

Argonne Training Program on Extreme-Scale Computing

ATPESC 2024 Application Overview w/Q&A


Ray Loy
ATPESC Program Director



Outline

 **Motivation**

 **Overview**

 **Curriculum**

 **Public training materials**

 **Applying for ATPESC 2024**

Motivation for ATPESC

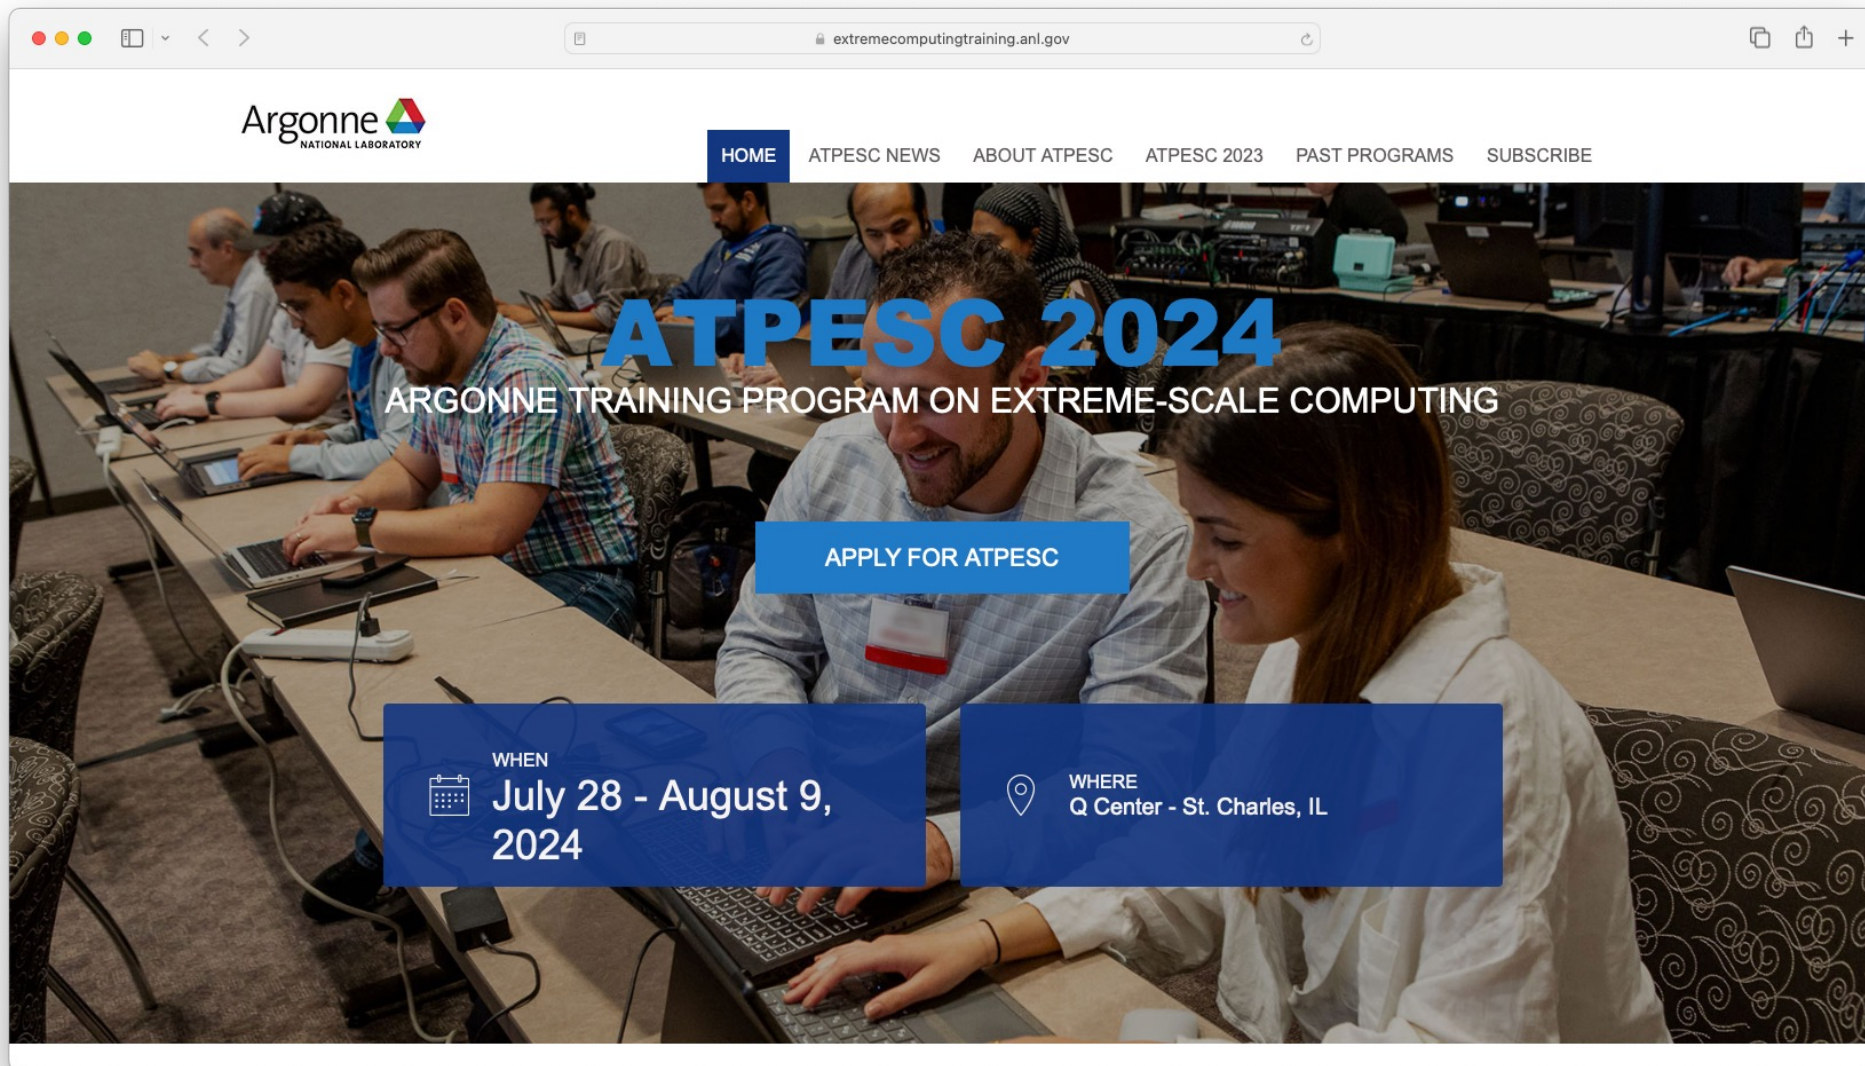
- Today's most **powerful supercomputers** have **complex hardware architectures** and **software environments**
 - and even greater complexity is on the horizon on next-generation and exascale systems
- The **scientific and engineering applications** developed for these systems are themselves **complex**
- There is a **critical need for specialized, in-depth training for the computational scientists** poised to facilitate breakthrough science and engineering using these systems

ATPESC Overview

- Founded by Paul Messina in 2013
- Conceived as a 2-week retreat
- Renowned computer scientists and HPC experts from US national laboratories, universities, and industry serve as lecturers and guide hands-on sessions.
- Target audience: advanced doctoral students, postdocs, and early career computational scientists
- No fee to participate. Domestic travel, meals, and lodging provided.
- Competitive application process reviewed by committee
 - Must have experience in MPI and/or OpenMP and/or Data Science frameworks
 - Experience with at least one HPC system
 - Concrete plans to conduct CSE research on large-scale computers

ATPESC Website

extremecomputingtraining.anl.gov



Argonne – a part of DOE National Laboratory System

Office of Science Laboratories

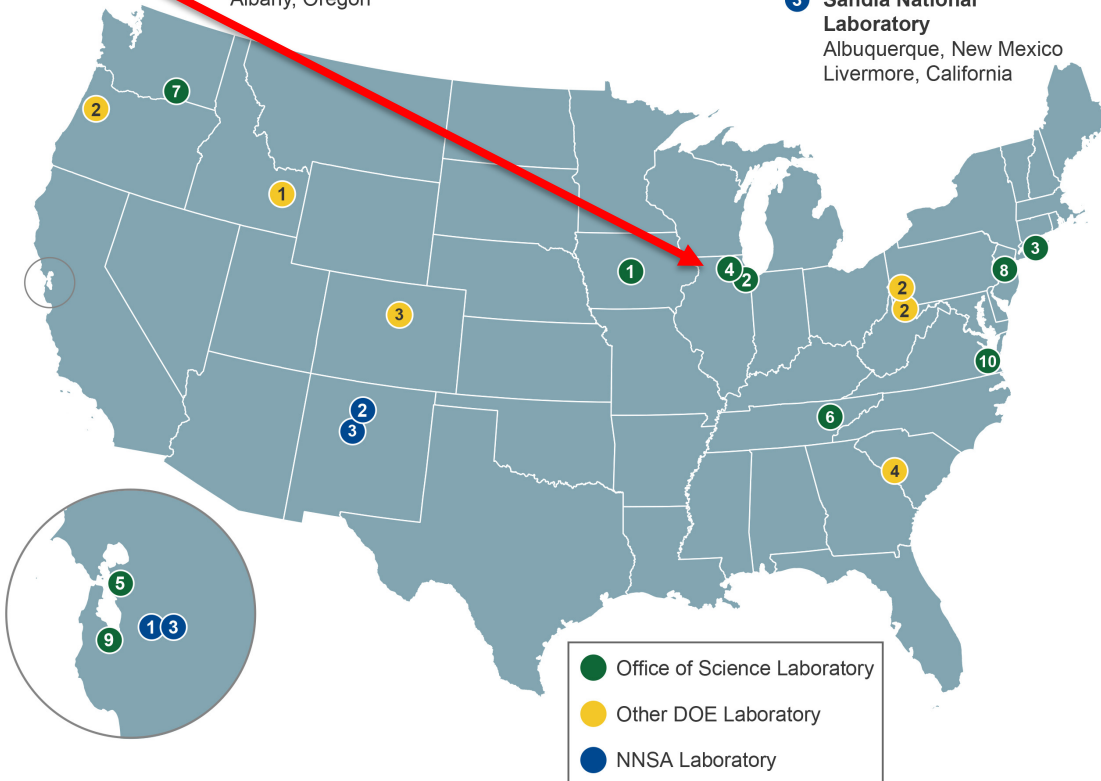
- 1 Ames Laboratory
Ames, Iowa
- 2 Argonne National Laboratory
Argonne, Illinois
- 3 Brookhaven National Laboratory
Upton, New York
- 4 Fermi National Accelerator Laboratory
Batavia, Illinois
- 5 Lawrence Berkeley National Laboratory
Berkeley, California
- 6 Oak Ridge National Laboratory
Oak Ridge, Tennessee
- 7 Pacific Northwest National Laboratory
Richland, Washington
- 8 Princeton Plasma Physics Laboratory
Princeton, New Jersey
- 9 SLAC National Accelerator Laboratory
Menlo Park, California
- 10 Thomas Jefferson National Accelerator Facility
Newport News, Virginia

Other DOE Laboratories

- | | |
|--|---|
| 1 Idaho National Laboratory
Idaho Falls, Idaho | 3 National Renewable Energy Laboratory
Golden, Colorado |
| 2 National Energy Technology Laboratory
Morgantown, West Virginia
Pittsburgh, Pennsylvania
Albany, Oregon | 4 Savannah River National Laboratory
Aiken, South Carolina |

NNSA Laboratories

- 1 Lawrence Livermore National Laboratory
Livermore, California
- 2 Los Alamos National Laboratory
Los Alamos, New Mexico
- 3 Sandia National Laboratory
Albuquerque, New Mexico
Livermore, California



Together, the **17 DOE laboratories** comprise a preeminent federal research system, providing the Nation with strategic scientific and technological capabilities. The laboratories:

- Execute long-term government scientific and technological missions, often with complex security, safety, project management, or other operational challenges;
- Develop unique, often multidisciplinary, scientific capabilities beyond the scope of academic and industrial institutions, to benefit the Nation's researchers and national strategic priorities; and
- Develop and sustain critical scientific and technical capabilities to which the government requires assured access.


Source: https://science.energy.gov/~media/ /images/laboratories/DOE_Laboratories_Map_2014_Hi-res.jpg

Major Scientific User Facilities at Argonne

**Advanced
Photon
Source**



**Argonne Tandem Linear
Accelerator System**



**Center for
Nanoscale
Materials**



**Argonne
Leadership
Computing
Facility
Aurora**

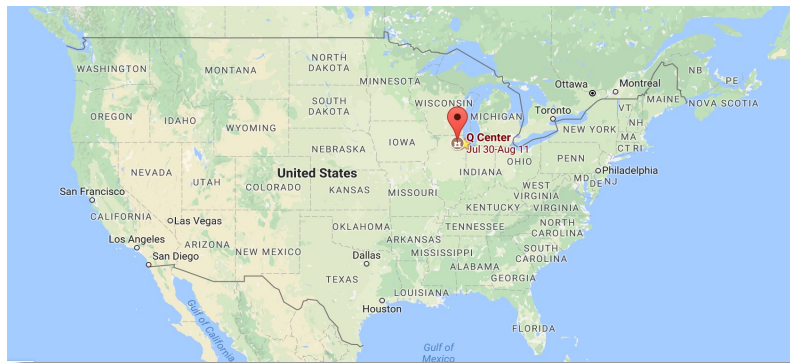
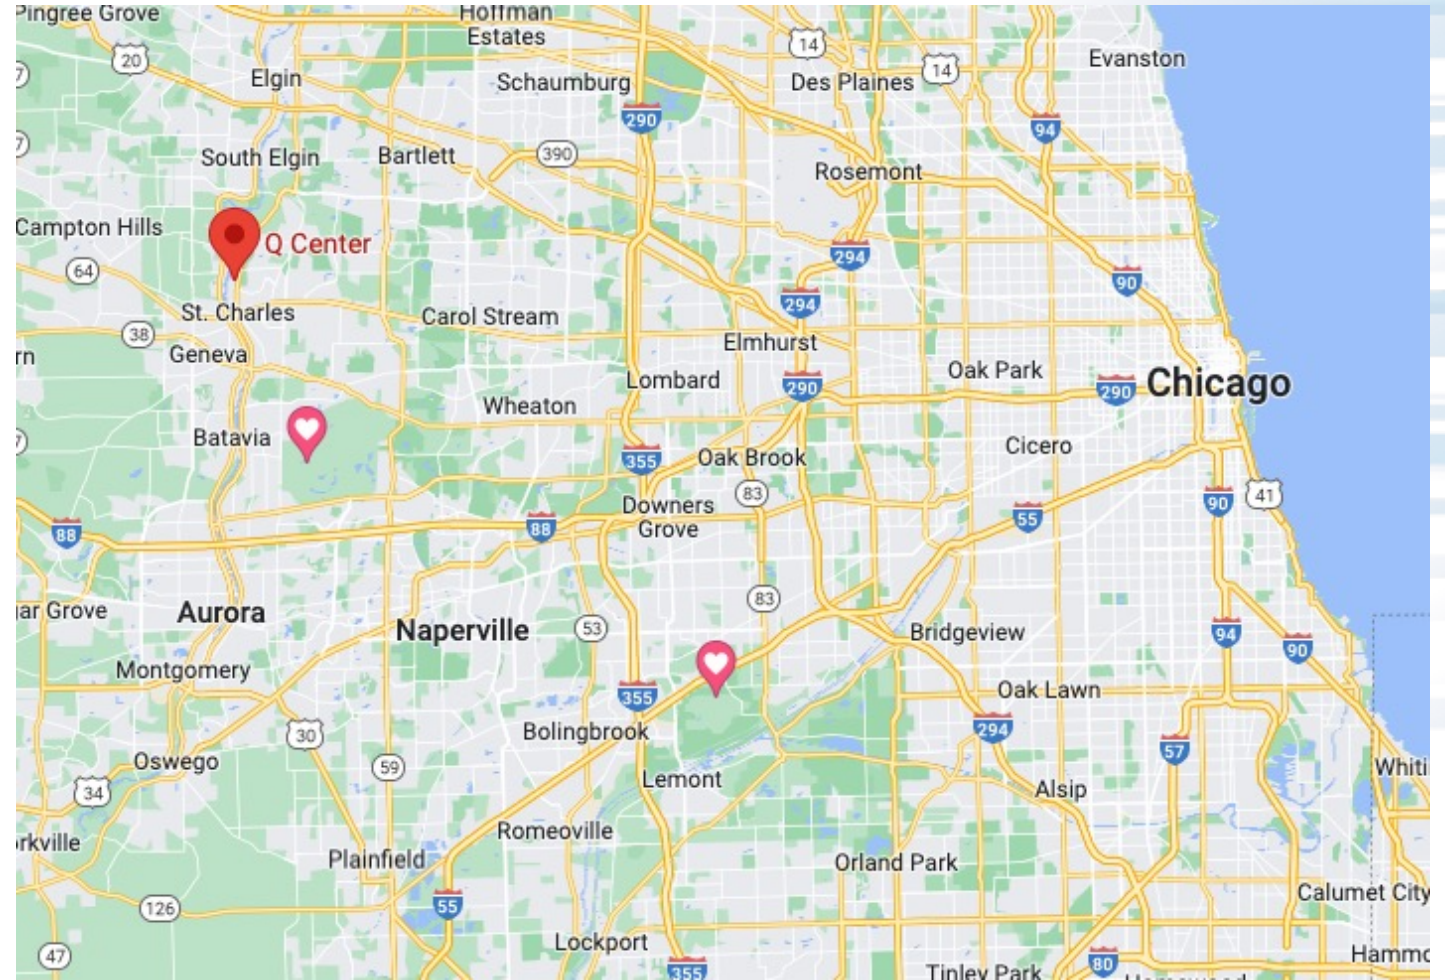


**Electron
Microscopy
Center**



ATPESC 2024 Venue

Q Center



ATPESC 2023



Curriculum Tracks (2023)

Over 100 hours of content

- **Track 1:** Hardware Architectures
- **Track 2:** Programming Models and Languages
- **Track 3:** Data-intensive Computing and I/O
- **Track 4:** Visualization and Data Analysis
- **Track 5:** Numerical Algorithms and Software for Extreme-Scale Science
- **Track 6:** Performance Tools and Debuggers
- **Track 7:** Software Sustainability
- **Track 8:** Machine Learning and Deep Learning for Science

ATPESC Computing Resources

The screenshot shows the website science.osti.gov/User-Facilities/User-Facilities-at-a-Glance/ASCR. The page is titled "ASCR User Facilities" and lists several national scientific user facilities. The left sidebar contains navigation links such as "Home", "About", "Programs", "User Facilities", "Laboratories", "Universities", "Funding", and "Initiatives". The main content area includes a search bar and a list of facilities with accompanying images and brief descriptions.

- ASCR User Facilities**
The Advanced Scientific Computer Research program supports the following national scientific user facilities:
- ALCF – Polaris and ThetaGPU**
Argonne Leadership Computing Facility (ALCF)
The ALCF provides the computational science community with a world-class computing capability dedicated to breakthrough science and engineering.
- OLCF - Ascent**
Oak Ridge Leadership Computing Facility (OLCF)
The OLCF provides the computational science community with world-class computing capability dedicated to breakthrough science and engineering.
- NERSC – Perlmutter, et al.**
National Energy Research Scientific Computing Center (NERSC)
The NERSC is the mission high performance computing facility for the Department of Energy's Office of Science, and is a world leader in accelerating scientific discovery through computation.
- Vendor resources – e.g. Intel Devcloud**

ALCF – Polaris and ThetaGPU

OLCF - Ascent

NERSC – Perlmutter, et al.

Vendor resources – e.g. Intel Devcloud


Goals for Attendees



Exposure to New ideas




Take advantage of ATPESC Resources



Talk F2F with Lecturers, other Participants, Support

Application Instructions

A screenshot of a web browser displaying the 'APPLICATION INSTRUCTIONS' page for the ATPESC training program. The browser's address bar shows 'extremecomputingtraining.anl.gov/application/'. The page has a navigation menu with 'HOME', 'ATPESC NEWS', 'ABOUT ATPESC', and 'PAST PROGRAMS'. The main heading is 'APPLICATION INSTRUCTIONS'. Below it, a section titled 'Your application must include:' lists three requirements: 1. Statement of Purpose (SOP), 2. Curriculum Vitae (CV), and 3. Letter of recommendation (LOR). Each requirement is followed by detailed instructions and a 'Note' section. A vertical scrollbar is visible on the right side of the page content, and a small 'UP' button is located near the scrollbar.

HOME ATPESC NEWS ABOUT ATPESC PAST PROGRAMS

APPLICATION INSTRUCTIONS

Your application must include:

- 1. Statement of Purpose (SOP)**

This document should include your purpose for attending the training program, a description of your computing experience, and your current/future CSE research plans. This part of the application should be no longer than two pages. Please include your First and Last Name, Job Title, and Institution at the top of the document.

 - Note:** the curriculum of the training program is geared for participants with substantial MPI, OpenMP, and/or Data Science framework programming experience who have used an HPC system for a reasonably complex application and have conducted or are preparing to conduct CSE research on large-scale computers. Please be prepared to describe in detail your previous experience with HPC systems (e.g., size of simulations, kind of systems, experience as user and/or developer, software and solvers used, etc.) and your future CSE plans on leadership-class supercomputers. **Specific details about your experience in relation to the Eligibility Criteria are essential to a successful application.**
- 2. Curriculum Vitae (CV)**

The submitted CV must provide all science and technology specialties, all work positions with no gaps in time, and current/accurate names of all academic institutions attended.
- 3. Letter of recommendation (LOR)**

One letter of recommendation from an advisor or supervisor. The recommendation should address your experience with respect to the eligibility requirements. It should include **specific details about your experience and personal contributions in relation to the Eligibility Criteria.**

 - Note:** if your advisor or supervisor wants to make a confidential recommendation, they should send it directly to support@extremecomputingtraining.anl.gov before the submission deadline. In that case, ATPESC support staff will attach it to your application. Please note that you are still responsible for ensuring final submission by your advisor or supervisor. Incomplete applications will not be reviewed.

Please do not include in your application any articles, additional letters, or other documents that are not among the three mentioned above.

When you have compiled all the documents listed above, log in to [EasyChair](#) to apply for ATPESC 2022

Important: on the [EasyChair](#) application form:

Eligibility

"Doctoral students, postdocs, and computational scientists interested in attending ATPESC can review eligibility and application details on the application instructions web page."

- These are the most common applicants but in general, **anyone who is qualified and who would benefit from the program may apply.**
- *If your situation is unusual, just address that in your Statement of Purpose.*

Eligibility

- Substantial experience in MPI and/or OpenMP programming. Equivalent experience with Data Science frameworks will also be considered.
- Experience using at least one HPC system for a reasonably complex application.
- Plans to conduct computational science and engineering (CSE) research on large-scale computers.

Application Components

- Statement of Purpose
- Curriculum Vitae (CV)
- Letter of Recommendation (LOR)

The Key to success

Provide details

The #1 problem a reviewer faces is trying to guess your qualifications when the materials are not clear.

1. Substantial experience in MPI and/or OpenMP programming. Equivalent experience with Data Science frameworks will also be considered.

- Weak statements
 - "I took a class on MPI"
 - "I ran a code that uses MPI"
- Stronger statements
 - "I work on the NNN code. I implemented and tested the message-passing routines for the boundary exchange of the MMM solver which has a communication pattern that exchanges between... taking into account.... Benchmarking of performance showed..."
- If you are part of a group-developed code or working on a community code, be specific about your individual contribution.

Experience using at least one HPC system for a reasonably complex application.

- We do not expect that you have access to a world-class system
 - University/corporate clusters are relevant
- Be clear if you are developing code and running it vs. running an existing code
- What types of runs have you accomplished?
 - Core/node counts are informative
 - Did you perform scaling studies or other performance inquiries?
- What type of queuing system(s) were involved?
 - Did you construct job scripts?
 - Did you run a large campaign and employ any workflow techniques to manage jobs?

Plans to conduct computational science and engineering (CSE) research on large-scale computers.

- The training should be put to good use
- What is your current situation?
 - Are you currently part of a research effort?
- Where are you headed in your career?
 - If you are a student, what positions are you applying to?
 - How do you expect your current work to carry over?

ATPESC Communications

- Subscribe to "info" mailing list linked from the homepage
 - <https://extremecomputingtraining.anl.gov>
 - *Menu tab at top right -or- button at bottom of page*
- Direct inquiries to
 - support@extremecomputingtraining.anl.gov

Acknowledgments

This training is supported by, and uses resources of, the Argonne Leadership Computing Facility, a U.S. Department of Energy (DOE) Office of Science user facility at Argonne National Laboratory and is based on research supported by the U.S. DOE Office of Science-Advanced Scientific Computing Research Program, under Contract No. DE-AC02-06CH11357.

Questions/Discussion

