# The Argonne Training Program on Extreme-Scale Computing


Paul Messina
Argonne Leadership Computing Facility
Argonne National Laboratory

# Argonne Training Program on Extreme-Scale Computing

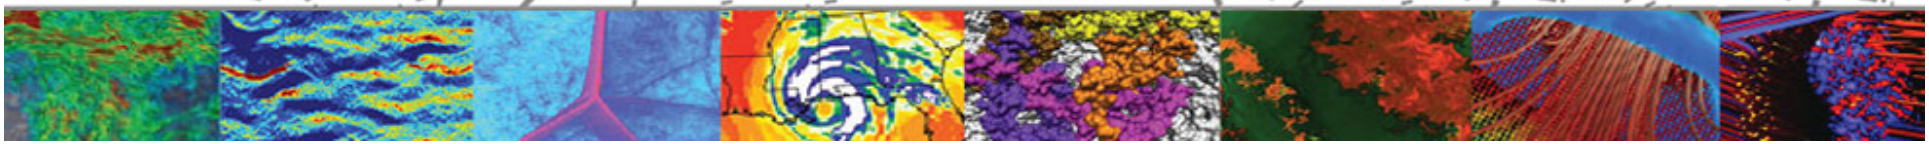

- **The Argonne Training Program in Extreme-Scale Computing (ATPESC) provides training on the key skills, approaches, and tools necessary to design, implement, and execute Computational Science and Engineering (CS&E) applications on current high-end systems**
  - and the leadership-class systems expected to be available in 2017 and beyond.
- **The training features lectures by leading computer and computational scientists in the topics covered and**
- **Extensive hands-on exercises on leadership-class computers.**

**http://extremecomputingtraining.anl.gov**

# Motivation for the ATPESC

- **With the challenges posed by the architecture and software environments of today's most powerful supercomputers -- and even greater complexity on the horizon from next-generation and exascale systems -- there is a critical need for specialized, in-depth training for the scientists poised to facilitate breakthrough science and engineering using these computational resources.**

# Computational research at the high end requires expertise in many areas

- The architecture and software environments of today's most powerful systems - including those at the Argonne Leadership Computing Facility and at the Oak Ridge Leadership Computing Facility - pose new and unique challenges to their users.
- Technology trends indicate that the next generation of systems, circa 2017, and the exascale systems projected for 2022 will feature far greater parallelism and complexity of architectures.
- Computational research projects tackle ever more complex projects.

# Scope of the Argonne Training Program

- **The ATPESC covers many of the key skills that computational scientists and engineers need to carry out their research on leadership-class systems**
- **In two weeks it is not possible to present detailed material on all topics; in some cases we aim to make participants of the existence of certain approaches, methods, and tools**
  - E.g., PGAS languages, software engineering tools

# High-level view of curriculum

- Computer architectures, mathematical models and numerical algorithms
- Programming methodologies that are effective across a variety of today's supercomputers and that are expected to be applicable to exascale systems
- Multiple approaches on unifying concepts and levels of abstraction that provide migration paths and performance portability among current and future architectures
- Approaches to building community codes for HPC systems, and methodologies and tools relevant for Big Data applications

# Curriculum tracks and their leaders

- Architectures – Pete Beckman
- Programming models and languages – Rusty Lusk, Rajeev Thakur
- Mathematical algorithms and software -- Lois McInnes, Lori Diachin
- Software engineering in scientific computing – Katherine Riley, Anshu Dubey
- Performance and debugging tools – Kalyan Kumaran, Scott Parker
- Visualization and Data Analysis – Mike Papka
- Data-intensive computing – Rob Ross and Rob Latham
- Community codes– Katherine Riley and Anshu Dubey
- Application case studies– Katherine Riley and Anshu Dubey

# Logistics

- **The training program is held at the Pheasant Run Resort in St. Charles, Illinois**
  - Lodging and all meeting rooms are at the same facility
  - Moderately priced: we use the old conference center
- **Participants pay no registration fee**
- **Domestic airfare, meals, and lodging are provided for the participants**

# Program Schedule

- **Daily schedule**
  - 08:30 – 10:30 Lectures
  - 10:30 – 11:00 Break
  - 11:00 – 12:00 Lectures
  - 12:00 – 13:00 Lunch and Hands-on exercises
  - 13:00 – 15:00 Lectures
  - 15:00 – 15:30 Break
  - 15:30 – 17:30 Lectures
  - 17:30 – 18:30 Dinner and dinner talk
  - 18:30 – 21:30 Hands-on exercises and lectures
- **Program begins mid-afternoon on a Sunday**
- **Free time in middle week-end: Saturday afternoon and Sunday**
- **Program ends noon second Friday**

# Qualifications of Applicants

- **Target participants are**
  - doctoral students,
  - postdocs, and
  - computational scientists.
- **Developers of programming models, algorithms, and tools for leading-edge systems may also apply.**
- **This program provides advanced training. Therefore, successful applicants must**
  - Have substantial experience in MPI and/or OpenMP programming
  - have used at least one HPC system for a reasonably complex application and
  - are planning to conduct CS&E research on large-scale computers.
- **~60 participants are selected.**

# Application and selection process

- **Applicants submit through a website:**
  - Letter written by applicant stating purpose for attending the training program and description of computing experience.
  - Curriculum Vitae.
  - One letter of recommendation from an advisor or supervisor.
- **In 2013 we received 160 applications**
- **A committee of five people selected 63 participants, based on candidates having**
  - Substantial MPI and/or OpenMP programming experience
  - Used an HPC system for a reasonably complex application
  - Having conducted or are preparing to conduct computational science and engineering research on large-scale computers

# ATPESC 2013 Participant Distribution

| | |
|---|---|
| Undergraduate student | 1 |
| Graduate student | 32 |
| Postdoc | 12 |
| Assistant professor | 2 |
| Research staff academia | 4 |
| Research staff US government labs | 7 |
| Research staff foreign labs | 2 |
| Industry | 3 |
| Total | 63 |

# Other ATPESC 2013 Participant Profiles

| | |
|---|---|
| US Government laboratory affiliation:<br>Argonne (3)<br>Army Research Laboratory<br>Geophysical Fluid Dynamics Laboratory<br>LANL (3)<br>LLNL<br>NCAR<br>National Energy Technology Laboratory | 7 research staff<br>4 postdocs |
| Foreign institution affiliation:<br>Brazil (Brazil Nat Lab)<br>Germany (Juelich, Aachen)<br>Italy (CMCC, University of Salento)<br>Singapore (National University Hospital)<br>Greece (National Technical U of Athens) | 2 graduate students<br>1 assistant professor<br>2 research staff<br>2 postdocs |
| Gender | 54 Male<br>9 Female (11 were selected but 2 did not accept the invitation due to changes in their schedules) |

# ATPESC main lecture hall

# Access to ATPESC 2013 Lectures

- **Slides for all lectures can be downloaded from the 2013 agenda page:**
- **http://extremecomputingtraining.anl.gov/curriculum/2013-agenda/**
- **Videos of the talks will be online soon**

# ATPESC 2013 Lectures
# Supercomputing Architecture Trends

- *Architecting for Exascale* -- Tryggve Fossum (Intel Corp.)
- *An Introduction to Parallel Supercomputing* -- Pete Beckman (ANL)
- *PIM & Memory: The Need for a Revolution in Architecture* -- Peter Kogge (U. of Notre Dame)
- *Macro Architecture for Exascale* -- Pavan Balaji (ANL)
- *Micro Architecture for Exascale* -- Tryggve Fossum (Intel Corp.)

# ATPESC 2013 Lectures
# Programming Models for HPC

- *Programming Models for HPC* -- Marc Snir (ANL)
- *MPI for Scalable Computing* -- Bill Gropp (U. of Illinois), Rusty Lusk (ANL), Rajeev Thakur (ANL)
- *Using OpenMP for Intranode Parallelism, OpenMP 4.0 and the Future of OpenMP* -- Bronis de Supinski (LLNL)
- *Writing Parallel Programs that Work* -- Paul Petersen (Intel)
- *Multi-resolution Global-View Programming in Chapel* -- Brad Chamberlain (Cray)
- *Partitioned Global Address Space Programming* -- Kathy Yelick (LBNL and UC Berkeley**)**
- *Programming with Parallel Migratable Objects* -- Sanjay Kale (U. of Illinois)

# ATPESC 2013 Lectures
# Accelerator and GPU Programming

- *High Performance with Accelerator Programming* – Jeff Vetter (ORNL) and Bronson Messer (ORNL)
- *An Introduction to OpenACC* -- Jeff Vetter (ORNL) and Bronson Messer (ORNL
- *The Evolution of GPU Accelerated Computing* -- Steve Parker (NVIDIA Corp.)
- *Scientific Computing on GPUs* --Nick Frontiere (ANL)
- *MPI and Hybrid Programming Models* – Bill Gropp (U. of Illinois)

# ATPESC 2013 Lectures
## Mathematical Software and Numerical Algorithms

- **Lecturers**
  - Mark Adams (LBNL)
  - Jim Demmel (UC Berkeley)
  - Lori Diachin (LLNL)
  - Jack Dongarra (UT and ORNL)
  - Robert Falgout (LLNL)
  - Glen Hansen (SNL)
  - David Keyes (KAUST)
  - Mark Miller (LLNL)
  - Mark Shepard (RPI)
  - Barry Smith (ANL)
  - Cameron Smith (RPI)
  - Tim Tautges (ANL)
  - Carol Woodward (LLNL)
- **Topics**
  - Architectural adaptation of algorithms
  - Algebraic solvers
  - PETSc toolkit
  - Preconditioners
  - Nonlinear and differential/ algebraic equation solvers
  - AMR technologies
  - Parallel adaptive unstructured mesh simulations
  - Communication-avoiding algorithms
  - Adaptive linear solvers and eigensolvers

# ATPESC 2013 Lectures
# Tookits and Frameworks

- Scott Parker (ANL)
- Mike Brim (ORNL)
- David Lecomber (Allinea **Software Ltd.**)
- Allen Malony (U. of Oregon)
- John Mellor-Crummey (Rice U.)
- Mike Wilde (ANL)

# ATPESC 2013 Lectures
## Community Codes and Application Case Studies

- Salman Habib (ANL)
- Sean Couch (U. of Chicago)
- Jeff Hammond (ANL)
- Hal Finkel (ANL)
- Anshu Dubey (LBNL)
- Jim Phillips (U. of Illinois)

# ATPESC 2013 Lectures
# Visualization and Data Analysis

- Mike Papka (ANL)
- Tom Peterka (ANL)
- Frank Childs (LBNL and U. of Oregon)
- Mark Hereld (ANL)
- Joseph Insley (ANL)
- Venkat Vishwanath (ANL)
- David DeMarle (Kitware, Inc.)

# ATPESC 2013 Lectures
# Data Intensive Computing and I/O

- Rob Ross (ANL)
- Rob Latham (ANL)
- Quincey Koziol (The HDF Group)
- Steve Tuecke (U. of Chicago)

# The Lecture and Hands-on Exercises Room

# Hands-on sessions

- **The ATPESC provides extensive access to leadership-class computers for extensive hands-on sessions**
- **The 2013 ATPESC provided access to**
  - Vesta: IBM Blue Gene/Q with 32K cores
  - Mira: IBM Blue Gene/Q with 768K cores
  - Tukey: visualization and analysis cluster at ALCF
  - Titan: Cray XK7 with 299,008 cores and 18,688 GPUs
  - Keeneland KIDS: a heterogeneous HP SL-390 (Ariston) cluster with 120 nodes, 240 CPUs and 360 GPUs.
- **In addition, during evening hands-on sessions, lecturers and their assistants provide individualized help and tutoring to the participants.**

# Evening hands-on exercises

# Rusty Lusk provides one-one training

# Working into the late evening

# Dinner talks

- **Purpose: present additional topics that will probably be relevant to the participants' research at some point in their career – or are interesting**

# Dinner Talks in 2013

- *Grid, Cloud, and Beyond: what have we learned about computing on demand?* -- Ian Foster (Argonne and U. of Chicago)
- *Petascale Post-doctoral Computing or: How I Learned to Stop Worrying and Blow Up Stars* -- Sean Couch (U. of Chicago)
- *Large-Scale Visual Analysis* -- Chris Johnson (U. of Utah)
- *Scientific Computing while Supercomputing: Software Carpentry* -- Aron Ahmadia (Continuum Analytics)
- *Perspectives on Teaming from the DOE National Labs* -- Lori Diachin (LLNL)
- *Dawn of the Era of Quantum Computation* -- Bob Lucas (USC-ISI)
- *Urban Center for Computation and Data* -- Charlie Catlett (Argonne)
- *Argonne National Laboratory: Innovative Research in the National Interest* -- Mark Peters (Argonne)
- *Big Data + Extreme-scale time to compute -> Actionable Insights* -- Alok Choudhary (Northwestern U.)

# Chris Johnson delivering an evening lecture

# ATPESC Student Snapshot, Jordan Musser

- Jordan works as research engineer in the DOE's National Energy Technology Laboratory's (NETL) Computational and Basic Sciences Division where he is a user and developer of MFIX, an open-source reacting multiphase CFD software.

- "My primary goal for applying to attend ATPESC was to gain knowledge that will help me refine MFIX so we can fully utilize the NETL-SBEUC, a 24K-core system used for energy-related simulations. I intend to incorporate many of the strategies I learned at ATPESC into my regular development activities."

# ATPESC Student Snapshot: Jeff Amelang

- Jeff is a fourth year Ph.D. candidate in Mechanical Engineering at the California Institute of Technology studying coarse-grained molecular dynamics, specifically the atomistic to continuum transition through the quasicontinuum method. His work includes a heavy emphasis on the application of high-performance computing to systems of engineering interest.
- "I study an area of physics that would be out of reach without HPC. I applied for the program because I love HPC, program HPC systems all day, taught an HPC class, and yet I hadn't had formal training in it for lack of opportunity and availability. The class was exactly what I was looking for: background and foundational knowledge, how to leverage the current state of the art, and a glimpse at what the future of HPC looks–all from leaders in the field. It was invaluable."

# ATPESC Student Snapshot: Fredrik Kjolstad

- **Fredrik is a second year Ph.D. student in Computer Science at MIT CSAIL , where he is a member of the D-TEC X-Stack project at work on declarative and operational programming model constructs and compiler techniques to support practical higher-level, tunable and performance-portable programming on extreme-scale machines.**

- **"As a computer scientist, the Programming Model and Application tracks were extremely helpful in giving me a greater understanding of how the algorithms used on extreme-scale machines fit together, and what the challenges of implementing them using contemporary programming models are. I plan to use what I learned at ATPESC to inform my research and to work to transfer the knowledge to the MIT X-Stack team. I hope this will help us develop technologies that solve real problems faced by the HPC community."**

# Argonne Training Program on Extreme-Scale Computing: 2013 participants

# Training

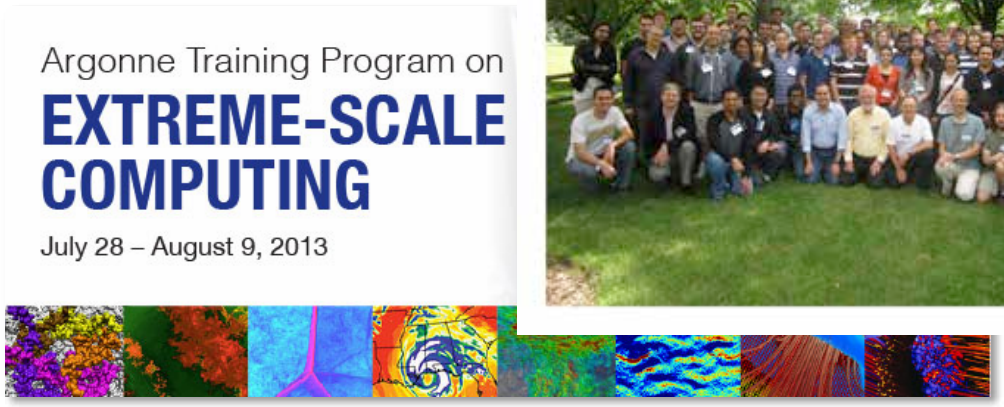


2014 dates:
August 3-15

- **ALCF Outreach**
  - Workshops
    - Tutorials
    - Lectures
    - Hands-on
    - One-on-one with ALCF experts
      - Benchmark, scale up, tune your application
  - Webinars
- *Mira* Performance Bootcamp
- Getting Started
- Leap to Petascale

# Acknowledgments


- The ATPESC is funded by DOE-SC/ASCR
- The FASTMath and SDAV SciDAC institutes presented major portions of the curriculum
- The Argonne and Oak Ridge Leadership Computing Facilities and the Keeneland National Institute for Experimental Computing provided access to their systems for hands-on exercises
- ALCF staff provided many lecturers and support for all aspects of the the program
- Rusty Lusk, Rajeev Thakur, and Pete Beckman (Argonne MCS Division) participated in the organizing committee