Present and Future Computing Requirements for NERSC repository m327:
"Parallel Simulation of Electron Cooling Physics and Beam Transport"

Presenter: D.L. Bruhwiler,[1]

Contributors: B.T. Schwartz,[1] V.H. Ranjbar,[1] G.I. Bell[1]
Other m327 users: J. Qiang,[4] S. White,[2] Y. Luo[2]
Collaborators: R. Ryne,[4] V.N. Litvinenko,[2] W. Fischer,[2]
G. Wang,[2] Y. Hao,[2] K. Paul,[1] I. Pogorelov[1]

1. Tech-X Corporation
2. Brookhaven National Lab
3. Thomas Jefferson National Lab
4. Lawrence Berkeley National Lab

Workshop: Large Scale Computing and Storage Requirements for Nuclear Physics

May 26, 2011

Work supported by the US DOE Office of Science, Office of Nuclear Physics, including grant No.'s DE-FC02-07ER41499 and DE-SC0000835. Resources of NERSC were used.

TECH-X CORPORATION

1. m327 Project Overview

- PI: David Bruhwiler (Tech-X Corporation)
- Summarize scientific objectives through 2014
  - Provide computational support to BNL and Jlab
  - Reduce technical risk for future Electron-Ion Collider
    - eRHIC (BNL concept) and ELIC (JLAB concept)
- Present focus is in three areas
  - electron cooling of relativistic hadron beams (increase luminosity)
  - beam-beam collisions (effect on beam dynamics, luminosity)
  - spin-tracking (how to keep polarized beam fraction high)
- In the next 3 years we expect to …
  - support CeC proof-of-principle experiment underway at BNL
  - larger-scale to support near-term RHIC efforts & also EIC
  - beam-beam to support RHIC, LHC, ELIC design

Coherent e- Cooling (CeC) is a priority for RHIC & the future Electron-Ion Collider

- 2007 Nuclear Science Advisory Committee (NSAC) Long Range Plan:
  - recommends "…the allocation of resources to develop accelerator and detector technology necessary to lay the foundation for a polarized Electron-Ion Collider."
  - NSAC website: http://www.er.doe.gov/np/nsac/index.shtml
- 2009 Electron-Ion-Collider Advisory Committee (EICAC):
  - selected CeC as one of the highest accelerator R&D priorities
  - EIC Collaboration website: http://web.mit.edu/eicc
- Alternative cooling approaches
  - stochastic cooling has shown great success with 100 GeV/n Au^{+79} in RHIC
    - Blaskiewicz, Brennan and Mernick, "3D stochastic cooling in RHIC," PRL **105**, 094801 (2010).
    - however, it will not work with 250 GeV protons in RHIC
  - high-energy unmagnetized electron cooling could be used for 100 GeV/n Au^{+79}
    - S. Nagaitsev et al., PRL 96, 044801 (2006). Fermilab, relativistic antiprotons, with γ~9
    - A.V. Fedotov, I. Ben-Zvi, D.L. Bruhwiler, V.N. Litvinenko, A.O. Sidorin, New J. Physics 8, 283 (2006).
    - Cooling rate decreases as $1/\gamma^2$; too slow for 250 GeV protons
  - CeC could yield six-fold luminosity increase for polarized proton collisions in RHIC
    - This would help in resolving the proton spin puzzle.
    - Breaks the $1/\gamma^2$ scaling of conventional e- cooling, because it does not depend on dynamical friction

Staging of all-in-tunnel e-RHIC

e^- energy increases from 5 to 30 GeV by building-up SRF linacs

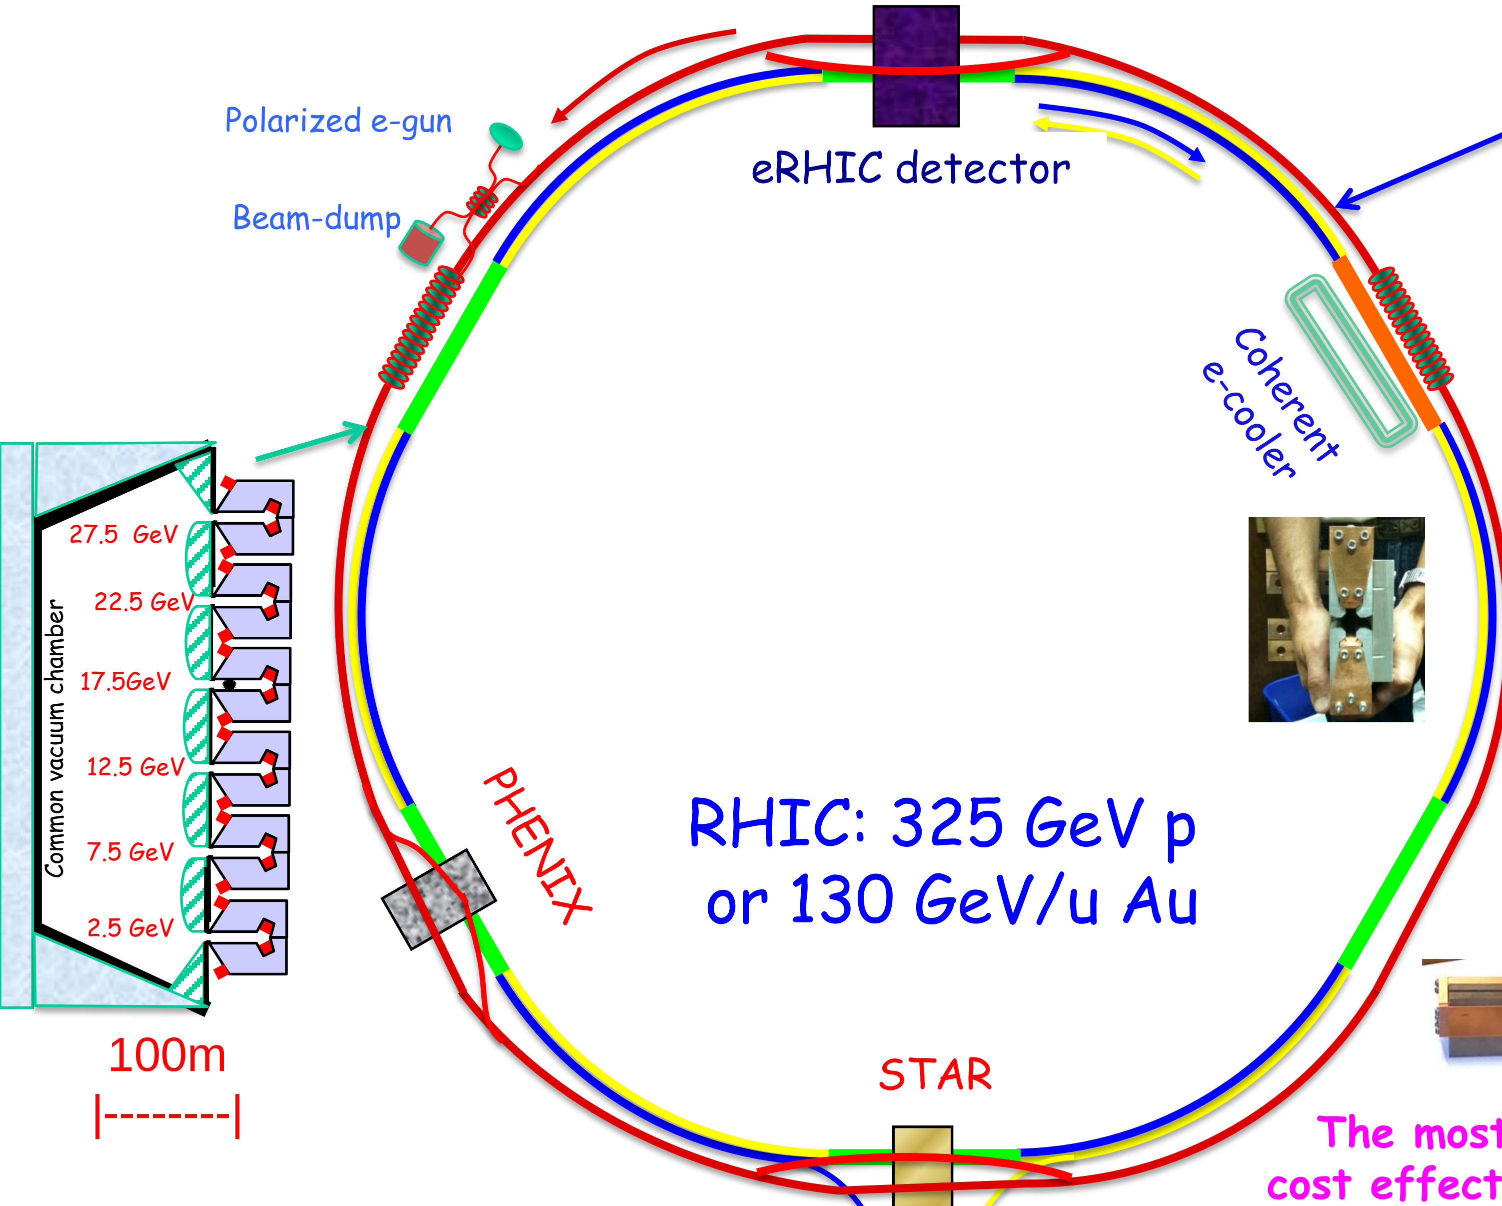

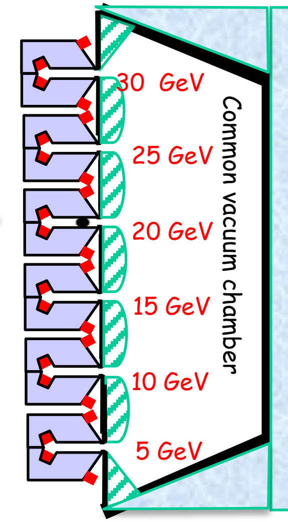

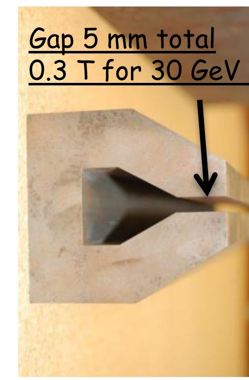

The most cost effective design

Schematic of a Coherent electron Cooling (CeC) system:

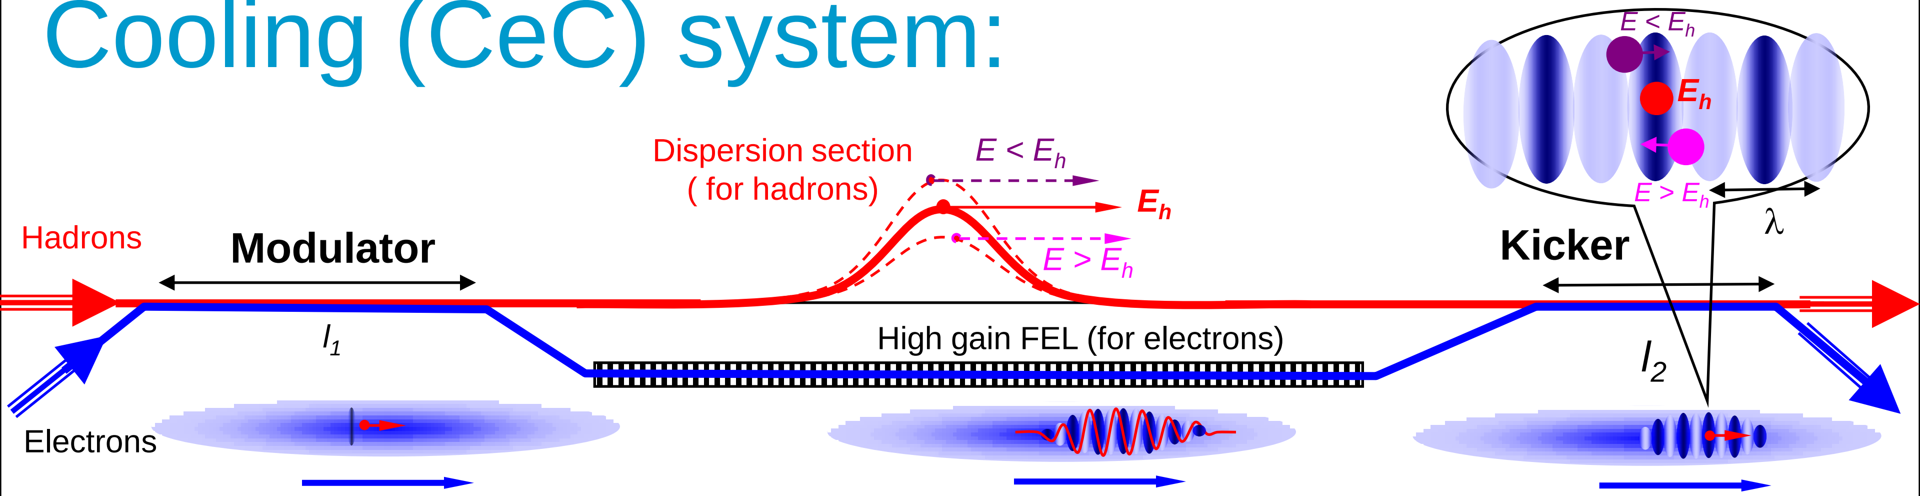

Litvinenko & Derbenev, "Coherent Electron Cooling," Phys. Rev. Lett. 102, 114801 (2009).

- Coherent Electron Cooling concept
  - uses FEL to combine electron & stochastic cooling concepts
  - a CEC system has three major subsystems
    - **modulator**: the ions imprint a "density bump" on e- distribution
    - **amplifier**: FEL interaction amplifies density bump by orders of magnitude
    - **kicker**: the amplified & phase-shifted e- charge distribution is used to correct the velocity offset of the ions

1.b. Limited Scope of this Presentation

- m327 is not the only repo supporting accelerator technology
- Other relevant NP activities in accelerator modeling and design:
  - SLAC
    - FEM modeling of SRF cavities at JLab, MSU/FRIB
  - LBL
    - parallel particle tracking & beamline design
    - additional beam-beam simulations for JLab, RHIC, LHC, EIC
  - ANL
    - FEM modeling of SRF cavities for FRIB
    - parallel particle tracking & beamline design for FRIB
    - Vlasov/Poisson algorithm development
  - Tech-X
    - FDTD modeling of SRF cavities for JLab
    - inverse cyclotron for light-ion stopping at FRIB
    - electron gun modeling for BNL (diamond amplifier project)

2.a. Current HPC Methods

- Algorithms used
  - coherent electron cooling (CeC)
    - ES PIC; δf PIC; Vlasov (all use FDTD, Poisson, unif. mesh)
  - beam-beam and spin-tracking
    - pushing particles through complicated external fields
    - Poisson solves used in some cases for "space charge kicks"
- Codes
  - The parallel VORPAL framework (Tech-X and collab's)
    - particle-in-cell; fluids; geometry; multi-physics; vlasov
    - electromagnetics, electrostatics
    - Trilinos, PETSc, parallel HDF5, new algorithm development
    - electron cooling, SRF cavities, laser-plasma, fusion, beams
    - DOE/NP, HEP, BES, OFES applications; also DOD
  - BeamBeam3D and IMPACT-T (LBL and collab's)
  - SimTrack (BNL and collab's)
  - Teapot-SpinTrack (part of UAL framework) (BNL and Tech-X)

2.b. Current HPC Methods

- Quantities that affect problem size, scale of simulations (electron cooling only)
- δf PIC uses macro-particles to represent deviation from assumed equilibrium distribution
  - much quieter for simulation of beam or plasma perturbations
  - implemented in VORPAL for Maxwellian & Lorentzian velocities
- Typical 3D simulation size
  - 3D domain, 40 λ_D on a side; 10 cells per λ_D → ~10^8 cells
  - 300 ptcls/cell to accurately model temp. effects → ~2 x 10^{10} ptcls
  - dt ~ $(dx/v_{th,x})$ / 5; $\omega_{pe} \sim v_{th} / \lambda_D$ → τ_{pe} ~ 300 time steps
  - 1 μs/ptcl/step → ~1,000 processor-hours for ½ plasma period

2.c. Current HPC Requirements (electron cooling only)

- Architectures currently used
  - Franklin, Hopper, small clusters
- Compute/memory load
  - 1,000 proc-hours per run; ~1,000 runs per year (param. scans)
  - 5 GB aggregate memory
- Data read/written
  - reading: input file (negligible size)
  - 20*5 = 100 GB (i.e. 20 restart dumps for movie generation)
- Necessary software, services or infrastructure
  - parallel i/o via HDF5; Trilinos; python; VisIt; IDL
- Known limitations/obstacles/bottlenecks
  - none at present or in next year; major problems are looming
- Hours requested/allocated/used in 2011
  - 2.2 million hours requested for FY 2011
  - 0.5 million allocated on Franklin; 0.7 million for Hopper
  - 0.35 million hours used so far (many free hours on Hopper)

3.a. HPC Usage & methods for next 3-5 years

- Upcoming changes to codes/methods/approaches to satisfy science goals (electron cooling only)
  - Large-scale Vlasov/Poisson sim's for e- cooling to benchmark PIC
  - Move beyond 10^5 cores for PIC, Vlasov and spin-tracking
    - in part via move towards effective use of GPUs
    - beginning exploration of OpenMP for hybrid parallelism
- Changes to Compute/memory load
  - more resolution & PPC needed in future to model realistic e- beams
    - 50,000 proc-hours per run; ~1,000 runs per year
    - 30 GB aggregate memory
  - full 3D3V Vlasov/Poisson (6D mesh) to benchmark/verify δf PIC
    - 300,000 proc-hours per run; ~100 runs per year
    - 150 GB aggregate memory
- Changes to Data read/written
  - δf PIC: 20*30 = 600 GB (i.e. 20 restart dumps for movie generation)
  - Vlasov: 20*150 = 3 TB

3.b. HPC Usage & methods for next 3-5 years

- Changes to necessary software, services or infrastructure
  - we may need assistance with visualizing 4D and 6D fields
  - assistance with the obstacles listed below may be necessary

- Anticipated limitations/obstacles/bottlenecks on 10K-1M PE system
  - I/O is not now a bottle neck, but it does not appear to scale, so…
  - dynamic load balancing may be required for good efficiency
  - must move to smaller surface-to-volume ratios for MPI domains
    - communication-related overhead will become a bottle neck
    - fault tolerance will become a major concern

VORPAL shows good efficiency at 4k Franklin cores for electrostatic PIC (strong scaling; production simulations of electron cooling)

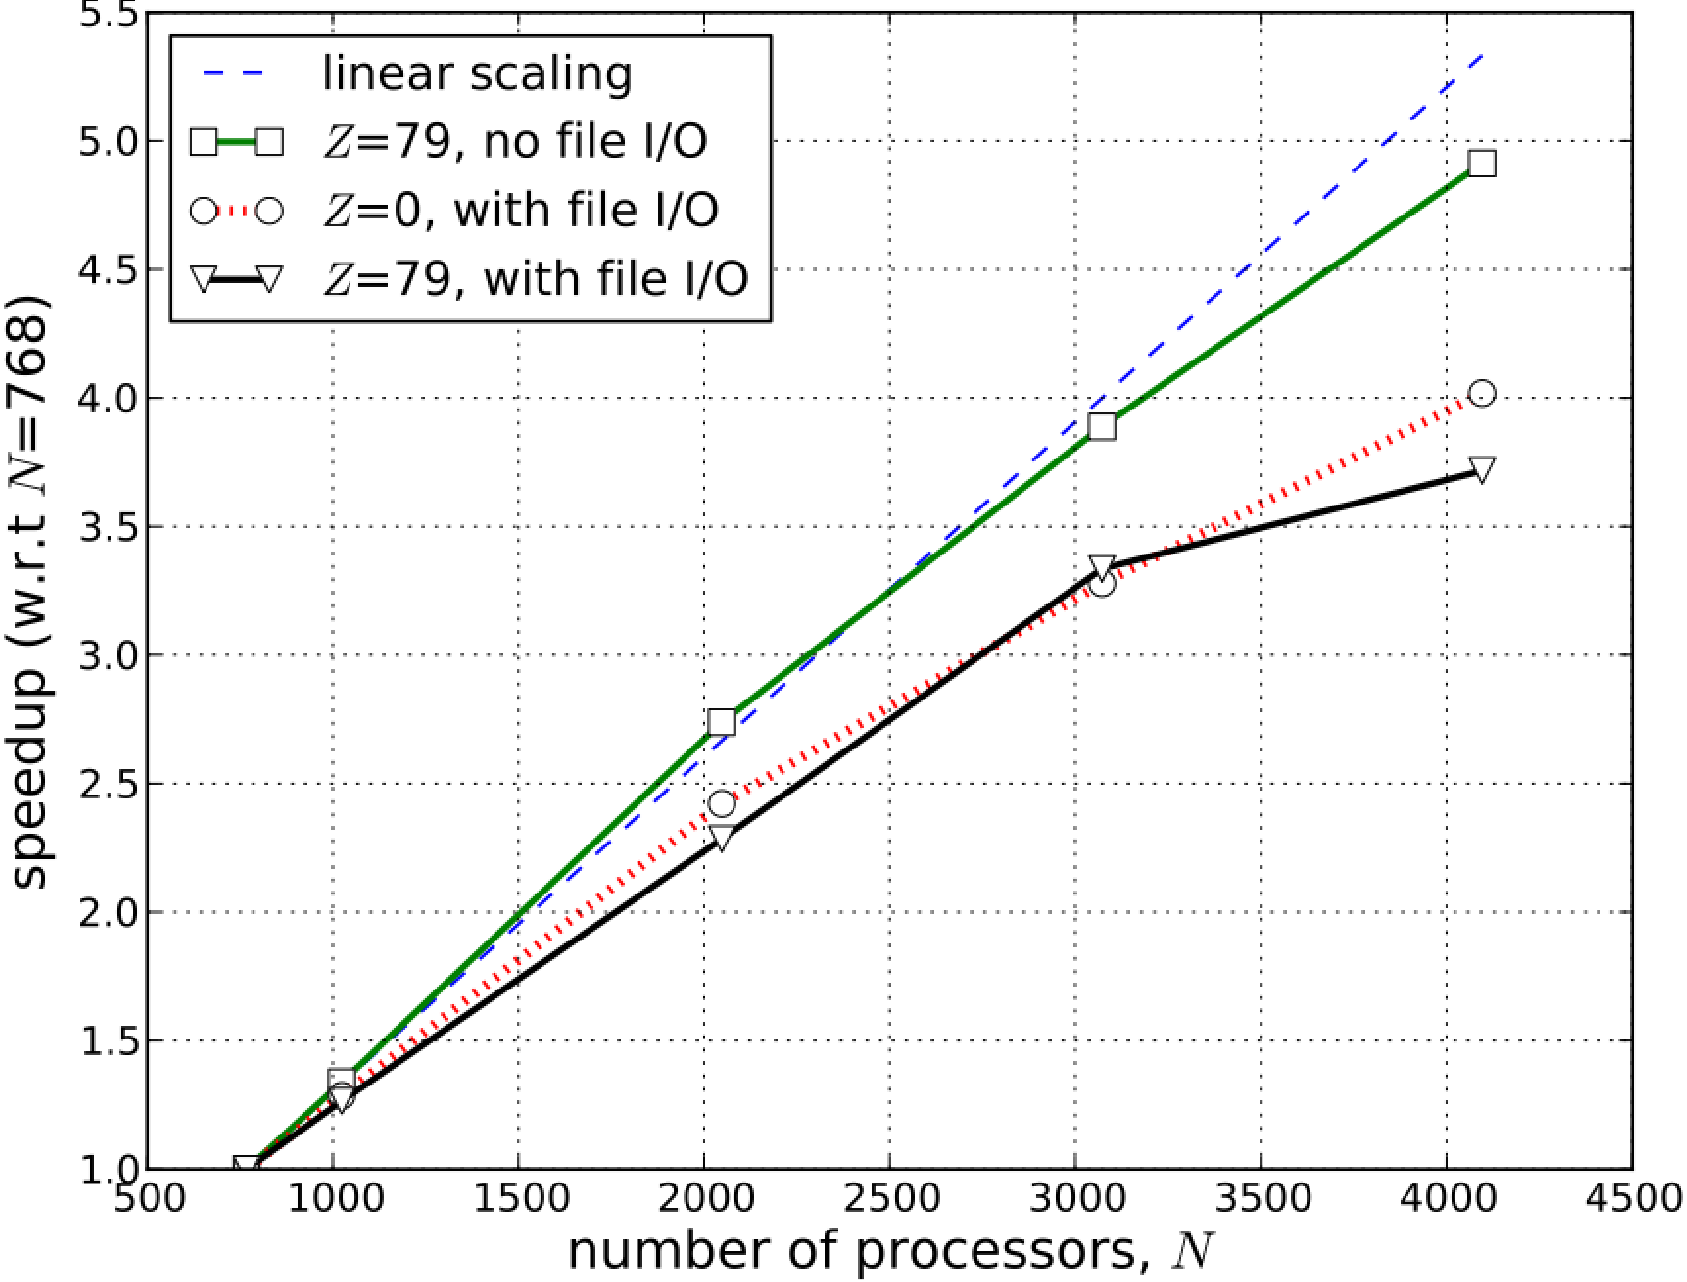

Strategy for New Architectures

- How are you dealing with, or planning to deal with, many-core systems that have dozens or hundreds of computational cores per node?
  - beginning to explore benefits of OpenMP for hybrid parallelism
  - hoping that MPI-3 will alleviate the problem (perhaps temporarily)

- How are you dealing with, or planning to deal with, systems that have a traditional processor augmented by some sort of accelerator such as a GPU or FPGA or similar?
  - VORPAL electromagnetics (w/ boundaries) is ported to multiple GPUs
  - electrostatic PIC has been prototyped on NVIDIA Fermi architecture
  - BNL codes TEAPOT and Spink were rewritten to use cuda
    - 100x speedup with 10,000 particles has enabled new physics

4.a. Summary I

- What new science results might be afforded by improvements in NERSC computing hardware, software and services?
  - faster beam-beam and spin-tracking simulations, with greater physical fidelity, could provide physical insight that points to beam dynamics changes that significantly increase the luminosity of RHIC, with greater polarization
    - this would reduce the time/cost required for obtaining important nuclear physics results, and perhaps enable new results
  - providing computational support to the CeC proof-of-principle experiment at BNL could help that effort succeed, resulting in a fundamentally new and important technique to increase the luminosity of RHIC or of any future EIC facility by orders of magnitude

4.b. Summary II

- Recommendations on NERSC architecture, system configuration and the associated service requirements needed for your science
  - present architecture and configuration should work well in the near future
  - Major changes in architecture (e.g. GPU or hybrid CPU/GPU) will require a great deal of additional software development
    - however, we are working to prepare for this transition
- NERSC generally acquires systems with roughly 10X performance every three years. What significant scientific progress could you achieve over the next 3 years with access to 50X NERSC resources?
  - would allow us to use full 3D3V (i.e. 6D mesh) Vlasov/Poisson to benchmark/verify our 3D δf PIC simulations for realistic e- distrib.'s
  - important; otherwise we have less confidence in our δf PIC results
- What "expanded HPC resources" are important for your project?
  - convenient 4D and 6D visualization of fields
  - GPU hardware, supporting libraries, consulting