XSCI: eXtreme Scale Computing LDRD Initiative: Driving Science Toward the Exascale

TP Straatsma, Darren Kerbyson

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

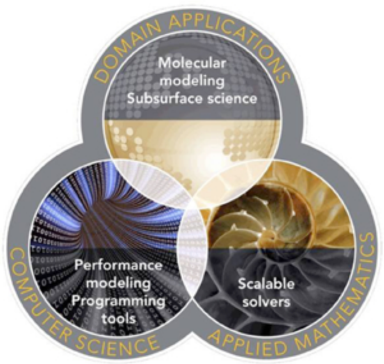

Strategy for PNNL's exascale software development

- Preparation for exascale computing applications (XSCI started in 2009)
- Design of next-generation programming models with fault tolerance and power considerations
- Investment in mission-critical domain science applications for future architectures
- Positioning for collaborative efforts in co-design centers
- Demonstrate new capabilities for PNNL's molecular and subsurface science applications

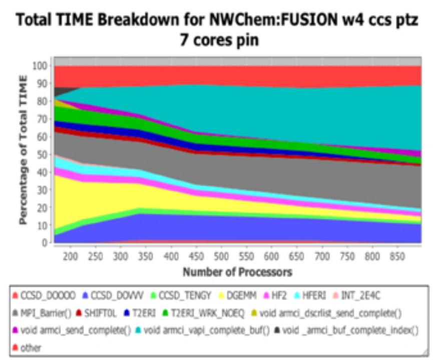

Global Arrays is under active development

- Unique one-sided asynchronous communication programming model
- Designed as part of a Molecular Science Software Suite (MS3) co-design effort
- Available for all current major HPC architectures
- Outreach activities to increase the application base
- Technical meetings: Seattle, May 2010; Anchorage, May 2011
- The GA/ARMCI framework is used in computational chemistry (NWChem, GAMESS UK, MOLPRO, MOLCASS), subsurface transport (STOMP), bioinformatics (ScalaBLAST) and computational fluid dynamics (Thethys) applications

Scalability

Removing the technical barriers for applications to use future systems with large processor counts

- CCSD(T) triples code scales to 225,000 cores and was a Gordon Bell finalist in 2008
- EOM-CCSD(T) triples codes scales to 120,000 cores
- e-STOMP fluid flow expected scaling to 200,000 cores by 2011 (now: 65,536), selected as ASCR Joule milestone
- ParaReal pwDFT scale to 25,000 cores per instance with 20-30 concurrent instances
- Ongoing efforts in multi-scale optimization, MD and MR-CCSD(T)

Scalability of the EOM-CCSD(T) triples

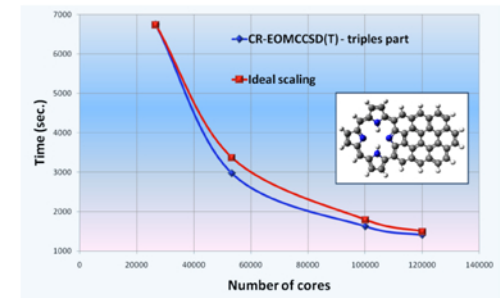

Architectures

Adapting to large processor count homogeneous and heterogeneous architectures

- Optimization strategies for multi- and many-core processors and accelerators
- Current test beds: Fermi GPUs, 48-core single nodes
- Leveraging work on IBM Power-7 (Blue Waters) and Cray Gemini networks
- Development of GPU versions of computational intensive kernels in MD, CCSD(T), e-STOMP Chemistry Modes

ParaReal Parallel in Time

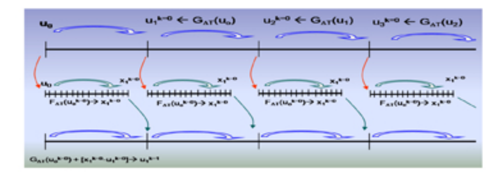

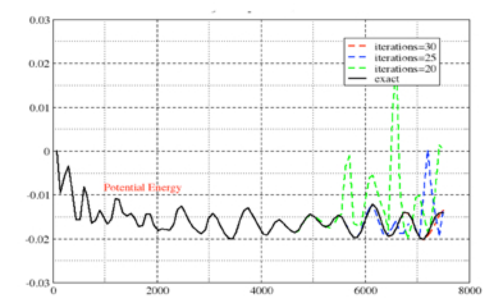

Fault Tolerance

Recovering from frequent hardware faults on many component systems

- Adding fault tolerance to GA involving computer science and application considerations
- Based on a combination of data replication and process management
- Demonstrated fault tolerance for CCSD(T) triples code

Productivity

Enabling efficient and effective development of highly scalable applications

- Providing performance diagnostics and analysis to the GA framework
- Scalable trace compression of GA/ARMCI communication calls
- Leveraging existing analysis and visualization capabilities in Tau, Scalasca

Scalability of the PW-DFT

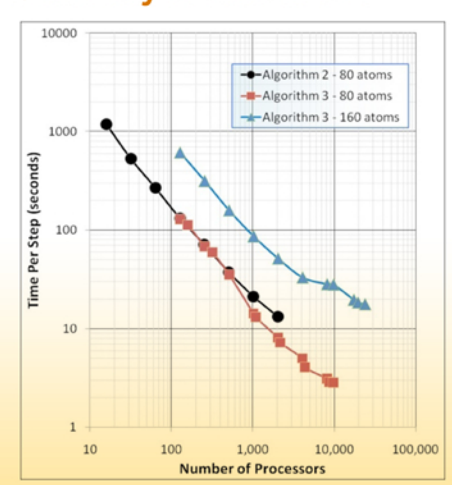

Power

Developing methods to deal with power as the measure of efficiency on future architectures

- Exploring energy-efficient one-sided communication libraries
- Exploring energy-efficient architectures
- Leveraging performance modeling technologies