

# Argonne National Laboratory

Marc Snir



# Department of Energy

- 3 weapon labs, managed by NNSA: LLNL, LANL, SNL
- 18 science labs, managed by the Office of Science
- Perform open, non classified research and have international collaborations
- 3 major CS and applied math divisions: ANL, LBNL, ORNL



# Direct descendent of Enrico Fermi's Metallurgical Laboratory



- Opened in Feb 1943 (as new site for Chicago's Metallurgical Laboratory)
- Became Argonne National Laboratory in July 1946 (first national laboratory)





# Argonne provides a vital ecosystems for major DOE scientific user facilities

APS

Protein  
facility

Nanoscience

Chemical  
engineering

Chemistry

Accelerator  
physics

Materials and  
energy  
sciences

Biology

High performance  
computing & data



# Argonne provides a vital ecosystems for major DOE scientific user facilities

## Employees

FTEs	3,375
Joint Faculty	149
Postdocs	273
Grad/Undergrad	874
Annual Facility	4,289
Visiting Scientists	477

FY11 Budget: \$680M

# Computing at ANL



# 60 Years of Computing at ANL

1953 – AVIDAC

1957 – Applied Mathematics Division spun off  
Numerical algorithms, math SW (...PACK)

1982 – Mathematics and Computer Science Division  
(same year as first CS department in the US) ~25 staff

Theorem proving, source-to-source, viz

1984 – Advanced Computing Research Facility

CM-2, DAP-510, BBN TC-2000, Intel iPSC/d5,  
Sequent Balance 21000, Encore Multimax, Alliant FX/8,  
Ardent Titan

1992 – High-Performance Computing Research Center  
IBM SP1

2005 – Argonne Leadership Computing Facility spun off



Math libraries (Lapack, PETSc, Nek5000, TAO),  
Meshing (MOAB) , Parallel software (MPICH, PVFS),  
Grid (Globus), Cloud (Nimbus)

# Production Systems at ALCF

## **Mira – BG/Q system**

- 49,152 nodes / 786,432 cores
- 786 TB of memory
- Peak flop rate: 10 PF

## **Vesta - BG/Q system**

- 4,096 nodes / 65,536 cores
- 64 TB of memory
- Peak flop rate: 832 TF

## **Cetus - BG/Q system**

- 2,048 nodes / 32,768 cores
- 32 TB of memory
- Peak flop rate: 416 TF

## **Tukey – NVIDIA system**

- 100 nodes / 1600 x86 cores
- 200 M2070 GPUs
- 6 TB x86 memory / 1.1TB GPU memory
- Peak flop rate: 220 TF

**Storage** - Scratch: 28.8 PB raw capacity, 240 GB/s bw (GPFS); Home: 1.8 PB raw capacity






# MCS

# MCS in Numbers

## People

Total: 209 (118 long term)





# MCS Division Areas

## Extreme Computing Data



*Pete Beckman*



*Rob Ross*

## Applied Math




*Paul Hovland*

## Applications



*Ray Bair*

## People (4/2013)



## Money (SC, FY 2012)



# Research Topics

## Applied Math

Paul Hovland

Optimization  
Mihai Anitescu /  
Sven Leyffer

Numerical PDEs  
Barry Smith

Comp Fluid Dynamics  
Paul Fischer

Meshes  
Tim Tautges

## Extreme Computing

Pete Beckman

Operating Systems  
Pete Beckman

Programming Models  
Pavan Balaji

Runtime Systems  
Rajeev Thakur

Performance Modeling  
Boyana Norris

Resilience  
Franck Cappello

## Data

Rob Ross

Storage Systems  
Rob Ross

Collaboration Systems  
Ian Foster

Visualization  
Mark Hereld

Data Analysis  
Stefan Wild

## Applications

Ray Bair

Molecular Science  
Ray Bair

Climate Modeling  
Rob Jacob

Reactor Simulation  
Andrew Siegel

Bioinformatics  
Folker Meyer

Cosmology  
Salman Habib





# Software Releases (2010-2012) [a]

1. **ADLB** –Asynchronous Dynamic Load Balancing MPI-Library (used by GFMC)
2. **ADIC2** – Tool for the automatic differentiation of programs written in ANSI C. (>1000 users)
3. **CGMA** – The Common Geometry Module, Argonne provides geometry functionality for mesh generation and other applications.
4. **CIFTS** – CIFTS provides coordinated and improved fault tolerance for high-performance computing systems.
5. **C-MPI** – Content-MPI is an implementation of a distributed hash table
6. **Darshan** – a scalable HPC I/O characterization tool.
7. **DCMP** – The Data Domain to Model Domain Conversion Package addresses evaluation of high-end climate models.
8. **ExM** – ExM provides support for extreme-scale, many-task applications.
9. **Falkon** – Fast and lightweight task execution framework.
10. **Globus Toolkit 5** –: security, communication, information infrastructure, fault detection, resource management, portability, and data management for the grid (> 15M transfers per day, 1 PB/day; 6000 registered users )
11. **IOFSL** – I/O forwarding scalability layer.
12. **MINOTAUR** – Open-source toolkit for solving mixed-integer nonlinear optimization problems.
13. **MPICH2 (soon MPICH3)** (>1500 downloads a month, deployed on all top supercomputers)



# Software Releases (2010-2012) [b]

14. **MOAB** – Component for representing and evaluating mesh data.
15. **Model Coupling Toolkit** – Coupler technology in the Community Earth System Model.
16. **NekCEM** –electromagnetic solver that uses the spectral-element discontinuous Galerkin method on a conformal spectral-element mesh.
17. **NekLBM** – NekLBM is a spectral-element discontinuous Galerkin lattice Boltzmann fluid solver.
18. **Nek5000** – Nek5000 simulates unsteady incompressible fluid flow with thermal and passive scalar transport.
19. **NEOS** – Network-enabled problem-solving environment for a wide class of applications in business, science, and engineering.
20. **Nimbus** – Integrated set of tools for scientific computing on clouds. (> 50,000 downloads/year, 200 subscribers to support list)
21. **OpenAD** – Tool for automatic differentiation (AD) of numerical computer programs.
22. **OpenAD/F** –Tool for automatic differentiation of Fortran codes.
23. **OpenAnalysis** – The OpenAnalysis toolkit seeks to separate analysis from the intermediate representation in a way that allows the orthogonal development of compiler infrastructures and program analysis.
24. **OSPRI** – OSPRI is an optimized one-sided communication runtime for leadership-class machines.





# Software Releases (2010-2012) [c]

- 25. **Parallel netCDF** – Library providing high-performance I/O while still maintaining file-format compatibility with Unidata's NetCDF.
- 26. **PETSc** – The Portable, Extensible Toolkit for Scientific computation is a suite of uni- and parallel-processor codes for solving large-scale problems modeled by partial differential equations. (>50,000 downloads a year)
- 27. **PMI** -- Scalable process-management interface for extreme-scale systems.
- 28. **PVFS** – The Parallel Virtual File System (PVFS) project brings state-of-the-art parallel I/O concepts to production parallel systems.
- 29. **RGG** – Reactor Geometry (and mesh) Generator is an open source tool to generate several types nuclear reactor assembly/core geometry and mesh.
- 30. **ROMIO** – ROMIO is a high-performance, portable implementation of MPI-IO.
- 31. **SISIPHUS** – SISIPHUS is a model and software infrastructure for climate simulations related to the dynamics of ice sheets.
- 32. **Swift** – Swift is a parallel scripting language that runs on multicores, clusters, clouds, and supercomputers.
- 33. **TACO** – The Toolkit for AMPL Control Optimization provides a set of extensions to the AMPL modeling language to conveniently model mixed integer optimal control problems for ODE or DAE dynamic processes.






## Software Releases (2010-2012) [d]

- 34. **TAO** – The Toolkit for Advanced Optimization focuses on the design and implementation of component-based optimization software for the solution of large-scale optimization applications. (> 1000 users)
- 35. **TUNE** – TUNE provides technology for compiler-directed automatic performance tuning, currently targeted at the Cray XT3 system.
- 36. **UNIC** – UNIC is an unstructured, deterministic neutron transport code that allows a highly detailed description of a nuclear reactor core in our numerical simulations.



# Recent Large Projects

- **CESAR** – Center for Exascale Simulation of Advanced Reactors [co-design center, Andrew Siegel, director]
- **M<sup>2</sup>ACS** -- Multifaceted Mathematics for Complex Energy Systems [ASCR, Mihai Anitescu, director]
- **ARGO\*** -- Exascale OS/R [ASCR, Pete Beckman, director]
- **Computation-Driven Discovery for the Dark Universe** – [SciDAC-3, Salman Habib]
- **SDAV** - Scalable Data Management, Analysis, and Visualization Institute [ASCR, Rob Ross, deputy director]
- **RDCEP** -- Center for Robust Decision Making on Climate and Energy [NSF, University of Chicago, Ian Foster, PI]
- Participation in SciDAC Institutes: **FASTMath, SUPER**
- Participation in SciDAC partnerships: **OSCon, PSI, NUCLEI**



# The Push-Pull Model

