Computational Challenges in High Intensity Ion Beam Physics

Robert D. Ryne

Los Alamos National Laboratory

presented at the
2nd ICFA Advanced Accelerator Workshop
on the Physics of High Brightness Beams
UCLA Faculty Center
November 9-12, 1999

Acknowledgements

- Special thanks to
  - Ji Qiang[1,2]
  - Kwok Ko[2,3]
  - Patrick McCormick[4]

[1]Los Alamos Accelerator Code Group, LANL
[2]Computational Accelerator Physics Grand Challenge, LANL
[3]Numerical Modeling Group, SLAC
[4]Vizualization Team, Advanced Computing Laboratory, LANL

Acknowledgements

- Support provided by
  - DOE Office of Science, High Energy Physics Division
    - David Sutter, Gerald Peters
  - DOE Office of Science, Mathematical, Information, and Computational Sciences Division
    - Daniel Hitchcock, Thomas Kitchens

Outline

- Future national accelerator projects:
  - What are the computational challenges in high-intensity ion beam simulation
- How is large-scale (i.e. *terascale*) computing different from ordinary computing?
- How is the accelerator community meeting these challenges?

What are the Computational Challenges in High-Intensity Ion Beam Simulation?

- Future accelerators : higher intensity, greater complexity, greater precision, new parameter regimes
  - high-intensity linacs & rings, muon-based systems, beam-beam in hadron colliders
- Modeling on parallel computers essential for design decisions
  - evaluate/reduce risk
  - reduce cost
  - optimize performance
- Example problems:
  - high intensity transport + beam cooling in a μ-system
  - accurate design/analysis of EM structures for SNS and APT

Example: Need for Large-Scale Intense Beam Simulations

- Future high-average power-accelerators will have to operate with ultra-low losses
- Excess losses ⟶ radioactivation
  - degrade components
  - hinder/prevent hands on maintenance
  - reduce reliability and availability

Allowed loss is ~ 1-10 particles per billion in certain apps

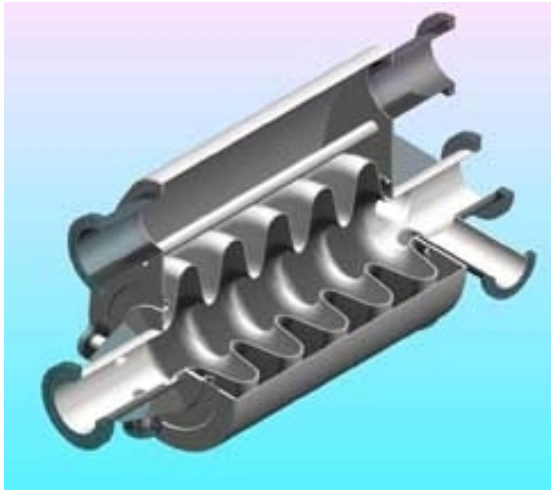

- Major loss mechanism: large amplitude halo particles striking the beam pipe

High-Resolution modeling using 0.1-1 billion particles is needed to make quantitative predictions of beam halo

Halo Formation from Mismatched Beam

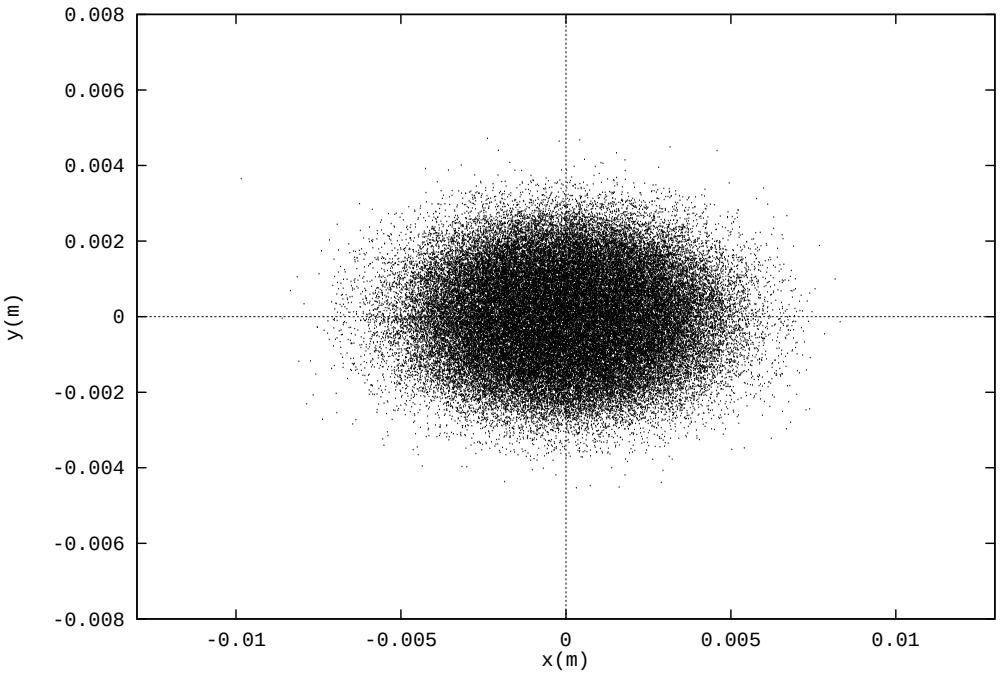

Matched beam

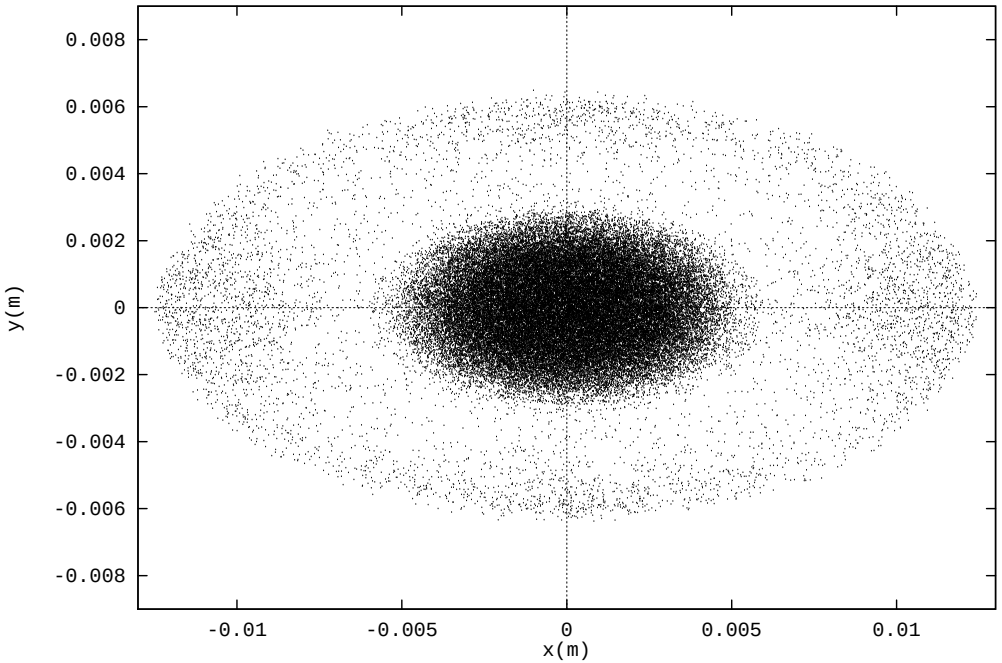

Mismatched beam

Large-Scale Simulations are Needed in order for Certain Calculated Quantities to Converge

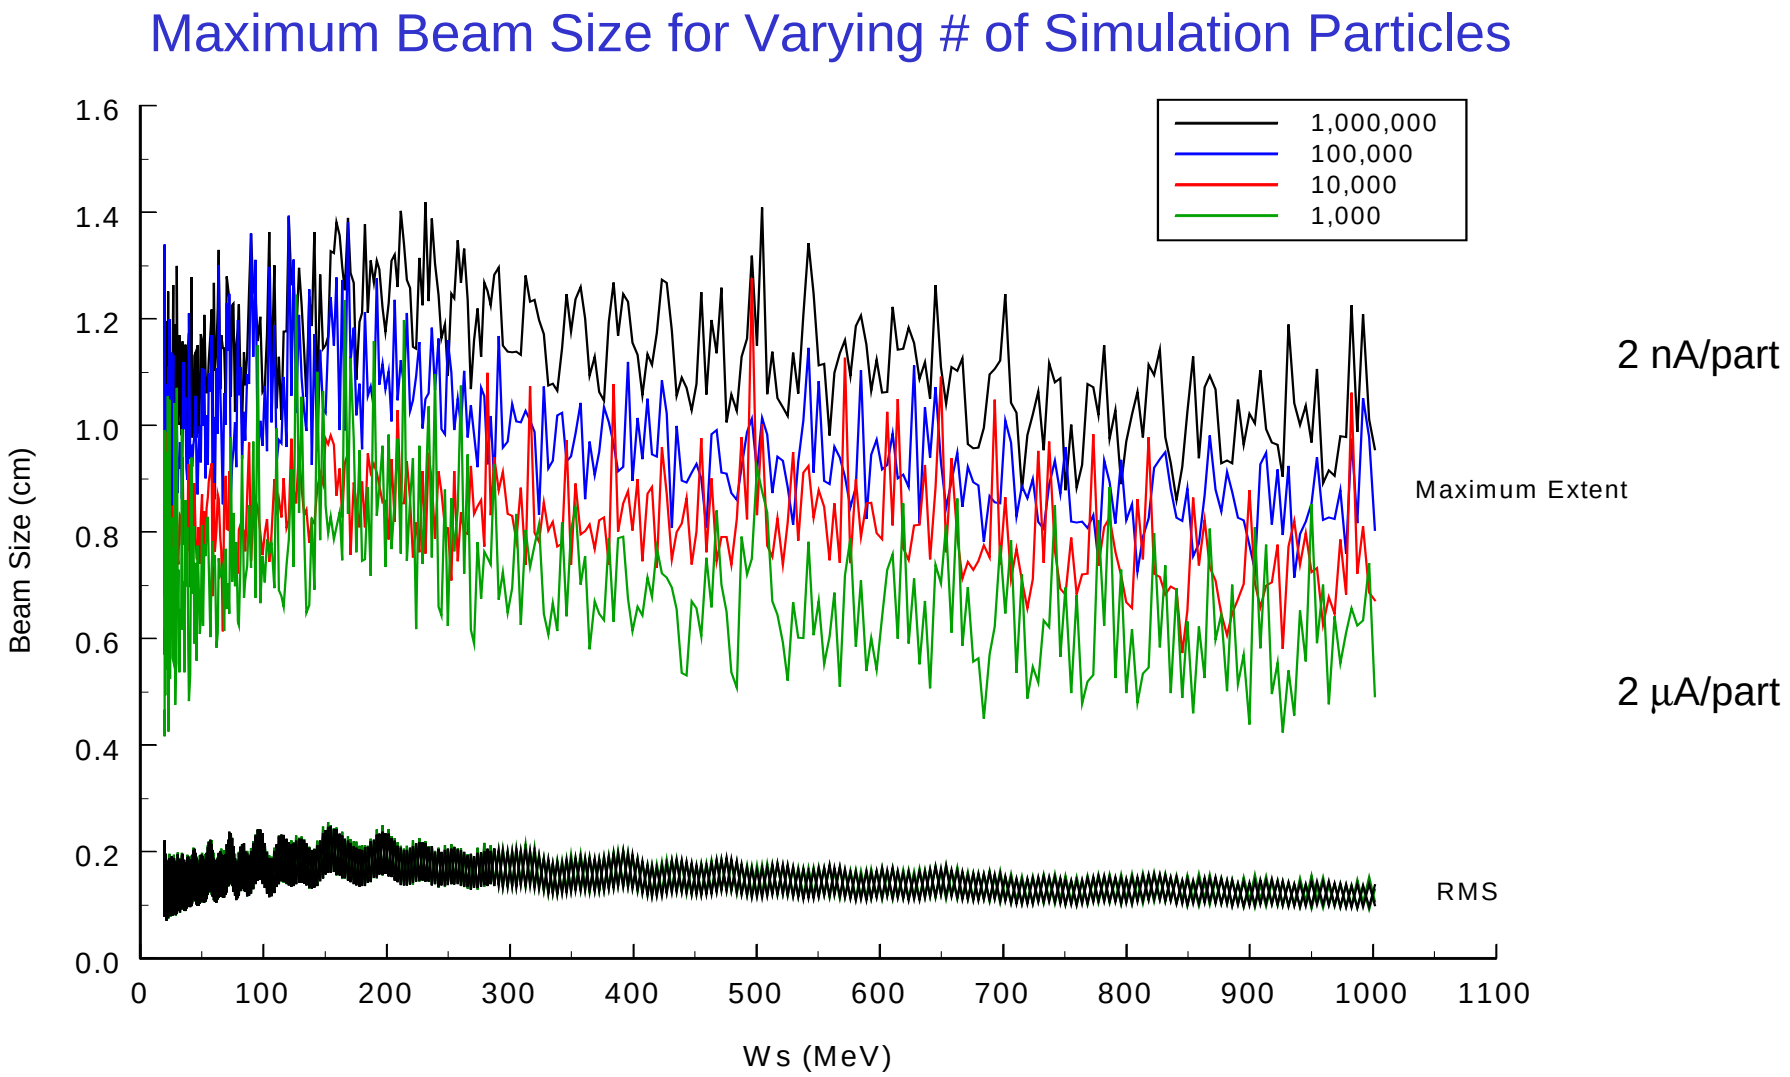

Example: Need for Large-Scale Electromagnetic Simulations

DDS for the NLC

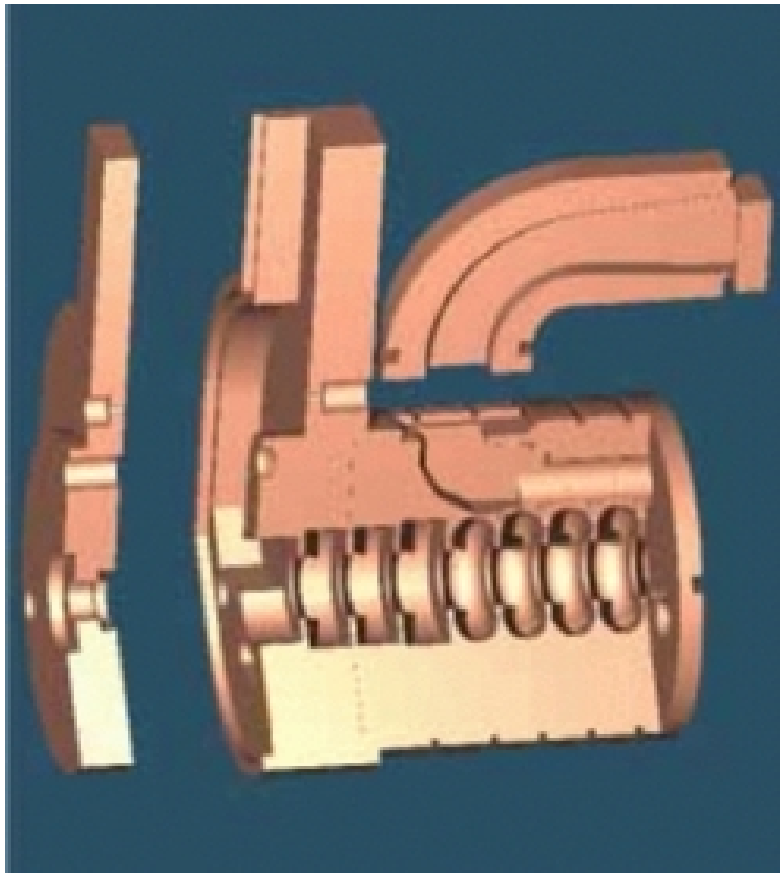

- Cell Design for NLC
  - frequency error of 1 part in 10,000.
  - mesh size close to fabrication tolerance (no tuning)
- RF Structure Design for SNS and APT
  - High-resolution, 3D modeling required for accurate field computation
  - Needed to compute frequencies, coupling

10^6 to 10^9 degrees of freedom

"High-Resolution" design, "System-Scale" analysis only possible on <u>very large scale</u> computers

Benefits of Large-Scale Simulation to Accelerator Science and Technology

- ***Tool of discovery*** to explore beams in novel configurations and under extreme conditions

Cooling for muon-based neutrino sources and colliders

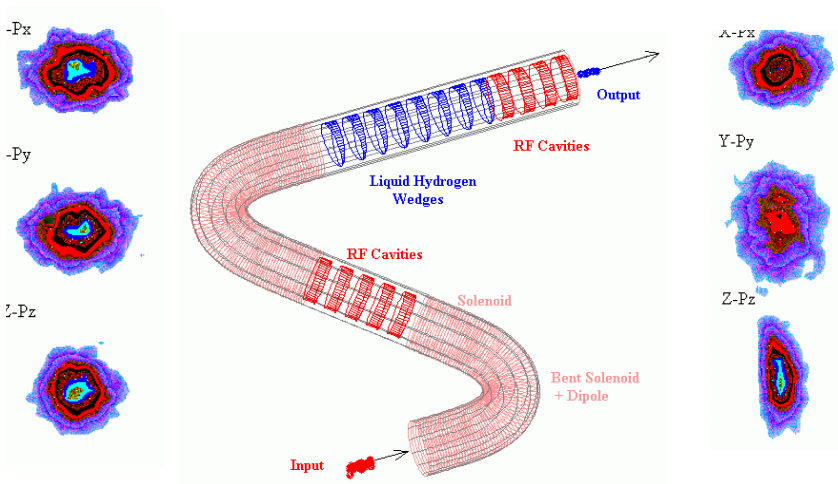

Plasma accelerators with gradients ~ 1-100 GeV/m

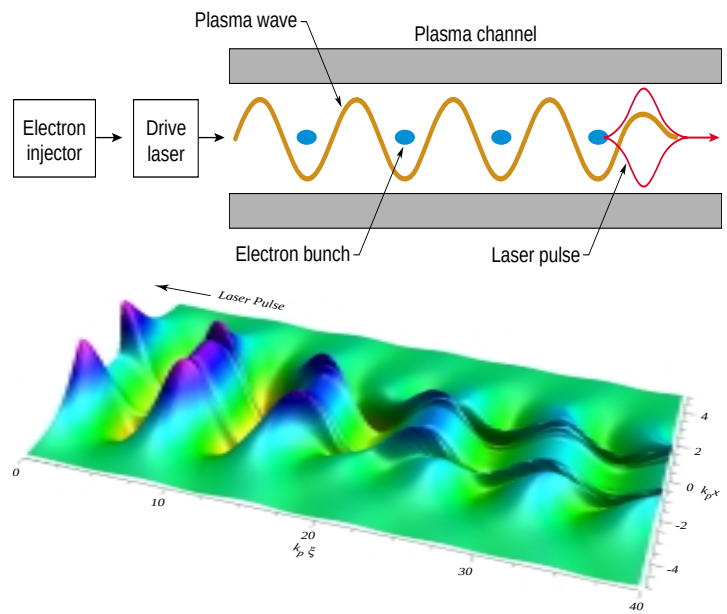

Large-scale simulation, coupled w/ theory & experiment, will help *advance the frontiers* of accelerator technology.

How is Large-Scale Computing Different from Ordinary Computing?

Architecture:

- Clusters of Symmetric Multiprocessors (SMP's)
  - Deep memory hierarchies
- PC clusters
  - Gaining popularity due to falling cost of networking hardware
  - Ideal for loosely coupled problems
  - Even cost effective for medium-scale tightly coupled problems

SGI Origin 2000 at ACL
IBM SP at NERSC

Programming:

- Complex issues associated w/ programming parallel computers, major differences compared with vector machines
  - *communication, cache*

Complexity & Scale:

- Highly three-dimensional
- 100's of millions of particles and grid points ⟶ data handling

Parallel Computing Activity is Increasing in the Accelerator Community

- Major NERSC allocations for accelerator science by year
  - FY99: 180K processor-hours
  - FY00: 532K processor-hours
- Large-scale accelerator science efforts at NERSC in FY00:
  - "Comp. Acc. Physics Grand Challenge," Ryne/Ko (LANL/SLAC)
  - "Continuing Studies of Plasma Beat Wave Accelerators," Mori (UCLA)
  - "3D and 2D Modeling of Intense Beams for HIF," Friedman (LBNL)
  - "Simulation of Advanced Acceleration Techniques," Esarey (LBNL)
  - "Effects of CSR on Accelerator Performance," Rui Li (TJNAF)
- Many others
  - John Galambos, Jeff Holmes (ORNL), Alfredo Luccio (BNL)
  - Andreas Adelmann (PSI)
  - ...

Development of large-scale parallel applications involves the integration of many pieces

Electromagnetics

- Parametric Geometry - CAD models to facilitate optimization
- Mesh Generation - impacts matrix conditioning/time-stepping
- Adaptive Refinement - improve accuracy/optimize resources

Beam Dynamics

- Frameworks - accelerator system geometry, models of beamline elements,...
- Transfer Map Methods -Symplectic maps corresponding to beamline elements
- Particle Managers - to minimize communication in parallel PIC codes

Both

- Parallel Solvers - scalable algorithms with fast convergence
- Domain Decomposition - partitioning tools for load balancing
- Visualization - efficient post-processing of large data sets; feature extraction and scientific discovery in simulation results
- Performance/Error Analysis - computer science/applied math

Success depends on integrating the combined efforts of a multi-disciplinary team including physicists, applied mathematicians, statisticians, computer scientists, software engineers, geometry builders, visualization experts,...

Brief Tour of Some Parallel Accelerator Simulations

- Beam Dynamics
  - IMPACT: Parallel PIC simulation for SNS
  - Ion Tracking: Parallel particle tracking for NLC
  - ORBIT: PSR comparison and SNS ring modeling
- Electromagnetics
  - Omega3P: NLC cavity design, PEP-II modeling
- Visualization
  - Parallel Volume Rendering: IMPACT results

Large-Scale Beam Dynamics Simulations using IMPACT

- 100M particle simulations require ~5 hours on 256 processors
- Simulations have been performed with 800M particles
  - Approaching real-world number of particles (900M for SNS)

Compare:

- 1M particle, 2D simulation using legacy code on a PC: weekend
- 100M particle, 3D IMPACT simulation on 256 processors : 5 hrs

- Parallel simulation is 100x larger and requires 1/10 the time
- Parallel computing results in a ***3 order-of-magnitude*** increase in simulation capability

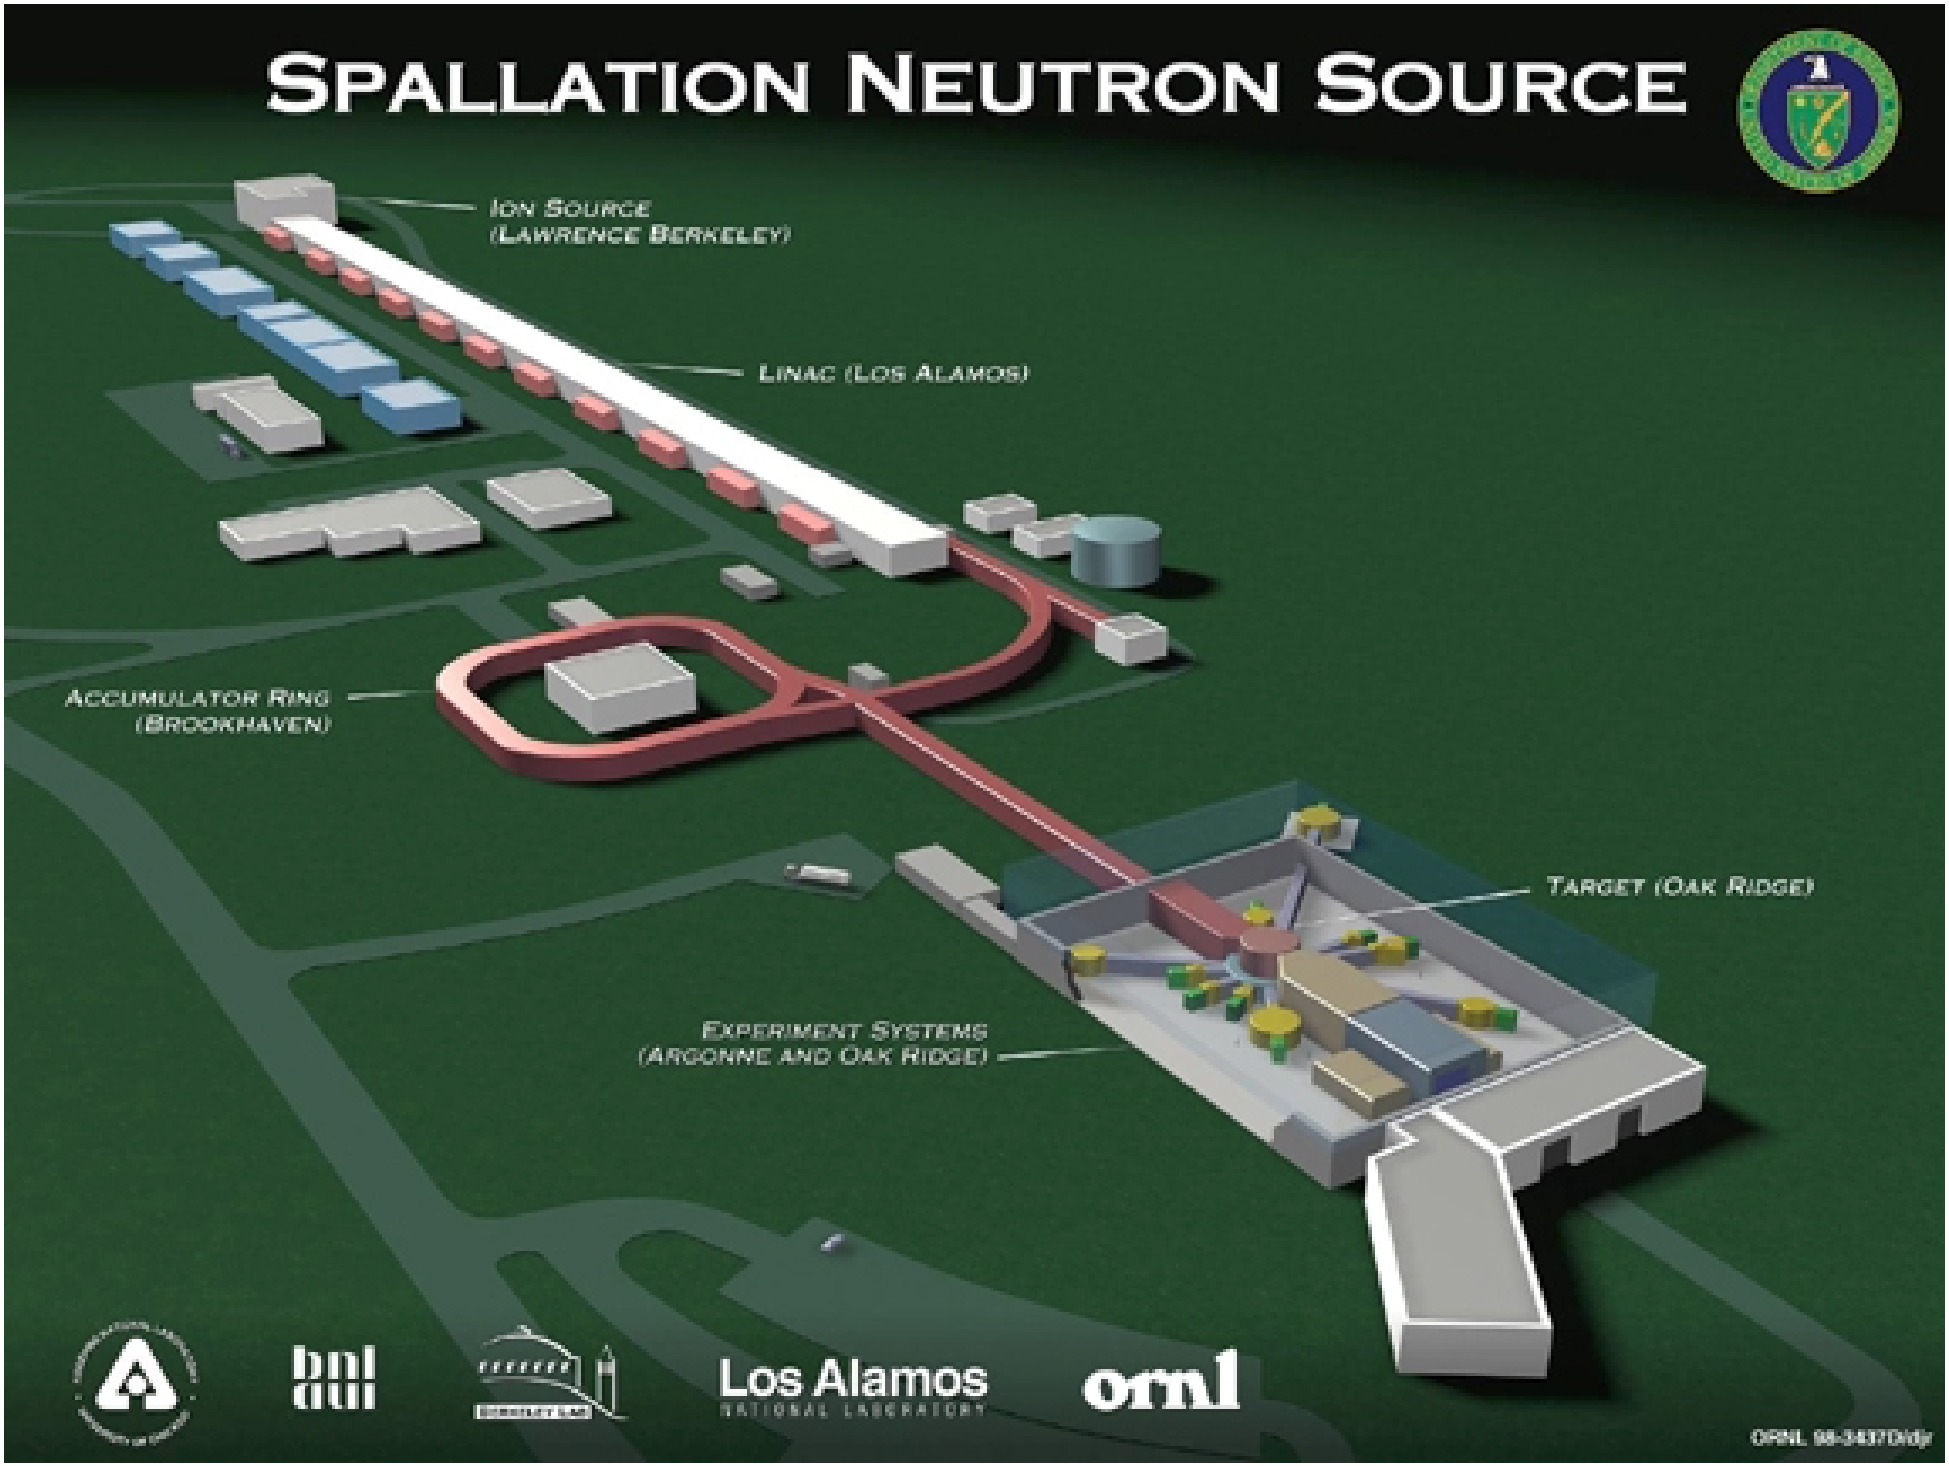
SPALLATION NEUTRON SOURCE
ION SOURCE
(LAWRENCE BERKELEY)
LINAC (LOS ALAMOS)
ACCUMULATOR RING
(BROOKHAVEN)
TARGET (OAK RIDGE)
EXPERIMENT SYSTEMS
(ARGONNE AND OAK RIDGE)
bnl
Los Alamos
NATIONAL LABORATORY
ornl

SNS Linac Configuration

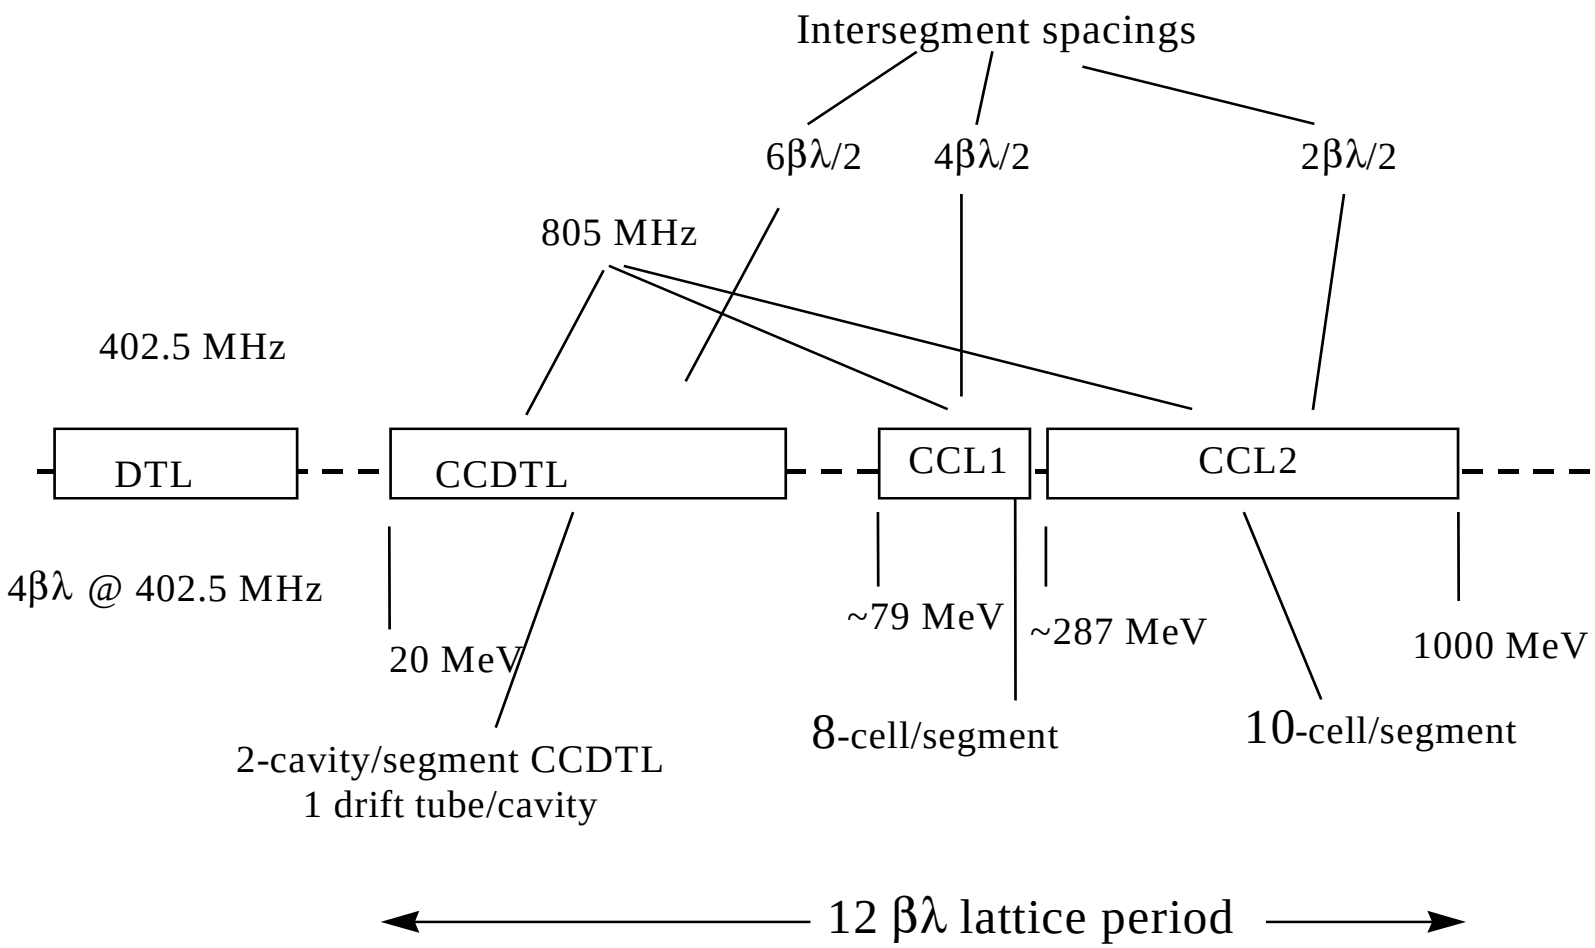

Design Energy and Accelerating Gradient vs. Distance

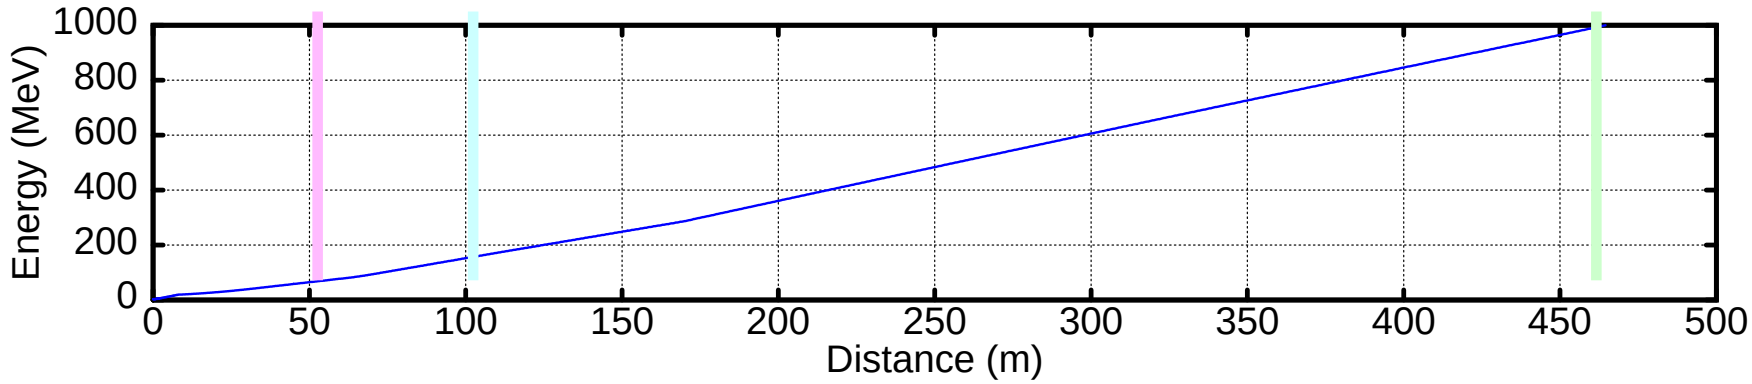

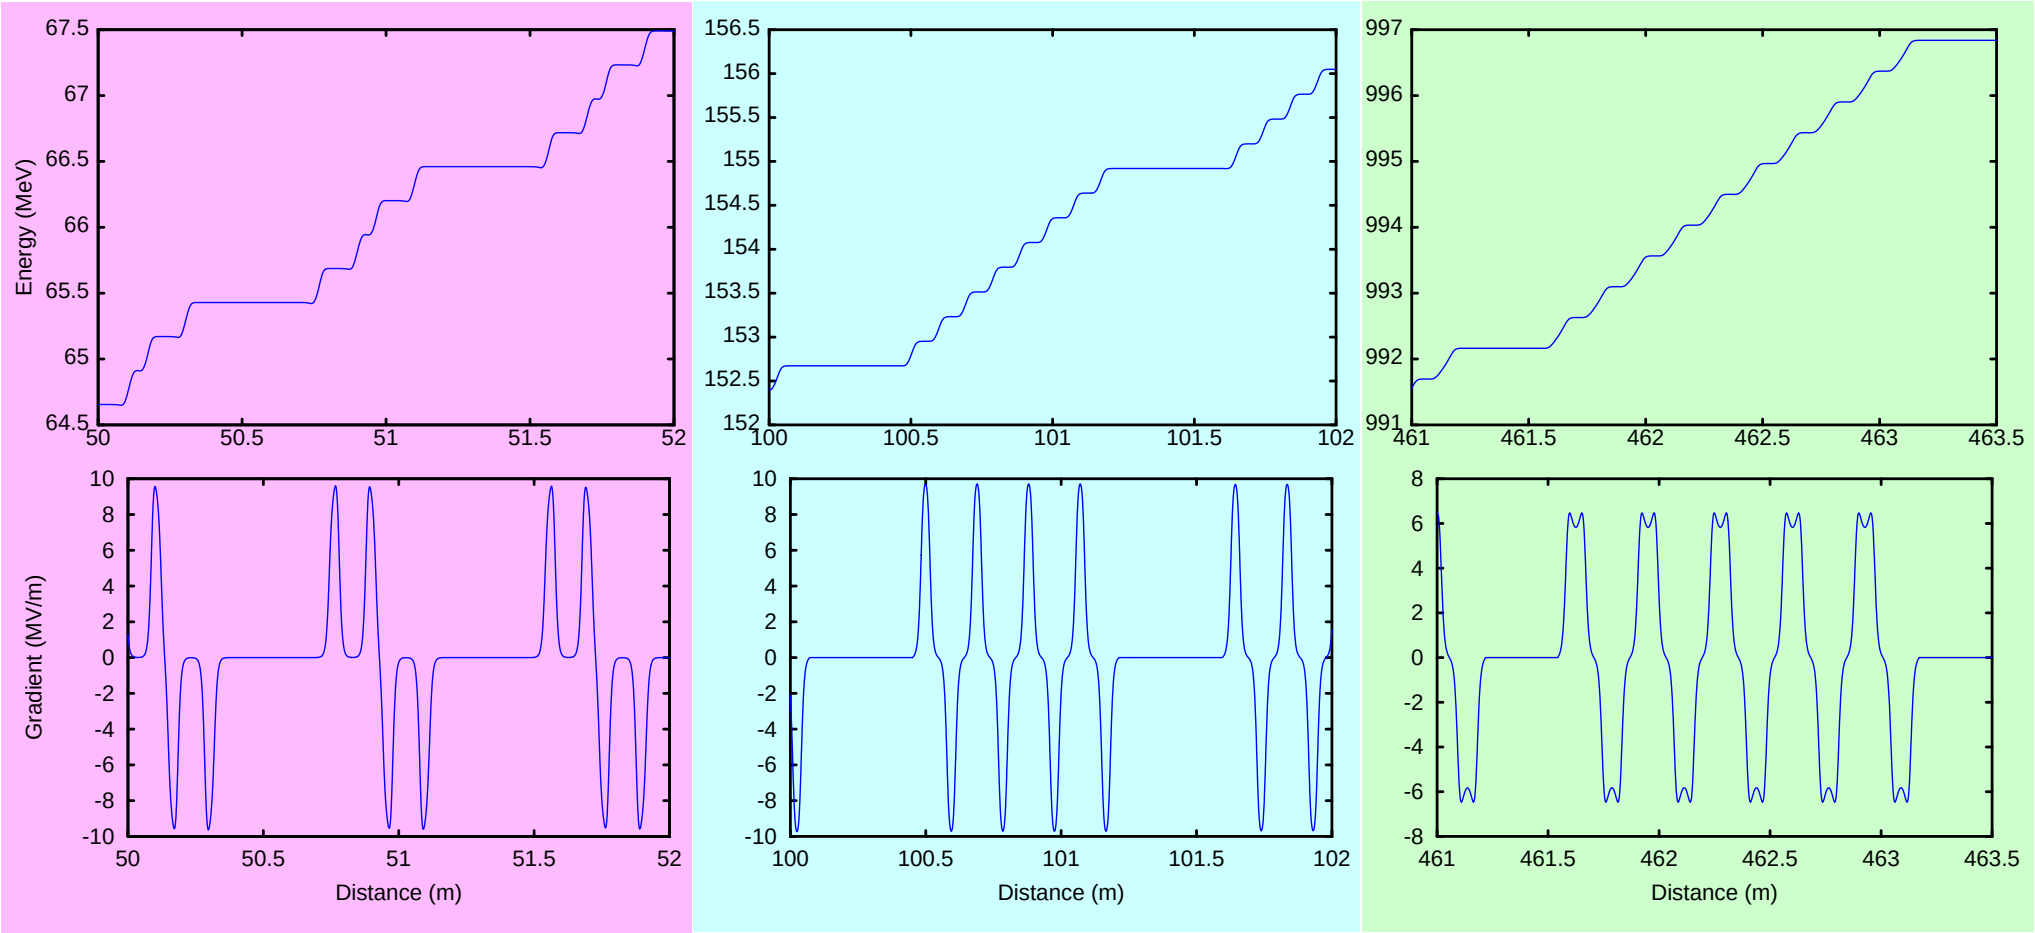

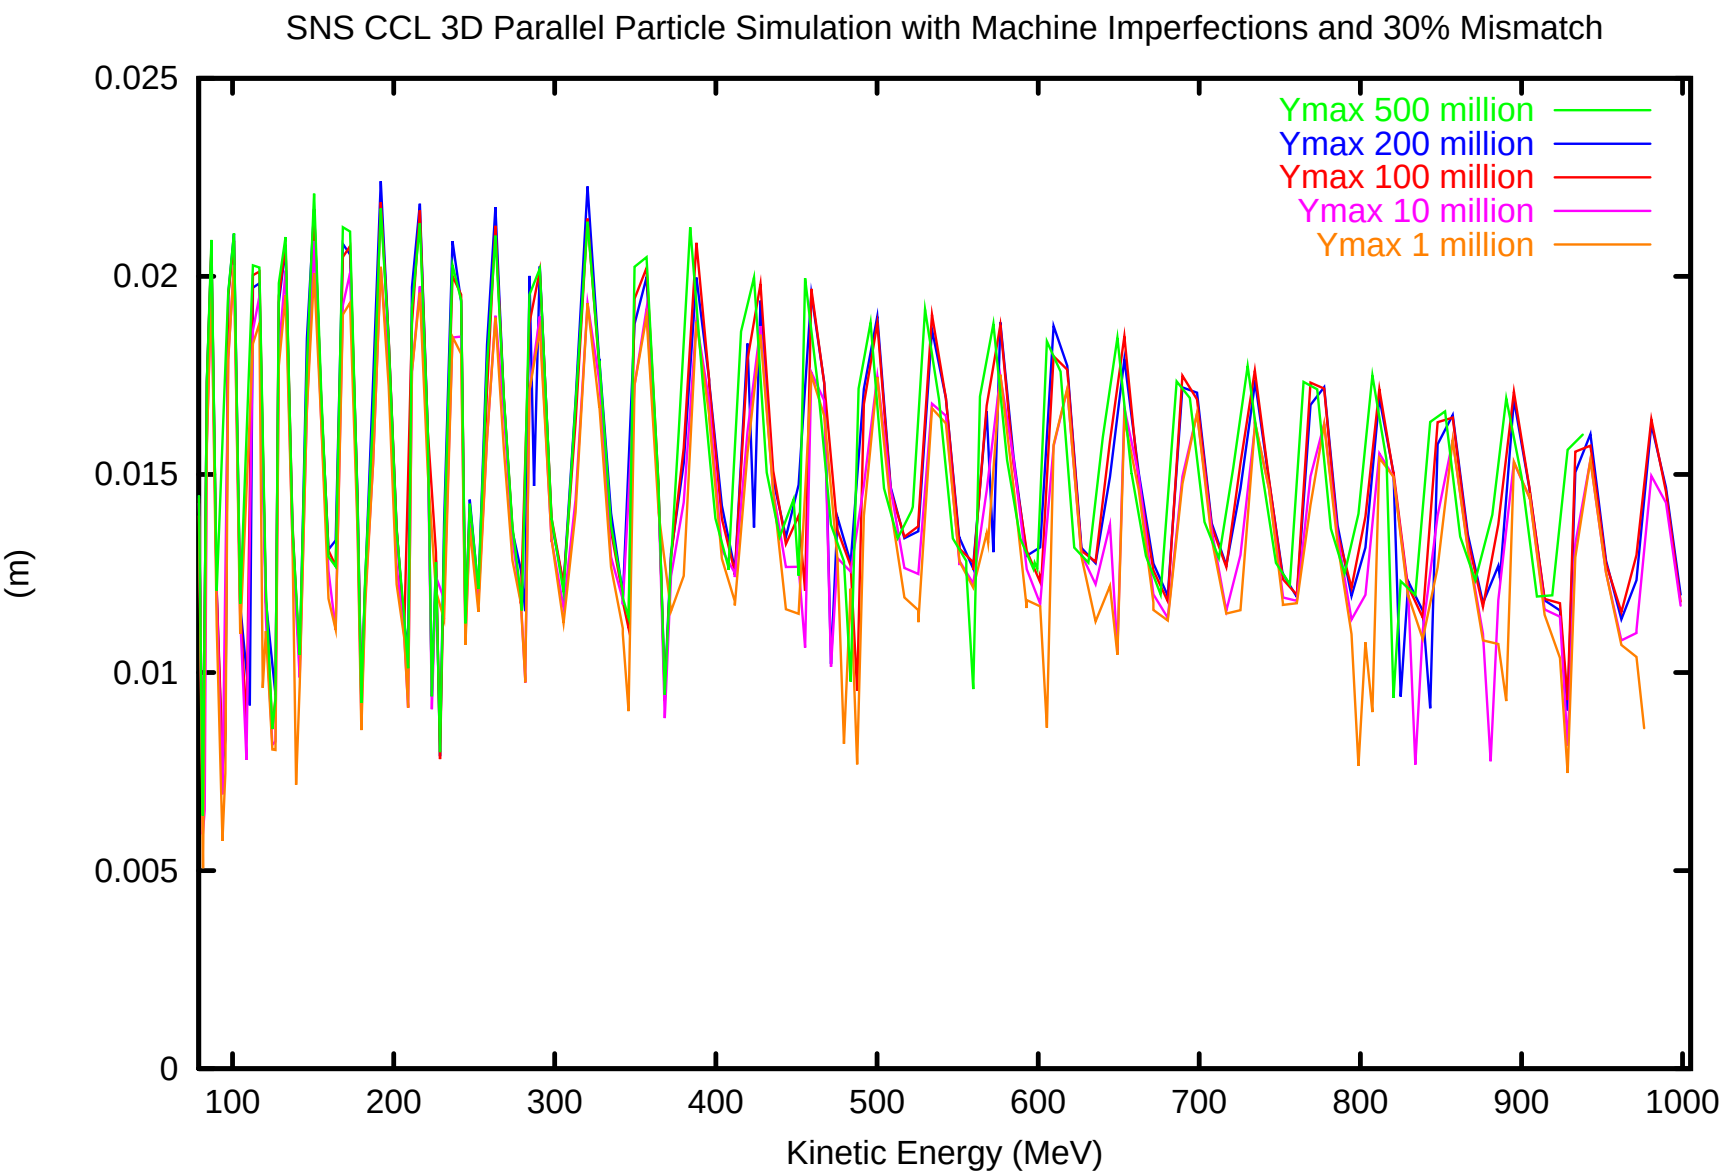

Maximum particle amplitude with varying # of simulation particles, run on the ACL Nirvana system using 32-1024 processors

Plots of "Integrated Density" from a 500 million particle linac simulation

X-P_X Plot

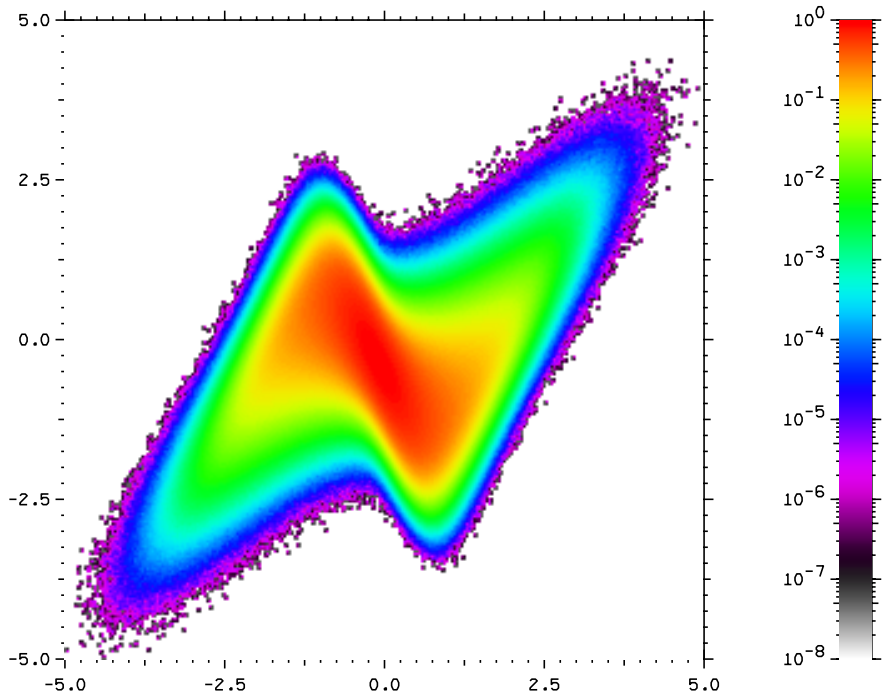

X-Y Plot

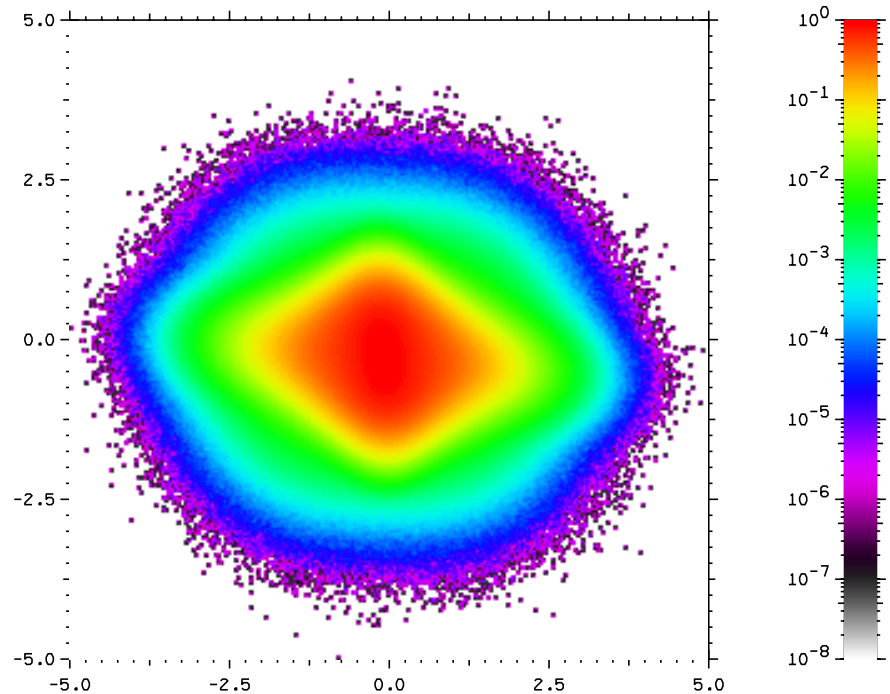

IMPACT: Integrated Map and Particle Accelerator Tracking Code

- Use of split-operator method with $H = H_{external} + H_{space\ charge}$
- Philosophy:

  Do not take tiny steps to push ~100 million particles

  Do take tiny steps to compute *maps*, then push w/ maps
- 3D space charge calculation
  - FFT-based with open boundary conditions
- Canonical variables (x, p_x, y, p_y, t, p_t)
- Dynamic load balance
- Choice of particle managers
- Parallel implementation using F90/MPI, POOMA, HPF

Split-Operator Approach

How to turn ***any*** magnetic optics code into a tracking code with space charge

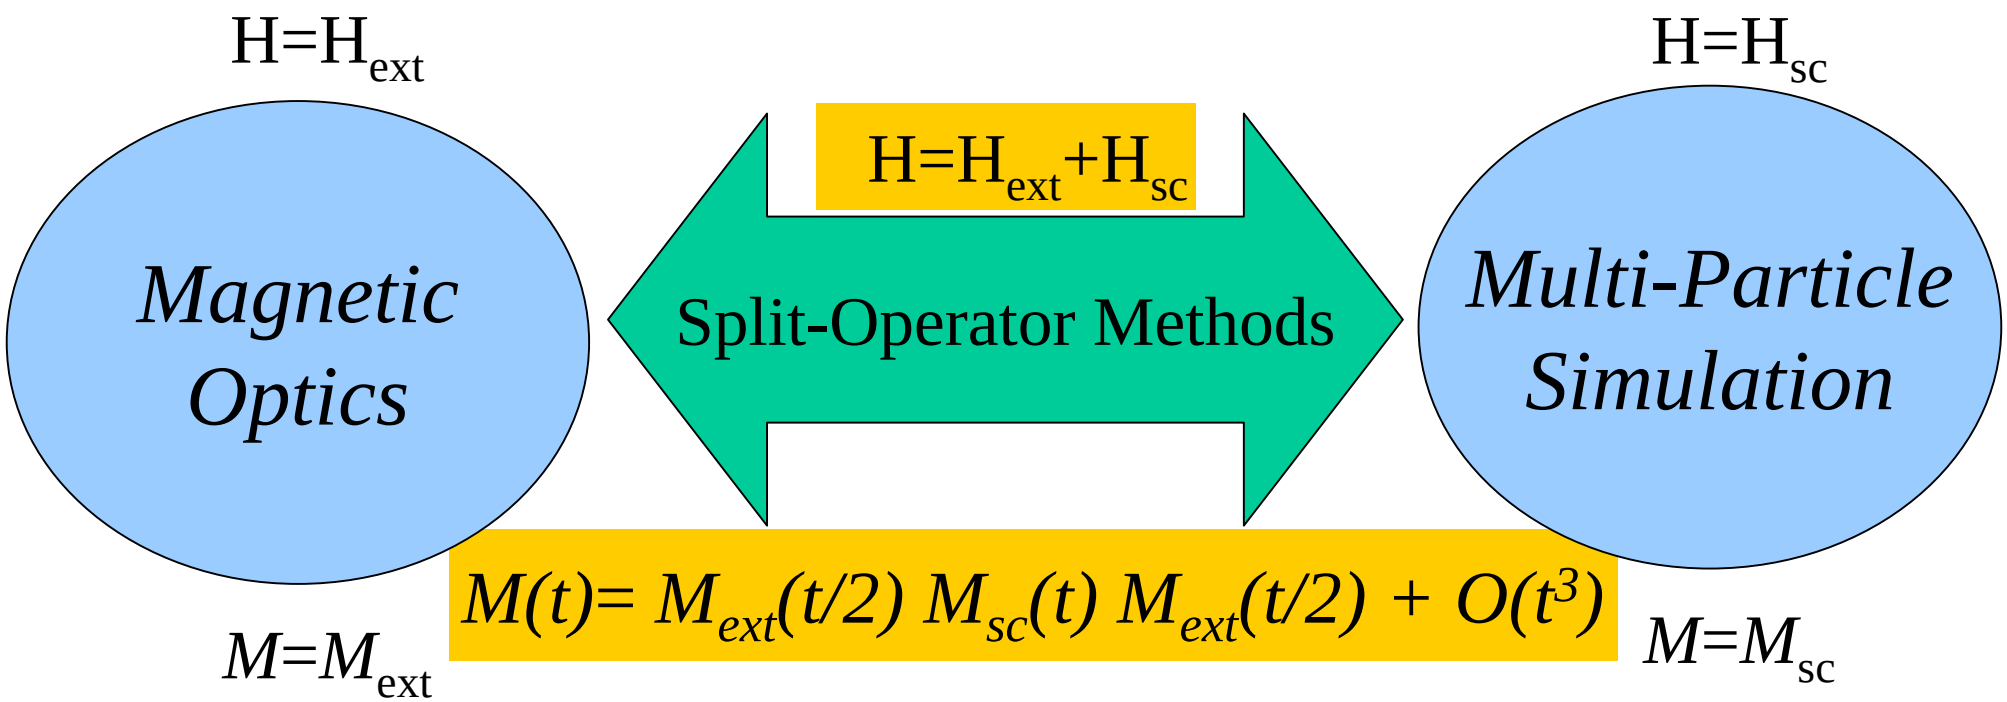

$$M(t)= M_{ext}(t/2)\, M_{sc}(t)\, M_{ext}(t/2) + O(t^3)$$

(arbitrary order possible via Yoshida)

IMPACT: Success of Dynamic Load Balancing

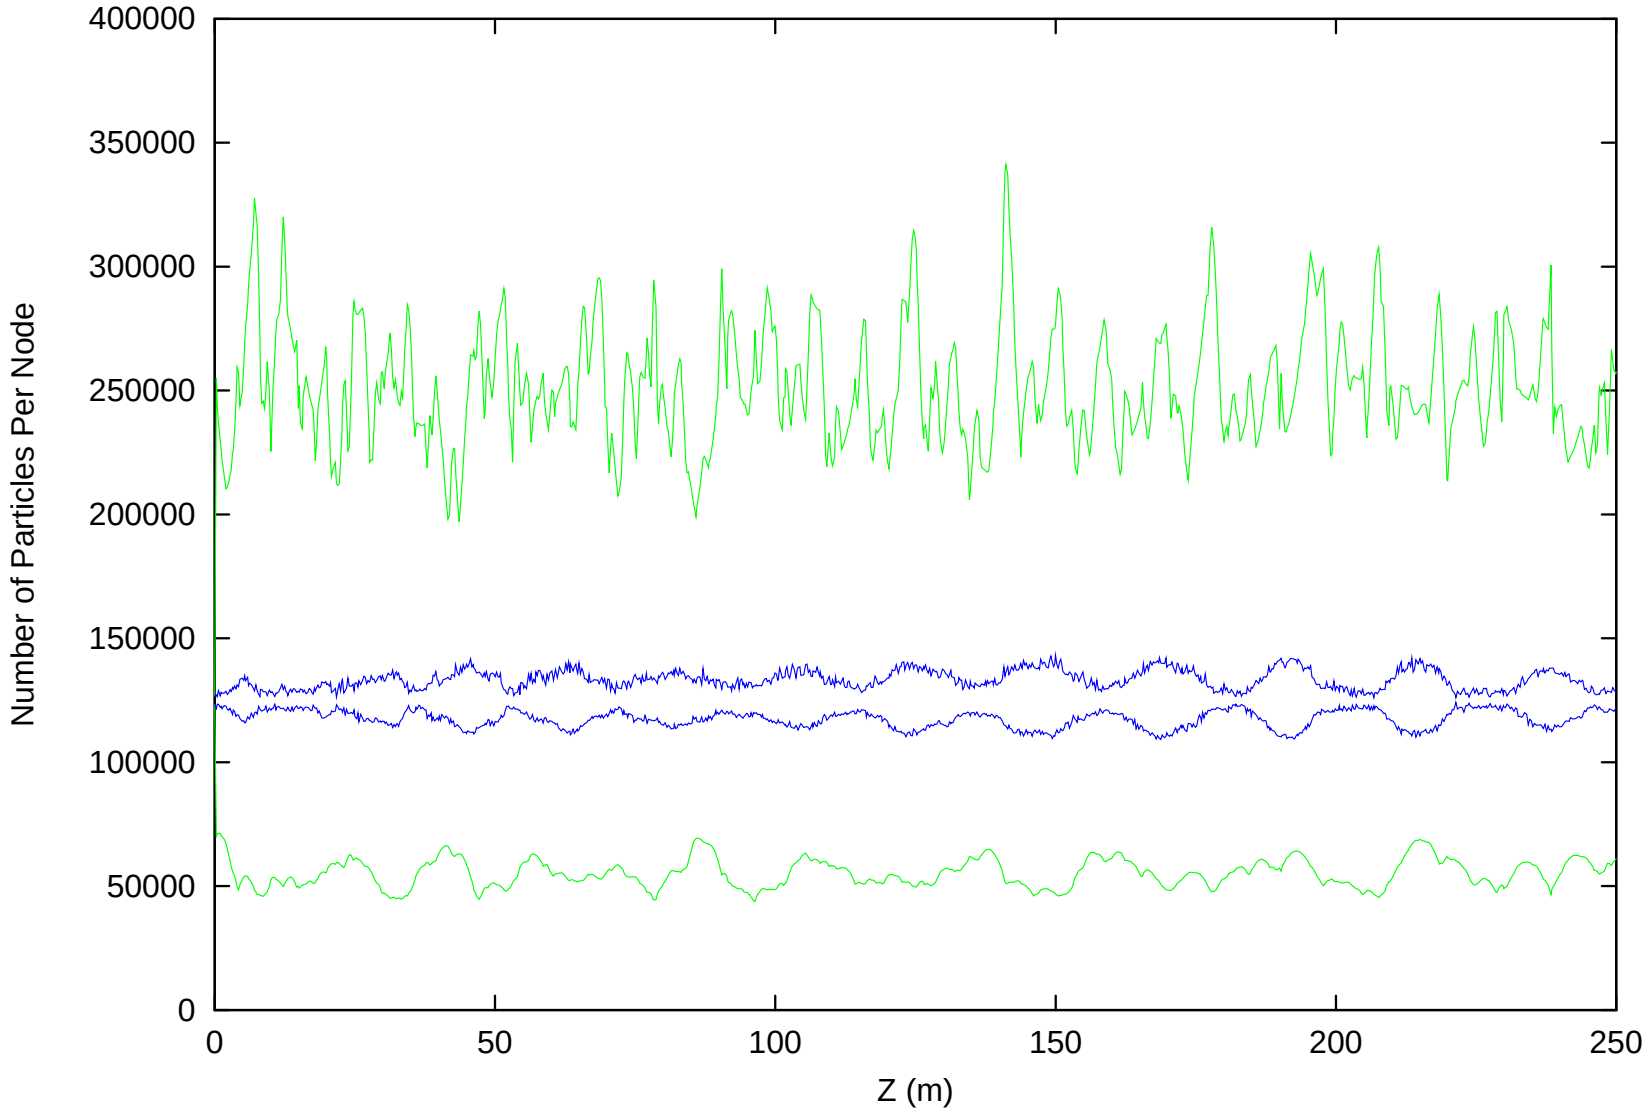

Summary of IMPACT Results for SNS Linac

- Largest, most detailed linac simulations ever performed
  - Physical model used 400 distinct accelerating structures
  - Computational model used 800M particles on a 512^3 grid
- 3 order-of-magnitude increase in simulation capability
  - Comparison w/ legacy PC code:
    - 100x larger simulations performed in 1/10 the time using 256 processors

Scaling of Ion Tracking Code on SLAC PC Cluster

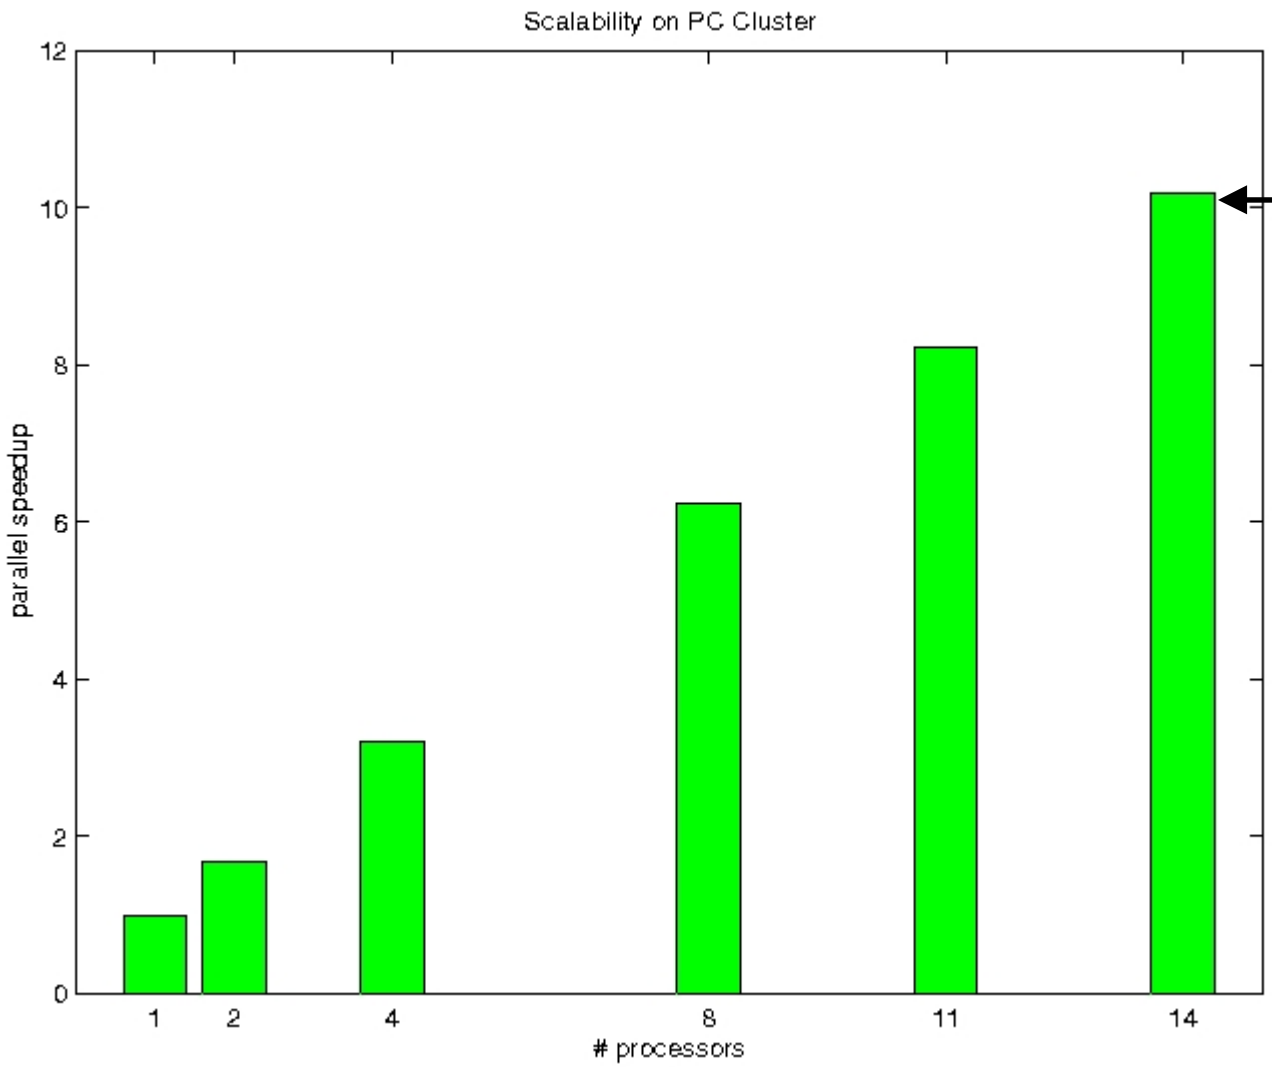

3-hour single processor job requires 18 min on 14 processors

10x speedup with 72% parallel efficiency

PSR Simulations using SAMBA: Comparison w/ Experiment

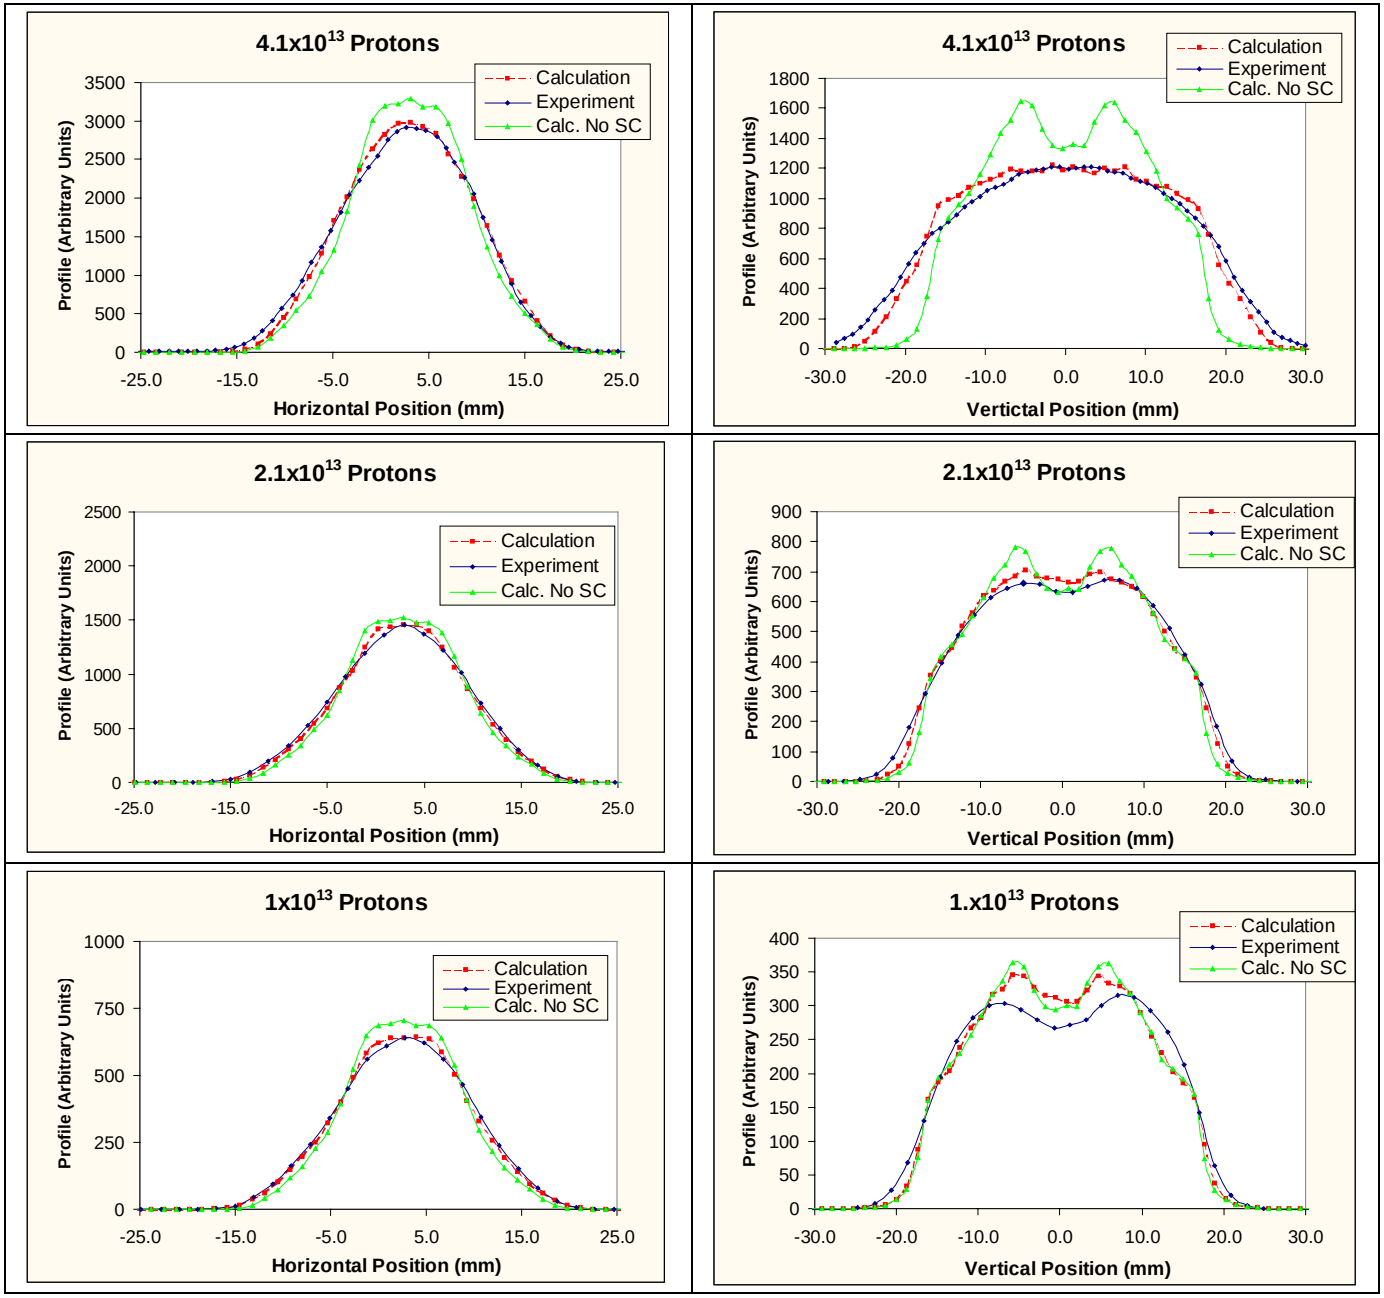

ORBIT: Parallel Simulation on ORNL Cluster

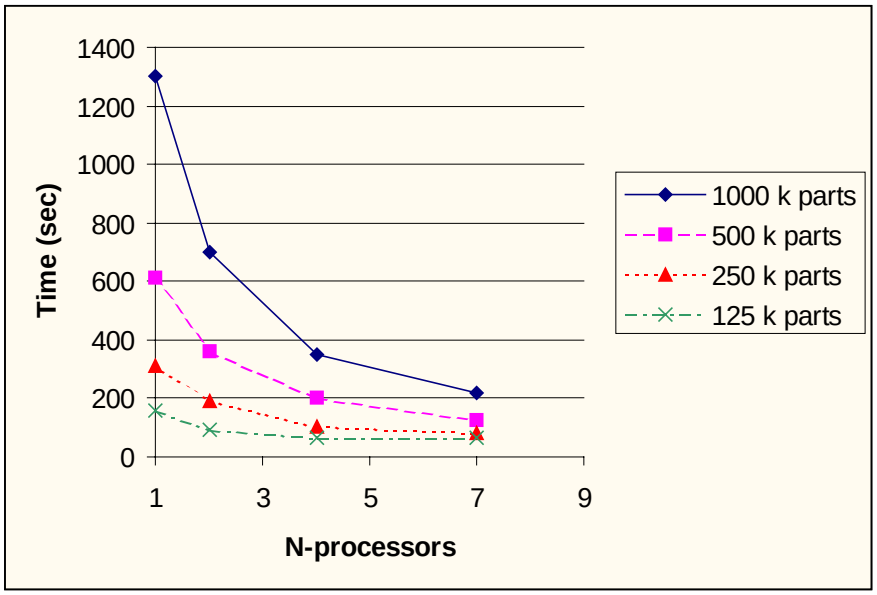

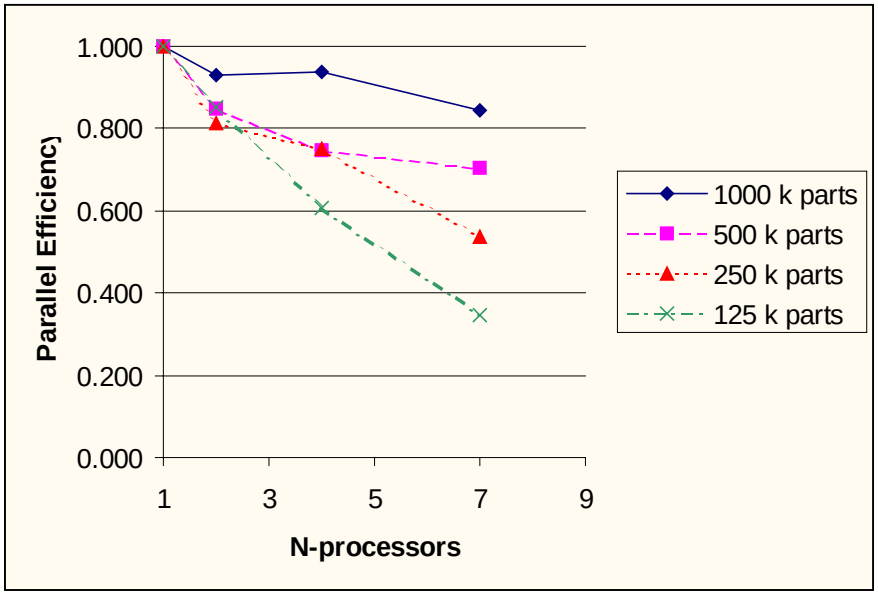

Timings and efficiencies for single turn using 128x128 grid and 480 steps/turn

```
//////////////////////////////////////////
// Add a non-accelerating RF Cavity //
//////////////////////////////////////////
  Real tFactor;
  Integer nRFHarms = 1;
  RealVector volts(nRFHarms),
    harmNum(nRFHarms), RFPhase(nRFHarms);
  harmNum(1) = 1; RFPhase(1) = 0.;

// make a new interpreted routine:
  Void PSRVolts()
  {
   tFactor = time/0.825;
   if(tFactor > 1.) tFactor = 1.;
   volts(1) = 8. + 9. * tFactor;
  }
//   call a compiled routine:
  addRampedRFCavity("RF 1", 75, nRFHarms,
      volts, harmNum, RFPhase, PSRVolts);

// change a "compiled" variable:
     nLongBins = 64;
```

Portion of an input script for ORBIT. Note mixing of interpreted and compiled code.

Parameterized solid model

– Analyze cavity with tuner using program Ω3

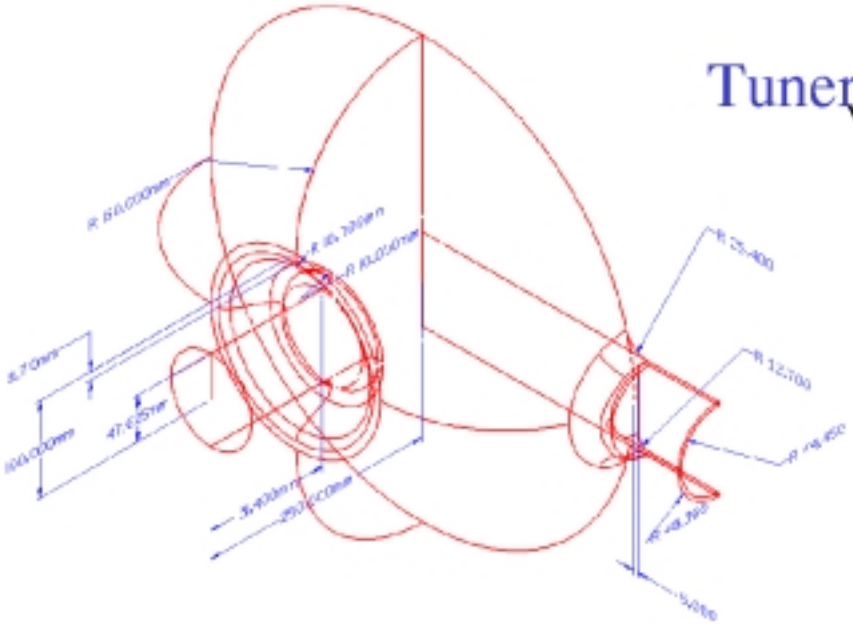

Engineering drawing

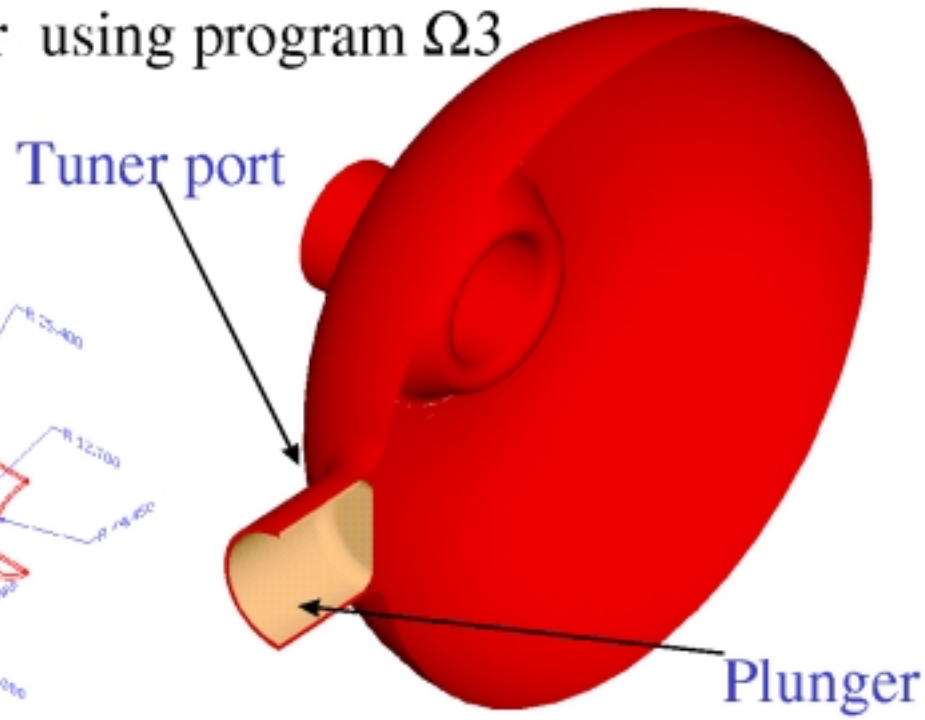

Parametric Solid model

- Interface to a variational solid modeler EMS/Intergraph
- Can resolve necessary features (e.g., small distance between the tuner port and the plunger)
- Easy to change plunger position

Mesh Generation for Omega3P and Tau3P Simulations

Size of meshes limited by workstation memory - <u>Parallel Mesh Generation</u>

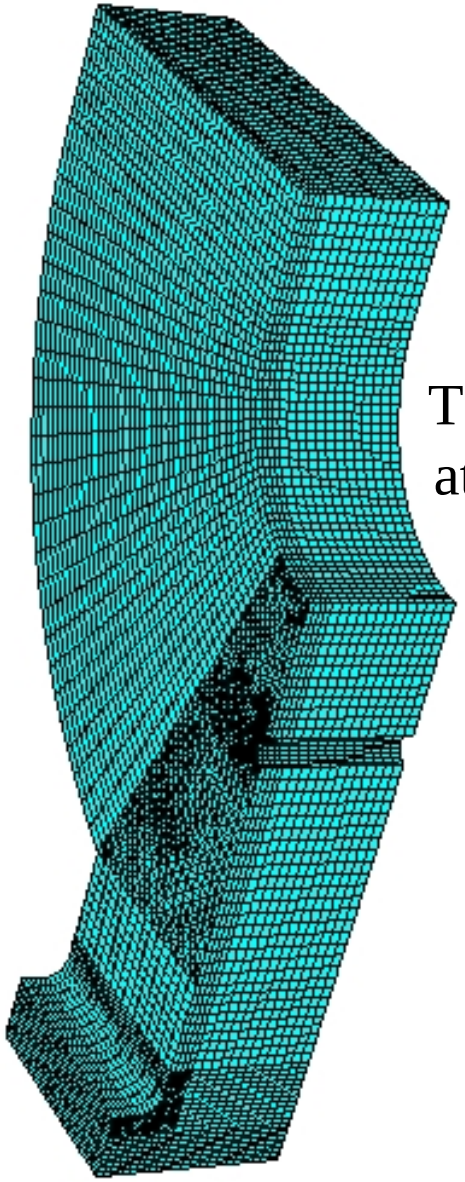

T-Junction
at X-Band

One octant of 1.5 DDS Cells
(500,000 elements)

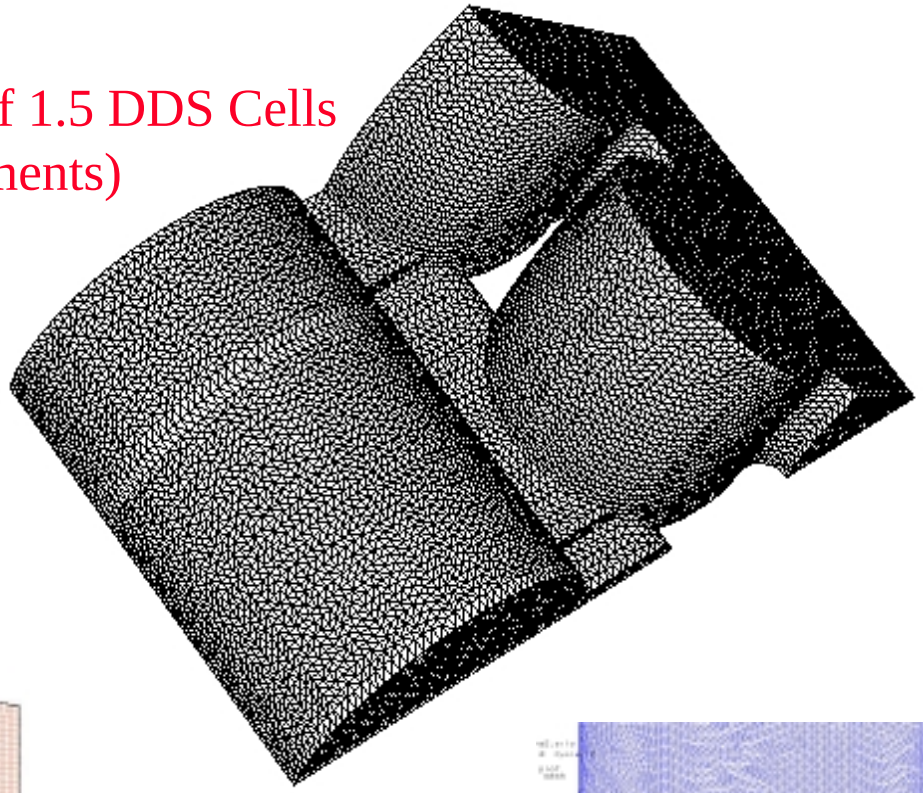

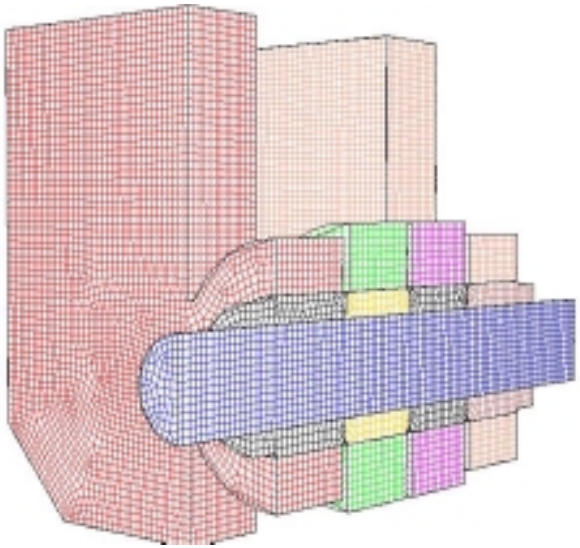

Input Coupler for NLC

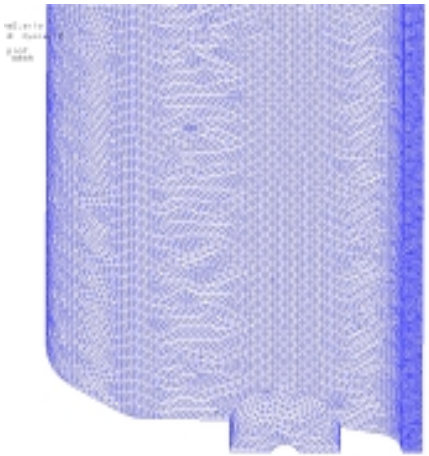

RFQ Cavity for SNS

Domain Decomposition

Partitioning to optimize load balance & minimize communication

Omega3P

Octant of 1.5 DDS cells in 16 domains

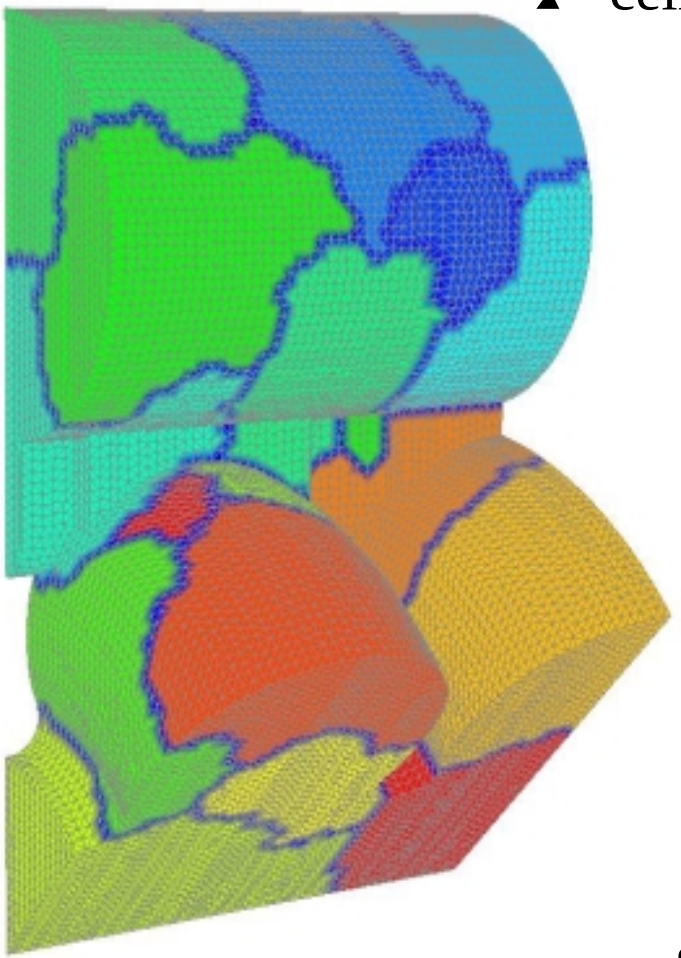

Tau3P

Quarter of RF choke in 4 domains

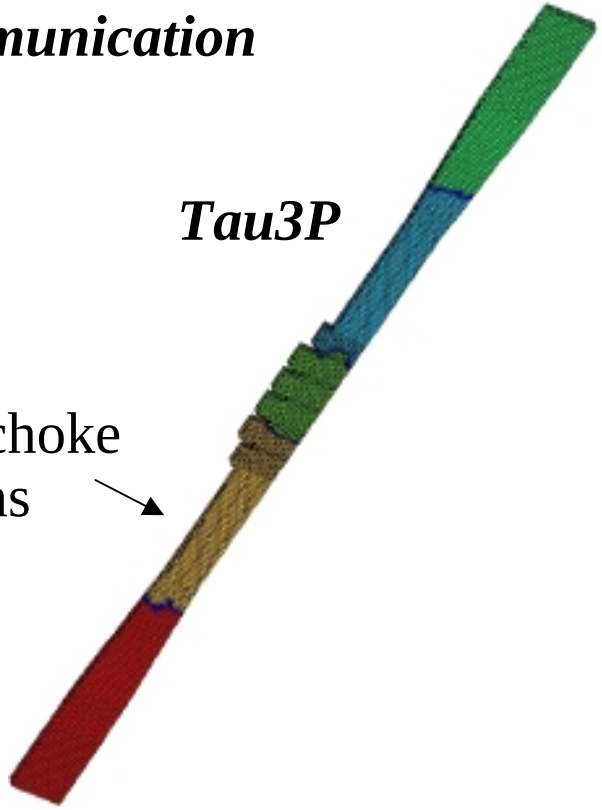

Omega3P

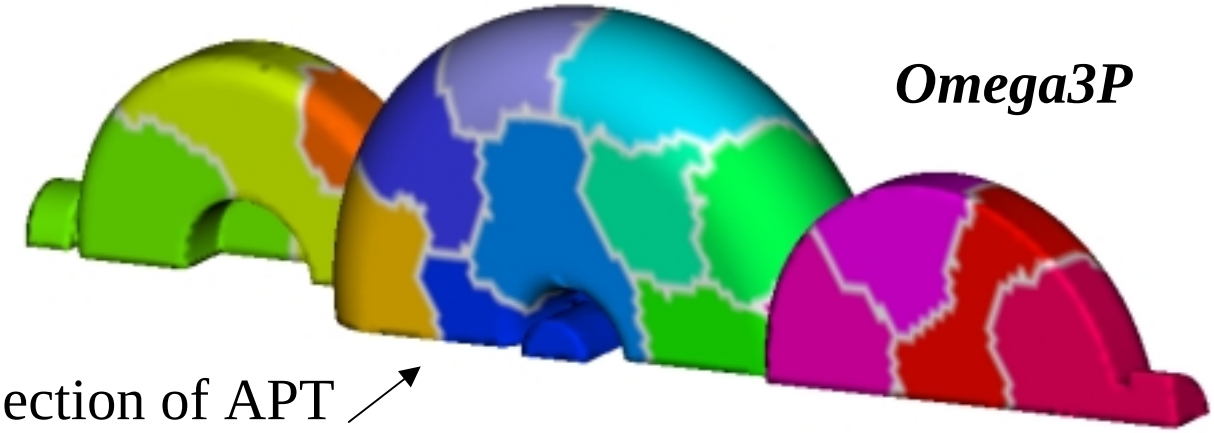

Section of APT side-coupled linac

Parallel Fields Solver

Omega3P *convergence on DDS example with 500K DOF's*

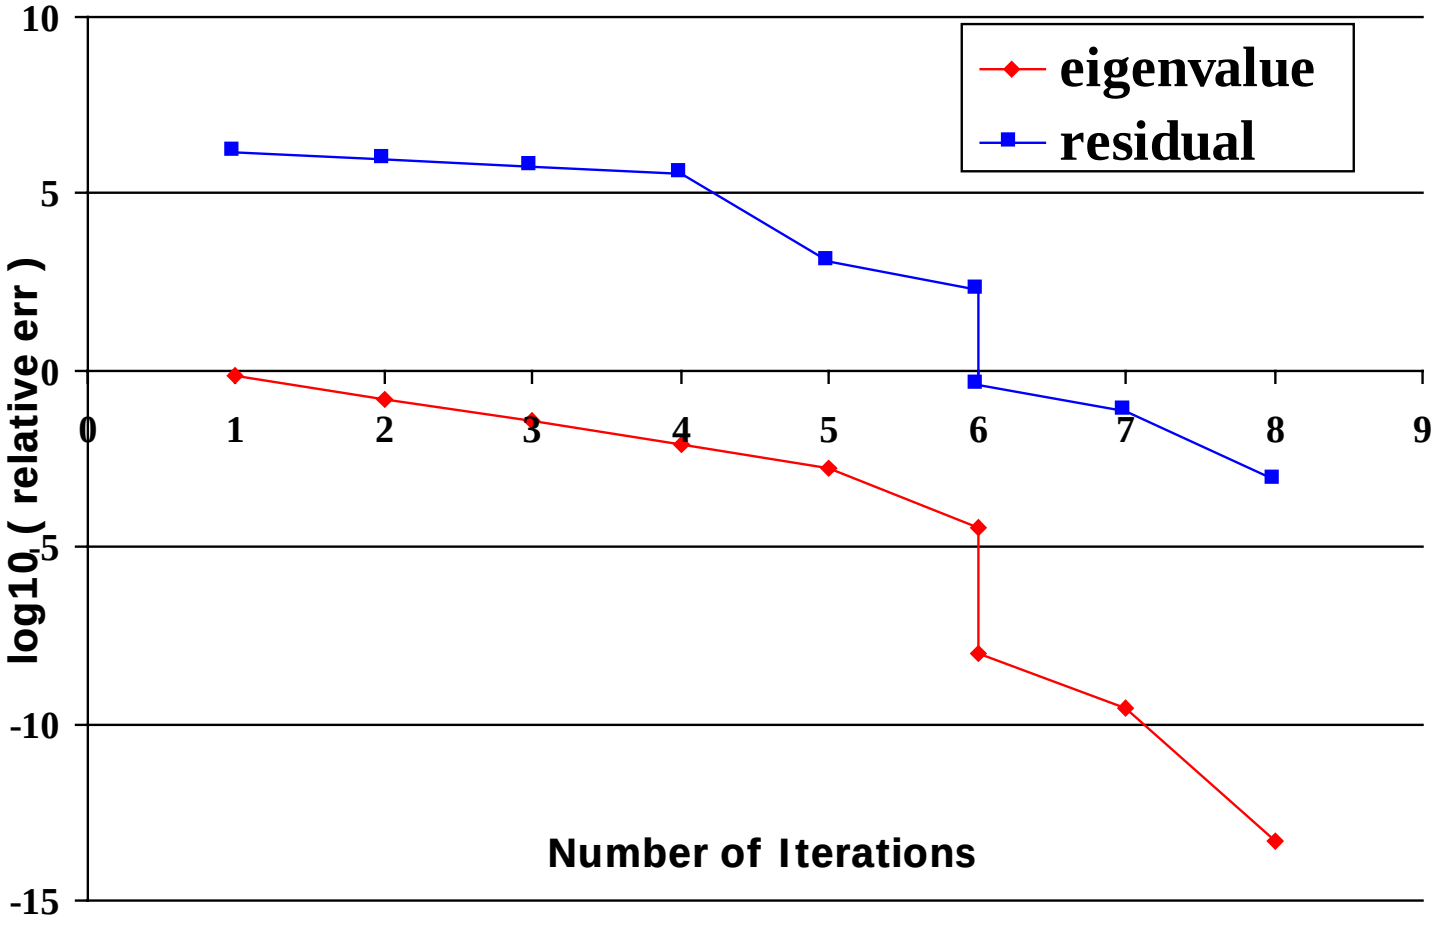

Example: PEP-II Cavity Simulation

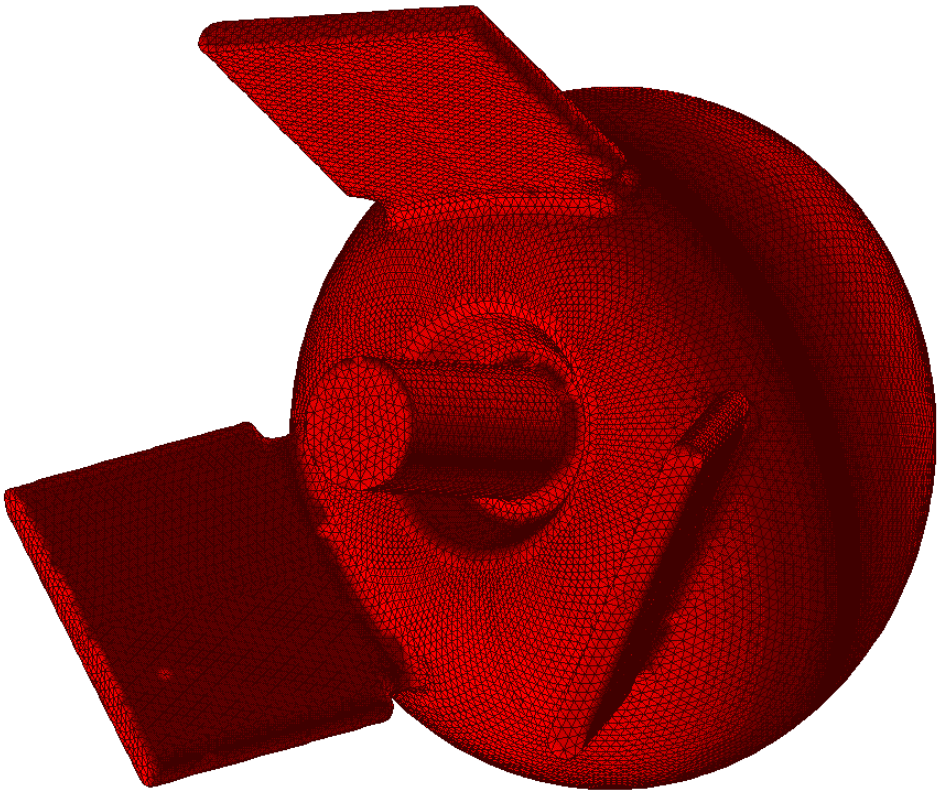

Computational Mesh

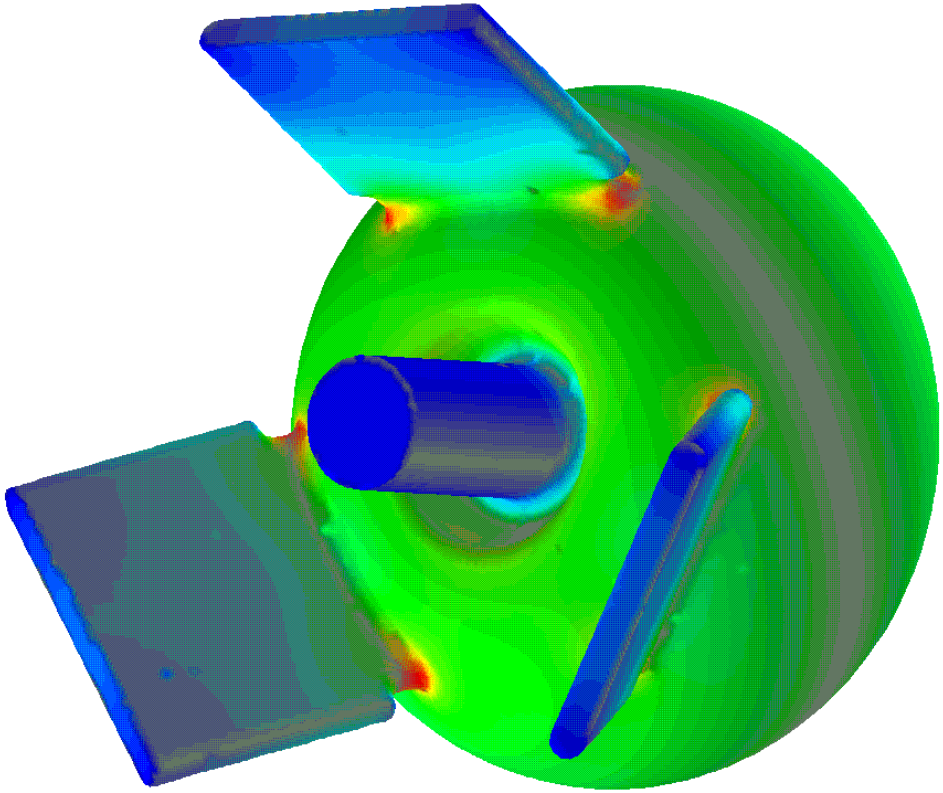

Wall Loss

LARGE SCALE VISUALIZATION

Why do we need large-scale visualization?

- **Huge** data sets
  - e.g. 0.1 - 1 billion particles, 5 - 50 GByte/timestep, TBytes per simulation
- Visualization is needed to:
  - provide physical picture of charge distribution
  - show time-evolution of the beam in phase space
  - explore very low density regions (beam halo)

High performance computing and visualization are needed to *reduce, render* & *explore* the data, and *provide guidance* & *validation* in the design process

Volume Rendering of APT Simulation Results

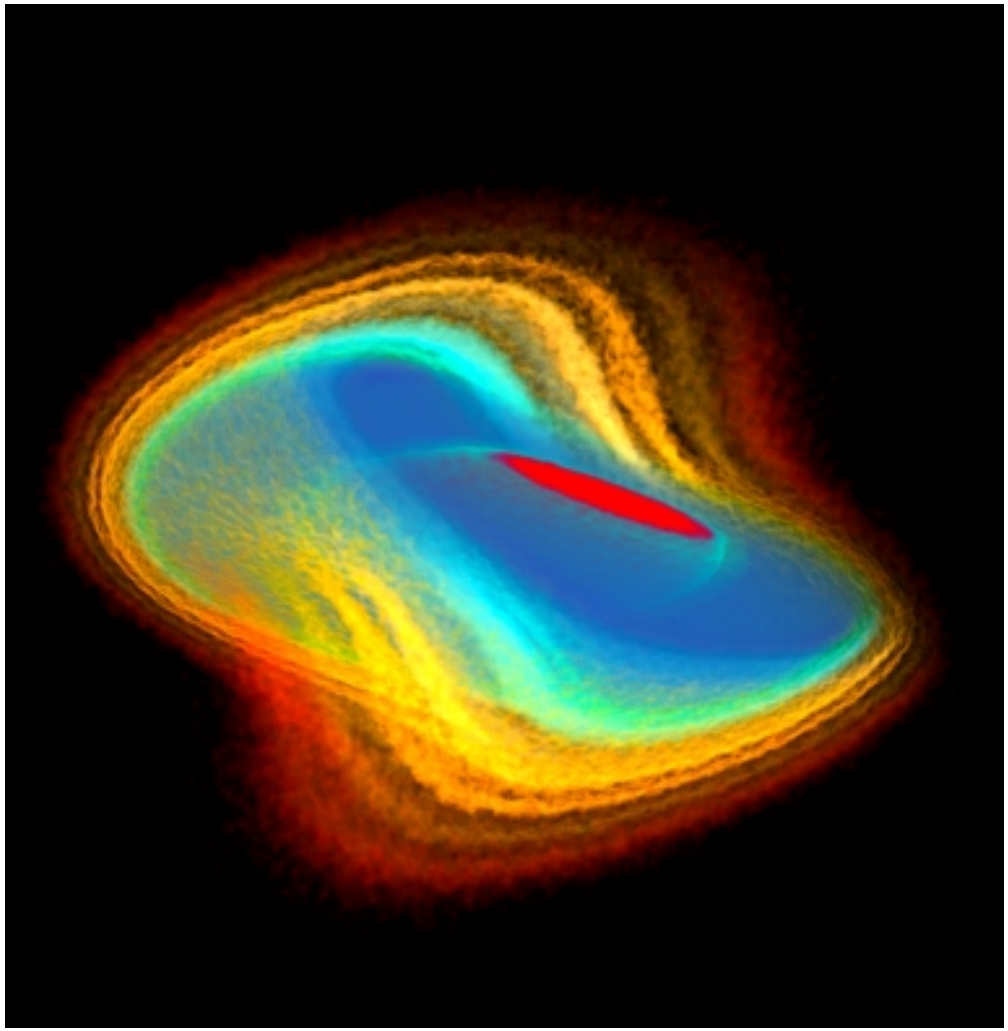

x-p_x-z space

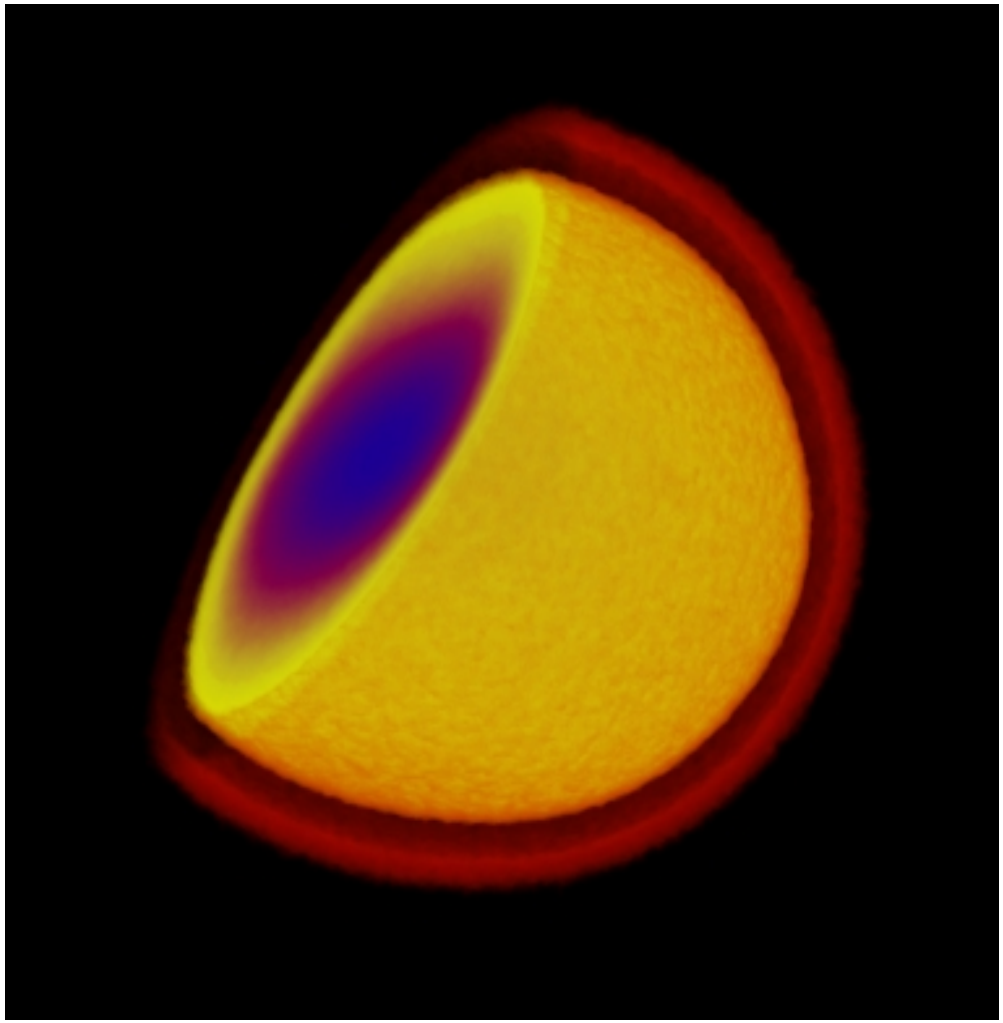

x-y-z space

High Performance Volume Rendering

- By equally subdividing a volume and assigning each subset to a pipe we can:
  - Support larger data sets -- texture memory scales with each pipe.
  - Maintain rendering performance -- fill rate scales with each pipe.
  - Disadvantage -- pipeline architecture causes latency to scale with each pipe.

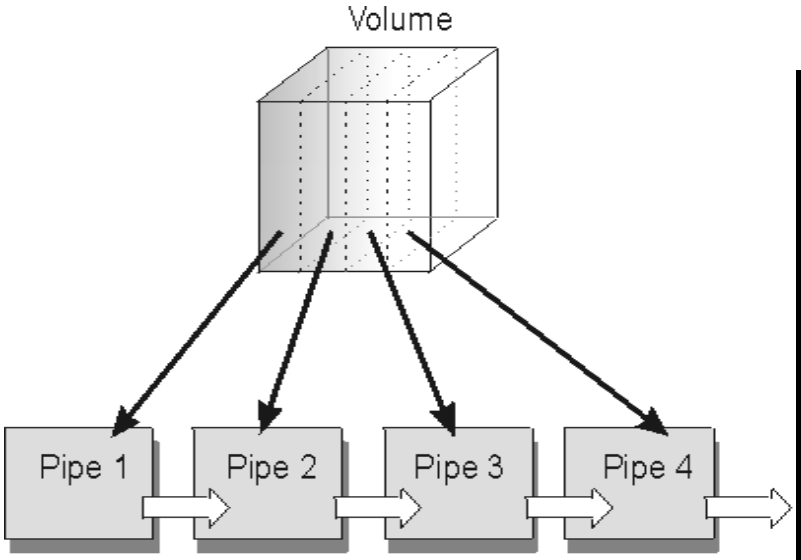

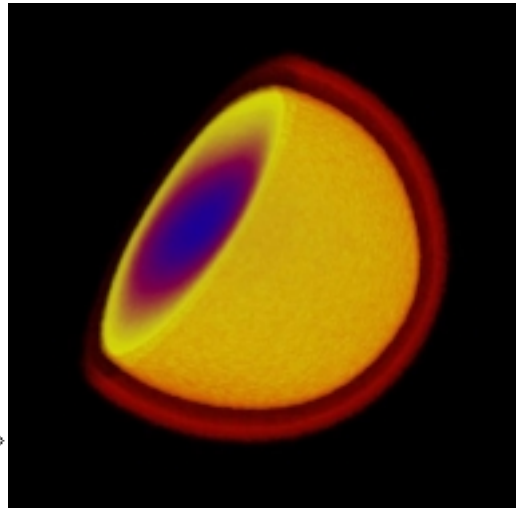

Final Image

Early work done in collaboration with SGI.

Volume Rendering of SNS Simulation Results

Transparency to see inside volume

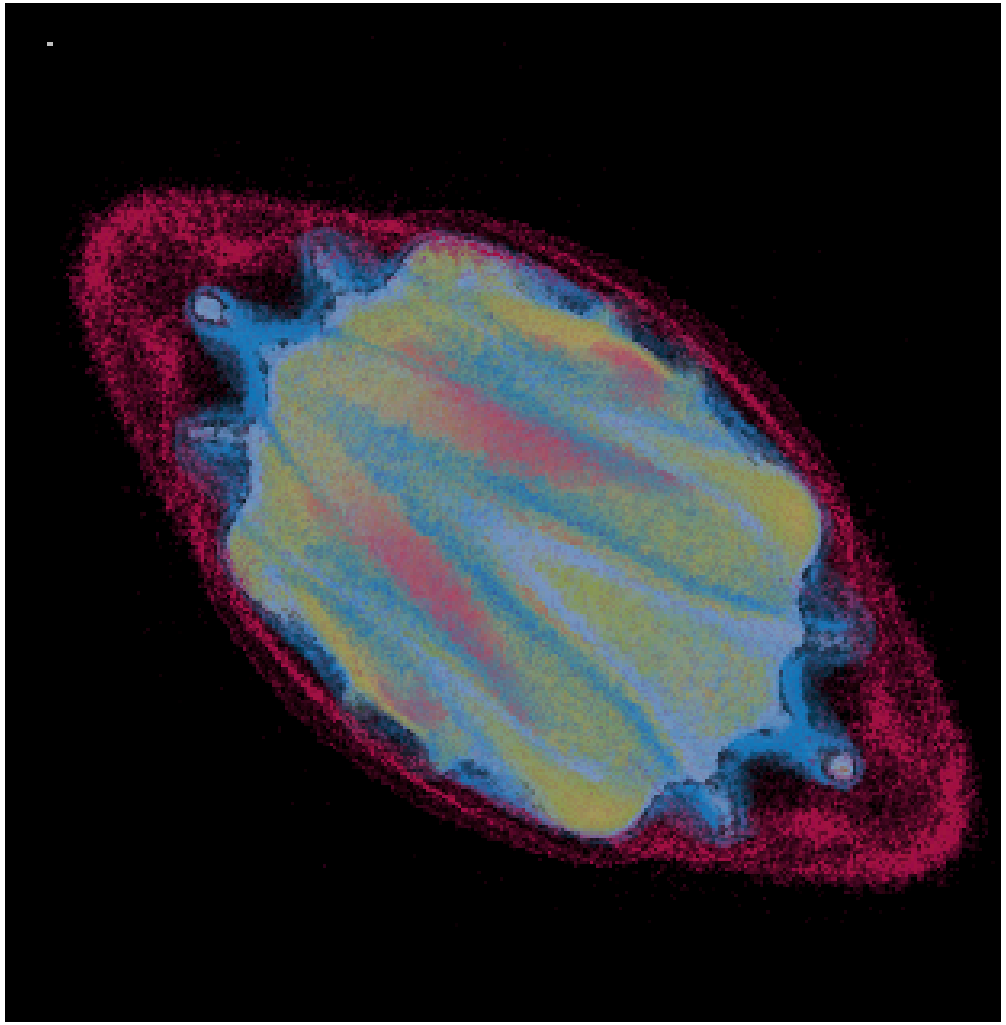

x-p_x-z space

Opacity to accentuate halo

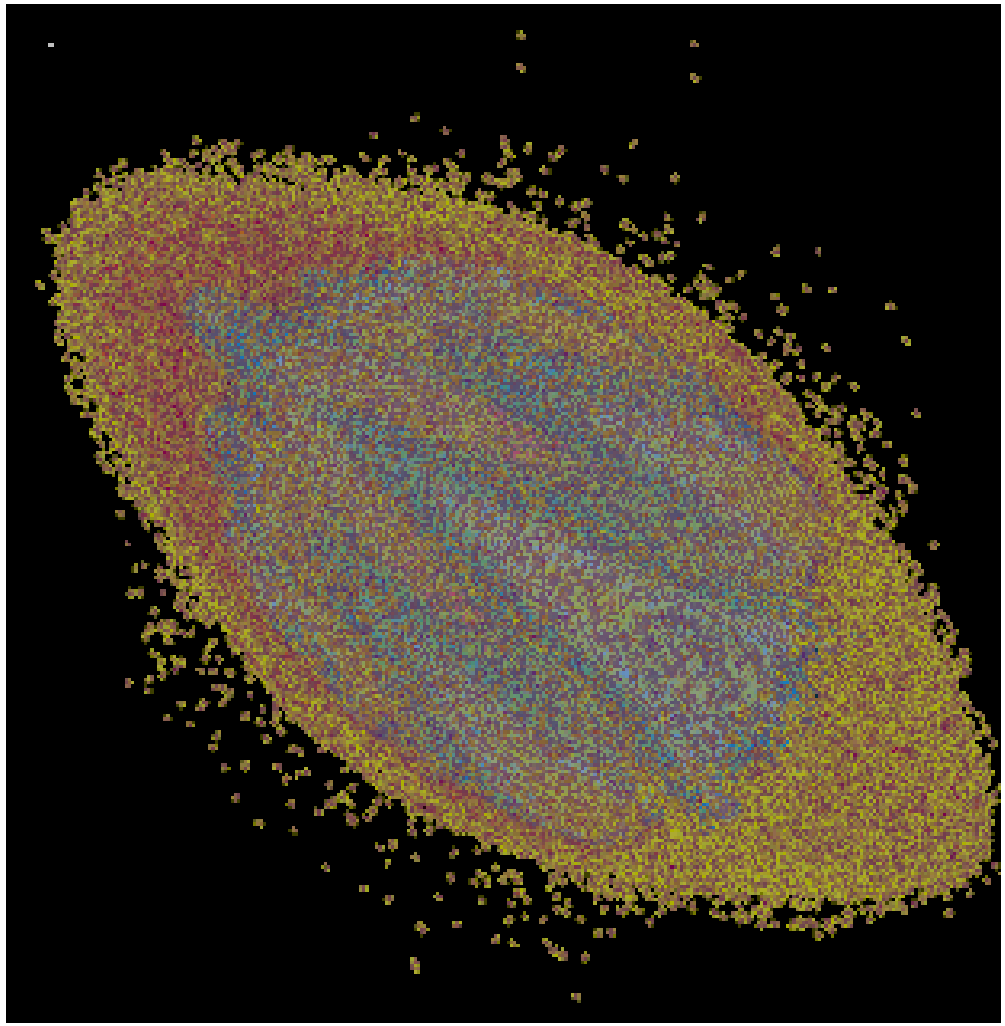

x-p_x-z space

Volume Rendering of SNS Simulation Results

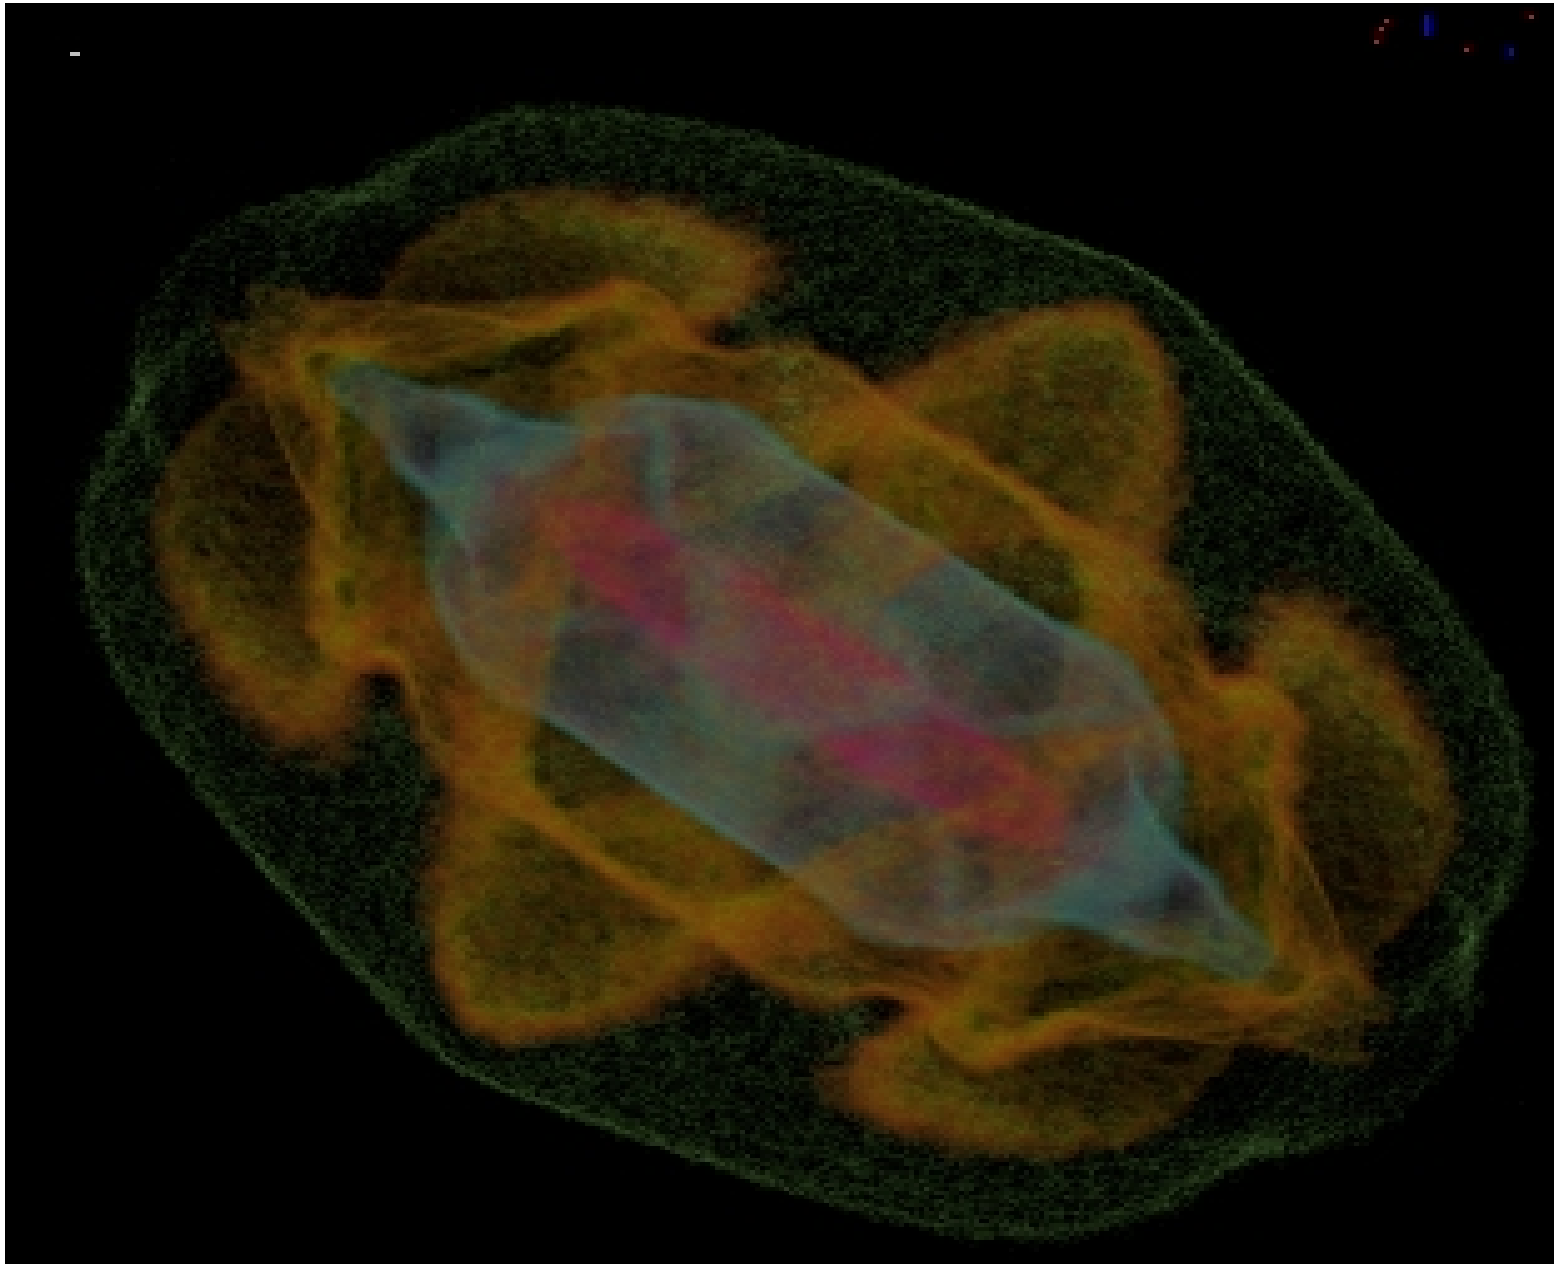

x-p_x-z space. Alpha & color maps set to show 4 isosurfaces

What is needed in a Next-Gen Beam Dynamics Code?

- Split-Operator based (high-order optics + self-fields)
- Modern computer science & software engineering concepts & tools
- Run on single- or multi-processor systems
- Choice of field solvers: structured grid, unstructured grid, tree, hybrid
- Image charges
- Canonical variables
- *Flexible* scaling to make canonical variables dimensionless
- Reference trajectory computed “on the fly” via numerical integration
- Wakes
- Collisions
- Multiple species
- Multiple beams
- Beam loading
-