Large Scale Computing and Storage Requirements For Accelerator Modeling

Finite Element Approach

Lie-Quan Lee

SLAC National Accelerator Laboratory

Large Scale Computing and Storage Requirements for High Energy Physics
NERSC/ASCR/HEP Workshop, Washington D.C., November 12-13, 2009

SLAC
NATIONAL ACCELERATOR LABORATORY

NERSC Project

* Project name: Advanced Modeling for Particle Accelerators
* Principle Investigator: Kwok Ko
* Participating institutions:
  - SLAC, BNL, FNAL, ORNL, TJNAF
  - CW09 Users
    - ANL
    - CERN
    - Cornell University
    - Los Alamos Lab
    - Michigan State University
    - Paul Scherrer Institut
    - Royal Holloway U London

Scientific Objectives

- Summarize your projects and its scientific objectives for the next 3-5 years
- Compact Linear Collider (CLIC)
  - Simulating wakefield and evaluate HOM damping in Accelerating Structures (AS) and Power Extract and Transfer Structures (PETS)
  - Evaluate the dark current issue
  - Simulating two-beam acceleration (AS coupled with PETS)
- LHC Collimator and Crab Cavity for upgrade
  - Determine the broadband impedance of collimators which dominate the LHC impedance budget
  - Optimize the couplers of the current crab cavity design
- ILC Electron Gun
  - New design to eliminate a damping ring by providing a flat beam.
  - study the emittance dependence on the quantum efficiency of the cathode.
- Project X (Cryomodule)

Scientific Objectives (Continue)

* Summarize your projects and its scientific objectives for the next 3-5 years (continue)

- High-Gradient Structure R&D
  - Dark current
  - RF breakdown
- Muon Collider
  - Optimize external magnetic field map and the cavity shape
  - Study dark current and multipacting
- Photonic Band Gap (PBG) Structures
  - MIT PBG higher-order-mode (HOM) calculation in the microwave regime
  - SLAC PBG coupler design in the optical regime

2. Current HPC Requirements

- Architectures
- Compute/memory load
- Data read/written
- Necessary software, services or infrastructure
- Current primary codes and their methods or algorithms
- Known limitations/obstacles/bottlenecks
- Anything else?

NATIONAL ACCELERATOR LABORATORY

Finite Element Discretization

- Curvelinear tetrahedral elements
  - High fidelity modeling

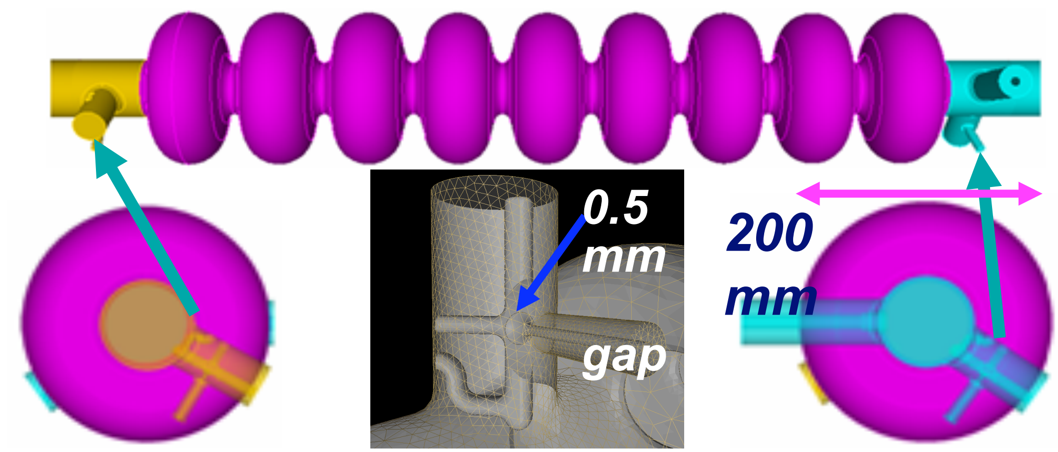

- High-order hierarchical vector basis functions (H(curl))
  - Provide tangential continuity required by physics
  - Easily set the boundary condition

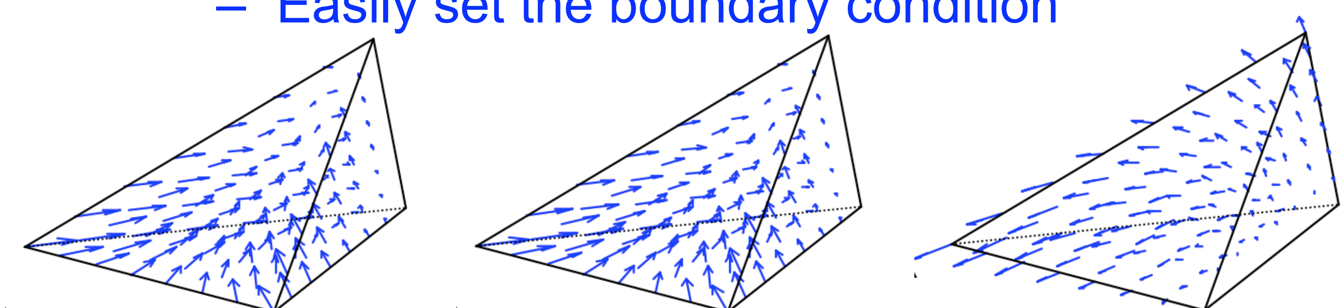

Shape functions

Entities	**# of bases**
Edge	p
Face	p(p-1)
Volume	p(p-1)(p-2)/2
Total	6E+4F+V

SLAC NATIONAL ACCELERATOR LABORATORY

Parallel Electromagnetic Suite ***ACE3P***

- **ACE3P** (**A**dvanced **C**omputational **E**lectromagnetics **3**D **P**arallel) is a comprehensive set of capabilities for the efficient and accurate *Modeling of complex accelerator structures* needed in Cavity design, Wakefield calculation, Dark current and multipacting simulation, RF Gun modeling, Multiphysics (RF, thermal, mechanical) analysis

ACE3P Modules		**– Accelerator physics application**
***Frequency Domain*:**	**Omega3P**	**– Eigensolver (nonlinear, damping)**
	S3P	**– S-Parameter**
***Time Domain*:**	**T3P**	**– Transients & Wakefields**
	Pic3P	**– EM Particle-In-Cell**
***Particle Tracking*:**	**Track3P**	**– Dark Current and Multipacting**
	Gun3P	**– Space-Charge Beam Optics**
***Multi-Physics*:**	**TEM3P**	**– EM-Thermal-Mechanical**
***Visualization*:**	**ParaView**	**– Meshes, Fields and Particles**

- User base of **ACE3P** formed to provide support to the worldwide accelerator community

(http://www-conf.slac.stanford.edu/CW09/)

SLAC NATIONAL ACCELERATOR LABORATORY

Eigenvalue Problem in Cavity Design using Omega3P

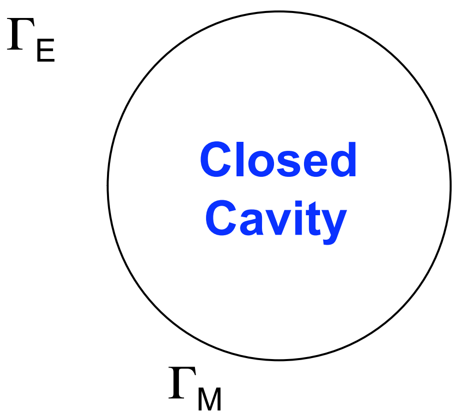

Time-Harmonic Maxwell's Equations

$$\nabla \times \left(\frac{1}{\mu}\nabla \times \vec{\mathbf{E}}\right) - k^2\epsilon\vec{\mathbf{E}} = 0 \ \ on \ \ \Omega$$

$$\vec{\mathbf{n}} \times \vec{\mathbf{E}} = 0 \ \ on \ \ \Gamma_E$$

$$\vec{\mathbf{n}} \times \frac{1}{\mu}\nabla \times \vec{\mathbf{E}} = 0 \ \ on \ \ \Gamma_M$$

H(curl) Conforming Element $\quad \mathbf{E} = \sum_i x_i \mathbf{N}_i$

$$\mathbf{Kx} = k^2\mathbf{Mx}$$

$$\mathbf{K}_{ij} = \int_\Omega (\nabla \times \mathbf{N}_i) \cdot \frac{1}{\mu}(\nabla \times \mathbf{N}_j)\, d\Omega$$

$$\mathbf{M}_{ij} = \int_\Omega \mathbf{N}_i \cdot \epsilon \mathbf{N}_j \, d\Omega$$

Find frequency and field vector of normal modes:

Eigenvalue: frequency

Eigenvector: field distribution

Cavity with Waveguide Coupling (Multiple Waveguide Modes)

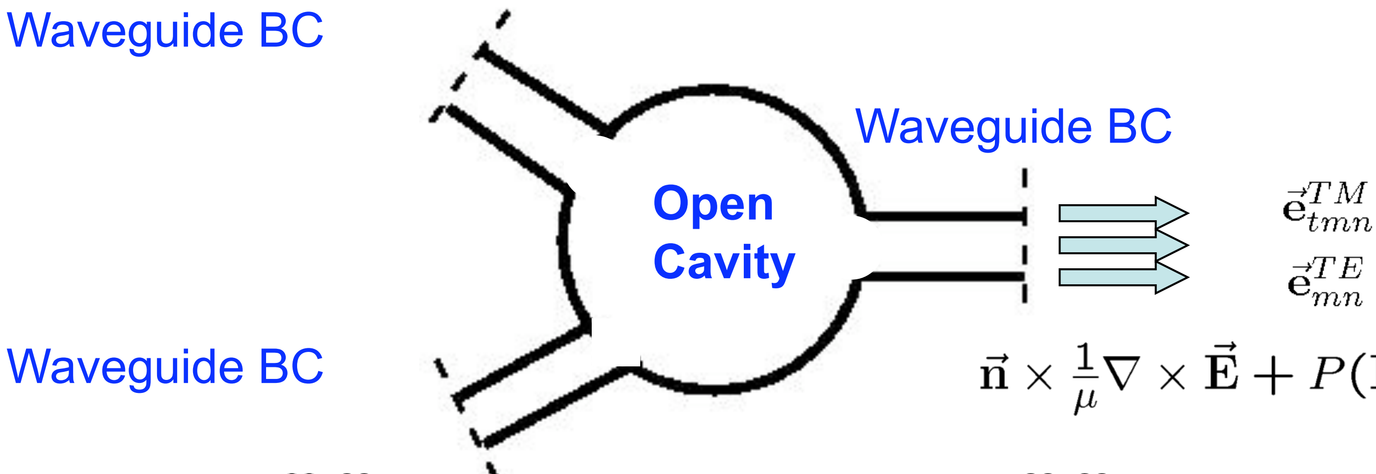

$\vec{\mathbf{n}} \times \frac{1}{\mu}\nabla \times \vec{\mathbf{E}} + P(\vec{\mathbf{E}}) = 0$ where

$$P(\vec{\mathbf{E}}) = \sum_m^\infty \sum_n^\infty \frac{k^2}{i\sqrt{k^2-k_{mn}^2}} \vec{\mathbf{e}}_{tmn}^{TM} \int_\Gamma \vec{\mathbf{e}}_{tmn}^{TM} \cdot \vec{\mathbf{E}}\, d\Gamma - \sum_m^\infty \sum_n^\infty i\sqrt{k^2-k_{mn}^2}\, \vec{\mathbf{e}}_{mn}^{TE} \int_\Gamma \vec{\mathbf{e}}_{mn}^{TE} \cdot \vec{\mathbf{E}}\, d\Gamma$$

* Vector wave equation with waveguide boundary conditions can be modeled by a **complex non-linear eigenvalue problem**

$$\mathbf{K}x + i\sum_{m,n}\sqrt{k^2-k_{mn}^2}\,\mathbf{W}_{mn}^{TE}x + i\sum_{m,n}\frac{k^2}{\sqrt{k^2-k_{mn}^2}}\mathbf{W}_{mn}^{TM}x = k^2\mathbf{M}x$$

where

$$(\mathbf{W}_{mn}^{TE})_{ij} = \int_\Gamma \vec{\mathbf{e}}_{mn}^{TE} \cdot \mathbf{N}_i\, d\Gamma \int_\Gamma \vec{\mathbf{e}}_{mn}^{TE} \cdot \mathbf{N}_j\, d\Gamma$$
$$(\mathbf{W}_{mn}^{TM})_{ij} = \int_\Gamma \vec{\mathbf{e}}_{tmn}^{TM} \cdot \mathbf{N}_i\, d\Gamma \int_\Gamma \vec{\mathbf{e}}_{tmn}^{TM} \cdot \mathbf{N}_j\, d\Gamma$$

$Q \equiv \frac{k_r}{2k_i}$: a measure for damping

SLAC
NATIONAL ACCELERATOR LABORATORY

Numerical Linear Algebra Algorithms Employed

* Generalized Real Eigenvalue Problems (closed cavity)
  – Shift-Invert Lanczos (with thick restart)
  – Implicit Restart Shift-Invert Lanczos (ARPACK)
* Complex Quadratic Eigenvalue Problems (open cavity)
  – Second Order Arnoldi
* Complex Nonlinear Eigenvalue Problems (open cavity)
  – Self-Consistent Iterations
  – Inverse Iterations
  – Nonlinear Rayleigh Ritz Iterations
  – Nonlinear Jacobi-Davidson
* Shifted Linear Systems (in solving EP)
  – Sparse direct solvers (MUMPS, SuperLU, WSMP)
  – KSP + spectral multilevel preconditioner
  – Hybrid linear solver

Implemented in Omega3P

SLAC NATIONAL ACCELERATOR LABORATORY

Omega3P on Cray XT (Franklin)

* Largest problems sizes
* Solver options
* Number of cores

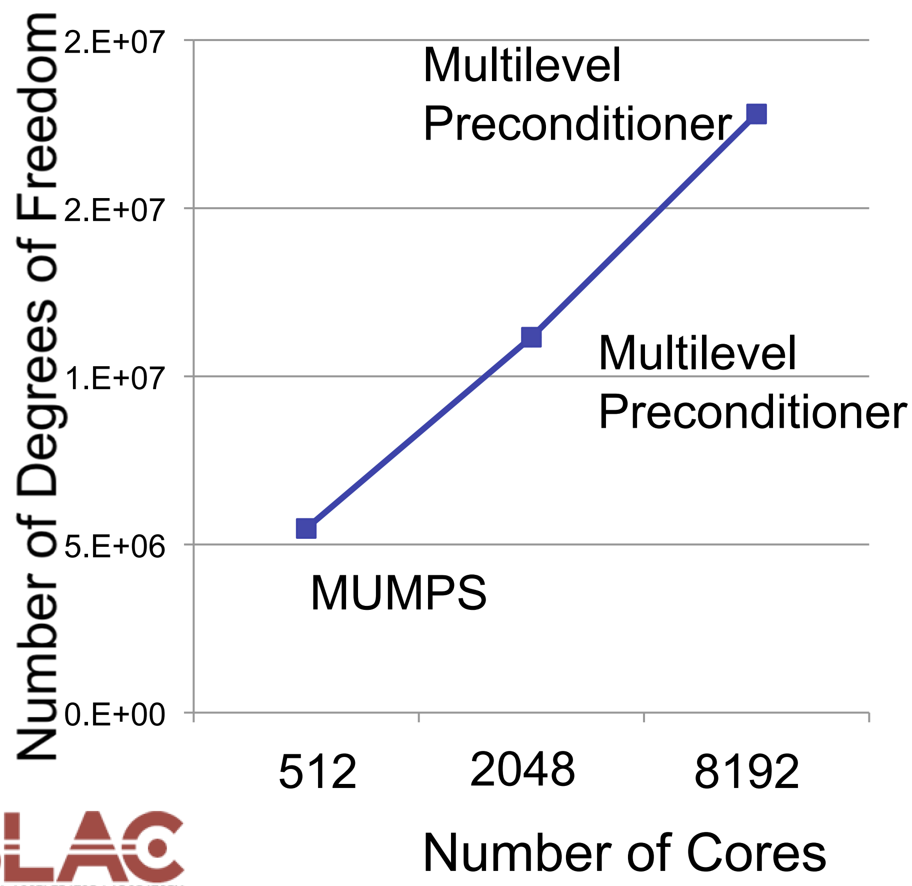

* Research activities on improving **scalability of per-node memory**:
  – Domain specific linear solvers and preconditioners
  – New scalable eigensolvers
  – Hybrid linear solver (TOPS/LBL)

Finite-Element Time-Domain Analysis using T3P

Compute transient and wakefield effects of beam inside cavities

Time-domain second-order vector wave equation:

$$\nabla \times \left(\frac{1}{\mu}\nabla \times \vec{\mathbf{E}}\right) + \sigma\frac{\partial \vec{\mathbf{E}}}{\partial t} + \epsilon\frac{\partial^2 \vec{\mathbf{E}}}{\partial t^2} = -\frac{\partial \vec{\mathbf{J}}}{\partial t} \;\; on \;\; \Omega$$

$$\vec{\mathbf{n}} \times \vec{\mathbf{E}} = 0 \;\; on \;\; \Gamma_E$$

$$\vec{\mathbf{n}} \times \frac{1}{\mu}\nabla \times \vec{\mathbf{E}} = 0 \;\; on \;\; \Gamma_M$$

H(Curl)-conforming Element $\vec{\mathbf{E}} = \sum_i x_i(t)\mathbf{N}_i(\vec{\mathbf{r}})$

$$\mathbf{M}\frac{1}{c^2}\frac{\partial^2 \mathbf{x}}{\partial t^2} + (\mathbf{R} + \mathbf{Q})\frac{1}{c}\frac{\partial \mathbf{x}}{\partial t} + \mathbf{K}\mathbf{x} = \mathbf{f}$$

$$\mathbf{K}_{ij} = \int_\Omega (\nabla \times \mathbf{N}_i) \cdot \frac{1}{\mu}(\nabla \times \mathbf{N}_j)\, d\Omega$$

$$\mathbf{M}_{ij} = \int_\Omega \mathbf{N}_i \cdot \epsilon \mathbf{N}_j \, d\Omega$$

$$\mathbf{R}_{ij} = \int_\Omega \mathbf{N}_i \cdot \sigma \mathbf{N}_j \, d\Omega$$

$$\mathbf{Q}_{ij} = \int_\Gamma (\mathbf{n} \times \mathbf{N}_i) \cdot \frac{1}{\mu}(\mathbf{n} \times \mathbf{N}_j)\, d\Gamma$$

$$\mathbf{f}_i = -\int_\Omega \mathbf{N}_i \cdot \frac{\partial \mathbf{J}}{\partial t}\, d\Omega$$

SLAC NATIONAL ACCELERATOR LABORATORY

Newmark-β Scheme for Time Stepping

$$\left(\mathbf{M} + \frac{c\Delta t}{2}(\mathbf{R}+\mathbf{Q}) + \beta(c\Delta t)^2\mathbf{K}\right) x^{n+1} = b$$

$$\begin{aligned} b = {} & (2\mathbf{M} - (1-2\beta)(c\Delta t)^2\mathbf{K})x^n \\ & - \left(\mathbf{M} - \tfrac{1}{2}c\Delta t(\mathbf{R}+\mathbf{Q}) + \beta(c\Delta t)^2\mathbf{K}\right) x^{n-1} \\ & -(c\Delta t)^2(\beta f_{n+1} + (1-2\beta)f^n + \beta f_{n-1}) \end{aligned}$$

* Unconditionally stable when β > 0.25
* For each time step, a linear system need to be solved
* Matrix in the linear system is SPD
* Conjugate Gradient + Block Jacobi / Incomplete Cholesky

SLAC
NATIONAL ACCELERATOR LABORATORY

T3P Scalability on Cray XT

Solution of linear system

- T3P uses a block-Jacobi preconditioner with incomplete factorization at every time step

Strong scaling studies

- Problem size with 56 million elements, 360 million DOFs

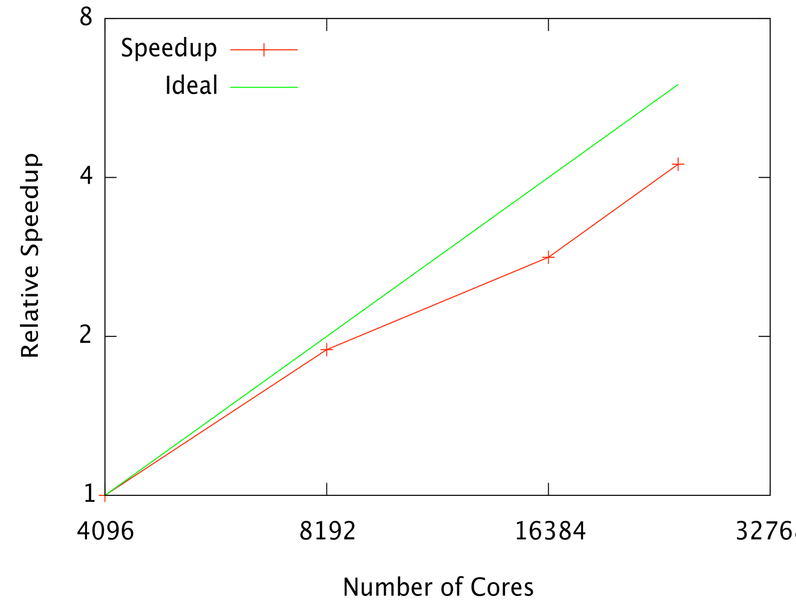

Performance of preconditioners

- Convergence of block-Jacobi preconditioner degrades as number of blocks (cores) increases

Cores	4096	8192	16384	24576
Iterations	251	257	261	268

SLAC
NATIONAL ACCELERATOR LABORATORY

Summary of Current HPC Requirements

Facilities Used or Using	☑ NERSC ☑ OLCF ☐ ACLF ☐
Architectures Used	☑ Cray XT ☑ IBM Power ☐ Blue
Total Computational Hours Used per Year	9,000,000 Core-Hours
NERSC Hours Used in 2009	1,174,091 Core-Hours
Number of Cores Used in Typical Production Run	4000
Wallclock Hours of Single Typical Production Run	12
Total Memory Used per Run	6000 GB
Minimum Memory Required per Core	1.5 GB
Total Data Read & Written per Run	1000 GB
Size of Checkpoint File(s)	50 GB

SLAC NATIONAL ACCELERATOR LABORATORY

3. HPC Usage and Methods for the Next 3-5 Years

- Upcoming changes to codes/methods/approaches
- Changes to Compute/memory load
- Changes to Data read/written
- Changes to necessary software, services or infrastructure
- Anticipated limitations/obstacles/bottlenecks on 10K-1000K PE system.
- Strategy for dealing with multi-core/many-core architectures

SLAC
NATIONAL ACCELERATOR LABORATORY

Strategy for Computing in Accelerator R&D

- Four closely interconnected efforts

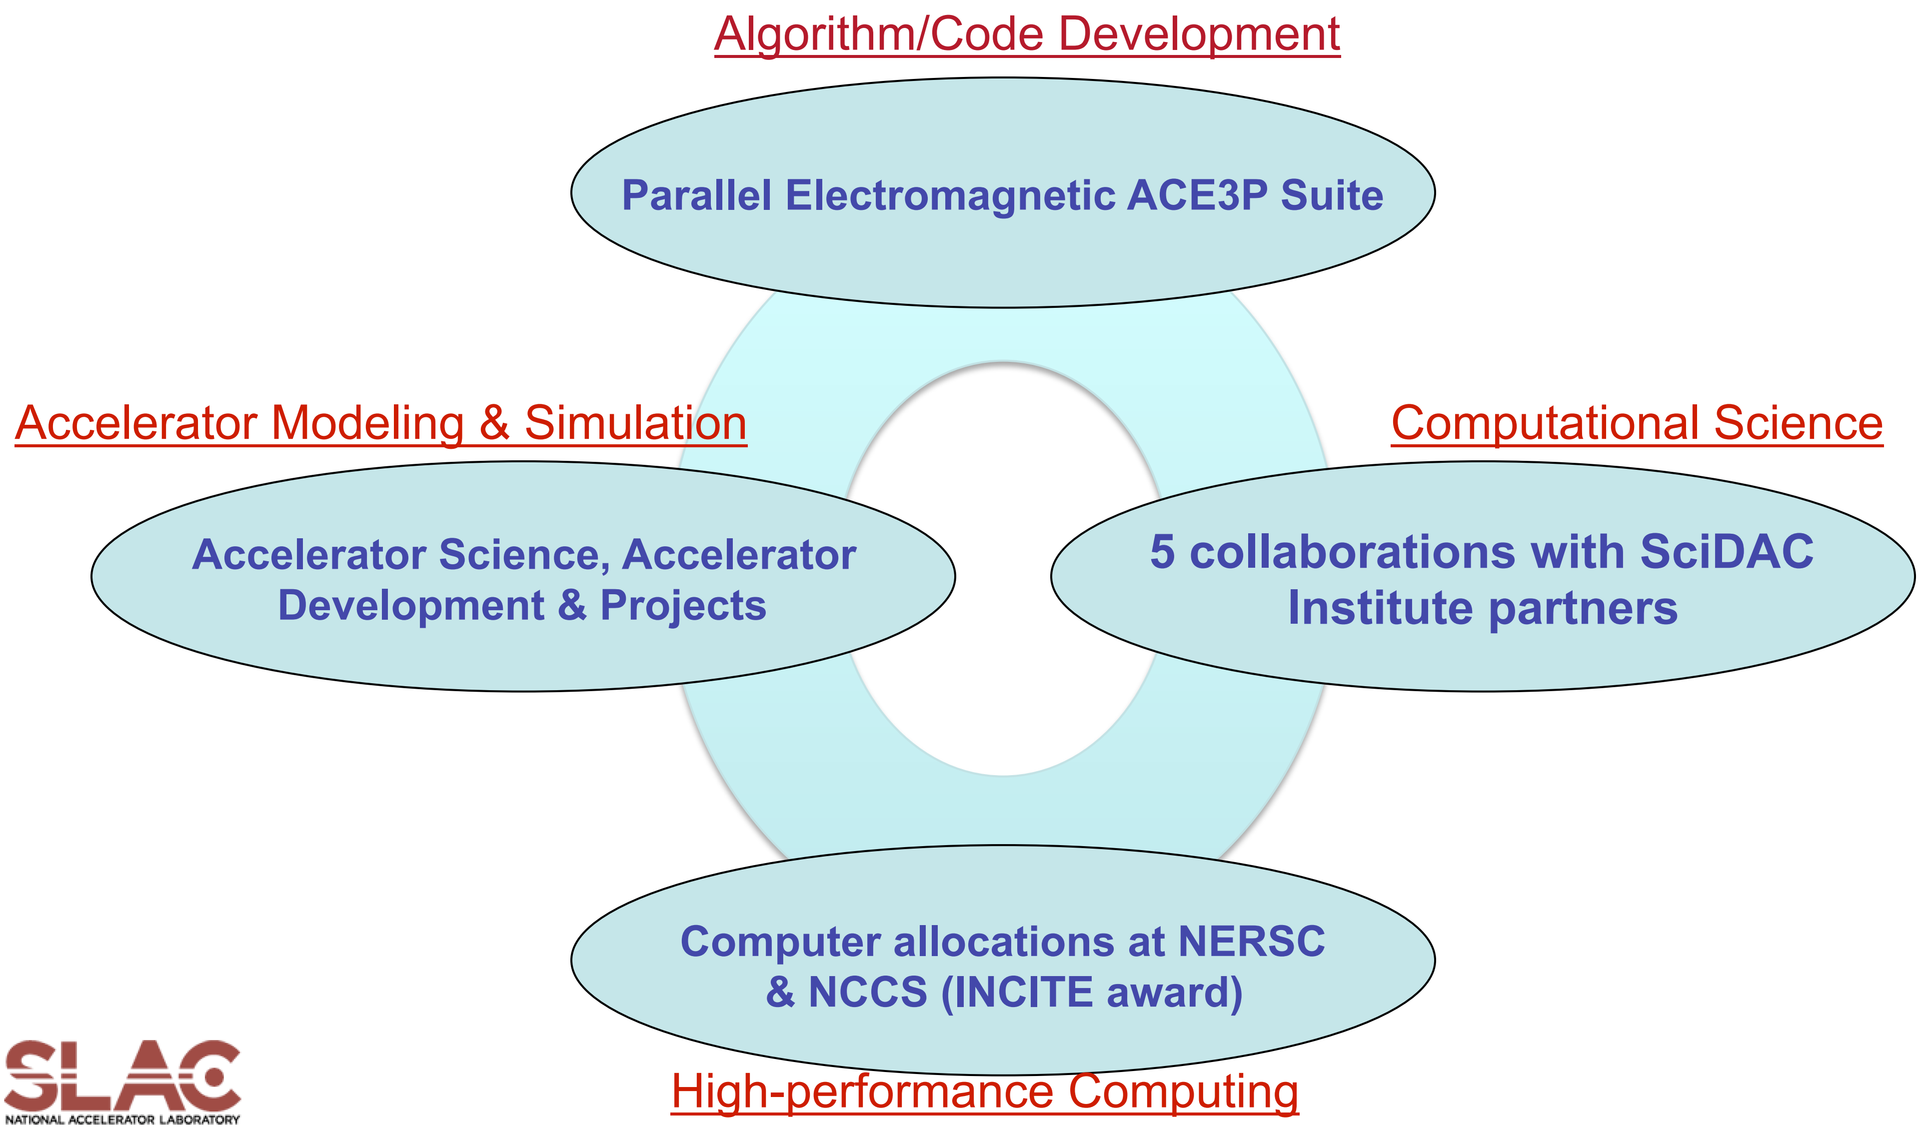

Developing New Capabilities

- RF Breakdown
  - Develop RF breakdown model including effects of plasma and particle-solid interaction
  - Verify and validate model with measurements
  - Quantify RF breakdown damage in structures

Image courtesy C. Adolphson

- Automatic accelerator prototyping (black box)

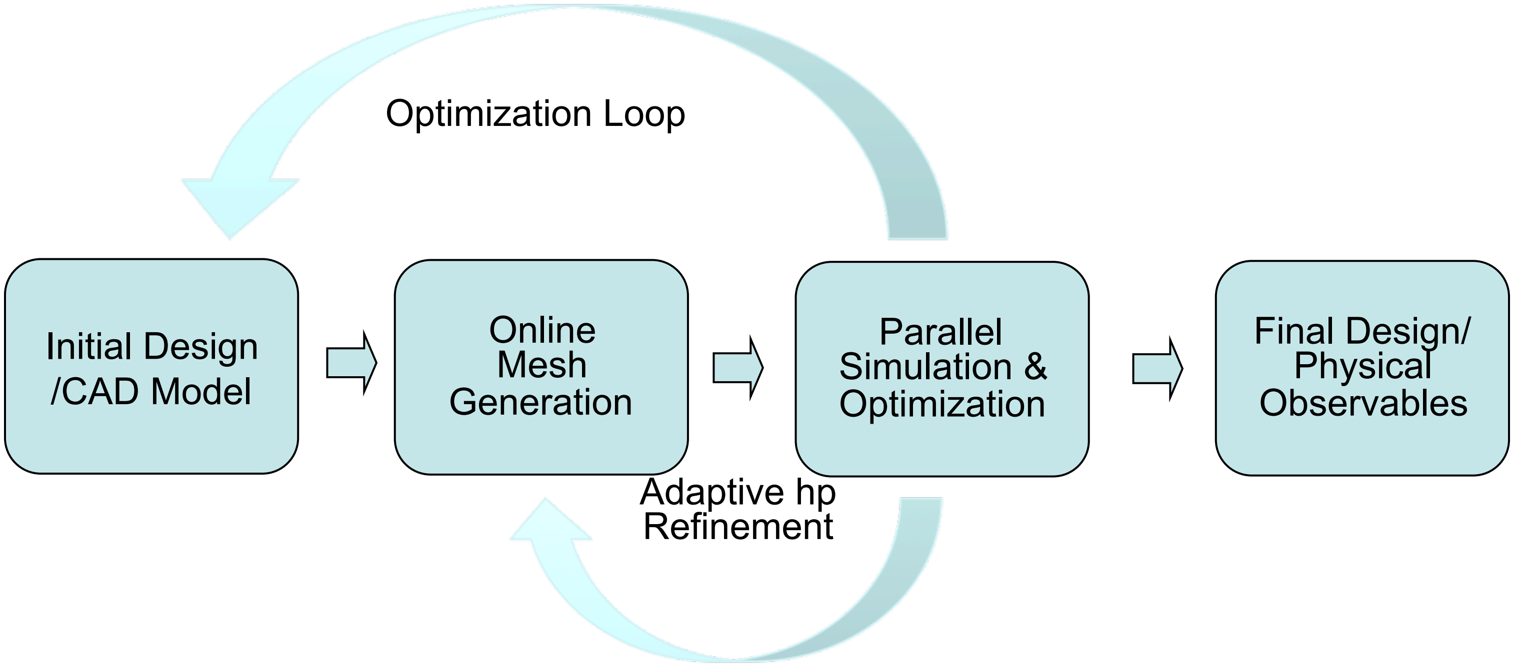

SLAC
NATIONAL ACCELERATOR LABORATORY

Towards 10K-1000K PE Systems

* More accurate workload and communication model in the partitioning scheme for better load balancing and domain decomposition
* More scalable linear solver which is the computational kernel of our finite element simulation suite
* Modular design of application codes
  – Swift adaptation of new components that can efficiently use those many-core or heterogeneous architectures
* Discontinuous Galerkin method
  – High-risk high-yield
  – Efficient, parallel method to solve very large problems
  – Expected to scale up to 1,000,000’s of cores
  – Allows GPU acceleration

Summary of HPC Requirements in 5 Years

Computational Hours Required per Year	50,000,000	
Anticipated Number of Cores to be Used in a Typical Production Run	50,000	
Anticipated Wallclock to be Used in a Typical Production Run Using the Number of Cores Given Above	24	
Anticipated Total Memory Used per Run	75,000	GB
Anticipated Minimum Memory Required per Core	1.5	GB
Anticipated total data read & written per run	10000	GB
Anticipated size of checkpoint file(s)	500	GB
Anticipated On-Line File Storage Required (For I/O from a Running Job)	10 100	GB and Files
Anticipated Amount of Data Moved In/Out of NERSC		GB per
Anticipated Off-Line Archival Storage Required		GB and Files

NATIONAL

4. Summary

- Recommendations on NERSC architecture, system configuration and associated service requirements needed for your science:
- What significant scientific progress could you achieve over the next 5 years with access to ~50X NERSC resources?
- What "expanded HPC resources" are important for your project?
- Any other special needs or NERSC wish lists?

SLAC NATIONAL ACCELERATOR LABORATORY

Recommendations

* A mid-range supercomputer with fat node
  - 64GB to 128GB per node
  - 1000 nodes
* A large supercomputer with
  - 2GB per core

SLAC NATIONAL ACCELERATOR LABORATORY

50X HPC Resources

* Predict **wakefield effects** of beam-environment interactions with **realistic bunch-size** in large complex accelerator structures
  – understand performance of the particle accelerator
  – provide information for further design optimization
* Model **self-consistent field-particle interactions** in space-charge dominated devices such as electron sources over **long time scales with high accuracy**
  – design high-quality and high-brightness beams for basic and applied scientific research
* Understand **dark-currents and RF breakdown** issues
  – provide insights for designing more efficient accelerating structures

SLAC
NATIONAL ACCELERATOR LABORATORY

Importance of "Expanded HPC Resources"

* More **per-node memory** will be extremely helpful in the foreseeable future in the frequency-domain analysis using Omega3P/S3P

* **More CPU hours** and the **ability to handle large jobs** are very important in time-domain analysis using T3P and particle tracking with Track3P

Thanks

Advancing Accelerator Science & Development

<u>Accelerator Science</u> – *PBG for Laser Acceleration, Muon Cavity Design, CERN high gradient structure (HGS)*

- Being tested at SLAC and KEK
- Dark current heating severe on disks towards output end – correlates well with measured breakdown rates
- *Work in progress on comparison with dark current measurements*

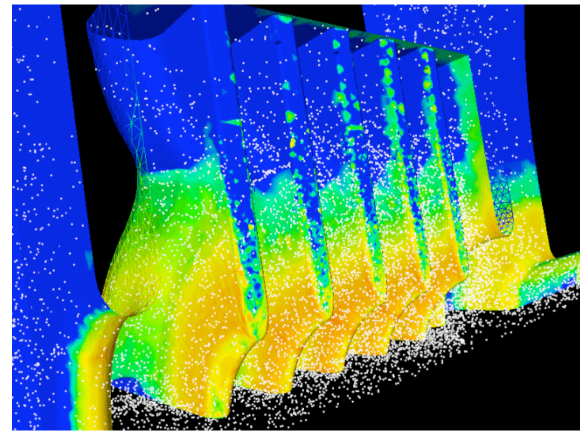

<u>Accelerator Development</u> – *LHC/LARP Crab Cavity & Collimator,*

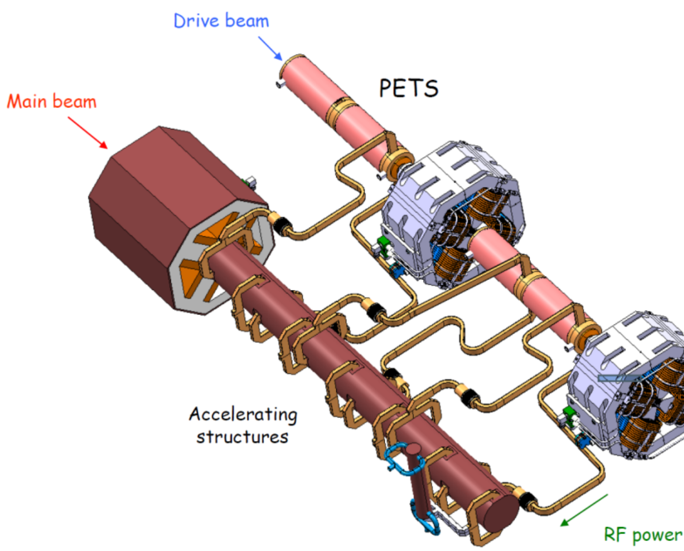

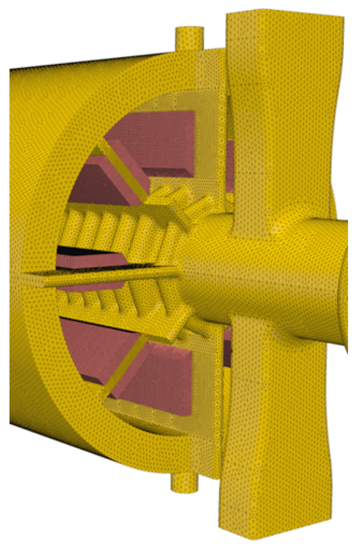

CERN CLIC Structure

- Wakefield damping in PETS
- RF power transfer from PETS to accelerating structures

Collaboration with CERN

SciDAC/ComPASS Collaborations under ASCR

Collaborations in Computational Science

E. Ng, X. Li, I. Yamazaki, C. Yang (**TOPS/**LBNL), L. Dianchin (**ITAPS**/LLNL), K. Devine, E. Boman, (**ITAPS/ CSCAPES**/SNL), D. Keyes (**TOPS/**Columbia), M. Shephard (**ITAPS**/RPI), W. Gropp (**CScADS**/UIUC), O. Ghattas (**TOPS/**UT Austin), Z. Bai (UC Davis), K. Ma (**ISUV**/UC Davis), A. Pothen (**CSCAPES**/Purdue), T. Tautges, H. Kim (**ITAPS**/ANL)

Contributed to advances in **Parallel Complex Nonlinear Eigensolver, Parallel Adaptive Mesh Refinement, Dynamic Load Balancing, Parallel and Interactive Visualization,** and **Shape Determination & Optimization** which helped solve the **CEBAF BBU problem**

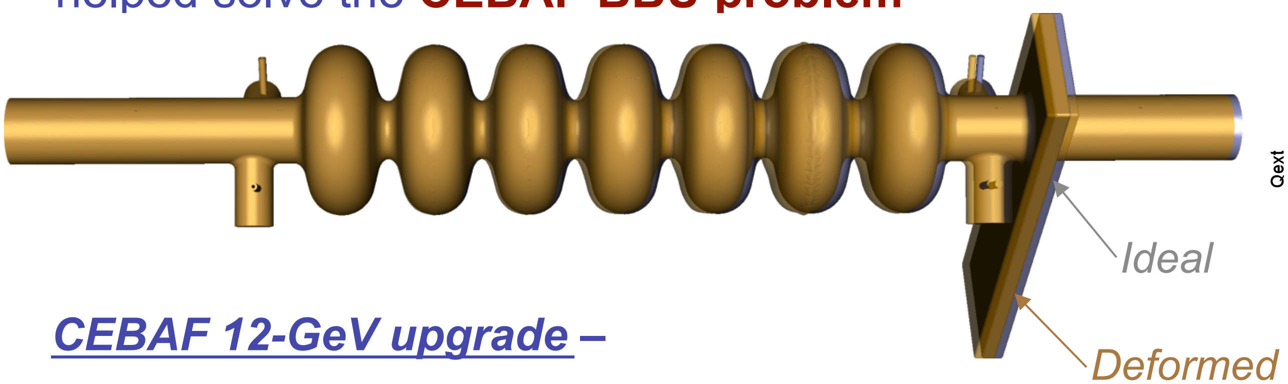

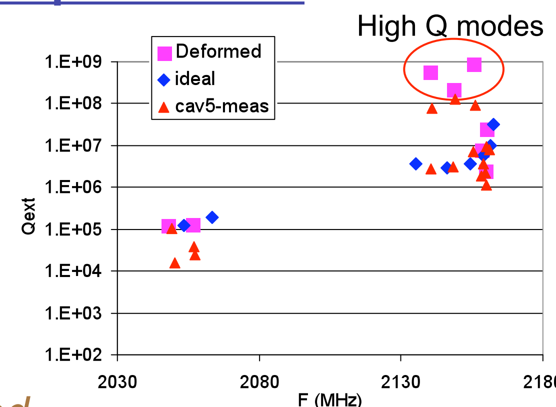

CEBAF 12-GeV upgrade **–**

- Beam breakup (BBU) observed at beam currents well below design threshold
- ***Solutions to the inverse problem*** identified the main cause of the BBU instability: **Cavity is 8 mm shorter** – predicted and confirmed later from measurements
- The fields of the **3 abnormally high Q modes** are shifted away from the coupler

SLAC NATIONAL ACCELERATOR LABORATORY