Center for Accelerated Application Readiness

Call for Proposals

Getting Applications Ready for Summit

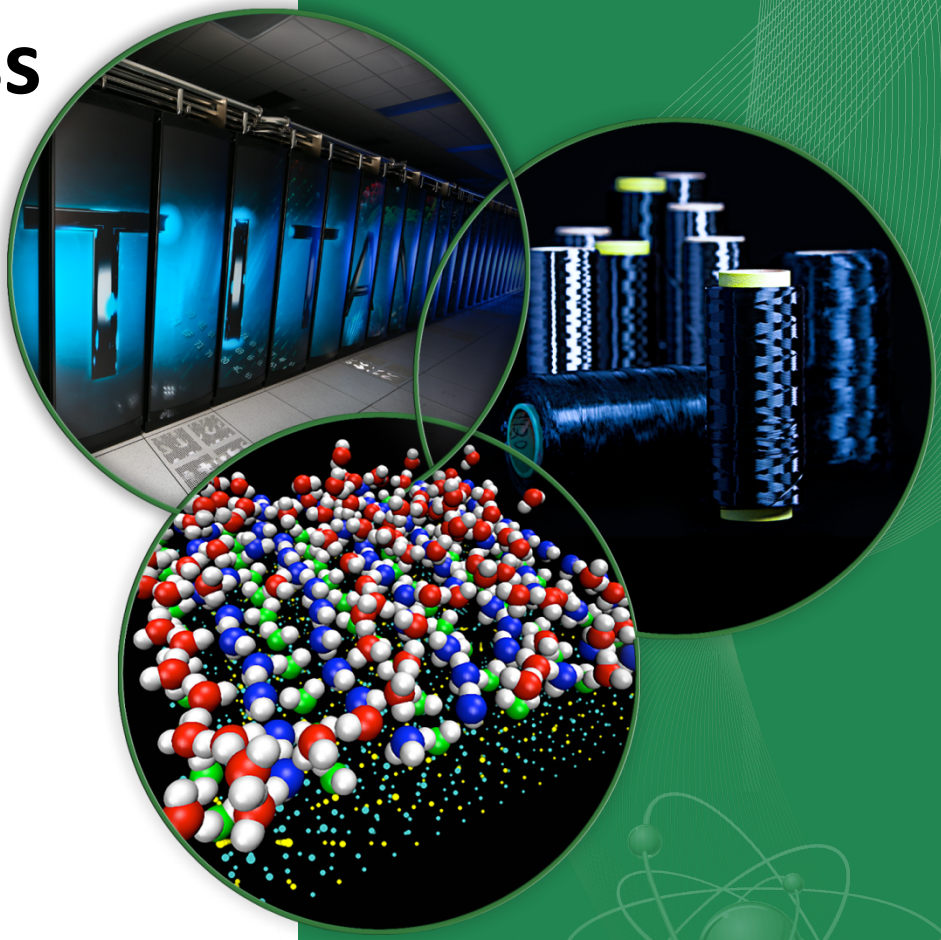

Tjerk Straatsma

OLCF Scientific Computing Group

OLCF User Group Presentation, December 3, 2014

ORNL is managed by UT-Battelle
for the US Department of Energy

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

Application Architecture and Performance Portability

Common Goals for ALCF, NERSC and OLCF

Two Tracks for Future Large Systems

Tianhe-2 (NUDT): TH-IVB-FEP
Intel Xeon E5-2692 12 C 2.2 GHz
TH Express-2
Intel Xeon Phi 31S1P

Titan (Cray): Cray XK7
AMD Opteron 6274 16C 2.2 GHz
Cray Gemini
NVIDIA K20x

Sequoia (IBM): BlueGene/Q
Power BQC 16C 1.6 GHz

K computer (Fujitsu)
SPARC64 VIIIfx 2.0 GHz
Tofu

Mira (IBM): BlueGene/Q
PowerPC A2 16C 1.6 GHz

Piz Daint (Cray): Cray XC30
Intel Xeon E5-2670 8C 2.6 GHz
Cray Aries
NVIDIA K20x

Edison (Cray): Cray XC30
Intel Xeon E%-2695v2 12C 2.4 GHz
Aries

Many Core

- 10's of thousands of nodes with millions of cores
- Homogeneous cores
- Multiple levels of memory – on package, DDR, and non-volatile
- Unlike prior generations, future products are likely to be self hosted

Hybrid Multi-Core

- CPU / GPU Hybrid systems
- Likely to have multiple CPUs and GPUs per node
- Small number of very fat nodes
- Expect data movement issues to be much easier than previous systems – coherent shared memory within a node
- Multiple levels of memory – on package, DDR, and non-volatile

Cori at NERSC

- Self-hosted many-core system
- Intel/Cray
- 9300 single-socket nodes
- Intel® Xeon Phi™ Knights Landing (KNL)
- 16GB HBM, 64-128 GB DDR4
- Cray Aries Interconnect
- 28 PB Lustre file system @ 430 GB/s
- Target delivery date: June, 2016

Summit at OLCF

- Hybrid CPU/GPU system
- IBM/NVIDIA
- 3400 multi-socket nodes
- POWER9/Volta
- More than 512 GB coherent memory per node
- Mellanox EDR Interconnect
- Target delivery date: 2017

ALCF-3 at ALCF

- TBA
- Target delivery date: 2017-18

Importance of Performance Portability

Application portability among NERSC, ALCF and OLCF architectures is critical concern of ASCR

- Application developers target wide range of architectures
- Maintaining multiple code version is difficult
- Porting to different architectures is time-consuming
- Many Principal Investigators have allocations on multiple resources
- Applications far outlive any computer system

Primary task is exposing parallelism and data locality

Challenge is to find the right abstraction:

- MPI + X (X=OpenMP, OpenACC)
- PGAS + X
- DSL
- …

Synergy Between Application Readiness Programs

NESAP at NERSC

NERSC Exascale Science Application Program

- Call for Proposals – June 2014
- 20 Projects selected
- Partner with Application Readiness Team and Intel IPCC
- 8 Postdoctoral Fellows

Criteria

- An application's computing usage within the DOE Office of Science
- Representation among all 6 Offices of Science
- Ability for application to produce scientific advancements
- Ability for code development and optimizations to be transferred to the broader community through libraries, algorithms, kernels or community codes
- Resources available from the application team to match NERSC/Vendor resources

CAAR at OLCF

Center for Accelerated Application Readiness

- Call for Proposals – November 2014
- 8 Projects to be selected
- Partner with Scientific Computing group and IBM/NVIDIA Center of Excellence
- 8 Postdoctoral Associates

Criteria

- Anticipated impact on the science and engineering fields
- Importance to the user programs of the OLCF
- Feasibility to achieve scalable performance on Summit
- Anticipated opportunity to achieve performance portability for other architectures
- Algorithmic and scientific diversity of the suite of CAAR applications.
- Optimizations incorporated in master repository
- Size of the application's user base

ESP at ALCF

Early Science Program

- Call for Proposals
- 10 Projects to be selected
- Partner with Catalyst group and ALCF Vendor Center of Excellence
- Postdoctoral Appointee per project

Criteria

- Science Impact
- Computational Readiness
- Proposed science problem of appropriate scale to exercise capability of new machine
- Confidence code will be ready in time
- Project code team appropriate
- Willing partner with ALCF & vendor
- Diversity of science and numerical methods
- Samples spectrum of ALCF production apps

Synergy Between Application Readiness Programs

- Application Developer Team involvement
  - Knowledge of the application
  - Work on application in development "moving target"
  - Optimizations included in application release
- Early Science Project
  - Demonstration of application on real problems at scale
  - Shake-down on the new system hardware and software
  - Large-scale science project is strong incentive to participate
- Vendor technical support is crucial
  - Programming environment often not mature
  - Best source of information on new hardware features
- Access to multiple resources, including early hardware
- Joint training activities

- Portability is a critical concern
- Experience benefits other developers and users
  - Coverage of scientific domains
  - Coverage of algorithmic methods and programming models
- Persistent culture of application readiness
  - More computational ready applications available
  - Experience of science liaisons and catalysts for user programs
  - Synergy with libraries and tools projects

OLCF Center for Accelerated Application Readiness

WL-LSMS
Illuminating the role of material disorder, statistics, and fluctuations in nanoscale materials and systems.

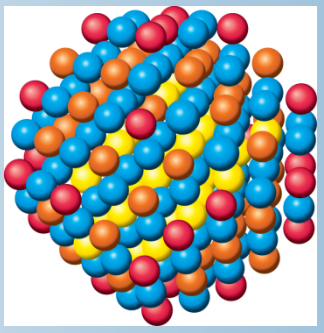

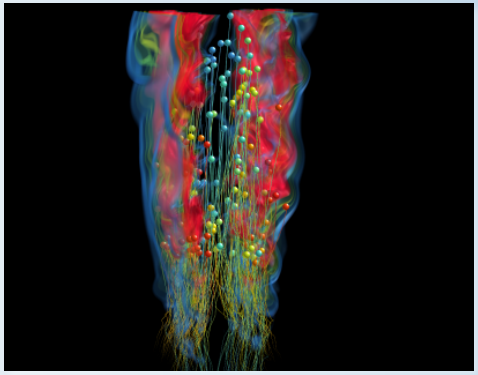

S3D
Understanding turbulent combustion through direct numerical simulation with complex chemistry.

NRDF
Radiation transport – important in astrophysics, laser fusion, combustion, atmospheric dynamics, and medical imaging – computed on AMR grids.

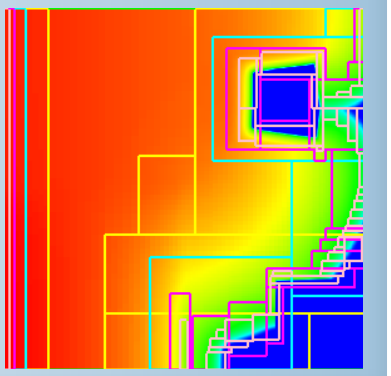

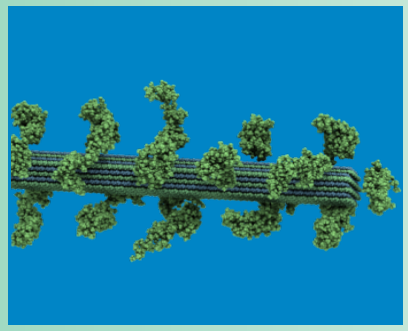

LAMMPS
A molecular description of membrane fusion, one of the most common ways for molecules to enter or exit living cells.

CAM-SE
Answering questions about specific climate change adaptation and mitigation scenarios; realistically represent features like precipitation patterns / statistics and tropical storms.

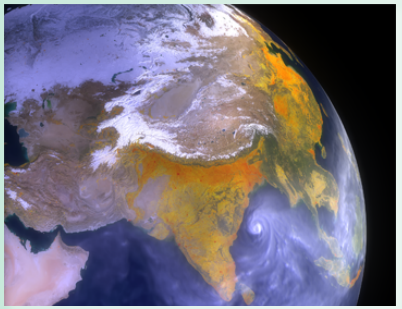

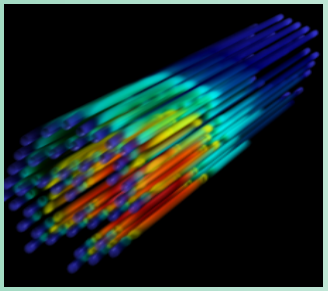

Denovo
Discrete ordinates radiation transport calculations that can be used in a variety of nuclear energy and technology applications.

OLCF Center for Accelerated Application Readiness

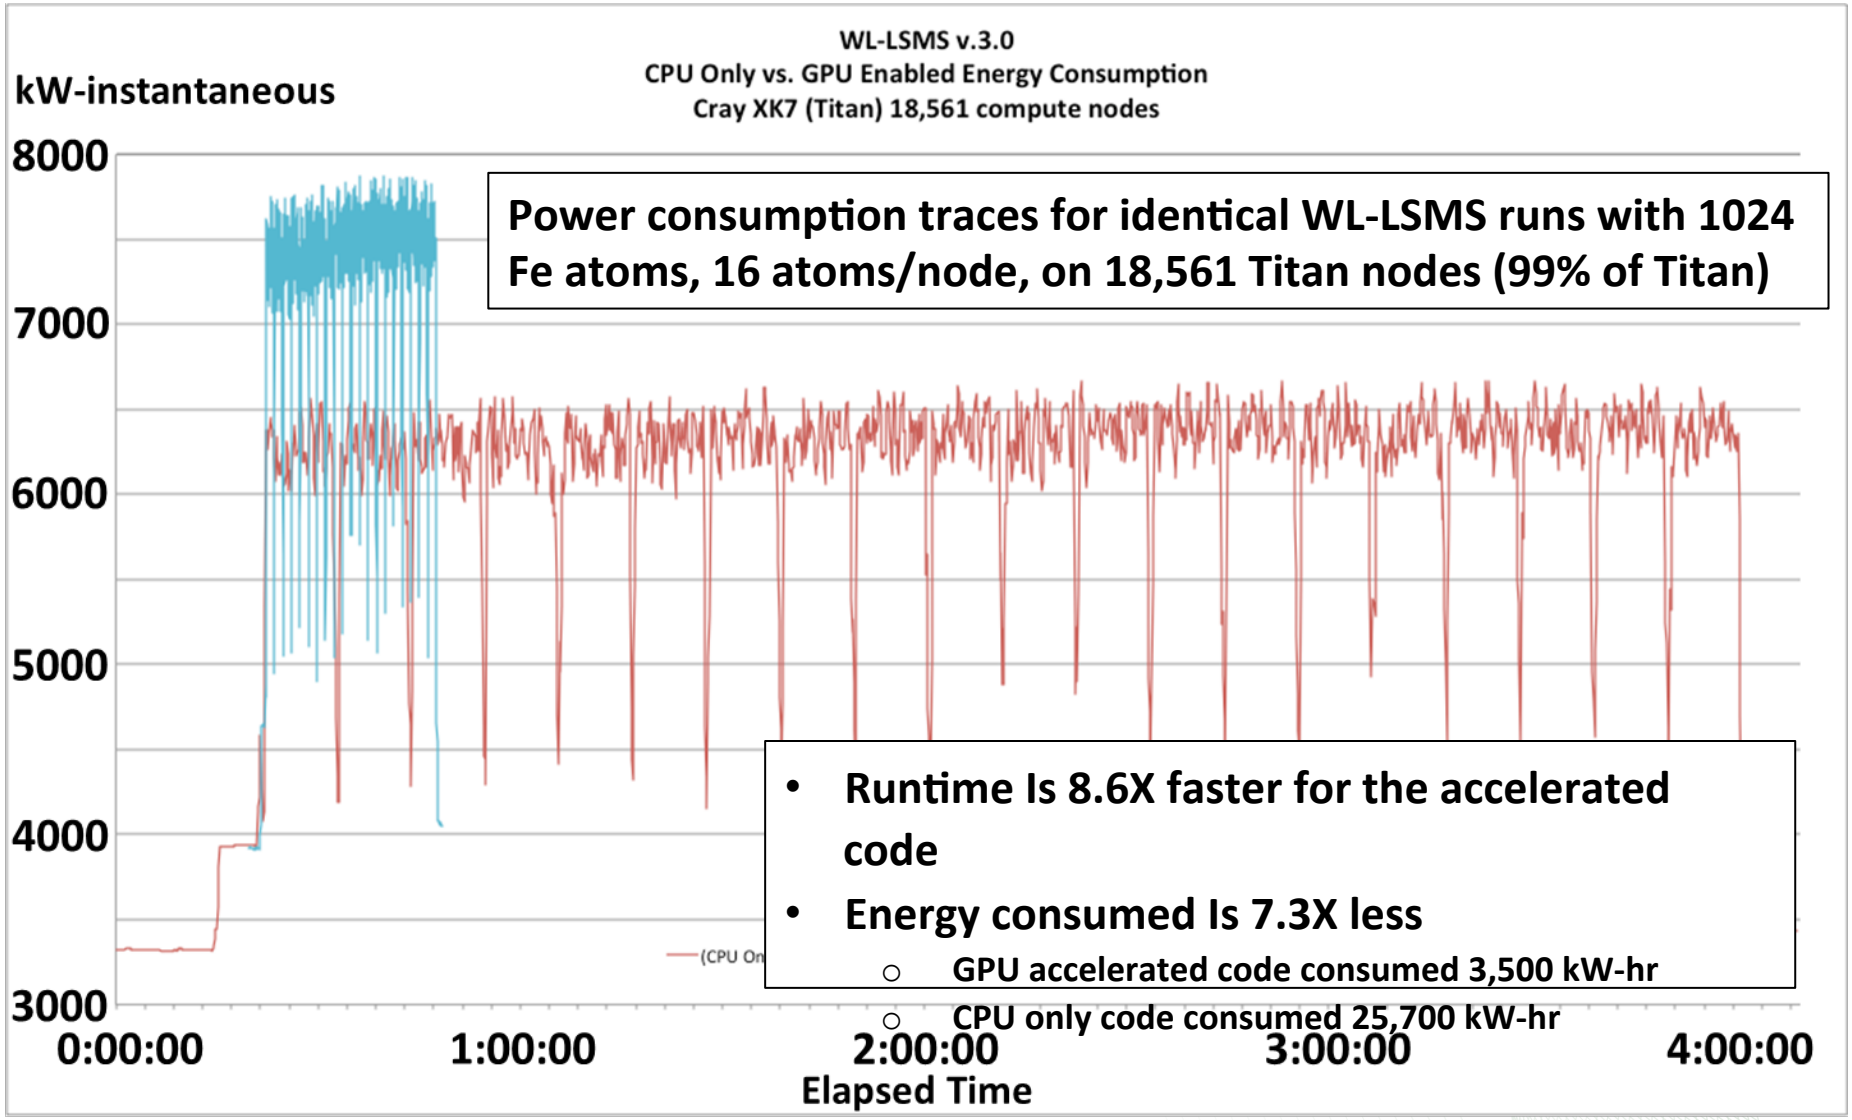

CAAR

Application Readiness Partnership Projects

CAAR Projects Overview

- Call for Proposals for eight Partnership Projects
- Partnership between
  - Application Development Team
  - OLCF Scientific Computing group
  - IBM/NVIDIA Center of Excellence
- Application Readiness phase for restructuring and optimization
- Early Science phase for grand-challenge scientific campaign
- OLCF Postdoctoral Associate per project
- Extensive training on hardware and software development environment
- Portability is critical concern
- Coordination with NERSC NESAP and ALCF ESP programs
- Allocations on Titan, early delivery systems and Summit
- Allocations on NERSC and ALCF

CAAR Partnership Responsibilities

- *Develop and execute a **Technical plan*** for application porting and performance improvement, developed and executed with reviewable milestones
- *Develop and work according to a **Management plan*** with clear description of responsibilities of the CAAR team
- *Develop and execute an **Early Science project*** for compelling scientific grand-challenge campaign
- *Assign an **Application Scientist*** to carry out the Early Science campaign together with the CAAR team
- *Provide **Documentation*** for semi-annual reviews of achieved milestones, and intermediate and final reports

CAAR Partnership Resources

- The core development team of the application, with a stated level of effort
- An ORNL Scientific Computing staff member, who will partner with the core application development team
- A full-time postdoctoral fellow, located and mentored at the OLCF
- Technical support from the IBM/NVidia Center of Excellence
- Allocation of resources on Titan
- Allocation of resources at ALCF & NERSC
- Access to early delivery systems and the *Summit*
- Allocation of compute resources on the full *Summit* system for the Early Science campaign

CAAR Partnership Activities

1. Common training of all Application Readiness teams

a. Architecture and performance portability

b. Avoidance of duplicate efforts

2. Application Readiness Technical Plan Development and Execution

a. Code analysis & benchmarking to understand application characteristics: code structure, code suitability for architecture port, algorithm structure, data structures and data movement patterns, code execution characteristics ("hot spots" or "flat" execution profile)

b. Develop parallelization and optimization approach to determine the algorithms and code components to port, how to map algorithmic parallelism to architectural features, how to manage data locality and motion

c. Decide on programming model such as compiler directives, libraries, explicit coding models

d. Execute technical plan– benchmarking, code rewrite or refactor, porting and testing, managing portability, managing inclusion in main code repository

3. Development and Execution of and Early Science Project, i.e., challenging science problem that demonstrates the performance and scientific impact of the developed application port

CAAR Timeline

1. November 2014: Call for CAAR applications
2. **February 20, 2015: CAAR proposal deadline**
3. March 2015: Selection of CAAR application teams
4. March 2015: CAAR application training workshop
5. April 2015: CAAR application teams start
6. June 2016: CAAR project review
7. September 2017: Call for Early Science projects
8. November 2017: Selection Early Science projects
9. December 2017: Early Science projects start
10. June 2019: Early Science project ends

Future Computational Scientists for Energy, Environment and National Security – Training Program

ASCR facilities host Distinguished Postdoctoral Associates programs with the objective of training the next generation of computational scientists

These programs provide:

1. Challenging scientific campaigns in critical science mission areas
2. Experience in using ASCR computing capabilities
3. Training in software development and engineering practices for current and future massively parallel computer architectures

Central to achieving these goals is access to leadership computing resources, availability of computational domain scientists to provide adequate mentoring and guidance, and facilities' association with universities with strong computational and computer science programs.

OLCF Distinguished Postdoctoral Associates Program

At the OLCF, **eight** Distinguished Postdoctoral Associate positions are available immediately for candidates interested in and capable of performing leading-edge computational science research and development.

Priorities include the development of methodologies and their efficient massively parallel implementation on current state-of-the-art accelerated computer architectures, as well as their application on large scientific challenge problems.

The center is specifically looking for candidates with strong computational expertise in the following scientific areas:

Astrophysics, Biophysics, Chemistry, Climate Science, Combustion, Fusion Energy Science, Materials Science, and Nuclear Physics.

For information about the positions or to apply, visit:
https://www.olcf.ornl.gov/summit/olcf-distinguished-postdoctoral-associates-program/

CAAR Selection Criteria

- *Composition, experience and commitment of application development team*
- *Assessment of anticipated scientific impact of ported application*
- *Assessment of porting feasibility based on provided porting plan and benchmarks*
- *Application user base*
- *Compelling vision of an Early Science project*

- *Technical domain expertise of OLCF liaison*
- *Technical expertise of IBM/NVIDIA Center of Excellence*
- *Coverage of science domains in CAAR portfolio and support for DOE and US mission*
- *Coverage of algorithms, programming approaches, languages, data models*

- *Consultation with NERSC and ALCF*
- *Consultation with DOE Office of Advanced Scientific Computing Research*

OLCF Scientific Computing

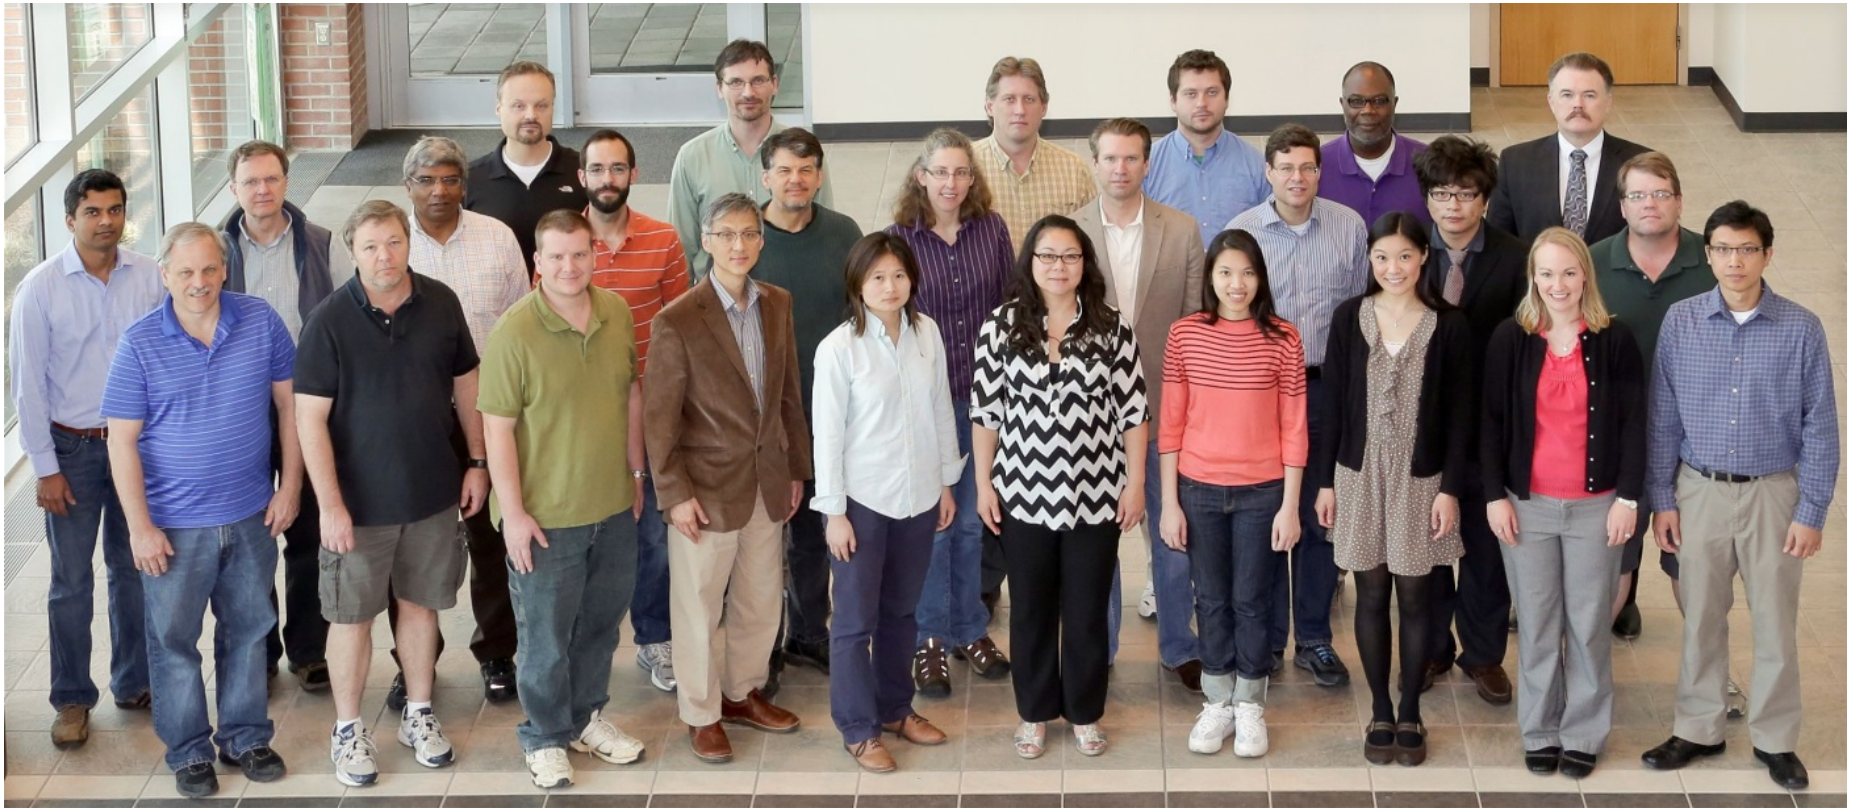

Ramanan Sankaran, Mike Matheson, George Ostrouchov, Duane Rosenberg, Valentine Anantharaj, Bronson Messer, Mark Berrill, Matt Norman, Ed D'Azevedo, Norbert Podhorski, Wayne Joubert, JJ Chai (postdoc, now in CSM), Judy Hill, Mark Fahey, Hai Ah Nam, Jamison Daniel, Dmitry Liakh, Supada Loasooksathit (postdoc), Markus Eisenbach, Arnold Tharrington, Ying Wai Li, Mingyang Chen (postdoc), Peyton Ticknor, Tjerk Straatsma, Dave Pugmire and Jan-Michael Carrillo (postdoc, now at SNS).

Questions & Discussion