SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

SciDAC-3 Institute for Sustained Performance, Energy, and Resilience

Bob Lucas
University of Southern California
July 24, 2013

Support for this work was provided through the Scientific Discovery through Advanced Computing (SciDAC) program funded by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research

Outline

SUPER Team

Research Directions

Application Engagement

New Directions

SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

SUPER Team

ANL	Paul Hovland Boyana Norris Stefan Wild
LBNL	Lenny Oliker Sam Williams
LLNL	Kathryn Mohror B. de Supinski Daniel Quinlan
Oregon	Kevin Huck Allen Malony Sameer Shende
ORNL	Philip Roth Patrick Worley
UCSD	Laura Carrington Ananta Tiwari
UMD	J. Hollingsworth
UNC	Rob Fowler Allan Porterfield
USC	Jacque Chame Pedro Diniz Bob Lucas (PI)
UTEP	Shirley Moore
UTK	George Bosilca Dan Terpstra
Utah	Mary Hall G. Gopalakrishnan

SUPER

SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

Broadly Based Effort

All PIs have independent research projects

- SUPER money alone isn't enough to support any of its investigators
- SUPER leverages other work and funding

SUPER contribution is integration, results beyond any one group

- Follows SciDAC-2 PERI model (tiger teams and autotuning)
- Collaboration extends to others having similar research goals
  - John Mellor-Crummey of Rice University hosted the last SUPER all-hands meeting

SUPER

Outline

SUPER Team

Research Directions

Application Engagement

New Directions

SUPER Research Objectives

Performance Optimization of DOE Applications

Performance tuning
Energy minimization
Resilient computing
Optimization of the above
Automation of the above

Collaboration with SciDAC-3 Institutes
Engagement with SciDAC-3 Applications

SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

SUPER Performance Engineering

- **Led by Mary Hall, University of Utah**
- **Automatic tuning for performance portability**
- **Measurement and monitoring**
  - **Adopting University of Oregon's TAU system**
  - **Building on UTK's PAPI measurement library**
  - **Also collaborating with Rice and its HPCToolkit**
- **Performance database**
  - **Extending TAUdb to enable online collection and analysis**
- **Performance modeling**
  - **PBound and Roofline models to bound performance expectations**
  - **PSINS to model communication**

SUPER

Performance Tuning of NWChem Texas Integrals

Bert de Jong, Lenny Oliker, Hongzhang Shan

Objectives

Improve the performance of TEXAS two-electron integral package in NWChem

Impact

Faster time-to-solution enabling larger and more accurate excited-state simulations with NWChem

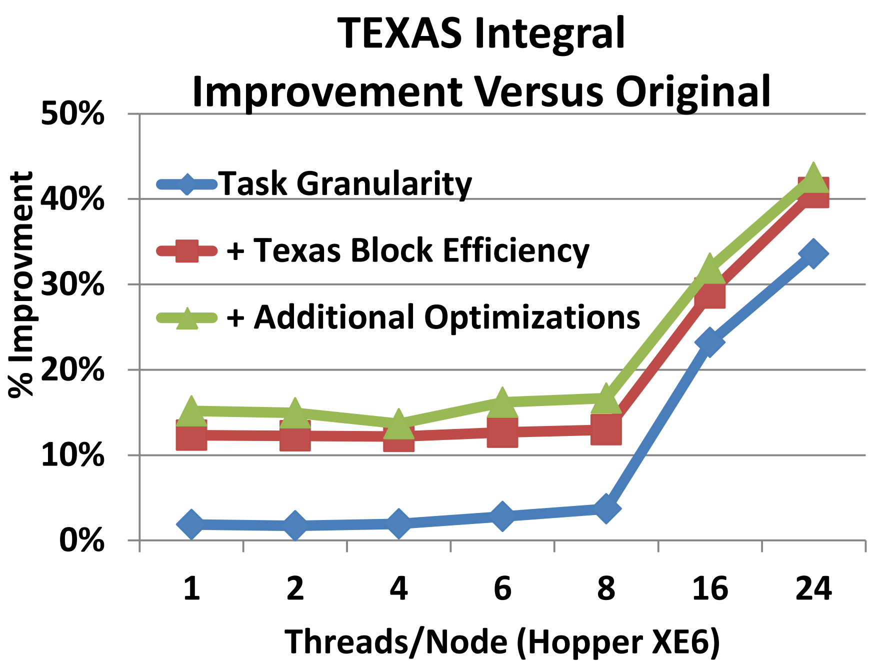

Progress & Accomplishments

Identified a load imbalance problem in dynamic task assignment

Up to 34% speedup on 24 cores (Task Granularity)

Improved the efficiency of the blocking structure

10% speedup for all configurations (+ Texas Block Efficiency)

Other minor improvements include using better sorting algorithms, removing redundant computations, unrolling loops, etc

Total overall speedup up to 15%-44% (+ Additional Optimizations)

Really automatic autotuning of PETSc with Orio

Enable autotuning by configuring PETSc with Orio as a compiler wrapper, e.g.

```
./configure --with-cc="orcc -r mpicc" --with-cxx="orcc -r mpicxx" --with-fc=0 --download-f2cblaslapack
```

During the installation, pragma-annotated code is empirically autotuned, for example in the function computing $w = y + \alpha x$ (w, y, x: vectors, α: scalar)

```
PetscErrorCode VecWAXPY_Seq(Vec win, PetscScalar alpha,Vec xin,Vec yin) {
  ...
#pragma Orio PerfTuning(import spec tune/waxpy;)
  ...
#pragma Orio Loop(transform Unroll(ufactor=UF, parallelize=PAR))
    for (i=0; i<=n-1; i++) ww[i] = yy[i] + xx[i];
  ...
#pragma Orio Loop(transform Unroll(ufactor=UF, parallelize=PAR))
    for (i=0; i<=n-1; i++) ww[i] = yy[i] + oalpha * xx[i];
  ...
}
```

Autotuning of PETSc library code with Orio

Example: Structured-grid PDE application using PETSc for GPUs (solid fuel ignition problem, SNES ex14). Comparison between library-based (MKL on CPU and Cusp on GPU) implementation and the Orio-tuned implementation (GPU).

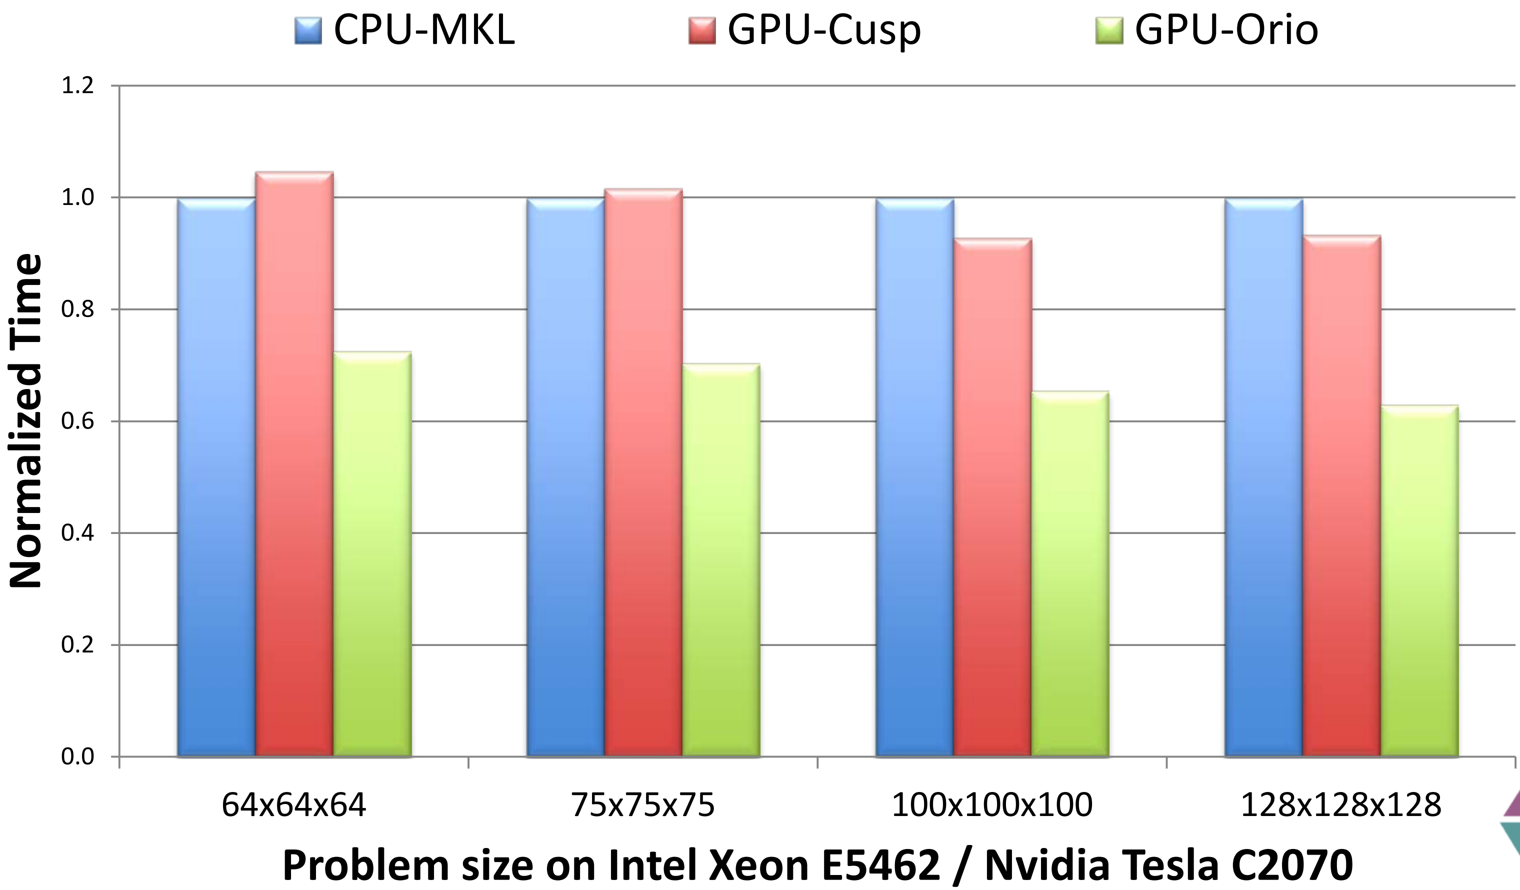

SUPER

SUPER Energy Minimization

- Led by Laura Carrington, University of California at San Diego
- Develop new energy aware APIs for users
  - I know the processor on the critical path in my multifrontal code
- Obtain more precise data regarding energy consumption
  - Extend PAPI to sample hardware power monitors
  - Build new generation of PowerMon devices
  - Extend performance models
- Transform codes to minimize energy consumption
- Inform systems to allow them to exploit DVFS

Towards Application-aware Model Driven Energy Optimization

Laura Carrington, Ananta Tiwari

SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

Green Queue: Application-aware Energy Optimization

- Develop a fully automated framework that utilizes fine-grained application characterizations and power and performance models to devise and deploy energy efficient policies

Goals of the SUPER Energy effort

- Understand how computation and communication patterns affect the overall energy requirements of HPC applications
- Leverage this understanding to design software- and hardware-aware optimization techniques that reduce the DOE's HPC energy footprint

Case Study: Modeling two computational phases of **GTC** (**1024** cores, Gordon Supercomputer)

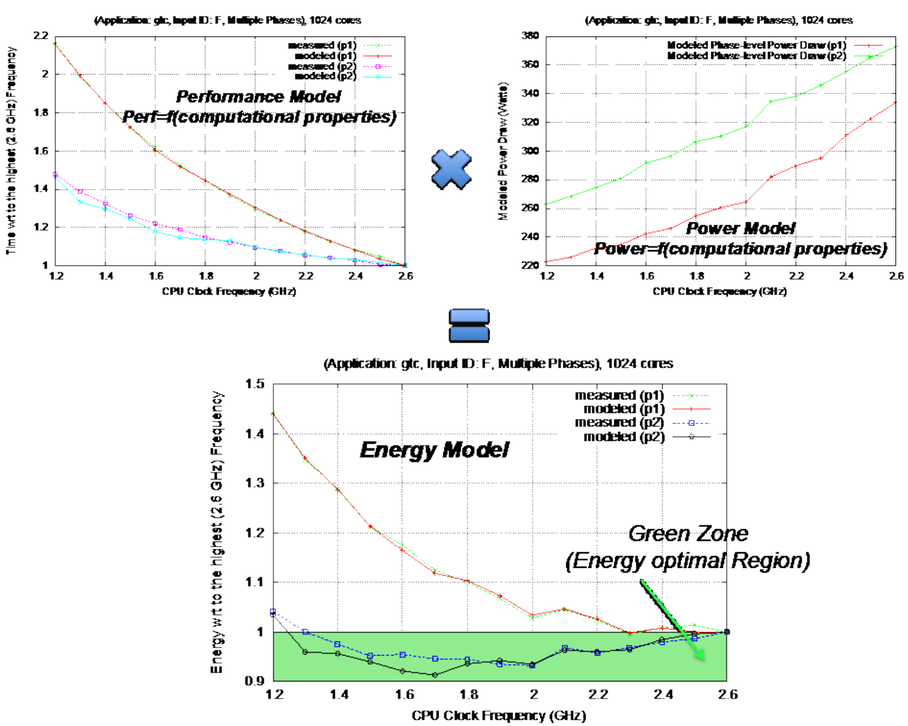

Progress & Accomplishments

- An application's computational behavior is captured by a series of characterization vectors; these vectors are inputs to power and performance models
- Developed flexible and highly customizable methodology to identify computational phases that can be targeted with appropriate energy optimization strategies
- Case study with one of DOE's high value HPC application, GTC. The green zone (in the figure) marks those frequency selections that provide energy savings and illustrate the accuracy of the models and how the models enable fine-grained customized DVFS settings for an application's individual computational phases

SUPER

INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

SUPER Resilient Computing

- Led by Bronis de Supinski, Livermore National Laboratory
- Investigate directive-based API for users
  - Enable users to express their knowledge w/r resilience
  - Not all faults are fatal errors
  - Those that can't be tolerated can often be ameliorated
- Automating vulnerability assessment
  - Modeled on success of PERI autotuning
  - Conduct fault injection experiments
  - Determine which code regions or data structures fail catastrophically
  - Determine what transformations enable them to survive
- In either event, extend ROSE compiler to implement transformations

SUPER

Programming Model Extensions for Resilience

(Saurabh Hukerikar, Robert F. Lucas and Pedro C. Diniz)

SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

Resiliency-oriented Programming Model

- Enable computational scientists to express knowledge about the resilience or fault tolerance of their codes.
- Use of code annotations such as directives and aspects.
- Evolutionary approach that extends familiar programming languages.

Goals of the SUPER resilience effort

- Identify vulnerable data and code regions.
- Design and implement simple and effective resilience strategies to improve vulnerability of sensitive pieces of code.
- Long term: develop a general methodology to automatically improve the reliability of generic HPC codes.

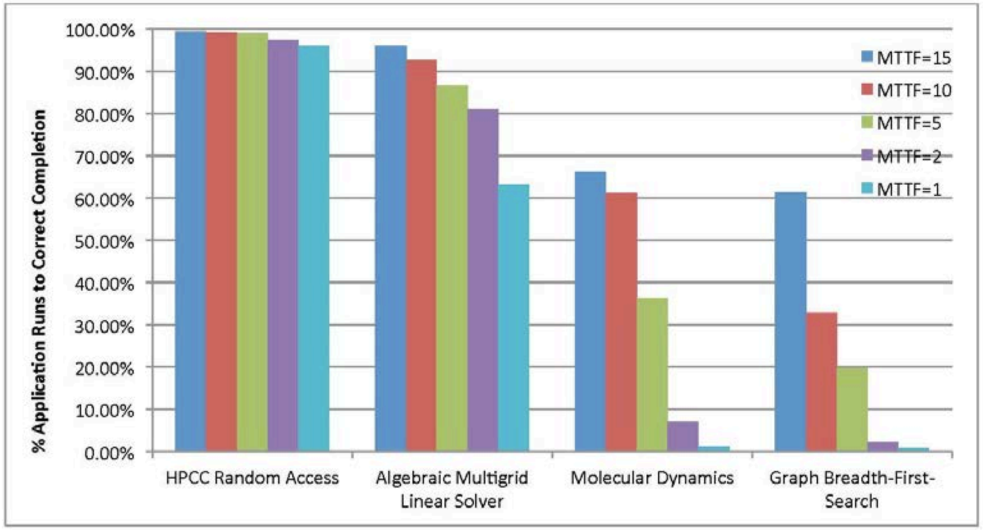

Probability of four computational kernels running successfully versus fault injection rate (mean time to fault in minute).

Progress & Accomplishments

- **Invented annotations for specifying regions of memory where uncorrected errors are tolerated.**
- **Developed of a methodology to automatically inject faults to assess the vulnerability of annotated codes to soft errors.**
- **Demonstrated that some computational kernels from scientific applications and benchmarks can tolerate uncorrected faults and run successfully to completion.**

SUPER

SUPER Optimization

Led by Paul Hovland, Argonne National Laboratory

Performance, energy, and resilience are implicitly related and require *simultaneous* optimization

- **E.g., Processor pairing covers soft errors, but halves throughput**

Results in a stochastic, mixed integer, nonlinear, multi-objective, optimization problem

Only sample small portion of search space:

- **Requires efficient derivative-free numerical optimization algorithms**
- **Need to adapt algorithms from continuous to discrete autotuning domain**

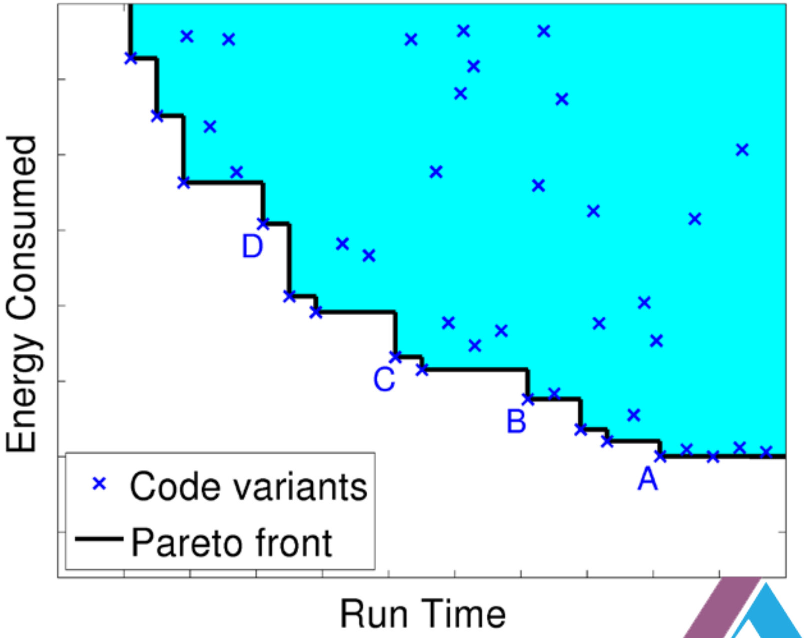

SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

Analyzing Multiple Objectives

(Prasanna Balaprakash, Paul Hovland, Ananta Tiwari, Stefan Wild)

Performance Tuning & Multiobjective Optimization

- Identify tradeoffs between potentially competing objectives such as execution time, system resiliency, power draw, and energy consumption
- Develop new search methods for use in autotuning tools

Motivation

- Autotuning essential as scientists move between generations of supercomputers
- Enable solution of decision problems facing systems administrators, hardware designers, & numerical library developers

Power-time tradeoffs for "jacobi" kernel on Intel Xeon E5530 CPU clocked at different frequencies

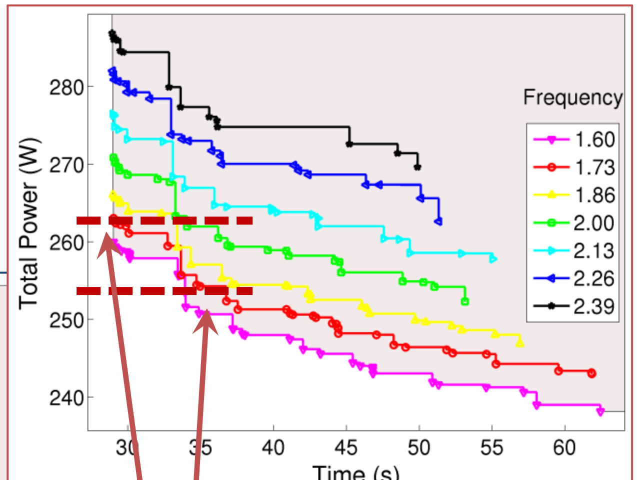

At 1.6Ghz: Increasing power cap from 250W to 260W yields 30% speedup

Energy-time tradeoffs for the mini-app "miniFE" on a 16,384-core rack of Blue Gene/Q

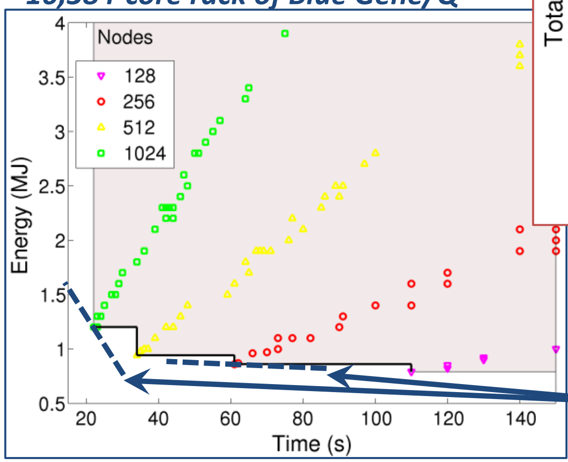

Different energy & time prices result in optimal codes with different parallelization levels

Progress & Accomplishments

- Mathematical formulation of multiobjective code optimization problem
- Mathematical analysis of necessary conditions for tradeoffs to exist
- First-of-its-kind empirical study of energy, power, and time on diverse platforms (Xeon Phi, Xeon E5530, BG/Q)
- Showed that tradeoffs exist for many codes, many tuning spaces, many platforms

INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

Outline

SUPER Team

Research Directions

Application Engagement

New Directions

SUPER

SciDAC-3 Application Partnerships

- **Collaboration with SciDAC Application Partnerships is expected**
  - **Yet SUPER funding is spread very thin**
- **SUPER investigators included in 12 Application Partnerships**
  - **Our time costs money, like everybody else**
- **Common features of our successful collaborations with SciDAC applications**
  - **A motivated member of the application team**
  - **A motivated member of the SUPER team**
  - **A specific need to improve performance of some aspect of execution, or a port to a new architecture**
  - **A computational kernel that represents the performance issue, including a representative input and output, validation and instructions for building and running it.**

SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

Participation in SciDAC-3 Application Partnerships

BER	**Applying Computationally Efficient Schemes for BioGeochemical Cycles (ORNL)**
BER	**MultiScale Methods for Accurate, Efficient, and Scale-Aware Models of the Earth Sys. (LBNL)**
BER	**Predicting Ice Sheet and Climate Evolution at Extreme Scales (LBNL & ORNL)**
BES	**Developing Advanced Methods for Excited State Chemistry in the NWChem S/W Suite (LBNL)**
BES	**Optimizing Superconductor Transport Properties through Large-scale Simulation (ANL)**
BES	**Simulating the Generation, Evolution and Fate of Electronic Excitations in Molecular and Nanoscale Materials with First Principles Methods (LBNL)**
FES	**Partnership for Edge Plasma Simulation (ORNL)**
FES	**Plasma Surface Interactions (ORNL)**
HEP	**Community Petascale Project for Accelerator Science and Simulation (ANL)**
NP	**A MultiScale Approach to Nuclear Structure and Reactions (LBNL)**
NP	**Computing Properties of Hadrons, Nuclei and Nuclear Matter from QCD (UNC)**
NP	**Nuclear Computational Low Energy Initiative (ANL)**

SUPER

INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

A Multi-Scale Approach to Nuclear Structure and Reactions: Forming the Computational Bridge between Lattice QCD and Nonrelativistic Many-Body Theory (CalLat)

PI: Wick Haxton, University of California-Berkeley

SUPER Component:

LBNL (Performance): Williams (5%) and 45% postodc

LLNL (Resilience): de Supinski (2.5%) and 37% postdoc

SUPER contingent at LBNL has spent significant time reviewing previous work (papers and online lectures) and having discussions with the physicists.

A new code focused on Wick Contractions is being developed…

- **Chroma-based**
- **Significant effort was required to get it up and running at LBL.**
- **Performance profiled with HPCToolkit (scales reasonably well on CPU-based machines).**
- **Future work will focus on analysis and optimization (complicated by the heavy C++ templating). Physics contingent must address a few key numerical/physics issues first.**
- **Ultimately, it will target CPU/BGQ, as well as accelerator (MIC/GPU) based systems**

LBL is also examining adding HDF5 support to QDP++

Finally, we have had discussions on optimizing the dslash operator for BGQ.

SUPER

SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

Computing Properties of Hadrons, Nuclei and Nuclear Matter from Quantum Chromodynamics

PI: Frithjof Karsch, Brookhaven National Laboratory

SUPER: Fowler (9%), Porterfield (15%), grad student (100%). (Year 2)

Goal: Use compiler technology and auto-tuning to help close the "Ninja programmer performance gap" between simple, high-level code and hand-optimized LQCD routines, especially for new computer architectures.

Problem statement: QDP++ is an embedded distributed-memory domain-specific language for LQCD computation implemented using C++ template meta-programming. QDP-JIT extends QDP++ to generate locally optimized Nvidia PTX code.

- **Code for each assignment expression (statement) is generated independently without global analysis.**
- **Memory management through run time software cache, no compiler analysis.**
  - **Excessive data movement, though most of this is to/from fast memory.**
  - **On Nvidia GPUs, memory bandwidth is the limiting resource.**
- **Per expression data movement (MPI) without global planning/coalescing/scheduling.**

Approach: With JLAB, add compiler analysis and code transformations to QDP-JIT.

- **Loop (expression) fusion to reduce memory traffic and generate tighter inner loops.**
- **Static analysis to improve memory management and reduce CPU↔GPU traffic.**
- **Optimize cluster-wide messaging operations across large units of code.**
- **Generate auto-tuning hooks to address hard problems, *e.g.*, performance portability.**
- **(LLVM chosen to leverage Intel and Nvidia infrastructures for accelerator devices.)**

SUPER

Nuclear Computational Low Energy Initiative

NUCLEI Nuclear Computational Low-Energy Initiative

SUPER INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

PI: Joseph Carlson, Los Alamos National Laboratory

SUPER Liason: Wild (15%)

Recent Progress: "Streamlining the nuclear force"

- **A decade of work has focused on hand-tuned potentials at next-to-next-to-next-to-leading order**
- **Apply state-of-the-art optimization methods for wide range of calibration problems in nuclear physics for scalable codes spanning *ab initio* to density functional theory approaches**
- **The derivative-free solver POUNDERS in TAO was used to systematically optimize potentials based solely on two-nucleon forces**

Correct placement of neutron dripline at ^{24}O

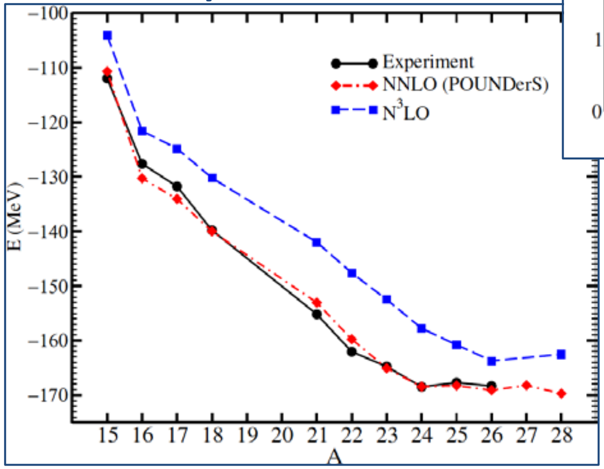

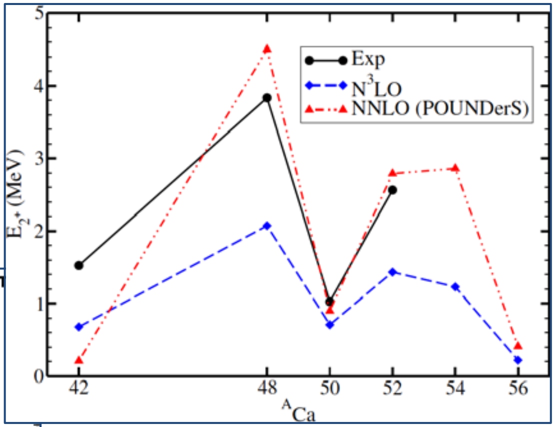

Magicity of ^{48}Ca and sub-shell closures in $^{52,54}Ca$ reproduced

TAO

- Physical Review Letters, 2013: "An *optimized chiral nucleon-nucleon interaction at next-to-next-to-leading order"*

- **Less pain and more gain through automated optimization: One can capture key aspects of nuclear structure from nucleon-nucleon interactions alone**

SUPER

Center for Edge Physics Simulation (EPSi)

PI: Choong-Seock Chang, Princeton Plasma Physics Laboratory

SUPER liaison: Patrick H. Worley, Oak Ridge National Laboratory

EPSi Science Team Lead for Performance: Worley (20%)

Recent Progress: Computer Performance Evaluation Study

GPU port and optimization of computational kernel, optimization of OpenMP parallelism to allow effective use of 16 threads per task, and optimization of MPI communication led to 4X performance improvement and good weak scaling out to full system size.

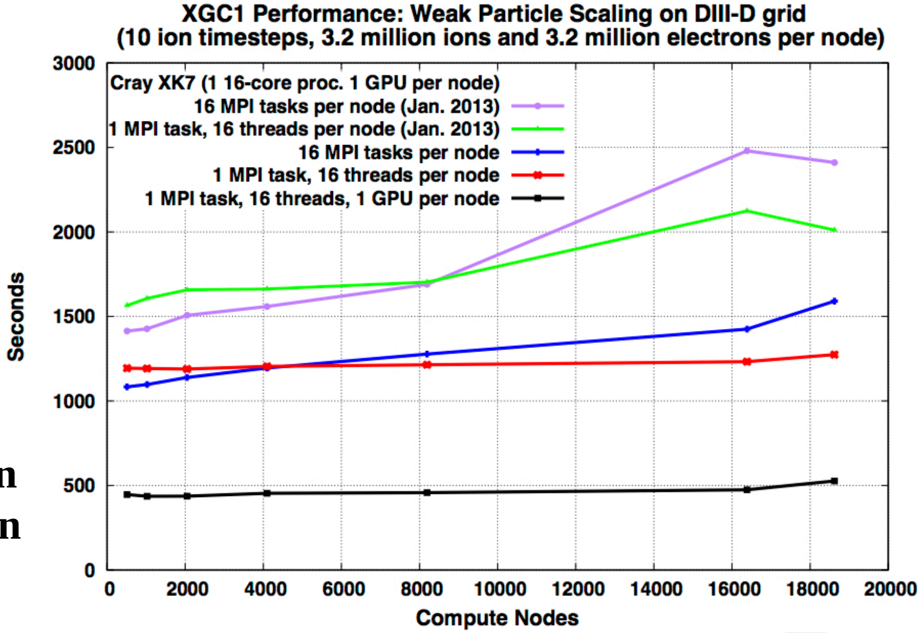

SUPER contributions included performance measurement and tracking, experimental design and performance diagnosis, and collaboration in parallel algorithm design and optimization.

Work far from over – new science capabilities will change performance characteristics dramatically. Performance tracking and diagnosis will continue to be a critical aspect of EPSi performance engineering and optimization, and new SUPER tools and techniques should accelerate the process.

Multiscale Methods for Accurate, Efficient, and Scale-Aware Models of the Earth System

INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

PI: William Collins, Lawrence Berkeley National Laboratory

SUPER: Oliker (10%), Williams (10%)

Two principal codes: MPAS-O (ocean model) and CAM-SE (atmosphere model)

LBNL has partnered with Doug Jacobsen (LANL) and Matt Cordery (NERSC) to study and optimize MPAS-O

Opportunity for FASTMath collaboration: Trilinos/implicit solvers (maybe Mesquite?)

Work has focused on restructuring the OpenMP implementation to be amenable to threading

- **Update: performance study with TAU shows MPI block decomposition + new OpenMP element decomposition approach performs faster than MPI alone for both RK4 and split integration solvers**
- **Identified opportunity for overlapping computation and communication**
- **Evaluating experimental weighted decomposition strategy with TAU**

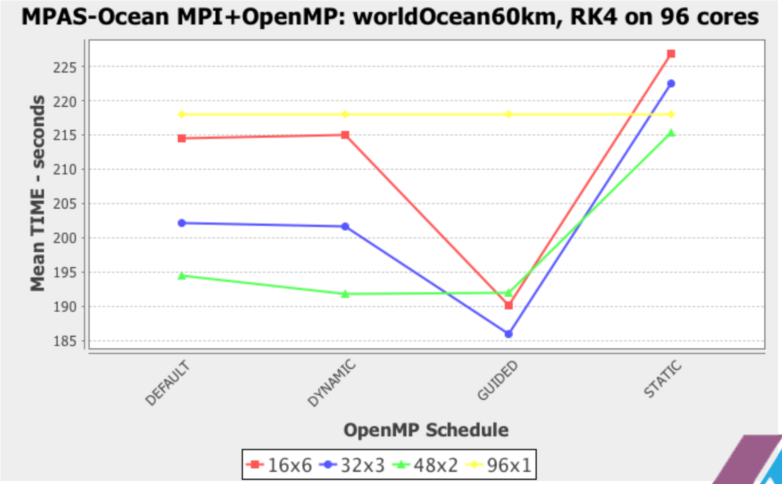

SUPER

INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

Outline

SUPER Team

Research Directions

Application Engagement

New Directions

SUPER

Collaboration with NNSA

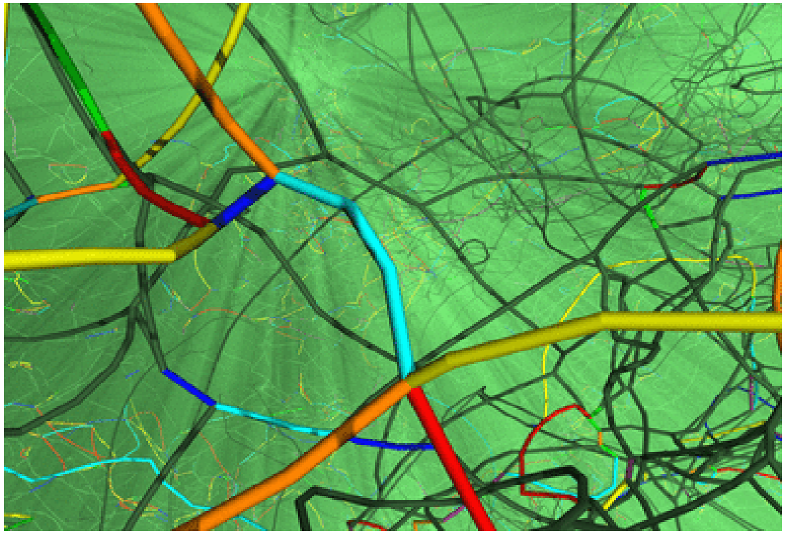

ParaDiS models dislocation lines in a physical volume

- **The volume is divided into several computational domains**

The dislocations interact and move in response to physical forces

- **External stress and inter-dislocation interactions determine the forces**

Evaluating accuracy and performance tradeoffs of the physical model

Reducing SMN execution frequency:

- **Improves performance**
- **Reduces the quality of the simulation**

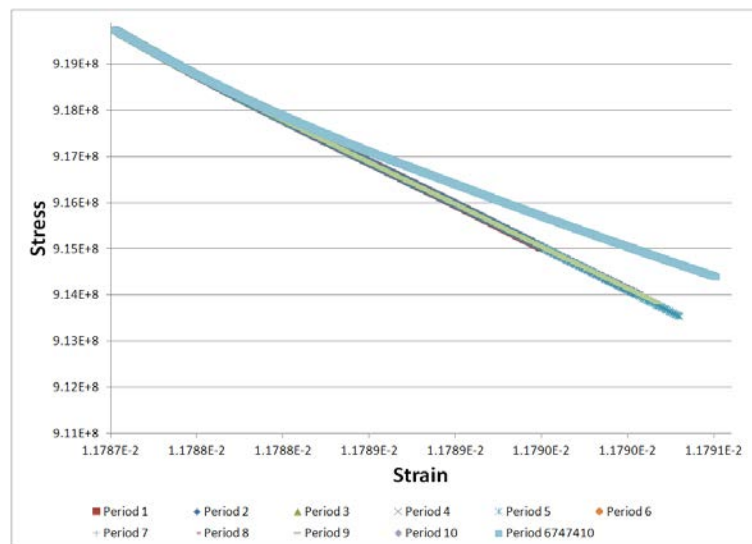

Geant4

Exploring the Transformation of Geant4 for the Future

- Geant4 is extremely important to the design and execution of HEP experiments
- How can it adapt to future multicore devices?

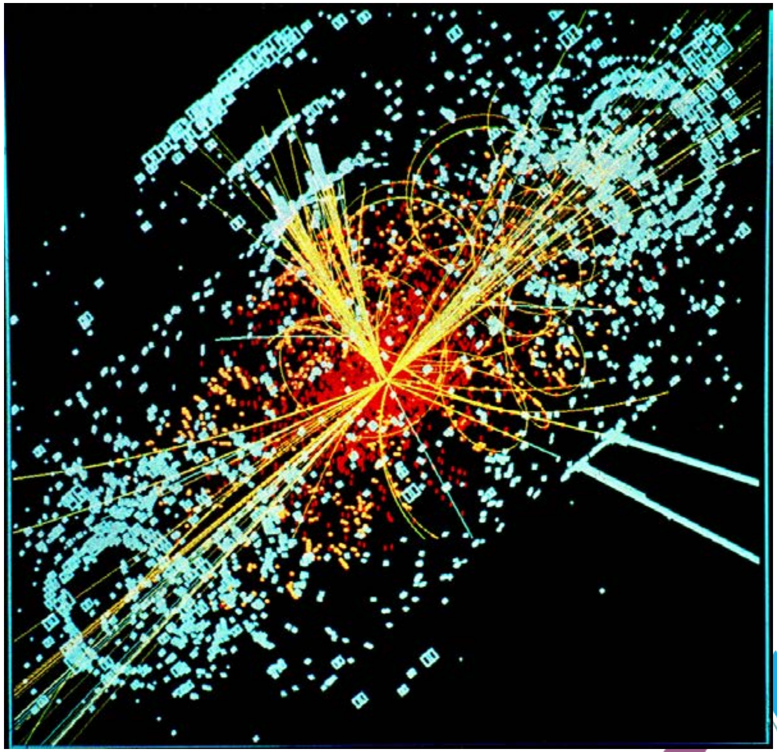

Partnership with HEP

- Lali Chatterjee is the HEP PM
- Ceren Suset is the ASCR PM

Modification to SUPER at ANL

- UNC and USC also contributing for ACSR
- FNAL and SLAC contributing for HEP

Please see the Geant4 poster

SUPER
INSTITUTE FOR SUSTAINED PERFORMANCE, ENERGY, AND RESILIENCE

Roofline Toolkit

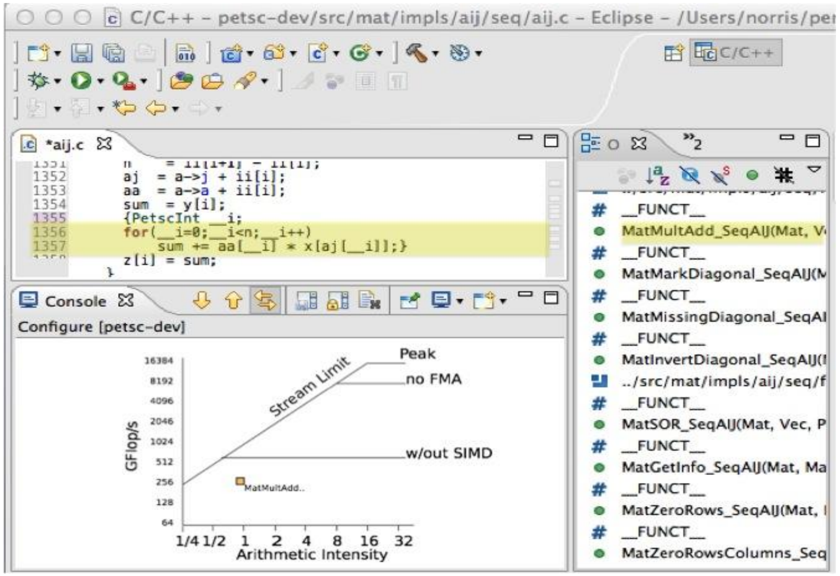

Enable construction, use, and interpretation of performance models by a broad range of DOE SciDAC scientists and beyond (figure shows tool sketch)

Modifications to ANL and LBNL & in partnership with FASTMath

- **Oregon also participating**

Roofline Toolkit is targeting downloadable collection of programs installable across a variety of systems, consisting of:

- **Hardware characterization, via portable, instrumented microbenchmarks and tools for many and multicore systems**
- **Software characterization via static analysis/modeling of source code, and performance counter execution instrumentation**
- **Data manipulation and visualization interface via TAU TAUdb database query system augmented for roofline analysis**

SUPER

SUPER Summary

SciDAC-3 Research

- Performance Engineering of Scientific Software
- Automatic performance tuning
  - New focus on portability
- Addressing the “known unknowns”
  - Energy minimization
  - Resilient computing
- Optimization of the above

Near-term impact on DOE computational science applications

- Application engagement coordinated with ALCF, NLCF, and NERSC
- Tool integration, making research artifacts more approachable
- Participation in SciDAC-3 Application Partnerships
- Outreach and tutorials