Advancing and Enabling Scientific Discovery Through Advanced Algorithms and Computing

Esmond G. Ng
Lawrence Berkeley National Laboratory

The Office of Science in the U.S. Department of Energy

- The Office of Science is one of the top offices in the U.S. Department of Energy (DOE). It is the major office that provides funding for basic research in energy and environment.

- **Computational science** is very important in the Office of Science.
- **SciDAC** (**Scientific Discovery Through Advanced Computing**) is a program that funds research and development in computational science.
  - It is quite unique and is widely recognized as a very successful program.
- This talk is about
  - the SciDAC program,
  - what makes it successful, and
  - a brief description of some of the SciDAC projects that I have been involved in.

Six Programs in the Office of Science

- **Advanced Scientific Computing Research** (ASCR)
  - Applied mathematics, computer science, computational science
- **Basic Energy Science**s (BES)
  - Chemical sciences, materials science, subsurface flow
- **Biological and Environmental Research** (BER)
  - Genomics, climate
- **Fusion Energy Sciences** (FES)
  - Fusion research
- **High Energy Physics** (HEP)
  - Particle physics, cosmology, accelerator physics
- **Nuclear Physics** (NP)
  - Low energy nuclear physics

DOE Office of Science – Computing Facilities

- Computation has always been an integral part of carrying out scientific research.
- Understanding the importance of computing in scientific research and recognizing the scale of computing needed for some of the scientific problems, DOE Office of Science, through ASCR, operates 3 high-performance computing facilities.
  - National Energy Research Scientific Computing Center (NERSC)
    + Cori: 29 petaflops
  - Argonne Leadership Computing Facility (ALCF)
    + Theta: 10 petaflops
  - Oak Ridge Leadership Computing Facility (OLCF)
    + Titan: 27 petaflops

Scientific Discovery Through Advanced Computing

- **Scientific Discovery Through Advanced Computing** (SciDAC)
  - A major computational science program in the Office of Science at DOE.
  - 5-year programs
    + SciDAC-1: 2001-2006
    + SciDAC-2: 2006-2011
    + SciDAC-3: 2012-2017
      – Some projects are 2011-2016
    + SciDAC-4: 2017-2022
      – Some projects are ending in 2020 and some in 2021.

Scientific Discovery Through Advanced Computing

- SciDAC is a very unique program …
  - *It is about enabling and accelerating* ***discoveries in scientific domains*** *that are of interest to DOE.*
  - *There is a* ***strong emphasis on partnership*** *among domain scientists, computational mathematicians, and computer scientists.*
  - *Also strong focus on scientific problems that require* ***high performance computing*** *(NERSC, ALCF, OLCF).*
- Recent SciDAC programs have 2 components.
  - Institutes
  - Partnerships.

SciDAC Partnerships

- Almost all partnership projects are large, multi-institutional projects in a number of DOE science domains.

- Each partnership project is funded **jointly** by the Advanced Scientific Computing Research (ASCR) Program and one of five science programs in the Office of Science:
  - Basic Energy Sciences (BES)
  - Biological and Environmental Research (BER)
  - Fusion Energy Sciences (FES)
  - High Energy Physics (HEP)
  - Nuclear Physics (NP)
  - *Nuclear Energy (NE)*

SciDAC Institutes

- SciDAC also funds large, multi-institutional Institutes in computational mathematics and computer science entirely through ASCR.
  - The main goal is the development of new high-performance scalable algorithms/tools for core components of scientific simulation, and the distribution of those algorithms/tools through portable high-performance libraries.
  - ***Strong sense of deployment, but also providers of math/computer science expertise.***
- Currently 2 Institutes under SciDAC-4:
  - **FASTMath** – Frameworks, Algorithms, and Scalable Technologies for Mathematics
  - **RAPIDS** – Resource and Application Productivity through Computation, Information, and Data Science

The FASTMath project brings leading edge computational mathematics technologies to the SciDAC Program

Develop advanced numerical techniques for DOE applications

- Eight focused topical areas based on application needs
- High level synergistic techniques

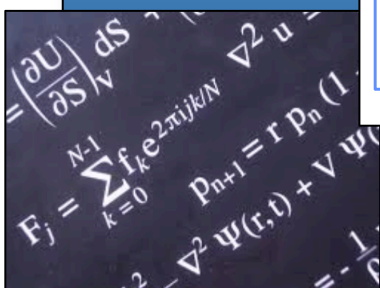

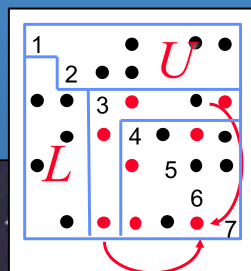

Deploy high-performance software on DOE supercomputers

- Algorithmic and implementation scalability
- Performance Portability
- Interoperability of libraries

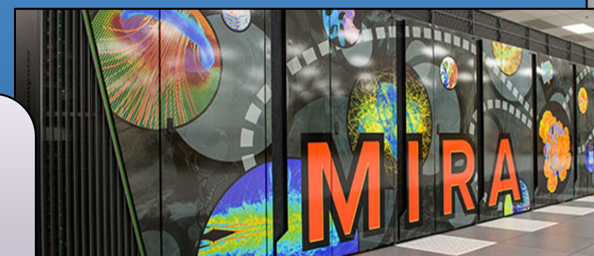

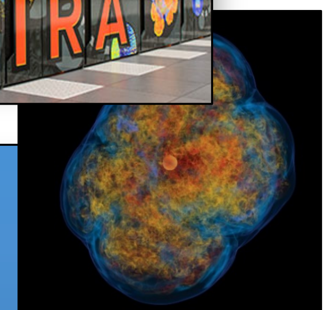

FASTMath Objective: Reduce the barriers facing application scientists

Demonstrate basic research technologies from applied mathematics

- Build from existing connections with basic research
- Focus on research results that are most likely to meet application needs

Engage and support of the computational science community

- Publications and presentations in highly visible venues
- Team tutorials
- Workforce pipeline and training
- Web presence

FASTMath is focused on eight core technology areas

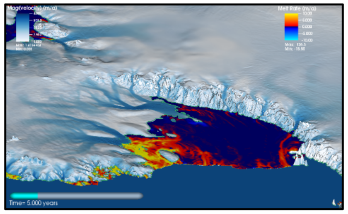

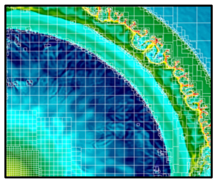

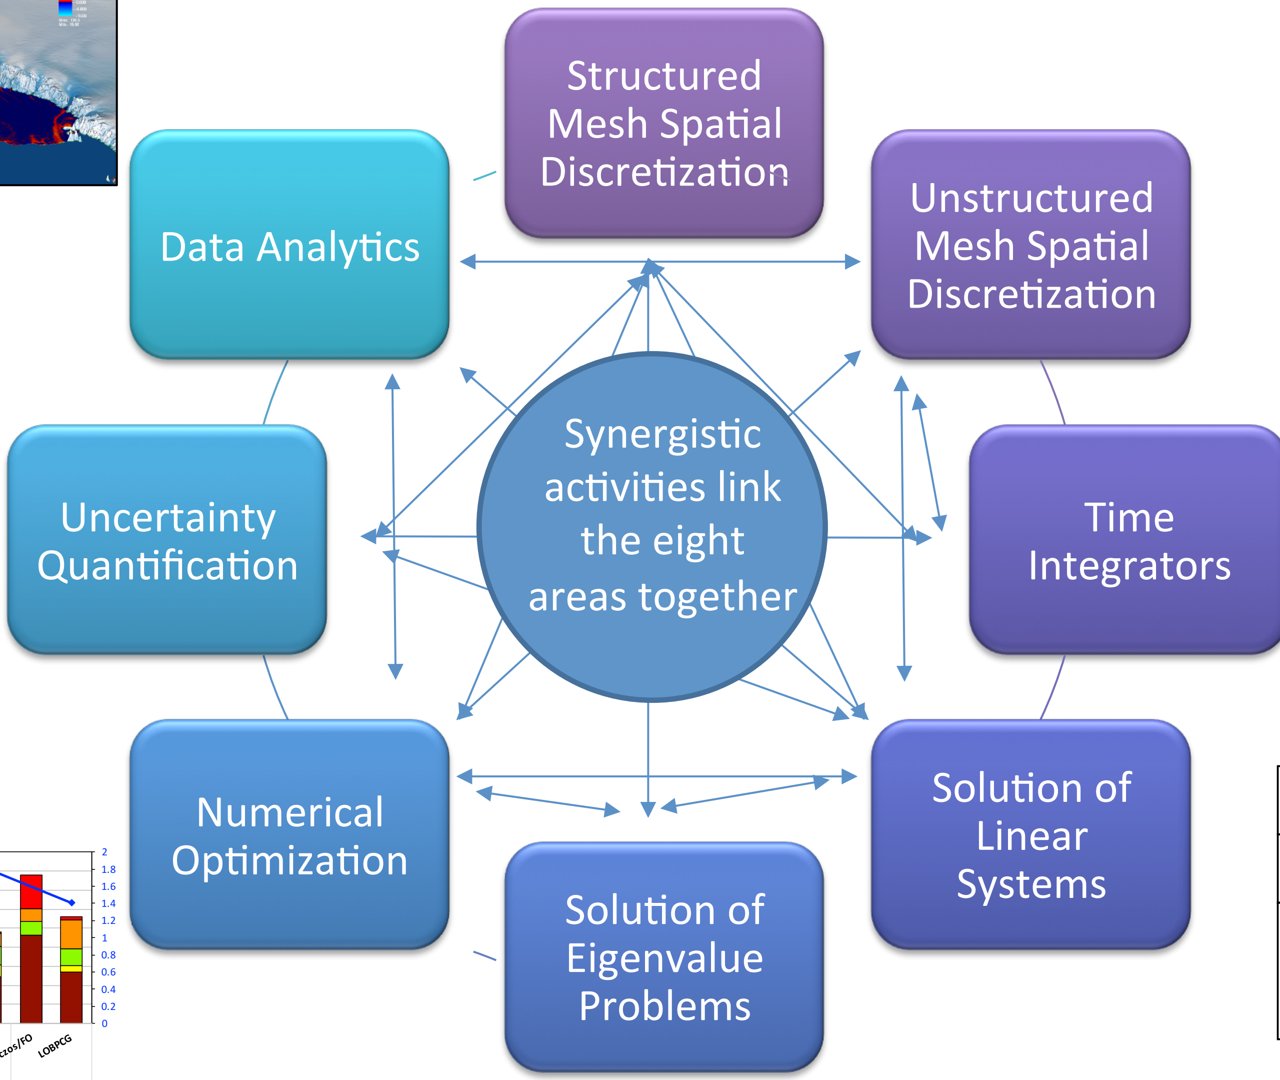

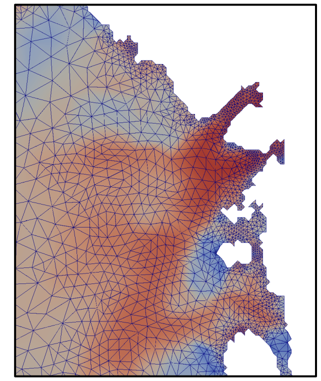

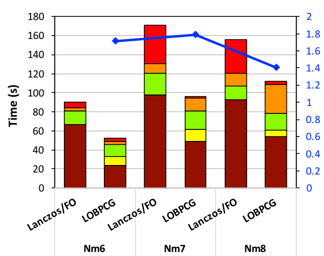

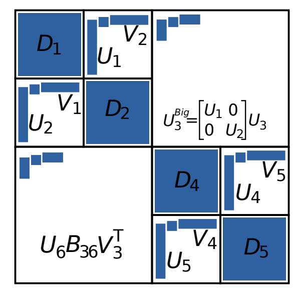

U.S. DEPARTMENT OF ENERGY | Office of Science

FASTMATH

FASTMath is actively engaged with 19 SciDAC-4 application partnerships

BER (5)

- Structured grid AMR
- Unstructured grid AMR
- Time integration
- Linear/Nonlinear solvers, Preconditioners
- Optimization
- Verification / UQ

FES (5)

- Unstructured meshes
- Discretization technologies
- Iterative Linear solvers
- UQ

HEP (3)

- Direct solvers
- Structured Grid AMR
- Optimization
- Sensitivity Analysis
- Inference and machine learning

NP (2)

- Structured grid AMR
- Eigenvalue problems
- Inference and Machine Learning

BES (2)

- Nonlinear and tensor eigenvalue problems
- Linear solvers and Preconditioners

NE (1)

- UQ

FASTMath is actively engaged with 19 SciDAC-4 application partnerships (BER (4/6))

Program Office	Project Lead PI/ Lead Institution	Project Title	FASTMath members in project	FASTMath Technologies
BER	Steve Price / LANL Esmond Ng / LBNL	Probabilistic Sea-Level Projections from Ice Sheet and Earth System Models	Esmond Ng, Dan Martin, Andy Salinger, Mauro Perego, John Jakeman, Juliane Mueller	optimization, preconditioners, structure-grid AMR, unstructured-grid, NGP performance
BER	Adrian Turner / LANL	Discrete Element Method for High-Resolution Sea Ice Models (DEMSI)	Dan Ibanez, Andy Salinger (consult)	performance portability
BER	Hui Wan / PNNL	Assessing and Improving the Numerical Solution of Atmospheric Physics in ACME	Carol Woodward, David Garnder	Time Integration, Verification
BER	Mark Taylor / SNL	A Non-hydrostatic Variable Resolution Atmospheric Model in ACME	Carol Woodward, David Gardner, Daniel Reynolds	Time Integration, nonlinear and linear solvers
BER	Dan Riccutio / ORNL	Optimization of Sensor Networks for Improving Climate Model Predictions	Khachik Sargsyan, Cosmin Safta, Youssef Marzouk	UQ

U.S. DEPARTMENT OF ENERGY | Office of Science

FASTMATH

FASTMath is actively engaged with 19 SciDAC-4 application partnerships (FES (5/8))

Program Office	Project Lead PI/ Lead Institution	Project Title	FASTMath members in project	FASTMath Technologies
FES	Xianzhu Tang / LANL	Tokamak Disruption Simulation Center	Barry Smith, Jonathan Hu (collaborating with ASCR-PI Shadid)	time integration, multi-level block preconditioners
FES	Steve Jardin / PPPL	Center for Tokamak Transient Simulations	Mark Shephard, Onkar Sahni, Sherry Li	Unstructured meshes
FES	C.S. Chang / PPPL	Center for High-Fidelity Plasma Simulations	Mark Shephard, Mark Adams, Onkar Sahni	Unstructured meshes, solvers
FES	Brian Wirth / ORNL	Plasma Surface Interactions: Predicting the Perfromance of PFC Surfaces	Mark Shephard, Barry Smith, Onkar Sahni, Habib Najm, Khachik Sargsyan	Solvers, Unstructured meshes, UQ
FES	Paul Bonoli / MIT	Center for Integrated Simulaions of Fusion Relevant RF Actuators	Mark Shephard, Tzanio Kolev, Onkar Sahni	Unstructured meshes
FES	Zhihong Lin / UCI	ISEP: Integrated Simulation of Energetic Particles in Burning Plasmas	Rob Falgout, Ulrike Yang	Parallel AMG

FASTMath is actively engaged with 19 SciDAC-4 application partnerships (HEP (3/5) / NP (2/3))

Program Office	Project Lead PI/ Lead Institution	Project Title	FASTMath members in project	FASTMath Technologies
HEP	Jim Amundson / FNAL	Community Project for Accelerator Science and Simulation 4 (ComPASS4)	Ann Almgren, Pieter Ghysels, Mathias Jacquelin, Esmond Ng; Stefan Wild	sparse symmetric direct solvers; structured-grid AMR; optimization
HEP	Jim Kowalkowski / FNAL	HEP Data Analytics on HPC	Juliane Mueller, Sven Leyffer	sensitivity analysis and large scale optimization, mixed-integer, derivative-free
HEP	Salman Habib / ANL	Accelerating HEP Science: Inference and Machine Learning at Extreme Scales	Juliane Mueller, Stefan Wild	multi-fidelity optimization algorithms
NP	William Raph Hix / ORNL	Towards Exascale Astrophysics of Mergers and Supernovae (TEAMS)	Ann Almgren	structured-grid AMR
NP	Joe Carlson / LANL	NUCLEI	Esmond Ng, Chao Yang, Stefan WIld	eigenvalue problems

FASTMath is actively engaged with 19 SciDAC-4 application partnerships (BES (2/4) / NE (1/1))

Program Office	Project Lead PI/ Lead Institution	Project Title	FASTMath members in project	FASTMath Technologies
BES	Martin Head-Gordon / LBNL	Advancing Catalysis Modeling: From Atomistic Chemistry to Whole System Simulation	Esmond Ng, Sherry Li, Chao Yang	nonlinear eigenvalue problem, linear solver, preconditioner
BES	Tom Devereaux / SLAC	Topological and Correlated Matter via Tensor Networks and Quantum Monte Carlo	Chao Yang, Roel Van Beeuman	eigenvalue problems with tensor structures

Program Office	Project Lead PI/ Lead Institution	Project Title	FASTMath members in project	FASTMath Technologies
NE	David Andersson / LANL	Simulation of Fission Gas in Uranium Oxide Nuclear Fuel	Habib Najm	UQ

Some SciDAC Partnership Success

Ice Sheet Modeling

- **Predicting Ice Sheet and Climate Evolution at Extreme Scales** (PISCEES)
- PISCEES aims to enable quantitative predictions of coupled ice-sheet/climate evolution using a new generation of high performance computers and computational tools.

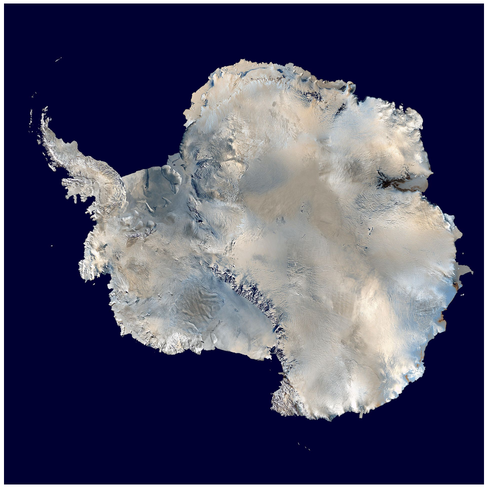

- Work at LBNL has focused on the Antarctic Ice Sheet.
  - Observation is that the “interior” of the ice sheet evolves much slower than the “grounding line”.
  - Ideal for deploying “adaptive mesh refinement”.
  - Developed the BISICLES code using the Chombo framework for structured grid adaptive mesh refinement.
    - + Able to model dynamics with high effective resolution but at a much lower cost.

"Anomalous Long Lifetime of Carbon-14"

Objectives

- Solve the puzzle of the long but useful lifetime of ^{14}C
- Determine the microscopic origin of the suppressed β-decay rate

Impact

- Establishes a major role for strong 3-nucleon forces in nuclei
- Verifies accuracy of *ab initio* microscopic nuclear theory
- Provides foundation for guiding DOE-supported experiments

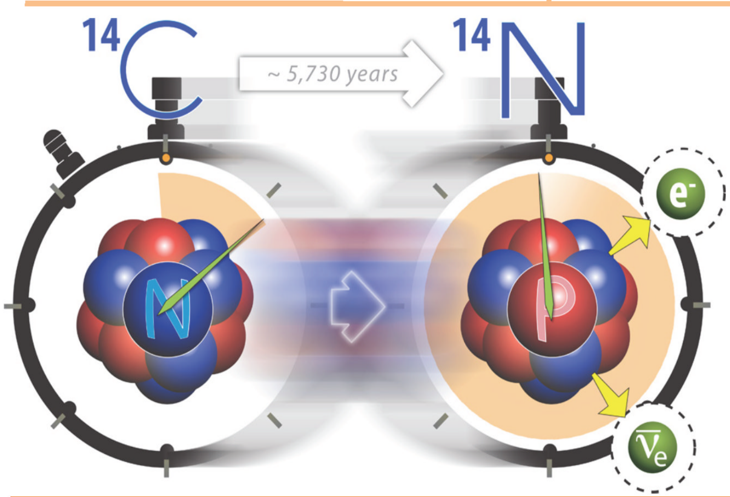

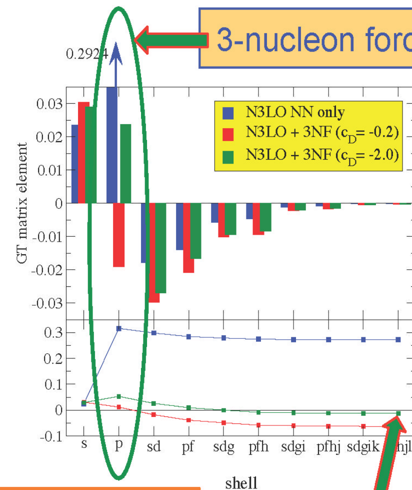

3-nucleon forces suppress critical component

net decay rate
Is very small

- Dimension of matrix solved for 8 lowest states ~ $1x10^9$
- Solution takes ~ 6 hours on 215,000 cores on Cray XT5 Jaguar at ORNL
- "Scaling of *ab initio* nuclear physics calculations on multicore computer architectures," P. Maris, M. Sosonkina, J. P. Vary, E. G. Ng and C. Yang, 2010 Intern. Conf. on Computer Science, Procedia Computer Science 1, 97 (2010)

PRL 106, 202502 (2011) PHYSICAL REVIEW LETTERS week ending 20 MAY 2011

Origin of the Anomalous Long Lifetime of ^{14}C

P. Maris,[1] J.P. Vary,[1] P. Navrátil,[2,3] W.E. Ormand,[3,4] H. Nam,[5] and D. J. Dean[5]

U.S. DEPARTMENT OF ENERGY | Office of Science | INCITE | IOWA STATE UNIVERSITY | TRIUMF

UNEDF SciDAC Collaboration — Universal Nuclear Energy Density Functional | NNSA | LAWRENCE LIVERMORE | OAK RIDGE NATIONAL LABORATORY — Managed by UT-Battelle for the Department of Energy

Solving CEBAF BBU Using Shape Uncertainty Quantification Method

SciDAC Success as a Collaboration between Accelerator Simulation, Computational Science and Experiment – Beam Breakup (BBU) instabilities at well below the designed beam current were observed in the CEBAF12 GeV upgrade of the Jefferson Lab (TJNAF) in which Higher Order Modes (HOM) with exceptionally *high* quality factor (Q) were measured. Using the shape uncertainty quantification tool developed under SciDAC, the problem was found to be a deformation of the cavity shape due to fabrication errors. This discovery was achieved as a team effort between SLAC, TOPS, and JLab which underscores the importance of the SciDAC multidisciplinary approach in tacking challenging applications.

Method of Solution - Using the measured cavity parameters as inputs, the deformed cavity shape was recovered by solving the ***inverse*** problem through an optimization method. The calculations showed that the cavity was 8 mm shorter than designed, which was subsequently confirmed by measurements. The result explains why the troublesome modes have high Qs because in the deformed cavity, the fields shift away from the HOM coupler where they can be damped. This shows that quality control in cavity fabrication can play an important role in accelerator performance. .

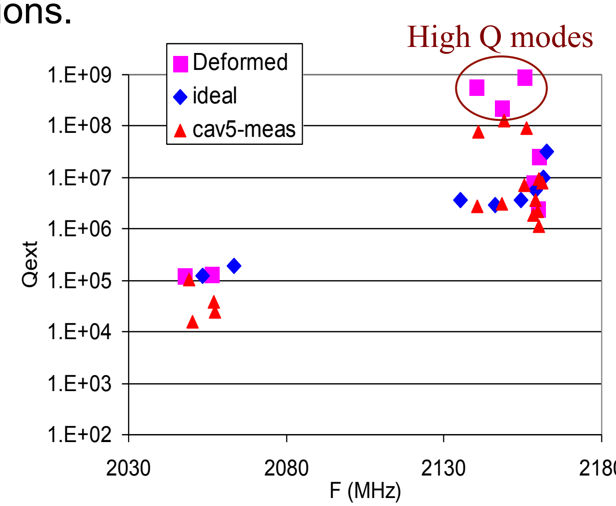

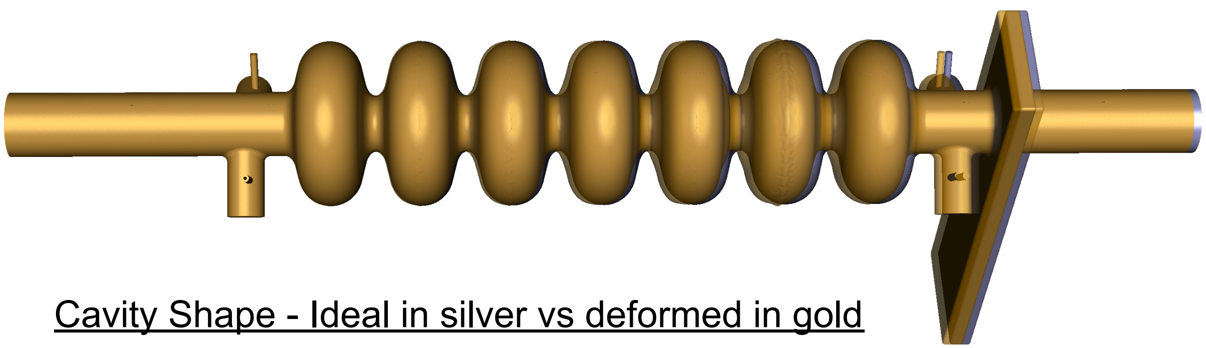

Cavity Shape - Ideal in silver vs deformed in gold

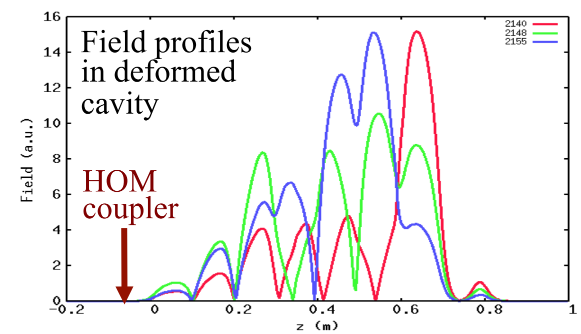

SciDAC Scientific Discovery through Advanced Computing
INCITE Advancing America's Science and Industrial Competitiveness
TOPS Terascale Optimal PDE Simulations
SLAC NATIONAL ACCELERATOR LABORATORY
Jefferson Lab

ACE3P: Multi-Physics Modeling of Superconducting Cavities

- **Electromagnetic simulation**

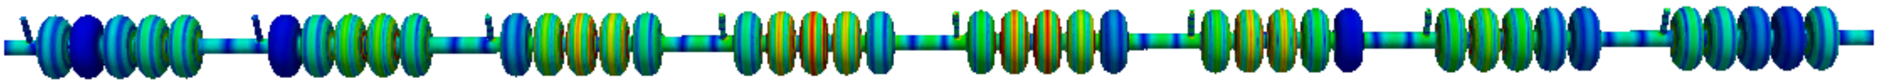

Trapped mode in Project X 650 MHz cryomodule's cavities

SciDAC collaboration in scalable solvers enables fast solution by significant reduction in ***memory*** footprint: 3 minutes per mode using 300 cores, 1.1 Tbytes of memory on NERSC Edison supercomputer.

- **Mechanical simulation**

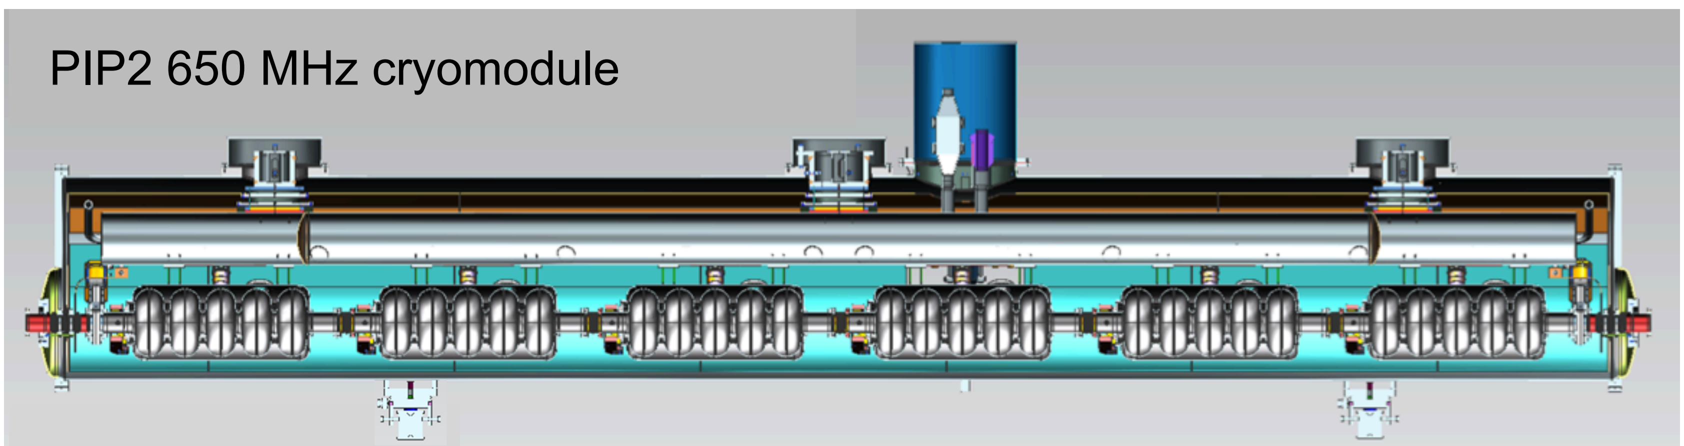

- Order of magnitude increase in problem size for modeling geometries with disparate spatial scales
- Linear solvers scalable in memory and on multi-core & AMR required

ACE3P: End-to-End Simulation of Dark Current and Radiation

- **Cryomodule simulation**

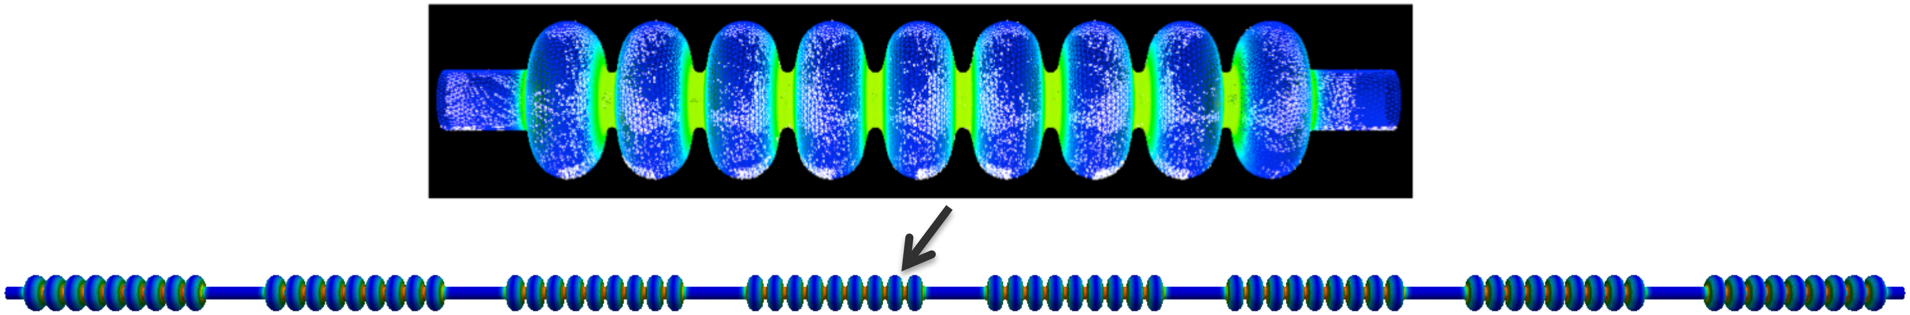

Dark current simulation for a single LCLS-II cromodule

  – 30 million particles, 20 minutes using 4800 cores on NERSC Edison

- **Linac simulation**
  - – Increase in problem size from modeling multiple cyomodules in superconducting linacs
    - • 25 cryomodules of different cavity designs in PIP-II
    - • 30 cryomodules in LCLS-II
  - – Integrated simulation of dark current and radiation effects (with Geant4) increases number of particles from secondary electrons.
  - – 2-3 orders more computing resources required

Wrapping up …

- SciDAC …
  - Enabling and accelerating **scientific discoveries** that require **high performance computing**.
  - Close **partnerships** among domain scientists, computational mathematicians, and computer scientists.

- Several forms of collaboration …
  - ~~Throw our algorithms/codes over the fence.~~
  - Build on what we have done/developed.
  - Develop and apply new applied/computational mathematical techniques.
  - Work with domain scientists and computer scientists often essential.
    + Mutual trust
    + Patience
    + Common language

Wrapping up …

- Partnership is key to the success of SciDAC.
  - Not just with domain scientists, but also between computational mathematicians and computer scientists.
  - This type of collaborations is particularly important and relevant in the forthcoming exascale computing era.
    - + Potentially drastic changes in computer architectures; deep memory hierarchy; limited memory per compute node, resilience, …

- Interesting and important scientific problems.
  - Both scientific curiosity and scientific satisfaction.

Thank you!