EMSL

NWChem: Enabling cutting-edge open-source science from the teraflop to the exaflop

W.A. de Jong, K. Kowalski, E.J. Bylaska, H.J.J. van Dam, N. Govind, T.P. Straatsma, K. Bhaskaran-Nair, O. Villa, A. Vishnu, S. Krishnamoorthy

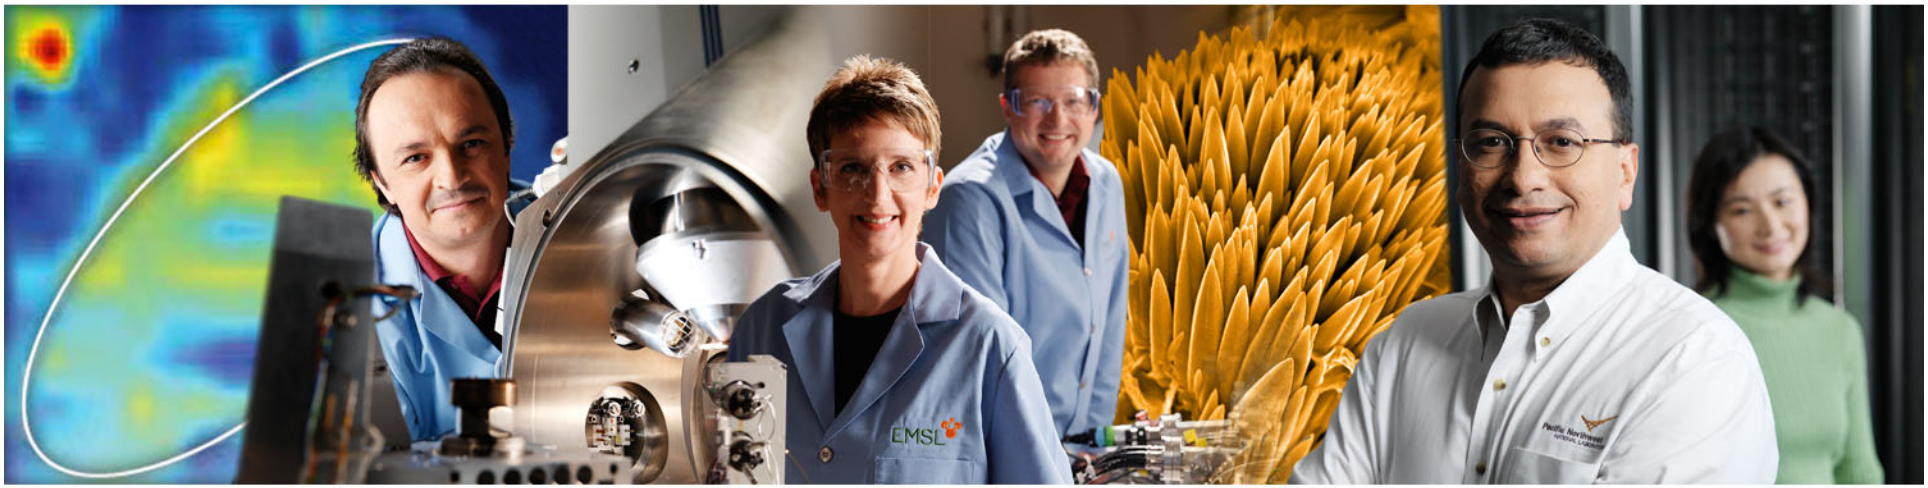

www.emsl.pnl.gov

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

U.S. DEPARTMENT OF ENERGY

Science driven software development

EMSL

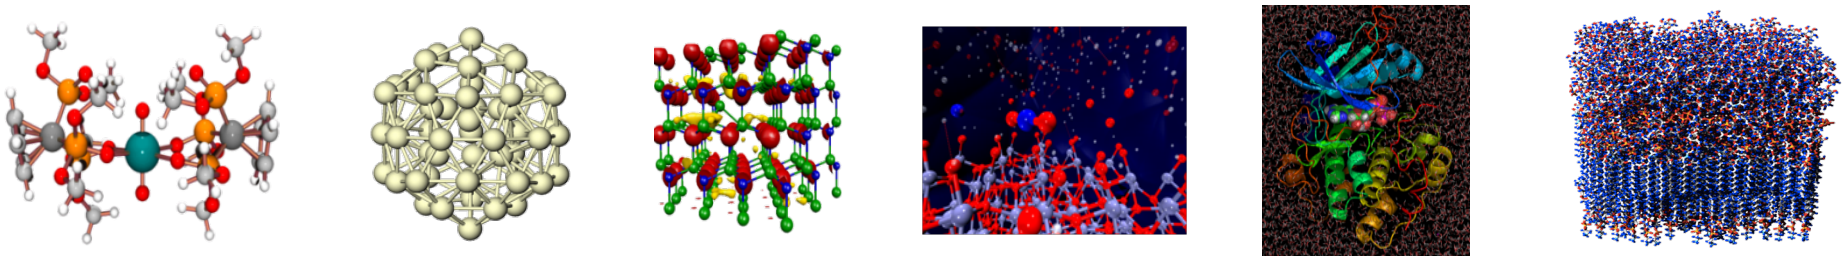

- DOE's Premier computational chemistry software
- One-of-a-kind solution scalable with respect to scientific challenge and compute platforms
- From molecules and nanoparticles to solid state and biomolecular systems
- Portable – runs on a wide range of computers
  - Supercomputer to Mac or PC with Windows

http://www.nwchem-sw.org

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

U.S. DEPARTMENT OF ENERGY

NWChem is Open-Source

- NWChem consortium delivers code and infrastructure for computational chemistry community to build upon
- Establish more collaborative development environment
- Close to 55K downloads in under 2 years
- License is Educational Community License (ECL 2.0)
  - Apache style license

NWChem is expanding its developers community

EMSL

Charge Transfer and Charge Transport in Photoactivated Systems

Developing Advanced Methods for Excited State Chemistry in the NWChem Software Suite

CAS/RASPT2 approaches

Approaches for ultrafast and excited state dynamics in solution, conformational sampling

Strong links with SciDAC SUPER and FASTMATH Institutes







BROOKHAVEN NATIONAL LABORATORY

UNIVERSITY OF MINNESOTA Driven to Discover℠

PENNSTATE 1855

U.S. DEPARTMENT OF ENERGY

Parallel scalability in NWChem through Global Array Toolkit

EMSL

- NWChem uses Global Array Toolkit to scale and perform on parallel architectures
- Programming model: Distributed dense arrays that can be addressed through a shared memory-like style and one-sided access

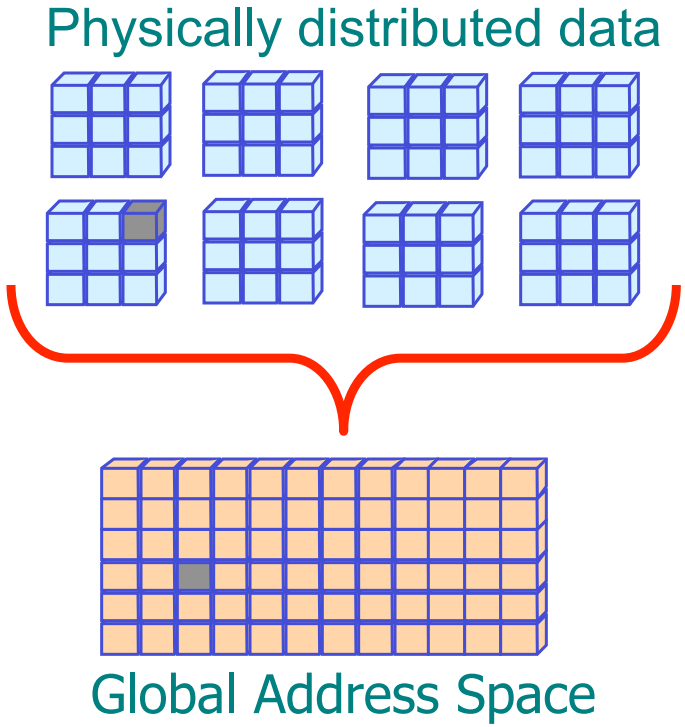

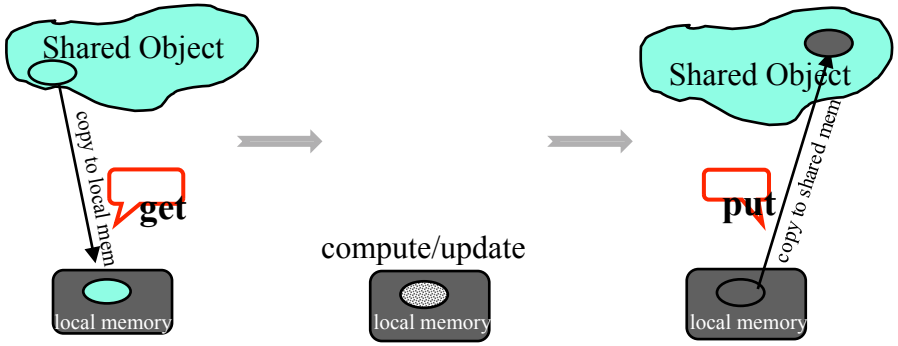

- Hides much of the necessary parallel infrastructure but awareness of locality of data for scalability
- Core-based model has to change change on exascale machines!!

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

U.S. DEPARTMENT OF ENERGY

Scalable single and multi-reference coupled cluster methods

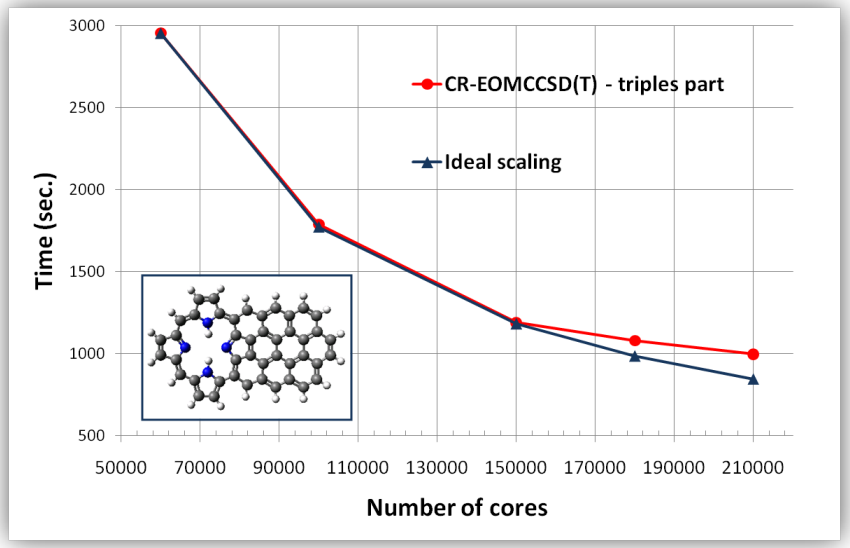

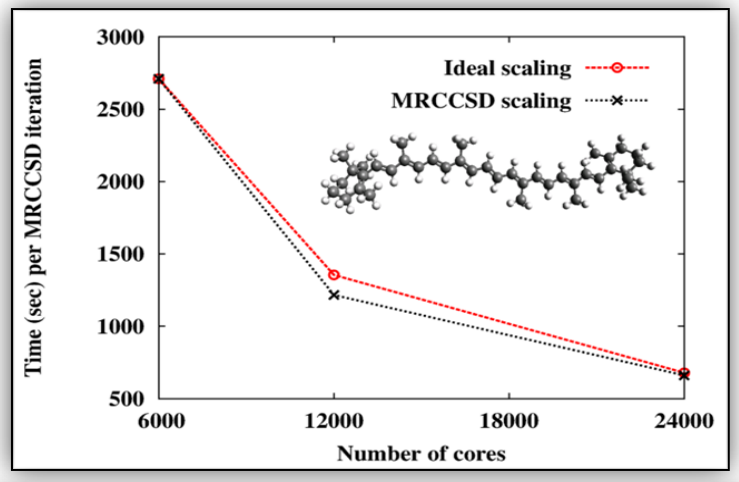

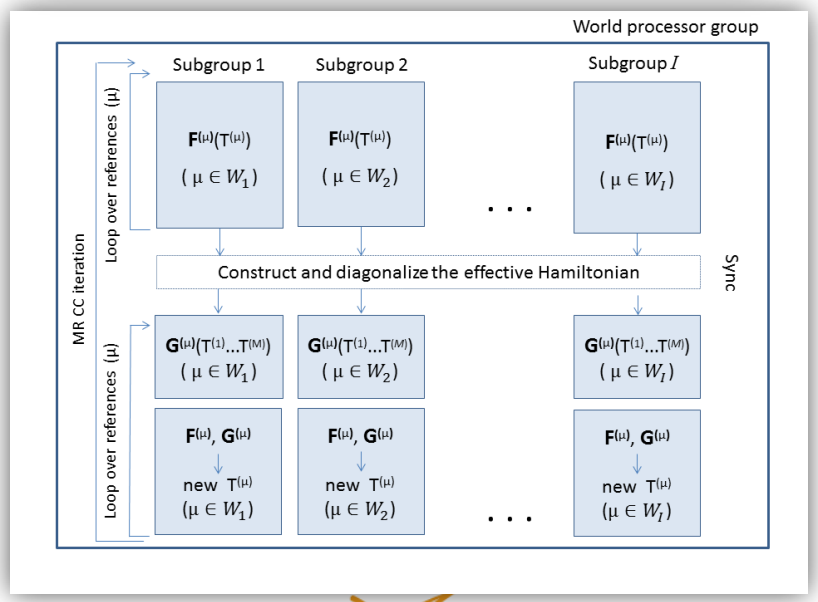

- MRCC in NWChem
  - Brillouin-Wigner MRCCSD
  - Mukherjee Mk-MRCCSD
  - State-Universal MRCCSD
  - Perturbative triples corrections MRCCSD(T)
  - Demonstrated scalability of MR-CCSD across 24,000 cores

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

U.S. DEPARTMENT OF ENERGY

Parallel Gaussian DFT performance

EMSL

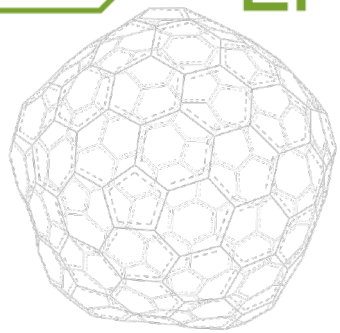

- Calculation on C_{240}
  - PBE0 functional, 6-31G*
  - Direct integral evaluation
  - Size 3600 basis functions
- Timings for different components of the Kohn-Sham matrix construction (scalar, vector, BLAS-2)
  - Fock 2e – two electron integrals
  - Fock xc – the DFT contribution
  - Diagonalization – eigenvector solver
- Scalability limited by diagonalization
  - Avoiding direct diagonalization, switched to quadratic convergence solver
- Other efforts are focused on
  - Reduced data movement
  - Reduced (linear scaling) approaches
  - Compute at lower precision

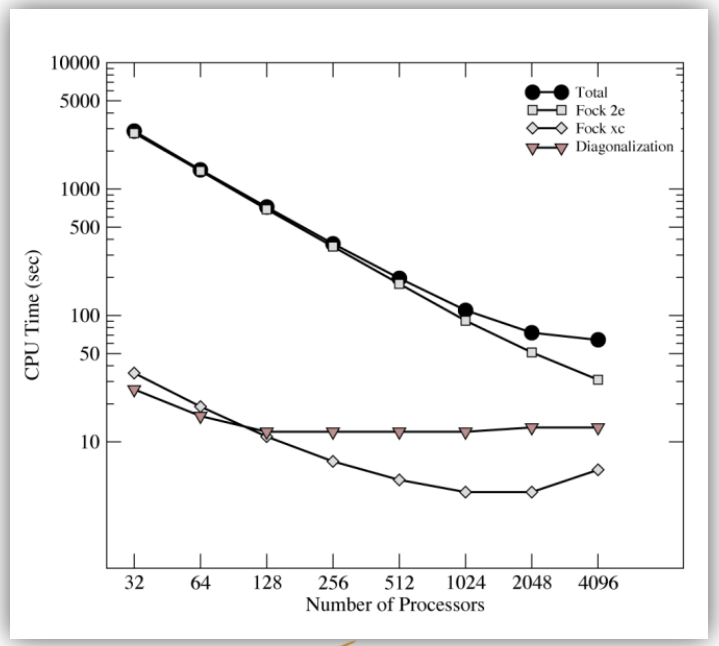

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

U.S. DEPARTMENT OF ENERGY

Plane wave DFT and dynamics for solution, surfaces, and materials

EMSL

- Plane wave density functional theory
  - Driven by FFTs and global DGEMMs
- Extensive dynamics functionality with Car-Parrinello
  - Parallel-in-time development
- Various exchange-correlation functionals
  - Exact exchange is very efficient
  - Lots of FFTs here

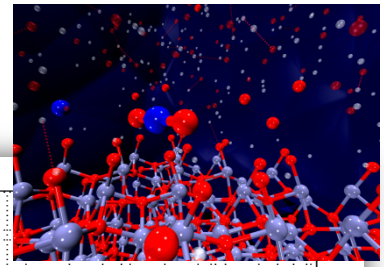

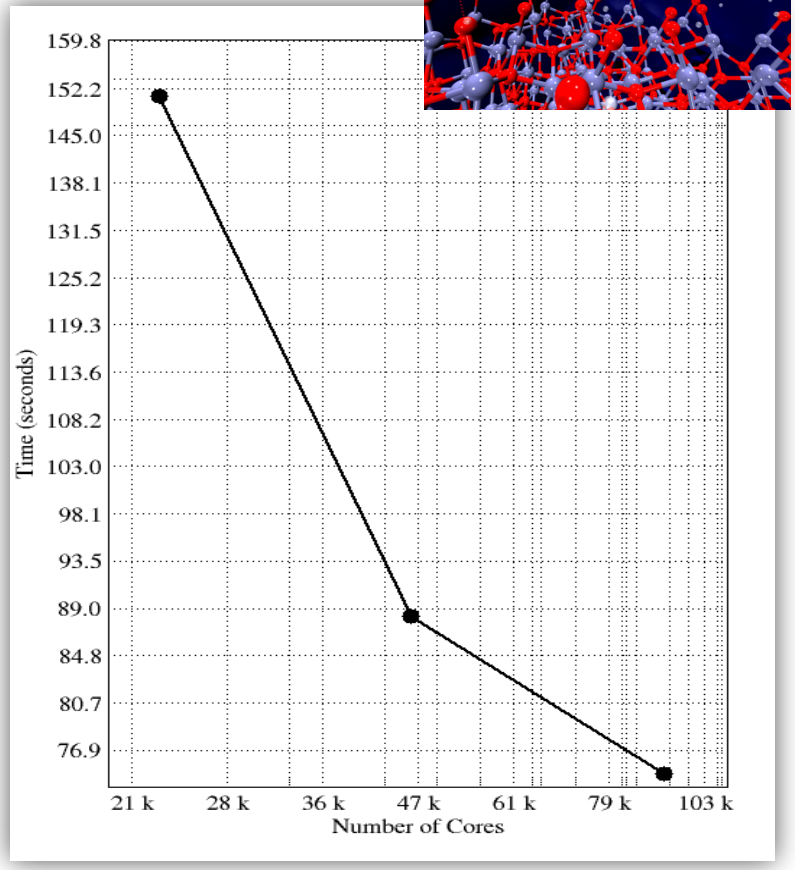

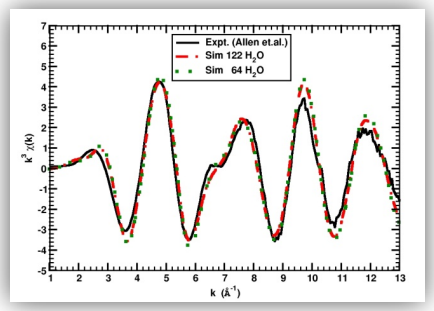

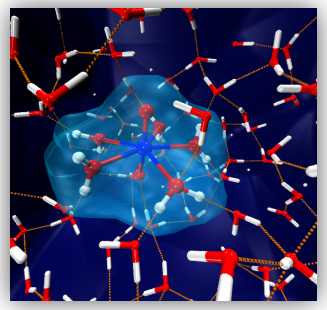

AIMD accurately models EXAFS of uranyl in water

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

U.S. DEPARTMENT OF ENERGY

Next-generation Molecular Dynamics

EMSL

- Removal of all explicit global synchronizations from the basic MD time steps
- More effective hiding of communication latency and bandwidth through computation scheduling
- Improved scalability through cell-cell pair instead of cell distribution
  - Minimizing communication
  - Improved load balancing

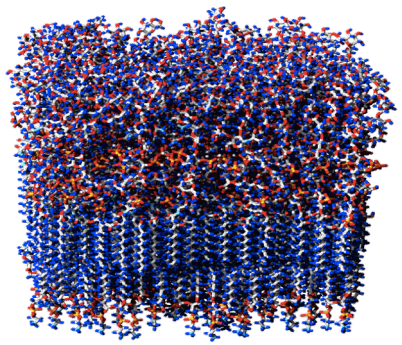

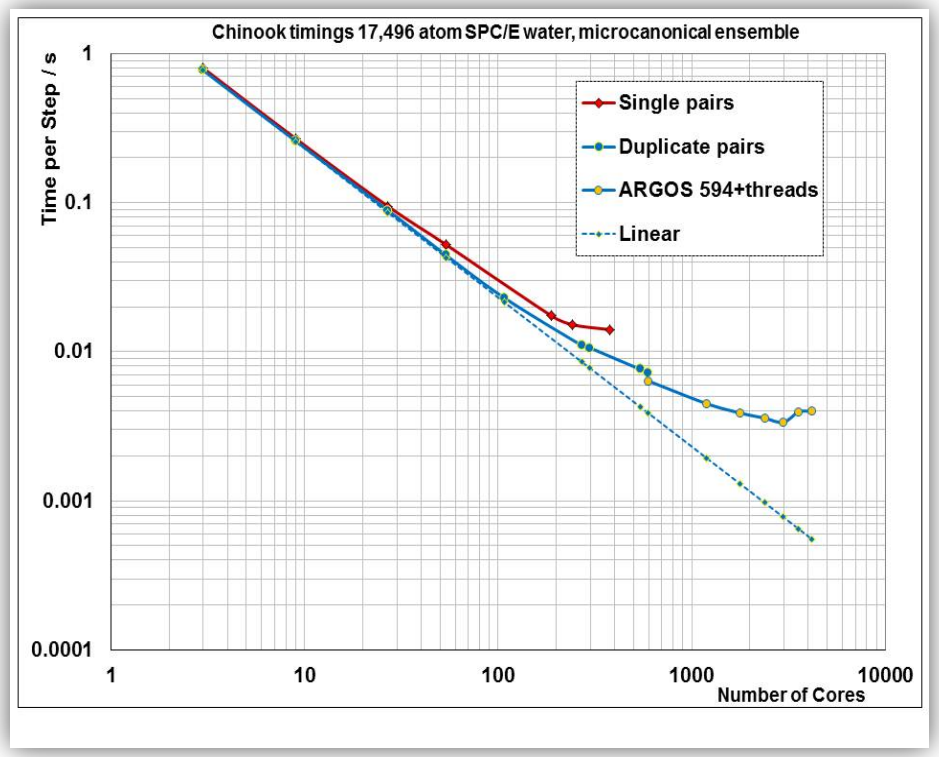

Scalability for 18,000 atoms on 3,000 cores

That is only 6 atoms per core!

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

U.S. DEPARTMENT OF ENERGY

Coupled-cluster on GPGPUs

- The CCSD(T)/Reg-CCSD(T) codes have been rewritten in order to take advantage of GPGPU accelerators
- 8X speedup obtained by moving coupled cluster DGEMMs to GPGPU
  - Efficient overlap of data movement and computation is crucial for performance

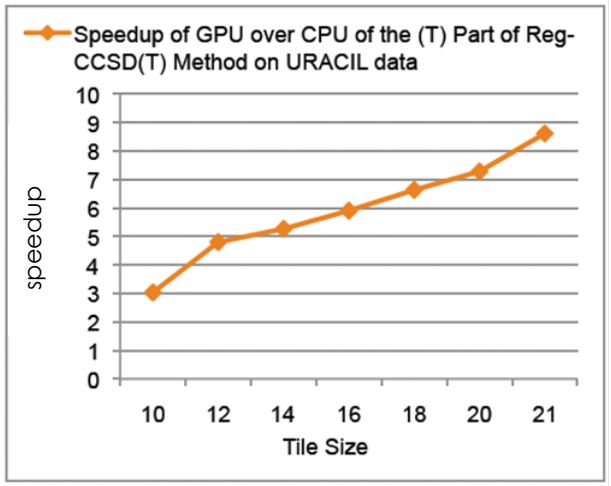

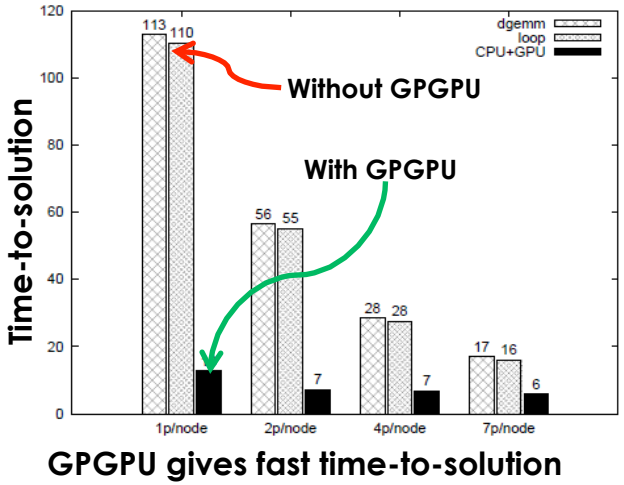

Oreste Villa
Sriram Krishnamoorthy
Karol Kowalski

- Ongoing work on GPGPUs: Cholesky decomposition methods

Ma, Krishnamoorthy, Villa, Kowalski, JCTC ***7****,1316 (2011)*

Handling faults in chemistry codes

EMSL

- NWChem's compute model is task based and data centric:
  - Continue with (N-1) after failure
  - Ensure access to critical data through redundancy
- Reliability Benefits
  - Significantly increased MTBF; handle decent number of failures
  - Recovery proportional to degree of failure
  - Continued execution with existing nodes
- Overheads
  - Replication with synced writes - Critical data structures are significantly smaller

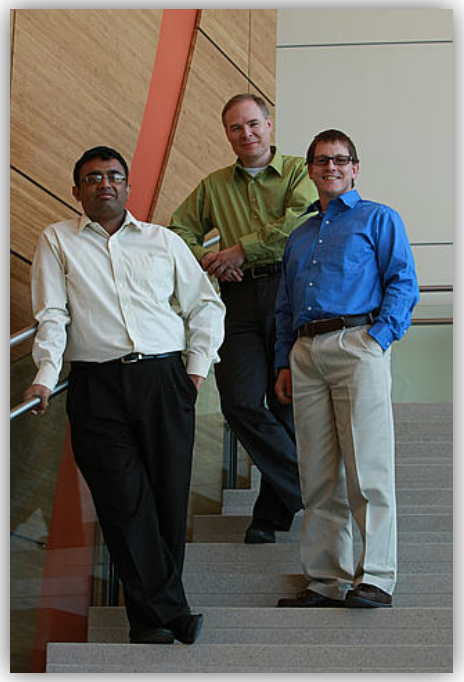

Huub van Dam
Abhinav Vishnu
Bert de Jong

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

U.S. DEPARTMENT OF ENERGY

NWChem CCSD(T) fault tolerant

EMSL

Design of Fault Tolerant Communication Runtime System

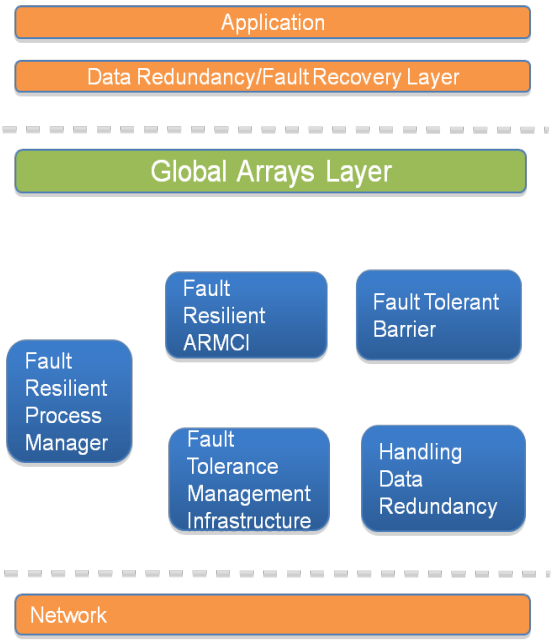

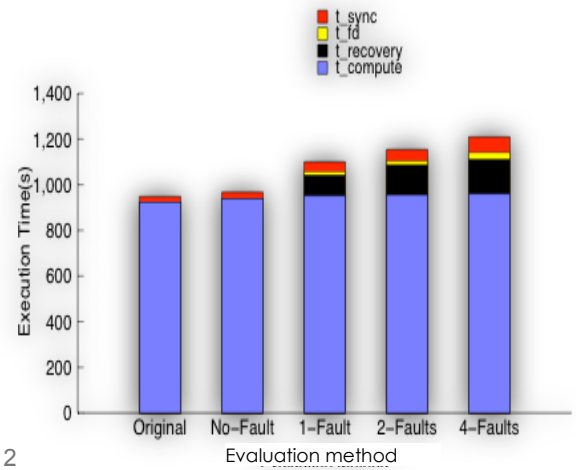

- Handling faults increases scientific productivity and discovery
- Unique Fault Tolerant Management Infrastructure Developed
- Fault tolerance to hardware errors in NWChem demonstrated on various platforms
- We're now addressing soft errors within the Hartree-Fock algorithm
  - Using algorithmic numerical bounds
  - Work also drove us to look at reduced numerical precision is needed

van Dam, Vishnu, de Jong, JCTC ***7****, 66 (2011)*

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

U.S. DEPARTMENT OF ENERGY

Where we want to go with computational chemistry programming models

EMSL

- Crucial to treat data and compute at the same footing
  - Moving from core to node as a basic data and task unit improves memory management
  - Cores within node performing different tasks
- Work load management
  - Next generation computational chemistry models will inherently be hierarchical
    - Dynamic task scheduling including futures
    - Tasks can spawn tasks (could include Python)
  - Avoid sync through data flow and task graphing
  - Move task to data and localization
  - Compute at minimally needed accuracy

EMSL

NWChem development is funded by:

- **US Department of Energy: BER, ASCR, BES**
- **PNNL LDRD**

EMSL: A national scientific user facility integrating experimental and computational resources for discovery and technological innovation