Large Scale Computing and Storage Requirements for High Energy Physics

Accelerator Physics

P. Spentzouris, Fermilab

SciDAC
Scientific Discovery through Advanced Computing

Motivation

- Accelerators enable many important applications, both in basic research and applied sciences
- Different machine attributes are emphasized for different applications
  - Different particle beams and operation principles
  - Different energies and intensities
- Accelerator science and technology objectives for all applications
  - Achieve higher energy and intensity, faster and cheaper machine design, more reliable operation

↳ a wide spectrum of requirements for very complex instruments. Assisting their design and operation requires an equally complex set of computational tools.

High Energy Physics Priorities

- High energy frontier
  - Use high-energy colliders to discover new particles and directly probe the properties of nature.
    - FNAL Tevatron, CERN LHC, design of future lepton collider
- High intensity frontier
  - Use intense beams to uncover the elusive properties of neutrinos and observe rare processes that probe physics beyond the Standard Model.
    - Future high intensity proton driver at FNAL: Project-X

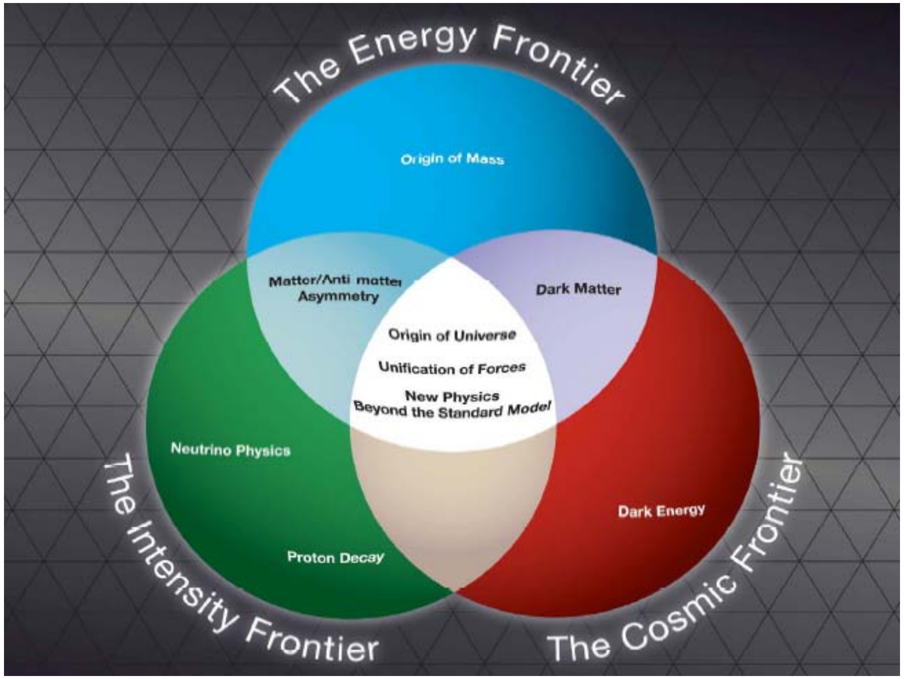

http://www.er.doe.gov/hep/HEPAP/reports/P5_Report%2006022008.pdf

SciDAC Scientific Discovery through Advanced Computing

ComPASS

- The SciDAC2 ComPASS project is developing HPC accelerator modeling tools for
  - Multi-physics, multi-scale for beam dynamics; “*virtual accelerator*”
  - Thermal, mechanical, and electromagnetic; “*virtual prototyping*”
  - Supporting and guiding R&D for new high-gradient acceleration techniques; “*Advanced Accelerator modeling*”

https://compass.fnal.gov/

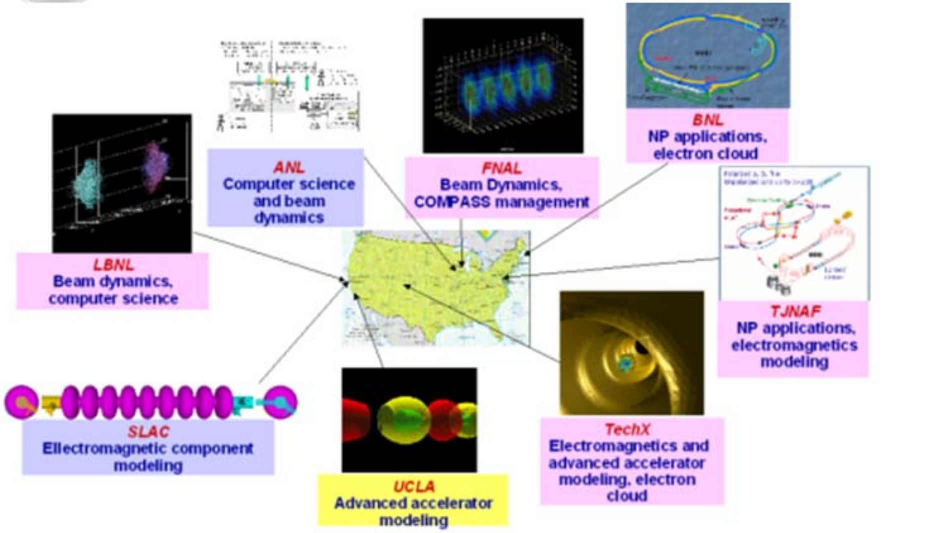

ComPASS focus areas

- Different areas call for a variety of approaches resulting to wide spectrum of computational requirements
  - Beam Dynamics *(this talk)*: electrostatic PIC
  - Electromagnetics: finite difference and finite element, time and frequency domain
  - Advanced Accelerators: full EM PIC and reduced PIC
- Application focus includes “discovery” and “design optimization” areas (need for both single and multi-physics capability development)
  - Results in both very-large-scale and very-large-volume medium scale computations
- The project utilizes both NERSC and LCF resources
  - NERSC HEP allocation: 2.2M core-hours (SciDAC), 4.2M core-hours (3 INCITE projects)

Accelerator design: multi-scale, multi-physics problem

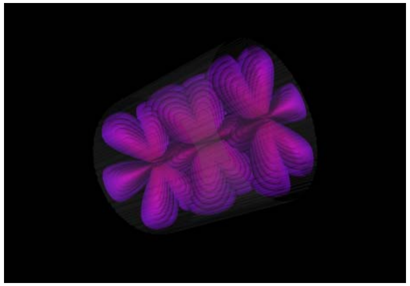

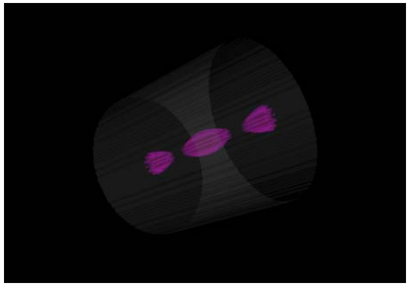

- Wide range of scales:
  - accelerator complex (10^3m) → EM wavelength (10^2-10 m) → component (10-1 m) → particle bunch (10^{-3} m) → PIC (10^{-12})
  - Simulations need to connect scales and allow inclusion of multiple physics effects at each level
  - Many parameters to understand and optimize, simulations have impact on engineering, and models can be tested with data

Beam Dynamics modeling tools example: Synergia

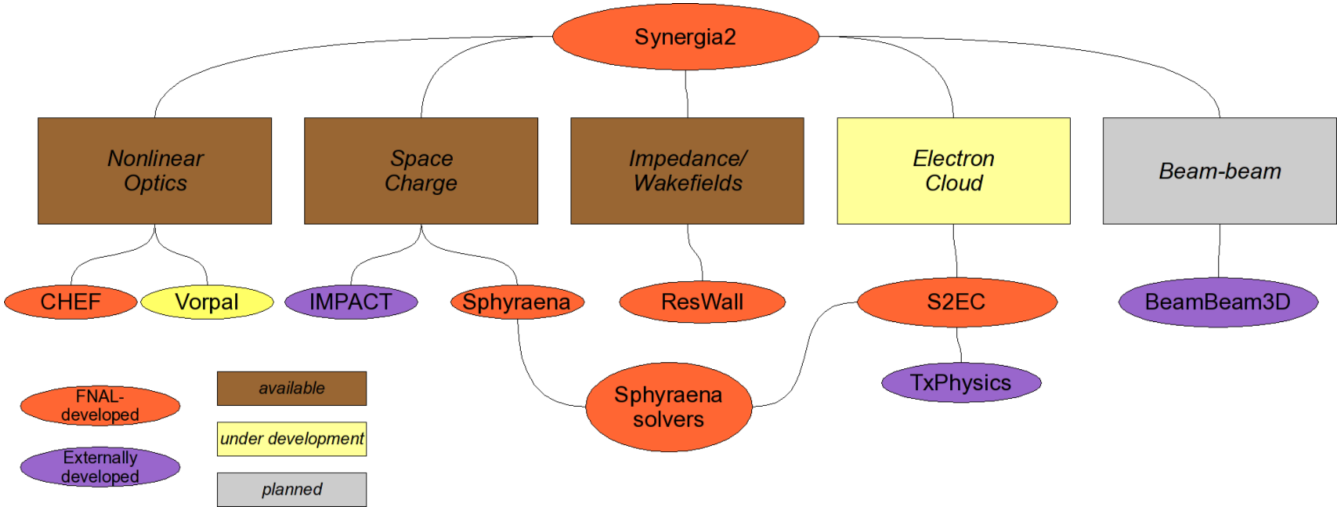

- Beam Dynamics framework with fully 3D PIC capabilities
  - Utilizes both native and external physics modules/algorithms
  - Includes space-charge & impedance (single and multi-bunch)
  - Includes non-linear optics to arbitrary order
- Runs on desktops, clusters and supercomputers, but requires support for shared libraries!
- Flexible framework allows for fully dynamic simulations including ramping, feedback, etc

SciDAC Scientific Discovery through Advanced Computing

ComPASS software

ComPASS

ComPASS
CET, Inst.

Facility

Accelerator modeling framework: geometry, physics models

Accelerator applications

Analysis, visualization infrastructure

Physics algorithms

CS, math algorithms, solvers

Scientific Software Libraries

Accelerator Science Applications: Tevatron

- Model beam-beam & impedance effects with 36 on 36 bunches at the Tevatron
    - Results used to optimize chromaticity settings at the squeeze, resulting in reduced losses

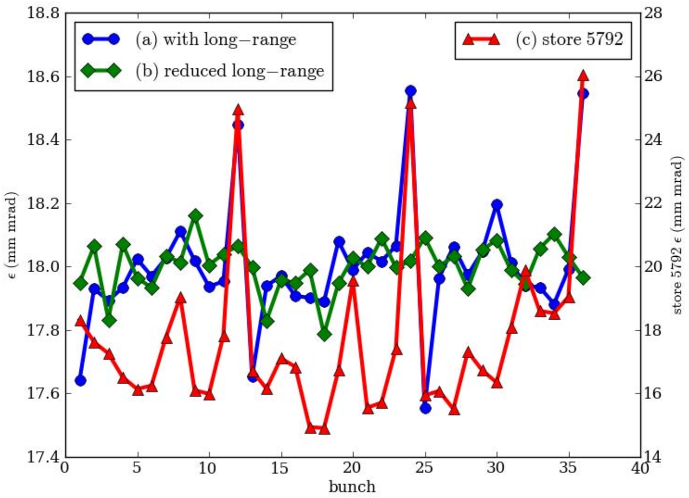

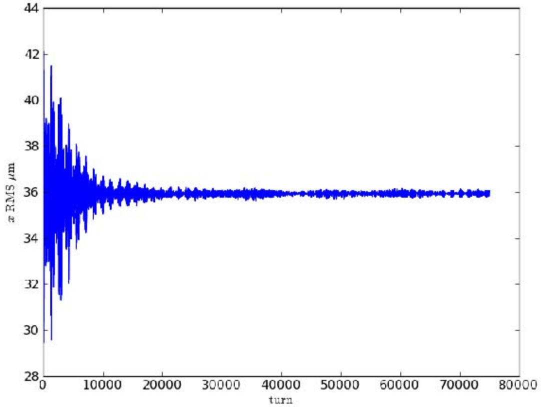

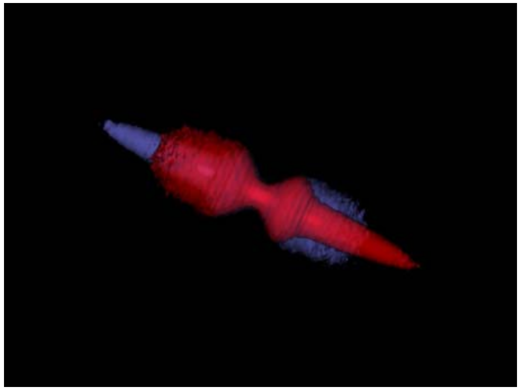

Tevatron, continued

- Success! Simulations support the use of lower chromaticity during injection!
  - Stern et al, “Fully 3D multiple beam dynamics processes simulation for the Fermilab Tevatron”, submitted to PRSTAB
- Just this one application required ***5M core-hours of running*** on intrepid @ ALCF (model validation & parameter optimization)!

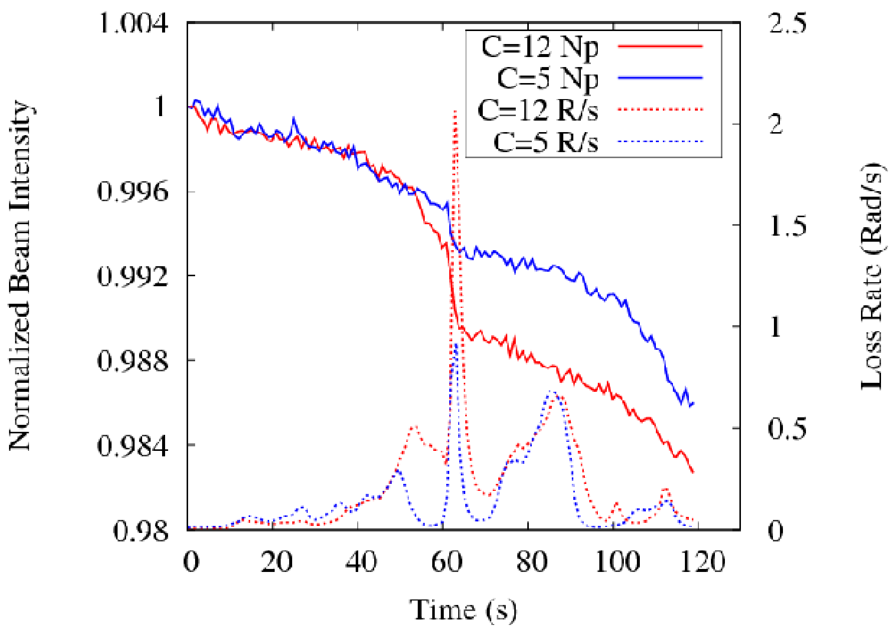

Synergia applications: Mu2e extraction design

- Model resonant extraction including space-charge at the FNAL Debuncher : Mu2e requires 10^5 more particles than current operations!
  - Optimize tune and resonant extraction parameters to minimize losses

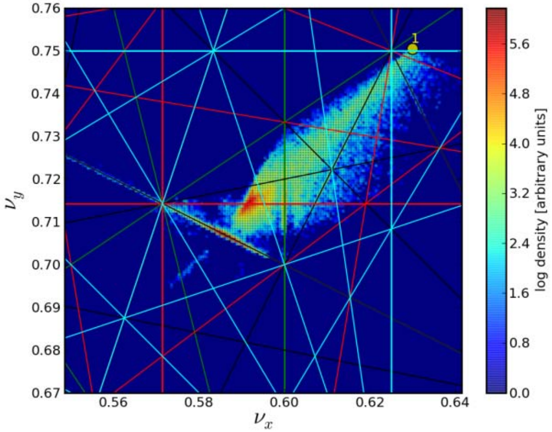

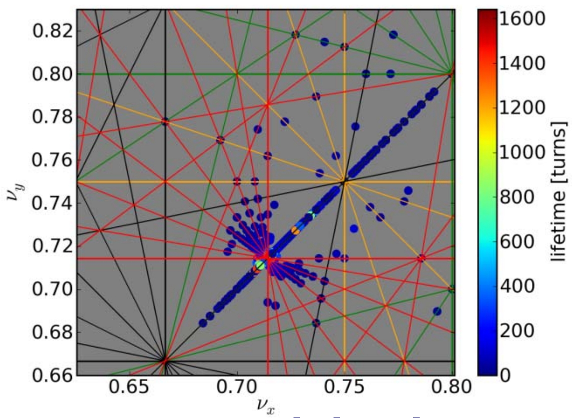

Particles lost (detailed tracking)

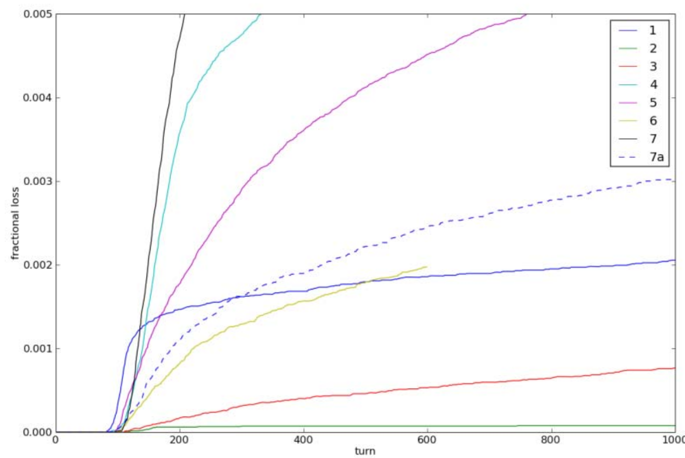

Losses vs time

Mu2e design parameter scans

Each point requires 3k core-hours, modeling only a small fraction of the cycle (extraction). No optimization algorithm or workflow is used (it is just a parameter scan)

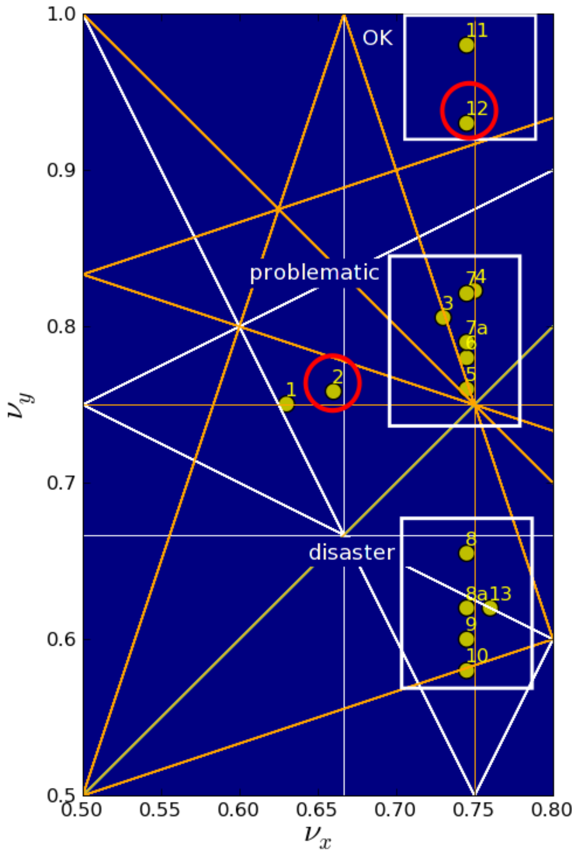

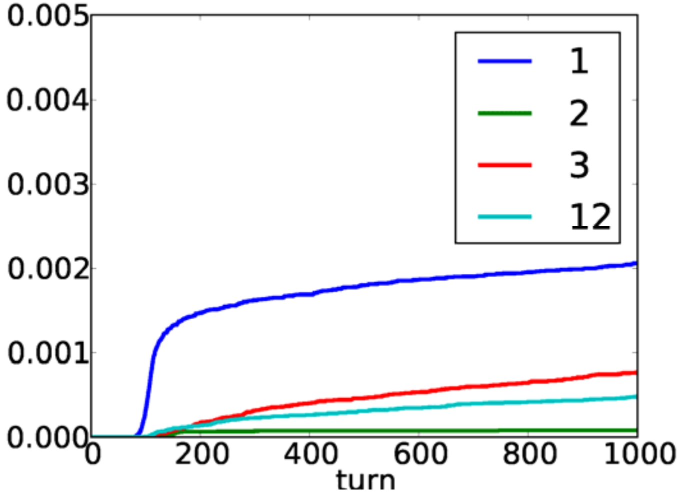

Accelerator Science Applications: Project-X

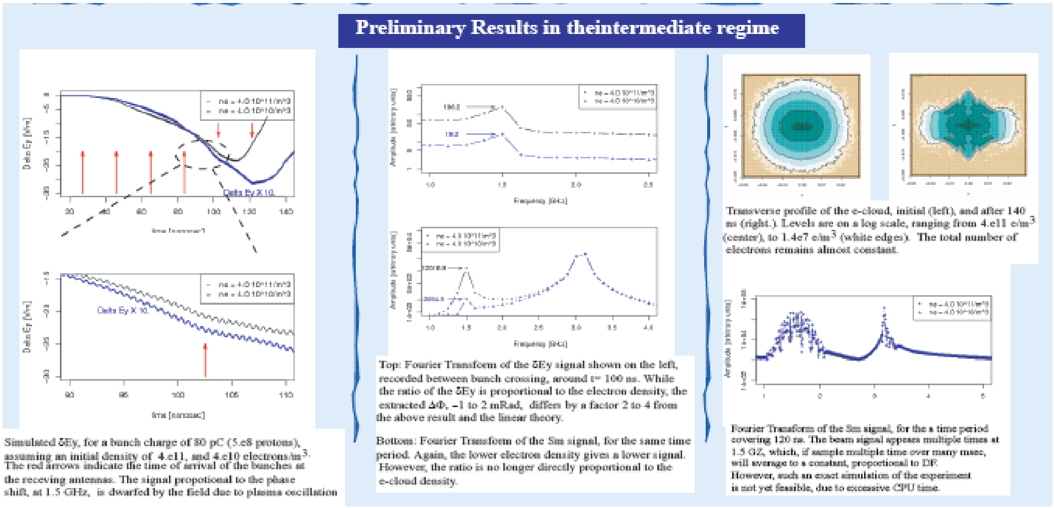

- Simulation of microwave transmission properties through an e-cloud at the Fermilab Main Injector (Project-X parameters)
  - Essential for utilizing e-cloud detector based on microwave propagation
  - Multi-scale, multi-physics problem
- Begin modeling space charge effects and mitigation techniques for Main Injector with Project-X beam parameters
- In the near future, we will need to model the dynamics of both space-charge and e-cloud

SciDAC
Scientific Discovery through Advanced Computing

Scaling for beam dynamics applications

- For a typical problem size, a single physics application will scale to a few thousand processors
- Multi-physics applications increase the size, but require more sophisticated workflows (example: bunch to bunch impedance plus any single-bunch multi-particle effect)
- Parameter optimization requires at least ensemble runs, preferably a parallel optimization scheme

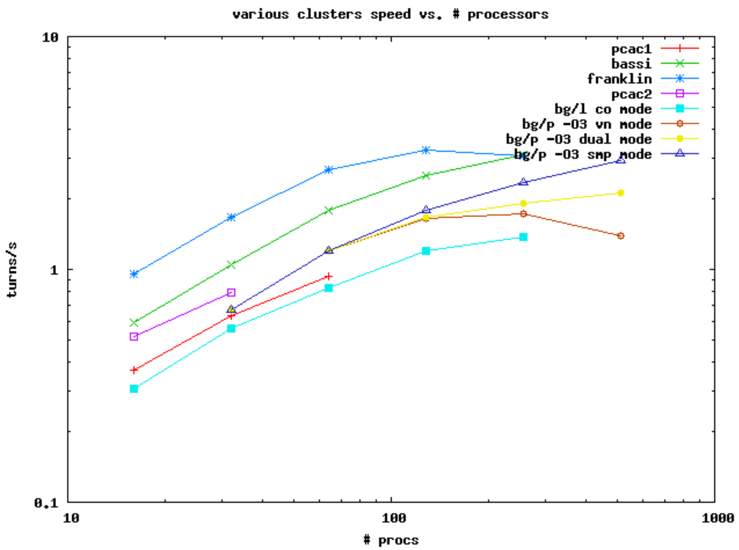

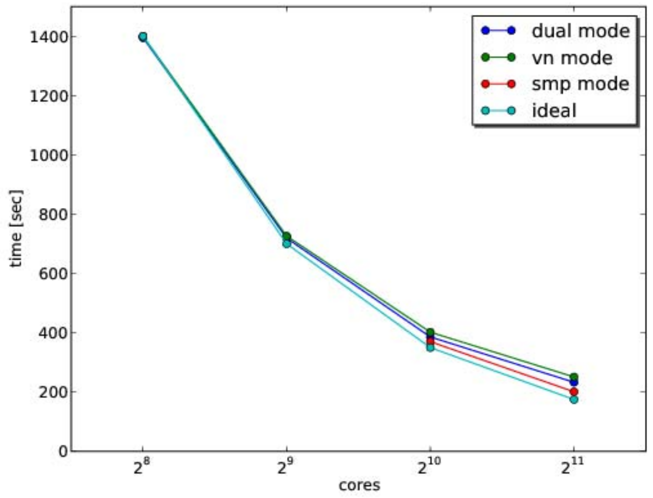

Other comments

- Stability, availability and reliability are essential. Example: the Tevatron optimization was performed on intrepid although it started on franklin because of the instability issues of the machine at the time the application was developed
- Better queue organization for algorithmic development and test jobs will be useful (this will reduce dependence to local development resources)
- The level of sophistication of simulation data analysis is increasing as is increasing the compute time for this step of a beam dynamics application (see example of tracking lost particles)
  - Current visualization and analysis strategy not homogeneous, but it is desirable that the heavy load is shifted on hpc facilities.
- Support of fully functional OS is essential for the deployment of multi-physics framework applications. The benefit of the additional physics capabilities is worth the performance loss.

HPC methods

- Main codes are Synergia, ML/Impact (both multi-physics frameworks), BeamBeam3D and NIMZOVICH (single purpose codes). The codes utilize electrostatic particle-in-cell model with structured grids, with different strategies and solver implementations:
  - Depending on the physics of the problem, the codes might use domain decomposition, particle decomposition, or hybrid decomposition. There may be communication of particle data, grid data, or both. Particle movement between Poisson solves may be slight or large, hence, some codes use a particle manager and some do not
  - Solvers. Our codes utilize spectral based, finite difference based, and hybrid descritisations with FFT and multi-grid based solvers.
    - Depending on the type of algorithm, we have different grid size limitations (memory): typical large grid 1024^3 for the first scheme (both particles and grids distributed), 256^3 for the second (only grid). This results to requirements for 10 to 100M macroparticles, depending on the application
    - Parallelization is implemented using MPI

New technologies

- We are currently starting our research program on understanding how to effectively utilize GPUs. Our applications (for machine design and optimization) have two main components: particle tracking and field soves. Our efforts to date have demonstrated that we can do efficient tracking with high-order-optics on GPUs. We are investigating field solves on GPUs and hybrid schemes involving a mixture of conventional cpus and GPUs.
  - We will need more information on the architecture of the future machines incorporating GPUs in order to design efficient multi-level parallelism schemes.
  - We will also need some guidance on the development environment (libraries, etc).

ComPASS at NERSC

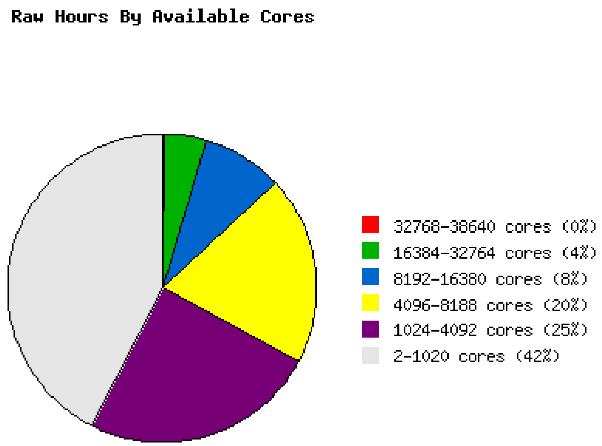

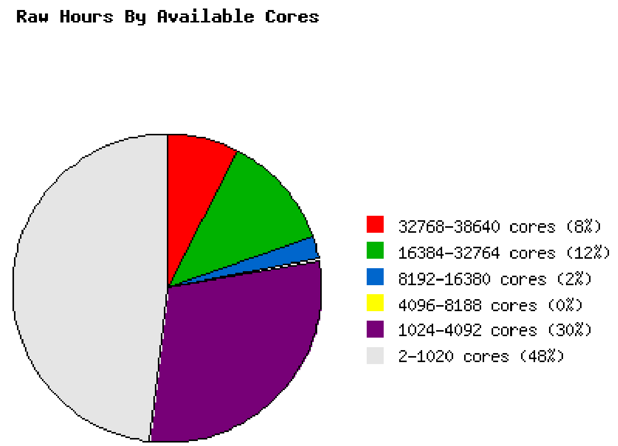

- ComPASS specific performance on franklin (repo m778) for calendar '09 (right), compared with all repos (left): ComPASS large jobs (>30 k cores) at ~10%, compared to 0.3% for all repos.
- High concurrency for average jobs (>8k cores) for all other repos (see next talks)