Fermilab

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

High Performance Computing Activities at Fermilab

James Amundson

Breakout Session 5C: Computing

February 11, 2015

High Performance vs. High Throughput Computing

- Much of the computing in HEP relies on High Throughput Computing (HTC)
  - One task per processor on many processors
  - Trivial to perform in parallel
- This talk is about High Performance Computing (HPC)
  - One large task on many processors
  - Tightly-coupled communications
    - Advanced networking
      - Low latency and high bandwidth
  - Includes Linux clusters with specialized networking and supercomputers

Cooperative Work in High Performance Computing (HPC)

- Lab facilities
  - HPC clusters
  - Next-generation HPC testbeds
- Lab competencies
  - Scientific
    - Computational Physics on HPC
  - Technical
    - HPC programming and optimization
    - HPC cluster support
- Relevant lab science topics
  - Lattice QCD
  - Cosmology
  - Accelerator Simulation

Much of this work Funded by DOE HEP CompHEP through SciDAC

Shared HPC Facilities

- Lattice QCD
  - Three clusters
    - ~25,000 CPU cores total
    - Two clusters include GPUs
    - Infiniband interconnects
- Accelerator Simulation and Cosmology
  - Two clusters
    - ~2,500 CPU cores total
    - Infiniband interconnects
- Next-generation research
  - Two clusters
    - 72 traditional cores

- Phi cluster
  - 16 Intel Xeon Phi 5110P accelerators
- GPU clusters
  - 2 NVIDIA Tesla Kepler K20m GPUs and 2 K40m GPUs
- Infiniband interconnects

Sharing HPC Facilities

- Facilities are operated by Computing (CS); use is shared between Particle Physics (PPD), Center for Particle Astrophysics (FCPA) and CS.
- All three subjects also utilize leadership class supercomputing facilities, for example
  - ALCF at Argonne (Mira: BlueGene/Q)
    - Resources obtained through the INCITE program.
  - NERSC at LBL (Edison: Cray XC30)
    - Resources obtained through the SciDAC program.

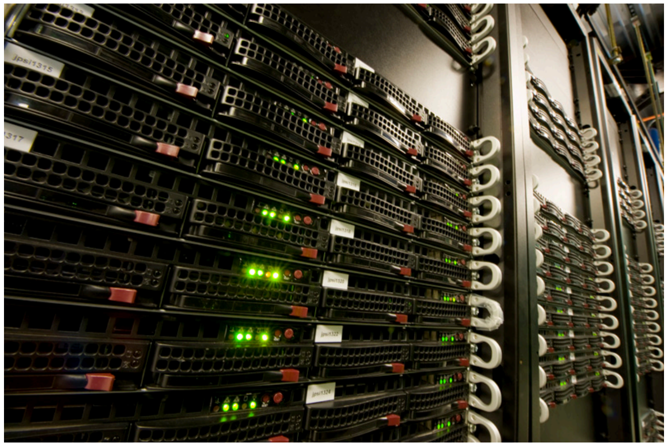

Shared Competencies

- HPC cluster support
  - LQCD group has developed highly specialized expertise in acquisition, deployment and support of HPC Linux clusters.
    - Difficulty of each is easy to underestimate.
      - Exotic hardware and drivers.
  - Expertise is shared with Accelerator Simulation and Cosmology through the support of specialized clusters.
- Computational Physics on HPC
  - Many overlaps between LQCD, Accelerator Simulation and Cosmology
    - Numerical Algorithms
      - Particle-in-cell techniques, linear algebra and spectral methods
    - Parallelization Algorithms
      - Load balancing, data layout, communication avoidance, etc.

Shared Competencies, continued

- HPC programming and optimization
  - Similar technical issues are faced by all HPC efforts at the lab.
  - Groups meet in a regular series of technical seminars ("NEAT Topics").
  - LQCD and Accelerator Simulation submitted SciDAC cross-cutting project.
  - Accelerator Simulation and Cosmology share technical personnel.

Lattice-QCD computing and USQCD

- Lattice-QCD applications cover all science cross-cuts, e.g. hadronic contributions to (g-2)μ, nucleon axial-vector form factor for CCQE X-section, quark masses & αs for Higgs predictions
  - Calculations require supercomputers (provided by DOE's LCFs) and even more flops on medium-scale computing clusters (naturally provided by DOE laboratories)
- Fermilab deploys and operates large computing clusters for U.S. lattice gauge theory community (organized by USQCD Collaboration), which is mostly based at universities.

- USQCD effort over 100 strong
- Hardware funded by DOE HEP & NP offices
- Hosts largest share of USQCD's computing hardware
- Highest user satisfaction among three USQCD facilities
- Fermilab plays leading roles in U.S. lattice effort including:
  - Paul Mackenzie USQCD spokesperson and project PI.
  - Bill Boroski (office of CIO) LQCD project manager.
- Computing and PPD interactions crucial for successful science!

Accelerator Simulation and ComPASS

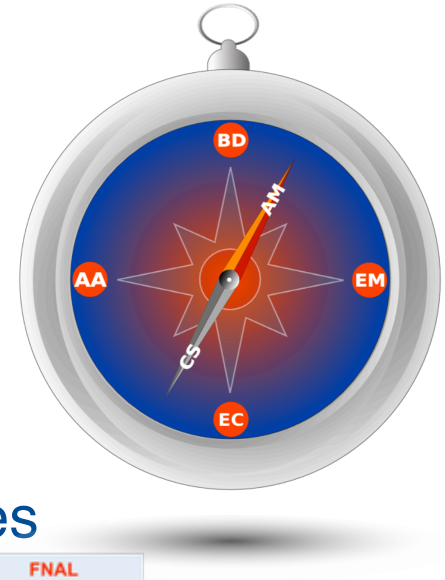

- Community Project for Accelerator Science and Simulation (ComPASS) collaboration
  - Mission is to develop and deploy the state-of-the-art accelerator modeling.
  - DOE CompHEP+ASCR funding through SciDAC
  - Collaboration includes national labs, universities and one company
- Fermilab is the lead institution on ComPASS
  - More discussion of accelerator simulation to follow…

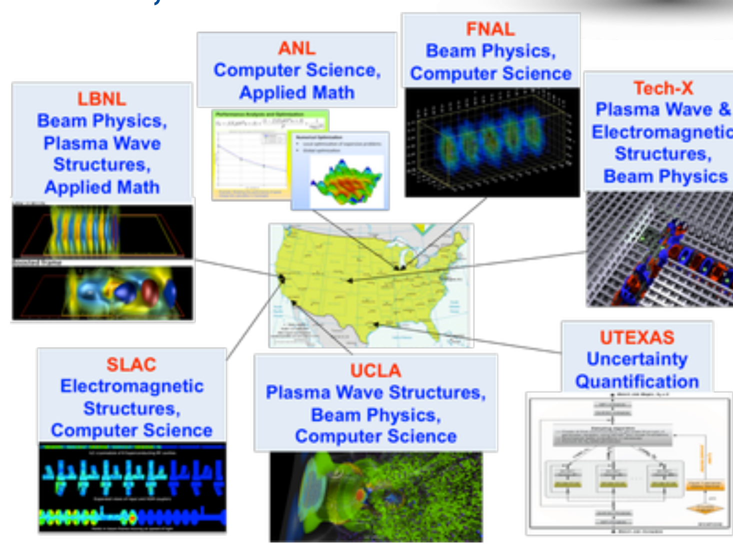

Fermilab

Science on HPC

- The point of these shared facilities and competencies is producing science.
  - Lattice QCD is covered in other cross-cutting sessions.
  - Cosmology is covered in other cross-cutting sessions.

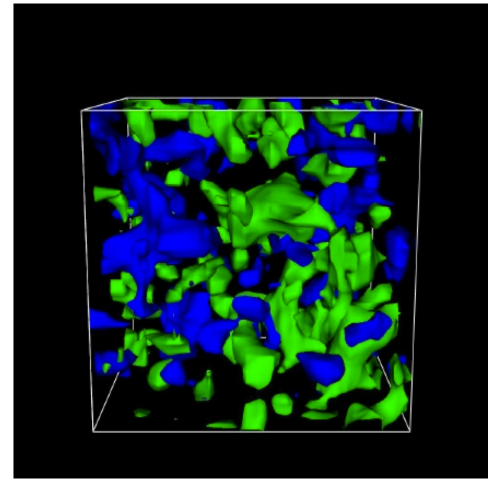

- I will focus on Accelerator Simulation.
  - Accelerator simulation serves HEP by enabling accelerator technology.
    - Directions must be in sync with HEP priorities.
    - Work is done in collaboration with Fermilab's Accelerator Division and other accelerator facilities, especially CERN.

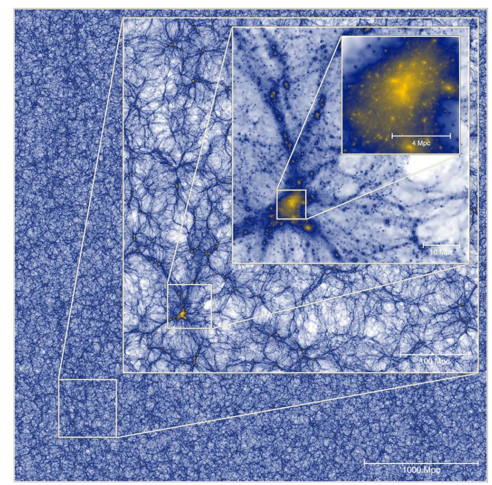

Accelerator Simulation Directions Set by Lab Goals

- We choose simulation topics based on Lab goals, which are ultimately driven by P5 goals.
  - Extend the scientific reach of existing accelerator facilities
    - Simulate existing Booster, Delivery Ring, Recycler and Main Injector.
  - Launch a test facility to enable "transformative" Accelerator Science
    - Simulate the Integrable Optics Test Accelerator (IOTA).
      - IOTA is an experimental machine to explore using intrinsically nonlinear dynamics to overcome intensity limitations inherent to ordinary linear machines.
  - Establish Fermilab as essential contributor to future large accelerators
    - Simulate LHC Injectors for HL-LHC.

Synergia

Beam Dynamics Framework

- Developed at Fermilab with support from SciDAC ComPASS Collaboration (CompHEP+ASCR)
  - Fermilab leads ComPASS
- Advanced capabilities
  - Collective effects
  - High precision through high statistics (many particles)
  - Single- and multi-bunch physics
    - Multi-bunch capability is unique to Synergia
- Runs on everything from laptops to supercomputers

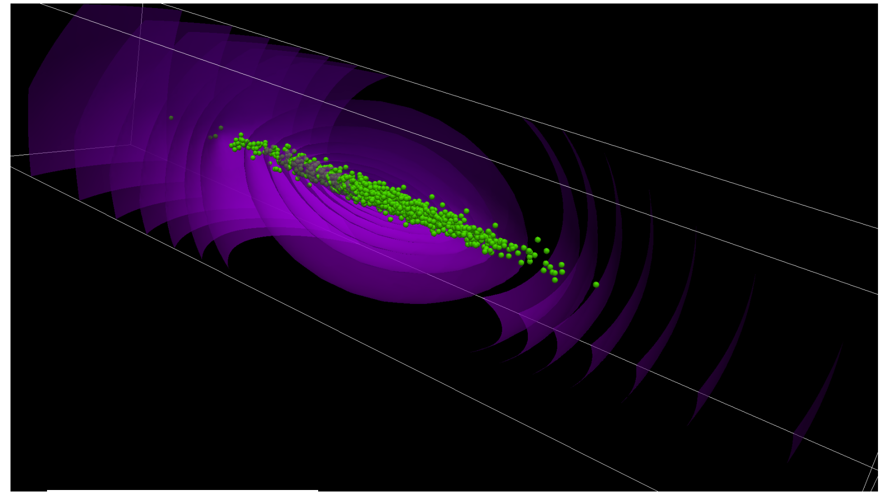

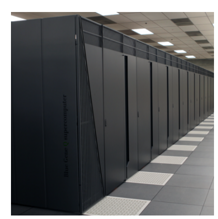

Synergia also used to teach students from around the world at the US Particle Accelerator School (e.g., two weeks ago)

Synergia and Current-Generation HPC

- Synergia was constructed for HPC
  - Core model is MPI + (optional) OpenMP
  - Strong scaling over 1000x
  - Weak scaling to 100,000+ cores

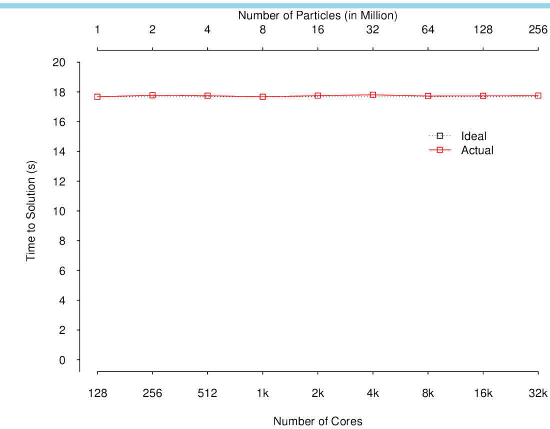

Weak scaling from 1M to 256M *particles*
128 to 32,768 cores

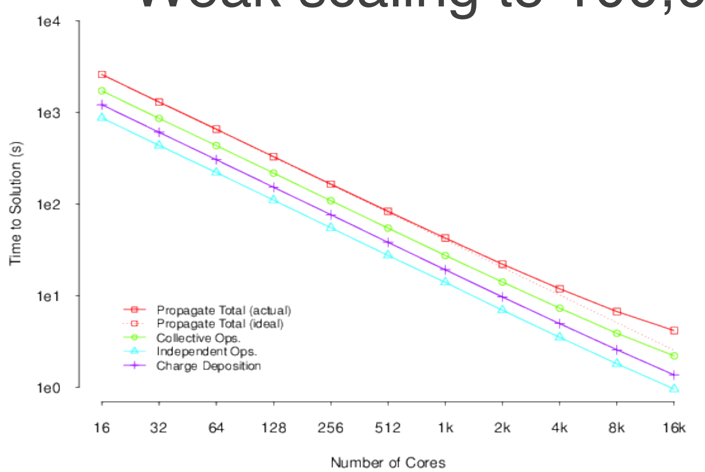

Single-bunch strong scaling from
16 to 16,384 cores
32x32x1024 grid, 105M particles

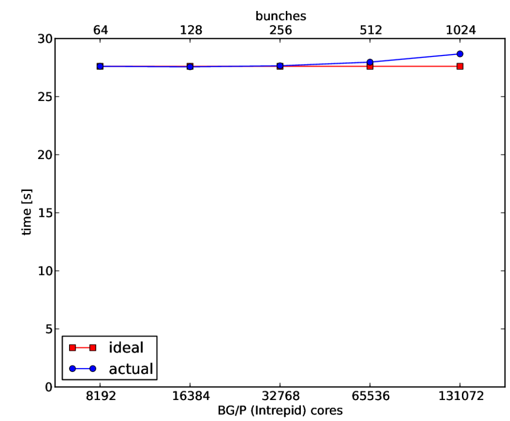

Weak scaling from 64 to 1024 *bunches*
8192 to 131,072 cores
Up to over 10^{10} particles

Synergia and Next-Generation HPC

- Dramatic changes in near future HPC architectures
  – Supercomputers and clusters alike
    - OLCF’s Summit: GPU + PowerPC
    - NERSC’s Cori: Intel MIC
- Synergia GPU and MIC ports in progress
  – Using shared expertise and facilities

First production-level GPU results

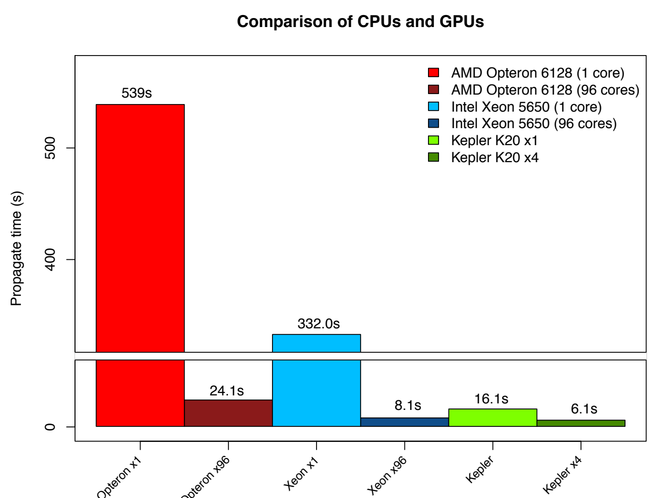

Vectorization benchmark for MIC

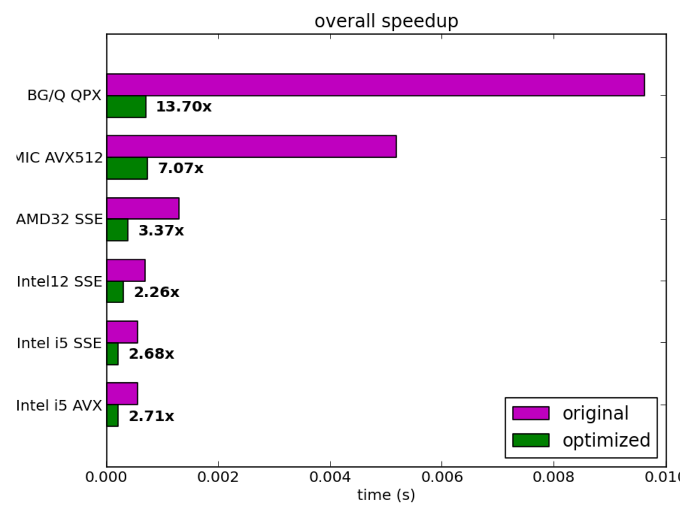

Synergia Applications

- Lab Goal: Extend the scientific reach of existing accelerator facilities
  - Simulate the Main Injector and Recycler under high intensities
    - Work done with Accelerator Division Main Injector personnel
    - Multi-bunch capabilities in Synergia used to simulate slip stacking
      - Slip stacking uses RF to combine multiple bunches to increase intensity

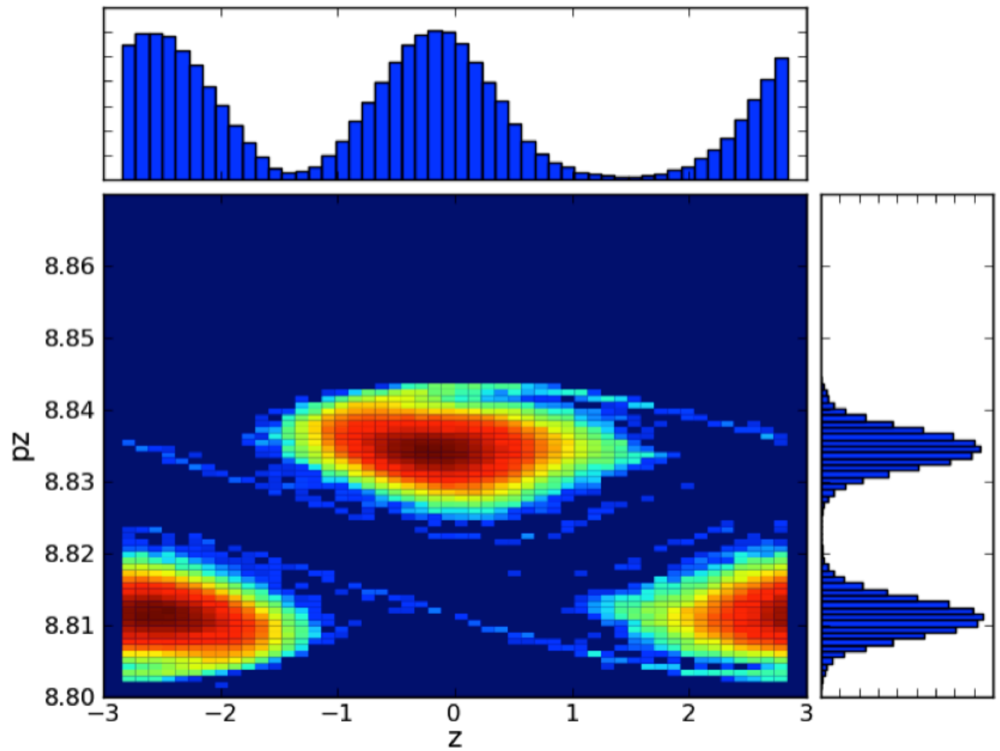

Fermilab

Synergia Applications, continued

- Resonant extraction in the Delivery Ring for Mu2e
  - Work done with Accelerator Division Muon personnel
  - Important collective effects

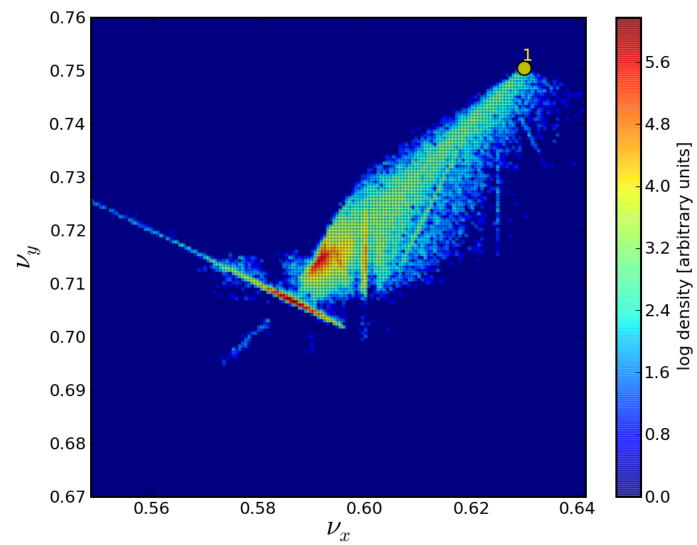

- Multi-bunch instability in the Booster
  - Work done with AD Proton Source personnel

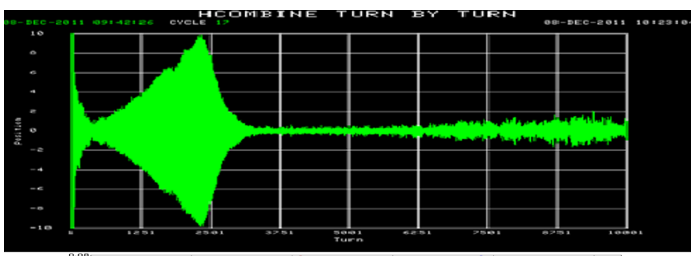

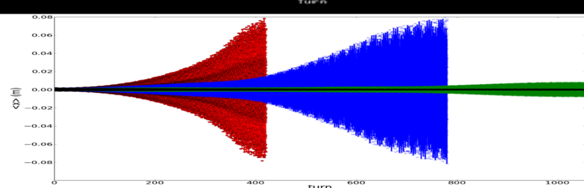

Requires multi-bunch physics with collective effects

Other Synergia Applications

- Lab Goal: Enable "transformative" Accelerator Science
  - Synergia extended to handle nonlinear lens for IOTA
  - Work done with Accelerator Physics Center personnel

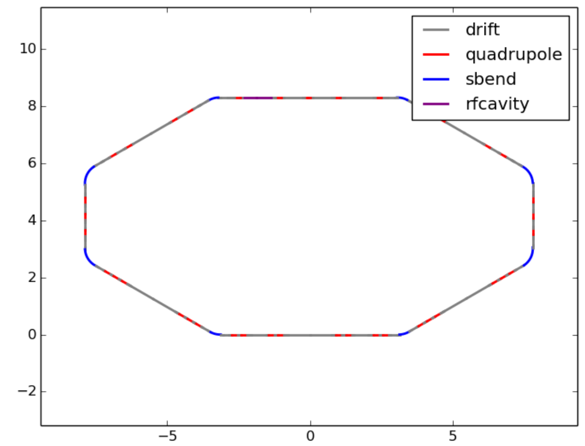

- Lab Goal: Establish Fermilab as essential contributor to future large accelerators
  - Synergia benchmarking exercise and simulations for HL-LHC
  - Work done with CERN Accelerator personnel

71 steps/turn
7,100,000 steps (~ 7 million)
4,194,304 particles (~ 4 million)
29,779,558,400,000 particle-steps (~ 30 trillion)
1,238,158,540,800,000 calls to "drift"
(~ 1 quadrillion)

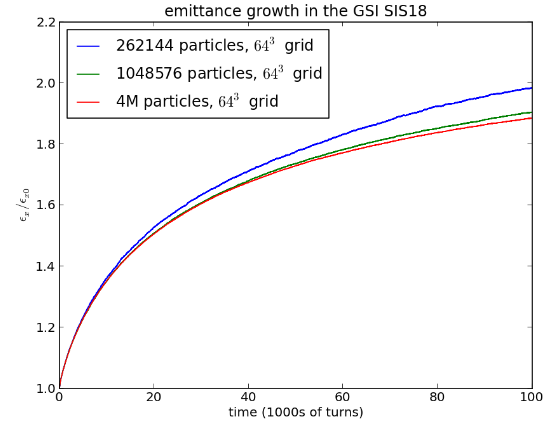

Conclusions

- Three scientific areas at the lab use significant HPC resources
  – Lattice QCD, Cosmology, Accelerator Simulation
- Share facilities
  – Acquired, deployed and operated by Computing
  – Used by all
- Share Competencies
  – Computational Physics
  – HPC programming and optimization
- Produce science aligned with lab priorities
  – Accelerator simulation in support of P5/Lab goals