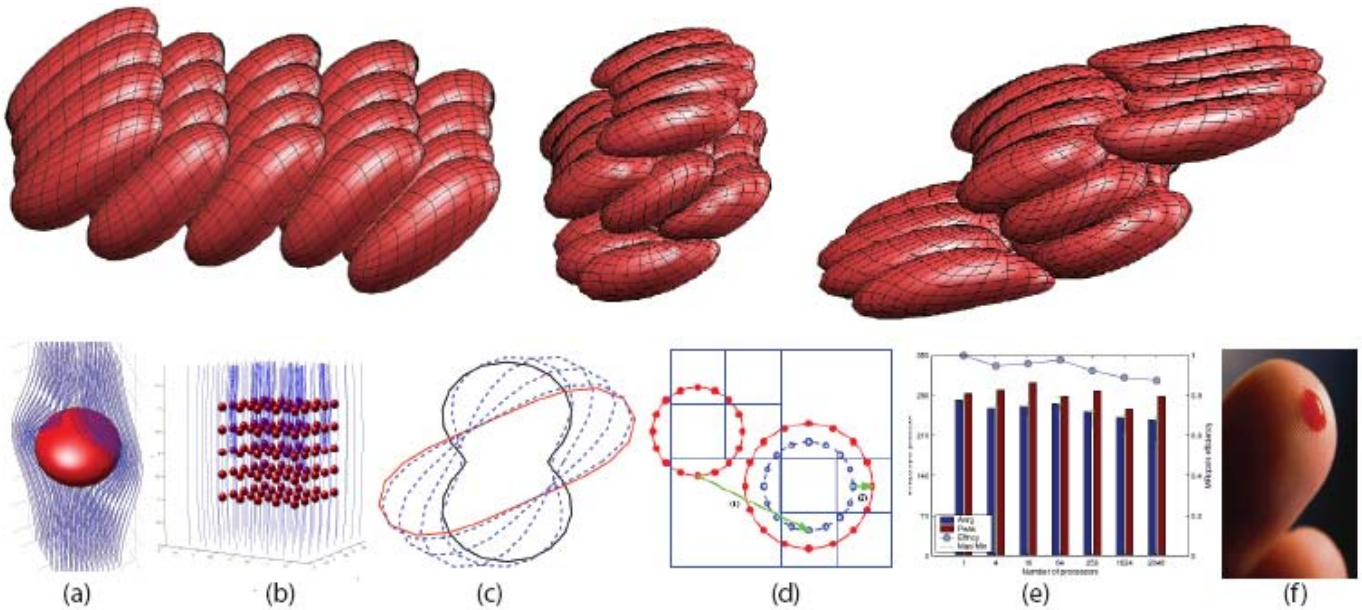

A. Rahimian, I. Lashuk, S. Veerapaneni *et al.*, *"Petascale Direct Numerical Simulation of Blood Flow on 200K Cores and Heterogeneous Architectures (Gordon Bell Winner)," in SC10: International Conf. High Performance Computing, Networking, Storage, and Analysis. New Orleans: ACM/IEEE, 2010.*

SEA Winners

- Release of version 0.9 of the AMP Nuclear Fuel Performance code for distribution through the Radiation Safety Information Computational Center (RSICC):
  - Kevin Clarno, James Banfield, **Bobby Philip**, **Srikanth Allu**, **Bill Cochran**, Gokhan Yesilyurt, **Srdjan Simunovic**, Jung Ho Lee, **Sreekanth Pannala**, **Pallab Barai**, **Phani Nukala**, **Jay Billings**, **John Turner**, Thomas Evans, and **Rahul Sampath**.
- Novel analyses of the simulation results from the CCSM 3.0 climate model:
  - Auroop Ganguly, Nagendra Singh, **David Erickson**, Karsten Steinhaeuser, **Marcia Branstetter**, Lawrence Buja, Esther Parish, and **John Drake**.

Other Recognitions

- **Pratul Agarwal** received the "Key Contributor Award" at the 2010 UT-Battelle Partnerships Award
  - commercializing Pratul's Enzyme Engineering technology.
- **Edoardo Apra** won HPCwire's Reader's Choice Award in supercomputing achievement at SC10
  - for his work on NWChem
- **Nagi Rao** best paper award at the 2010 IEEE International Conference in Multisensor Fusion and Integration
  - Cyber-Physical Trade-Offs in Distributed Detection Networks

Other

Directorate	Division	Total found in WoS	Total found in PTS	% found
Computing & Computational Sciences Dir	Center for Computational Sciences	38	18	47%
	Computational Sciences & Engineering Div	92	64	70%
	Computer Science and Mathematics Div	330	158	48%
	Computing & Computational Sciences Dir	5	0	0%
Energy & Environmental Sciences Dir	Biosciences Division	267	132	49%
	Energy & Environmental Sciences Dir	22	14	64%
	Energy & Transportation Science Division	123	49	40%
	Environmental Sciences Division	371	216	58%
	Measurement Science & Systems Engr Div	69	34	49%
ESH&Q Directorate	Nuclear & Radiological Protection Div	4	1	25%
	Quality Systems and Services Division	5	0	0%
Global Security Directorate	Global Security Directorate	12	3	25%
Neutron Sciences Directorate	Neutron Sciences Directorate	8	5	63%
	NScD Neutron Facilities Development Div	31	14	45%
	NScD Neutron Scattering Science Division	314	194	62%
	NScD Research Accelerator Division	36	21	58%
Nuclear Science & Engineering Dir	Fuel Cycle & Isotopes Division	12	2	17%
	Fusion Energy Division	249	94	38%
	Global Nuclear Security Technology Div	21	6	29%
	Nuclear Science & Engineering Dir	21	0	0%
	Reactor & Nuclear Systems Division	83	49	59%
Physical Sciences Directorate	Center for Nanophase Matls Sciences Div	319	196	61%
	Chemical Sciences Division	474	265	56%
	Materials Science and Technology Div	1358	868	64%
	Physics Division	373	288	77%
US ITER Project Office	US ITER Project Office	20	7	35%

Safety

5X,000,000 hours worked without a DART

DART = Days Away, Restricted or Transferred

The students are coming.....

- Opportunity to review safety issues in your immediate area and issues across the laboratory
  - Tripping and ergonomic issues are the biggest issues for the division
  - Be aware of your surroundings and the ergonomics of your environment
- Parking lots will be full, and “backing” issues are likely to rise
- Pay attentions to posted speed limits!

Sharing the Road

Sharing the Road: Motorists

Drive Cautiously:
•Reduce speed when encountering cyclists
•Don't tailgate, especially in bad weather
•Recognize hazards cyclists may face and give them space
Yield to Cyclists:
•Bicycles are considered vehicles and in most jurisdictions are prohibited from riding on the sidewalk
•Cyclists should be given the appropriate right of way
•Allow extra time for cyclists to get through intersections
Be Considerate:
•Scan for cyclists in traffic and at intersections
•Do not blast your horn in close proximity to cyclists
•Look for cyclists when opening doors
Pass with Care:
•When passing, leave at least three feet between you and a cyclist
•Wait for safe road and traffic conditions before you pass
•Check over your shoulder before moving back
Watch for Children:
•Children on bicycles are often unpredictable
•Expect the unexpected and slow down
•Don't expect children to know traffic laws
•Because of their size children can be harder to see
•Watch for children on bicycles crossing intersections from sidewalks

Sharing the Road: Cyclists

On the Road:
•The same laws that apply to motorists apply to cyclists
•Obey all traffic control devices
•Use hand signals to indicate stops and turns other users
•Wear a helmet, no matter how short the trip
Ride on the Right:
•Always ride in the same direction as traffic
•Use the rightmost right lane that heads to your destination
•Slower moving cyclists and motorists stay to the right
Ride Predictably:
•Ride in a straight line
•Don't swerve in the road or between parked cars
•Check for traffic before entering street or intersection
•Anticipate hazards and adjust your position accordingly
Be Visible:
•Wear brightly colored clothing that provides contrast
•Use a white front light in low light conditions
•Use a red rear light in low light conditions
•Use a reflector or reflective tape or clothing anytime
•Announce yourself by making eye contact with motorists

CSMD

OAK RIDGE National Laboratory

The Project Summaries one more time

Thanks

Project Title

John Smith, ORNL
Jane Doe, University of Tennessee

FWP #

Objectives

- *What will you accomplish?*

Approach

- *How are you going to accomplish your objectives*

Image Caption

Impact

- *Why should people care about this work?*
- *Highlights*

Project Title

John Smith, ORNL
Jane Doe, University of Tennessee

Challenges

- *Why is this hard to do?*

Formal Problem Statement

- *Speak to your peers*

Image Caption

Research Products/Artifacts

- *What are you building? What is the product of this research (beyond theorems and papers)*

Project Title

John Smith, ORNL
Jane Doe, University of Tennessee

Progress

- *How far along are you toward realizing your objectives?*

Accomplishments

- List of publications and code releases

CSMD Staff Supported

Name	Group	Support Level
John Smith	Comp Bio	75%

Scalable Graph Decompositions & Algorithms to Support the Analysis of Petascale Data

Blair Sullivan, Chris Groer, ORNL

ERKJU04

Objectives

- Leverage theoretical work on graph decompositions to create efficient computational framework for graph-based data.
- Enhance existing scientific computing applications by improving the efficiency of combinatorial kernels,
- Enable direct analysis of petascale data that has an underlying graph structure.

Approach

- Use tree decompositions to transform algorithm complexity so that it is exponential in the width, but polynomial in number of nodes.
- Develop efficient, scalable algorithms for computing low-width decompositions for large graphs.
- Integrate parallel computing with decomposition algorithms and dynamic programming.

A visualization of part of the internet graph

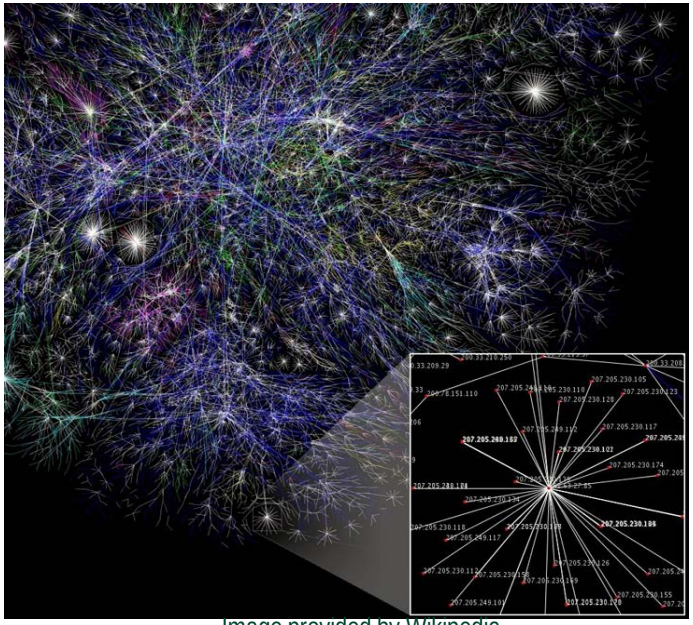

Image provided by Wikipedia

Impact

- Reduction of time required for key repeated computations on graphs that arise when modeling the discrete optimization problems in scientific computing applications.
- Ability to analyze petascale graphs that have recently emerged in areas such as the development of a more resilient electric grid and the classification of protein interactions in biofuel development which was previously prohibited by excessive time and memory requirements.

Scalable Graph Decompositions & Algorithms to Support the Analysis of Petascale Data

Blair Sullivan, Chris Groer, ORNL

Challenges

- For parallel computing, it is not sufficient to just have a low width tree decomposition – we need balanced bag sizes and a tree structure which is far from being a path, so we must develop techniques to control the structure of the decomposition tree.
- Branch decompositions are less well-studied and utility for large graphs is unknown.
- Given a decomposition, we still need to run a dynamic programming algorithm to solve a specific graph problem. These algorithms must be designed to run in parallel using only localized information at each node.

Formal Problem Statement

- The proof of Waring's conjecture led to the notions of tree/branch decompositions & the reduction of NP-hard problems to polynomial ones on graphs of fixed width.
- We must modify the methods for generating and utilizing these decompositions to transform them into a computational tool for the analysis of large, sparse graphs via parallel computing.

"Nice" Tree-Decomposition: bag sizes k & k-1, many leaves

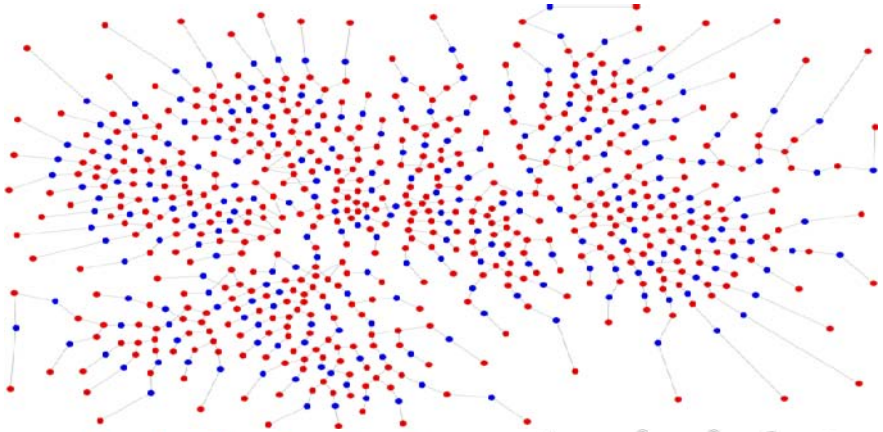

Research Artifacts

- Computational algorithms available as an open source library.

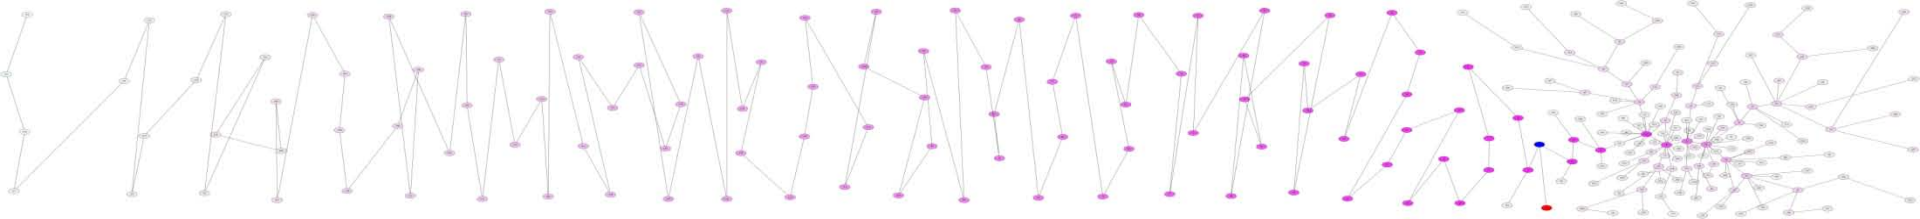

Traditional (Gavril) Tree-Decomposition – bag sizes vary widely, most nodes on long path-like structures (few leaves)

Scalable Graph Decompositions & Algorithms to Support the Analysis of Petascale Data

Blair Sullivan, Chris Groer, ORNL

Blair Sullivan	Complex Systems	50%
Chris Groer	Comp. Math	50%

Progress

- Development and implementation of serial algorithms in a C++ library, including:
- Heuristics for finding low-width tree decompositions
- Routines for generating "nice" tree decompositions (forces some structure & balances bag sizes)
- Implementation of a dynamic programming algorithm to solve an NP-hard problem (max weighted independent set) given a tree decomposition
- Randomized generator for graphs with a known upper bound on tree-width
- Integration with sparse matrix reordering routines (Metis, AMD) for additional elimination ordering algorithms

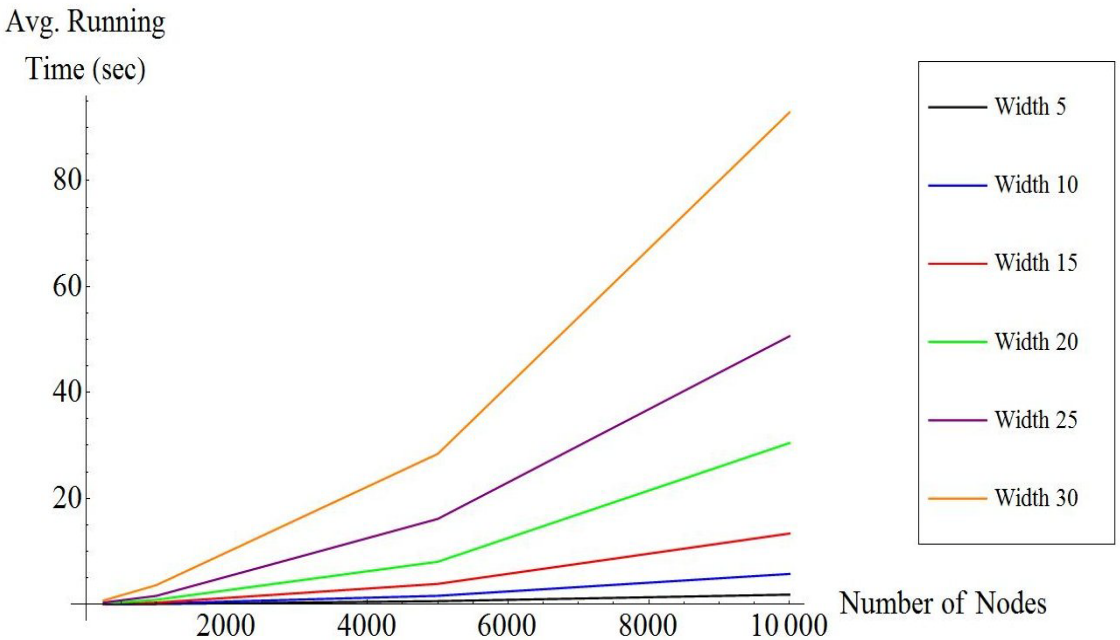

Accomplishments

- Generated preliminary timing data showing effect of width and graph size on solution time (as shown at left).
- Established connections with other researchers in field, including A. Koster (developed TWLib software) and I. Hicks (branch width expert).
- Tested new heuristics for finding elimination orderings with lower time complexity than traditional heuristics.
- Discovered that low-fill does not imply low-width for ordering heuristics.