Accelerator Modeling&&Future Experiments

P. Spentzouris
Fermilab

Computationally challenging accelerator physics, common to most HPC applications

- Machine design/operation: particles affected by machine components, other beam particles or beams **[Beam Dynamics]**
  - Space charge
  - Beam-beam
  - Electron cloud
  - Electron cooling
  - Intrabeam scattering
  - Geometry (1000’s of 3D elements)
- Component design **[Electromagentics]**
  - Impedance
  - Wakefields
  - multipacting
  - Thermal, mechanical
- New accelerator technologies **[Advanced Accelerators]**
  - Laser and plasma wakefields
  - Ionization cooling
- Multi-physics simulations, multi-parameter optimization runs necessary

Many particles

Many DOF

Many particles

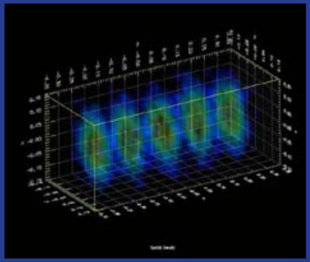

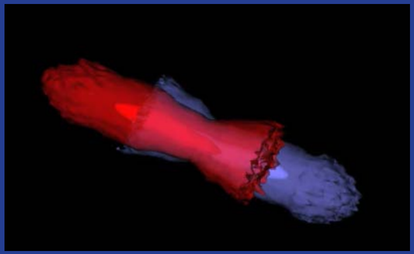

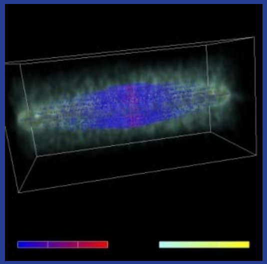

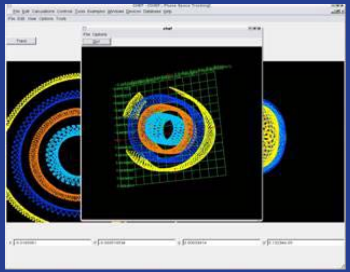

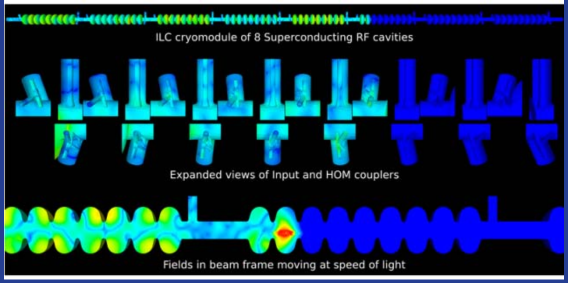

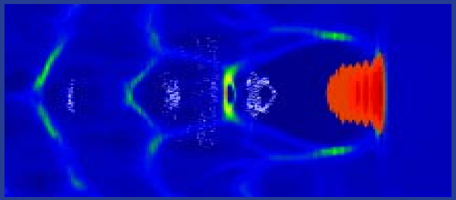

HPC accelerator modeling approaches

- Different physics call for a variety of approaches resulting to wide spectrum of computational requirements
  - Beam Dynamics: electrostatic PIC
  - Electromagnetics: finite difference and finite element, time and frequency domain
  - Advanced Accelerators: full EM PIC and reduced PIC
- Codes are spectral based, finite difference based, and hybrids, with both FFT and multi-grid based solvers.
- Depending on the physics of the problem, domain, particle, or hybrid decomposition is used. There is communication & storage of particle data, computational grid data, and geometry
- Application focus includes “discovery” and “design optimization” areas
  - Results in either very-large-scale and very-large-volume medium scale computations, thus large data sets per run

Current state of the art (simulation size)

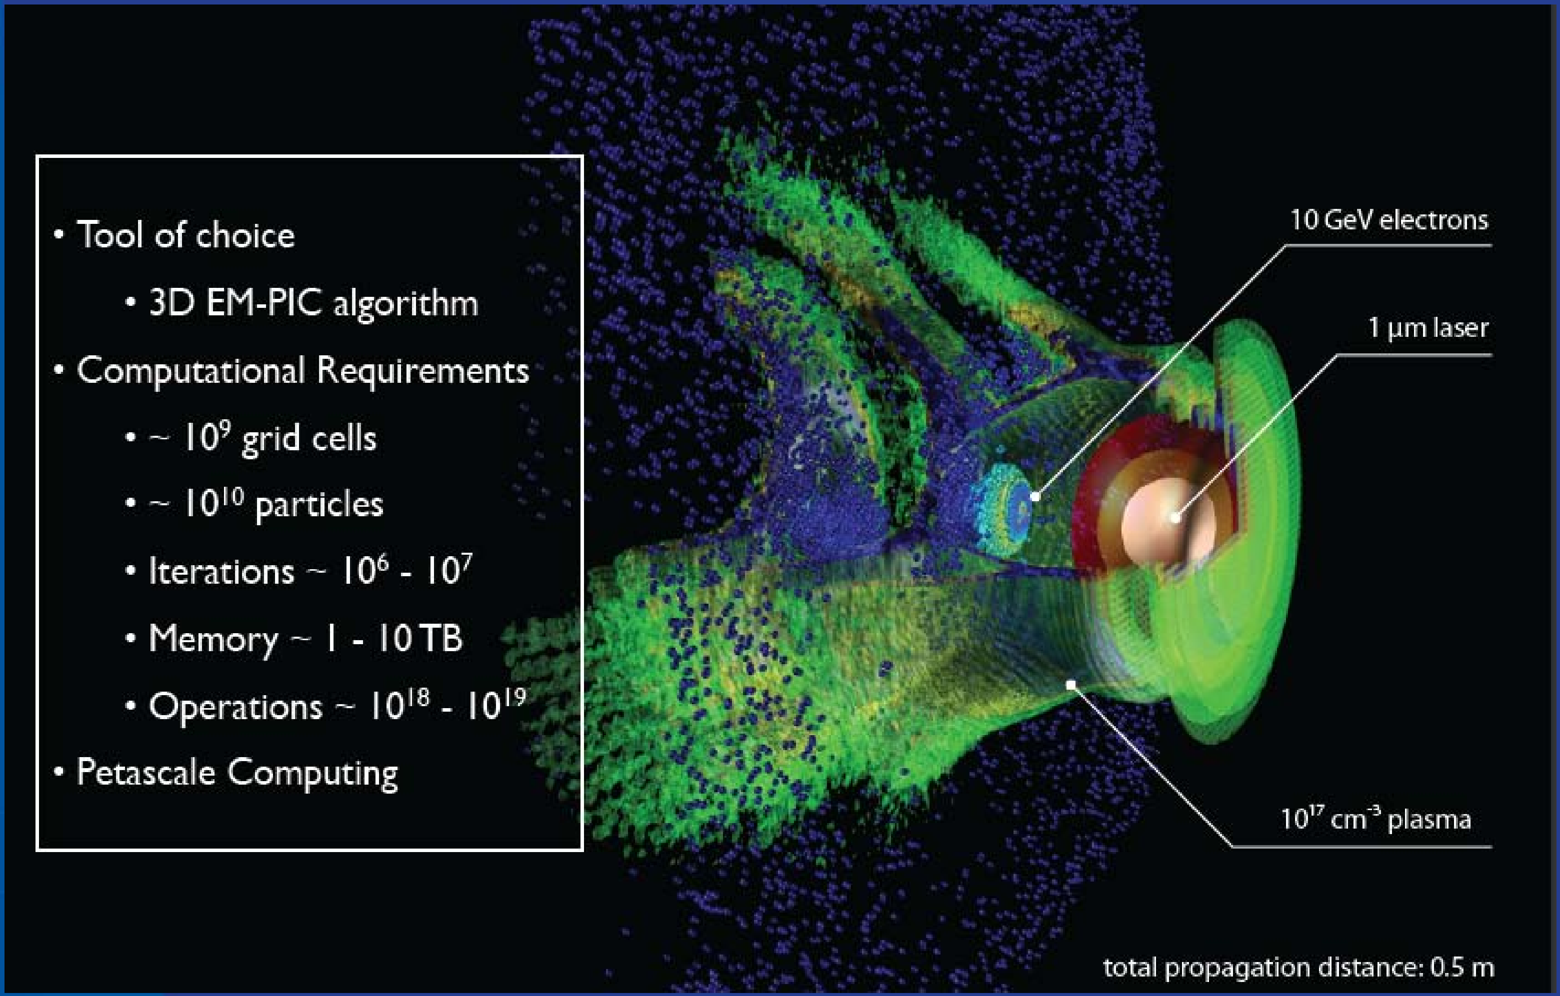

Exascale

➢Already getting good efficiency at large number of cores

Exascale provides opportunity for even larger simulations, in the case of plasma based acceleration the opportunity for full scale ab-initio simulations of experiments and real-time steering of these experiments, if the simulated data could be made available at the research facility.

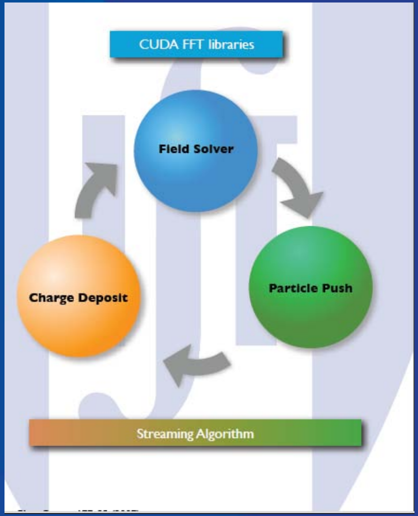

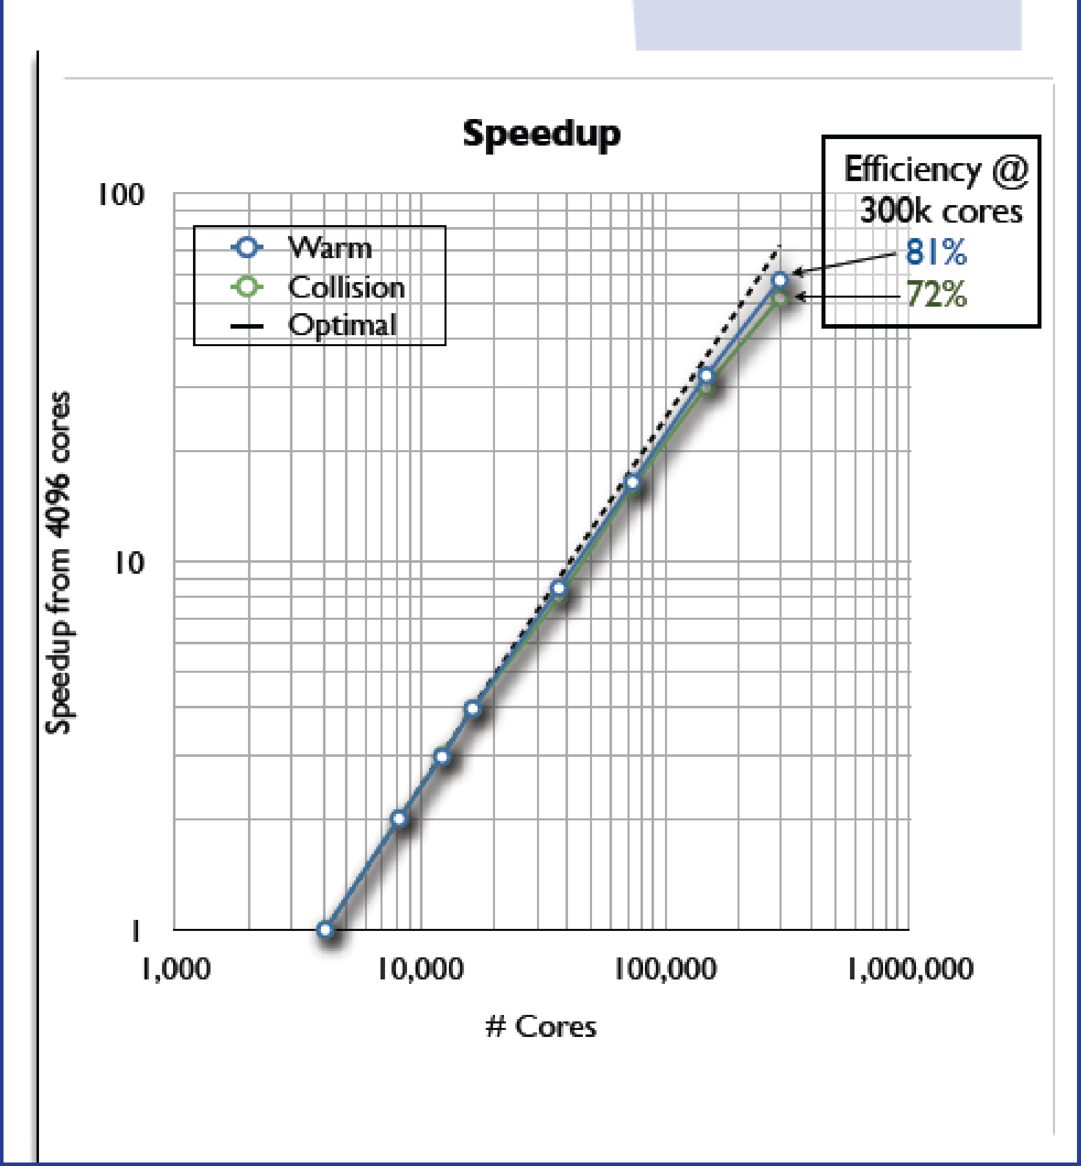

Plans for HPC accelerator simulation utilization

- Near-real-time steering of experiments (advanced concept R&D)
  - Beam facilities (for example BELLA, FACET, such utilization already an activity within the collaboration)
  - Component design (for example CLIC)
- Near-real-time steering of operating accelerators
- Requires data mining and interactive exploration of TB simulated data sets in the control room
  - Current paradigm has all analysis done at the SC centers, with high level info communicated to researchers and stakeholders

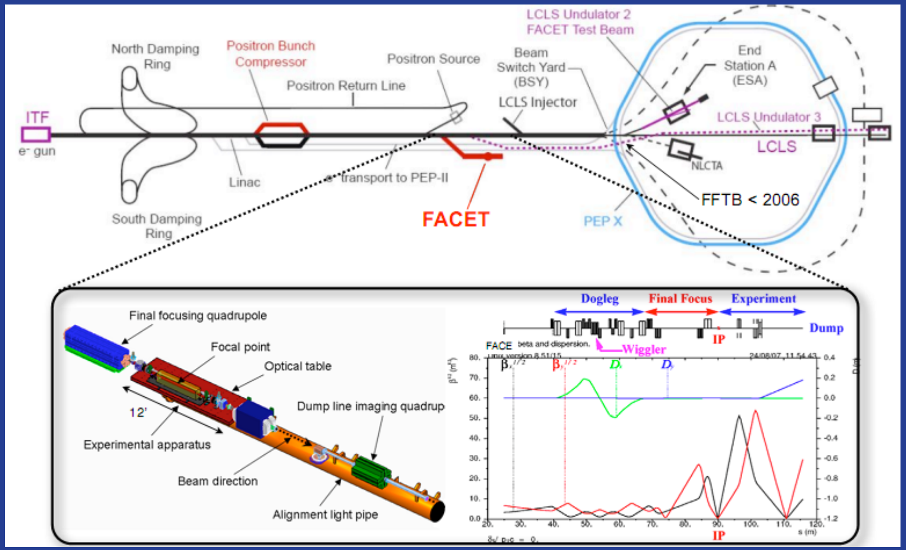

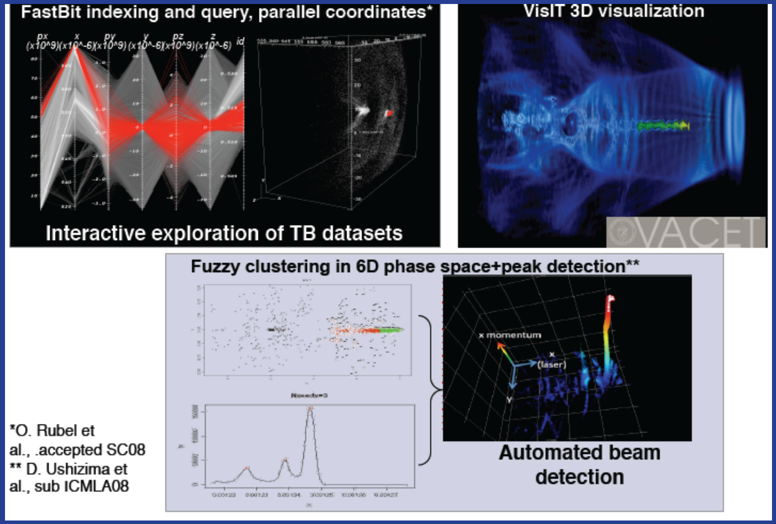

Fermilab

Component design

CAD model of LCLS RF gun

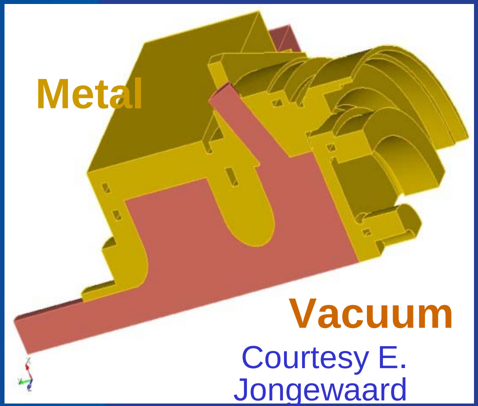

Numerical Model

Electromagnetics

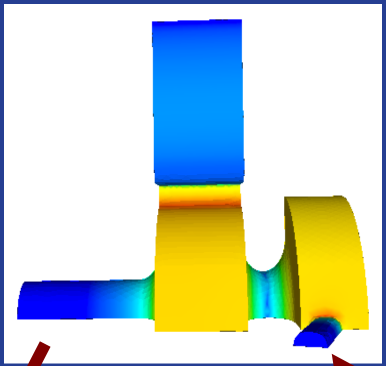

Thermal

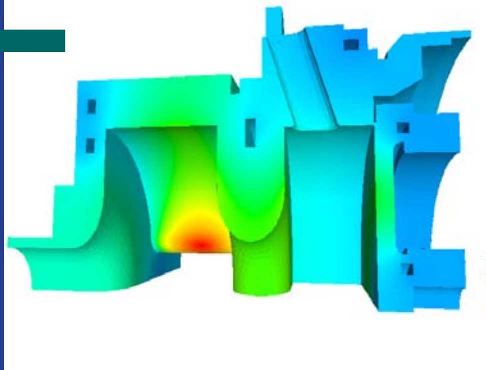

Mechanical

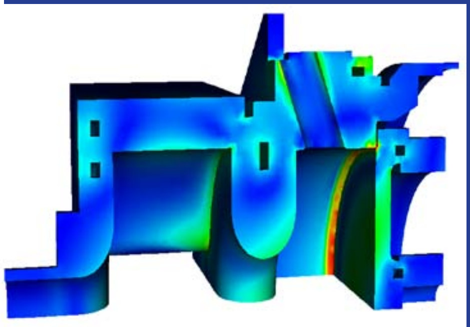

Engineering prototype

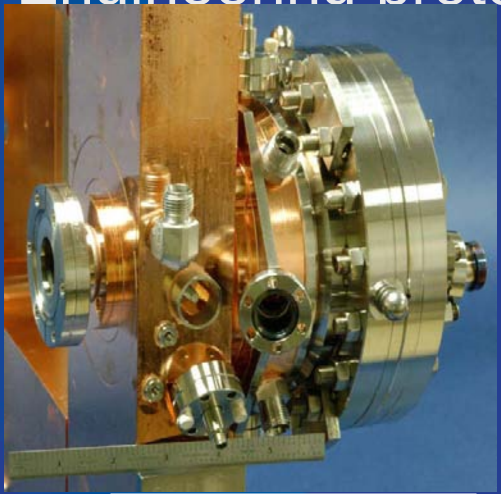

Courtesy D. Dowell

Requirements from NERSC planning workshop

BD

	Current	Next 3-5 Years
Computational Hours	12M	40M
Parallel Concurrency	2k-10k	10k-100k
Wall Hours per Run	48	72
Aggregate Memory	2 TB	10 TB
Memory per Core	0.5 GB	1 GB
I/O per Run	100 GB-1TB	500GB-5TB
On-Line Storage Needed	2 TB	15 TB
Data Transfer	500GB/week	5TB/run

EM

	Current	Next 3-5 Years
Computational Hours	9.2 million	100 million
Parallel Concurrency	8192	75,000
Wall Hours per Run	12	24
Aggregate Memory	12 TB	100 TB
Memory per Core	1.5 GB	1.5 GB
I/O per Run	1 TB	10 TB
On-Line Storage Needed	500 GB	5 TB
Data Transfer	100 GB/run	10TB/run

AA

	Current	Next 3-5 Years
Computational Hours	3.5M	150M
Parallel Concurrency	11k	500k (large) / 50k (sml)
Wall Hours per Run	24	12 / 12
Aggregate Memory	100GB	100TB / 10 TB
Memory per Core	<0.1 GB	0.5GB / 0.5 GB
I/O per Run	2 TB	50TB / 5 TB
On-Line Storage Needed	2 TB	50TB
Data Transfer	5 TB/run	50TB/run

Accelerator Modeling requirements (strawman)

Feature	Scientific instruments	Scientific Process	Requirements	
Time Scale			Network	Middleware
Near	• Large scale multi-physics simulations • Large data sets per run (~100 TB), many runs (parameter optimization) • Operating accelerators • Component prototyping and testing	•Distributed collaboration • Remote operation, steering • Remote visualization • Sharing of data and metadata, data mining	•Robustness •Reliability • Quality of service •Interoperability for security • Large-scale data storage	Remote I/O Collaborative data access
5+ years	Could provide info if this format useful			

Future Experiments

It seems that as far as data volume and rate demands nothing that "known" near future experiments require can surpass LHC experiment requirements. Maybe an interesting case would be the evolution of the idea of a trigerless streaming DAQ, where the data is not streaming from the detector to the experimental hall, but to computing resources residing in remote areas (collaborating institutions)

- Mu2E, for example, ~100GBytes/sec

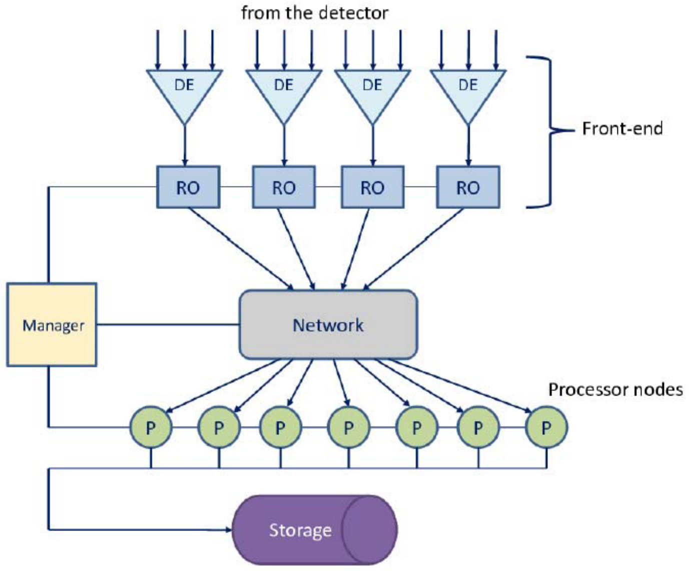

Figure 8.1: A scalable triggerless data acquisition architecture. Detector electronics (DE) sends hit information to readout modules (RO), which transmit data over high-speed optical data links to a network switch that routes data from different parts of the detector to individual processor nodes (P). Events are selected and prepared for storage in the processor nodes. The number of front-end data links and/or the number of processor nodes can be increased to add throughput to the system.