Faster! Faster! Highlights in Particle Accelerator Research @ NERSC

Jean-Luc Vay*
Advanced Modeling Program (Head)
Accelerator Technology and Applied Physics Division

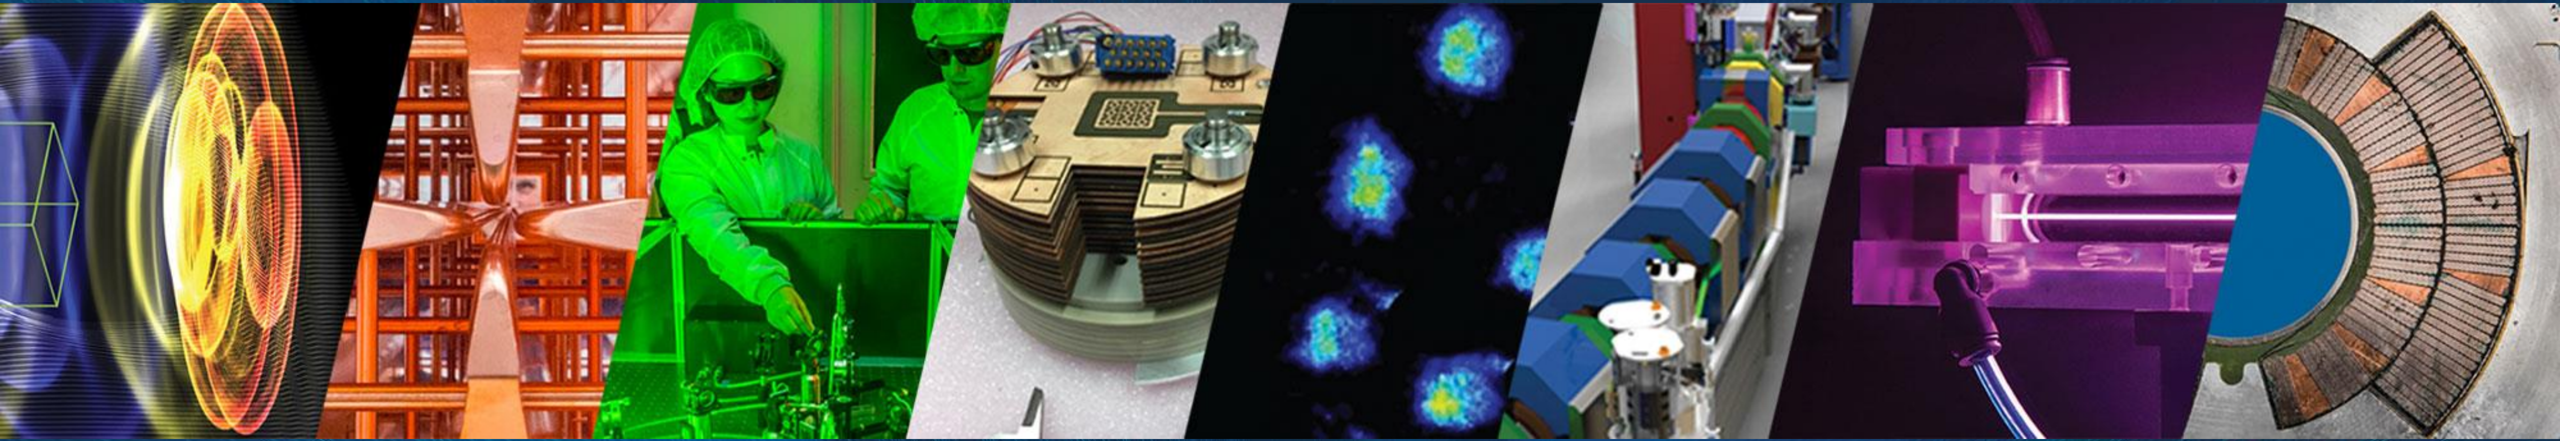

NERSC@50 Seminar Series
August 12, 2024

*Thanks for material provided by/borrowed from A. Formenti, M. Garten, C. Geddes, A. Huebl, R. Lehe & C. Ng

Faster! Faster!

A parallel between particle acceleration and computing

Particle accelerators

Goals (some):

- as many particles as possible
- going as fast as possible
- in an orderly fashion
- as cheaply as possible

Supercomputers

Goals (some):

- as many calculations as possible
- performed as fast as possible
- in an orderly fashion
- as cheaply as possible

Largest accelerators

Name	Cost ($B)	Power consumption (MW)
LHC (CERN)	~5	~120
FCC-ee* (CERN)	12-18	~290
FCC-hh* (CERN)	30-50	~560
HE-ILC* (Japan)	18-30	~400
CEPC* (China)	12-18	~340

*Planned

Largest supercomputers

Name	Cost ($B)	Power consumption (MW)
Perlmutter (NERSC)	~0.146**	~3
Frontier (OLCF)	~0.6	~23
Aurora (ALCF)	~0.5	~39
Fugaku (Japan)	~1.2	~30
El Capitan (LLNL)	~0.6	~30

**Construction + operation

➔ can be very expensive and power consumption has become a limiting factor

Quick (very incomplete) intro on particle accelerators & colliders

Particle Accelerators are Essential Tools in Modern Life

Medicine

- **~9,000 medical accelerators** in operation worldwide
- 10’s of millions of patients treated/yr
- 50 medical isotopes, routinely produced with accelerators

Industry

- **~20,000 industrial accelerators** in use
  - Semiconductor manufacturing
  - cross-linking/ polymerization
  - Sterilization/ irradiation
  - Welding/cutting
- Annual value of all products that use accel. Tech.: $500B

National Security

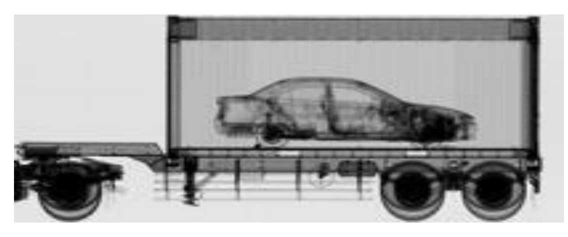

- **Cargo** scanning
- Active interrogation
- **Stockpile stewardship**: materials characterization, radiography, support of non-proliferation

Discovery Science

- **~30% of Nobel Prizes** in Physics since 1939 enabled by accelerators
- 4 of last 14 Nobel Prizes in Chemistry for research utilizing accelerator facilities

Particle accelerators: the basics

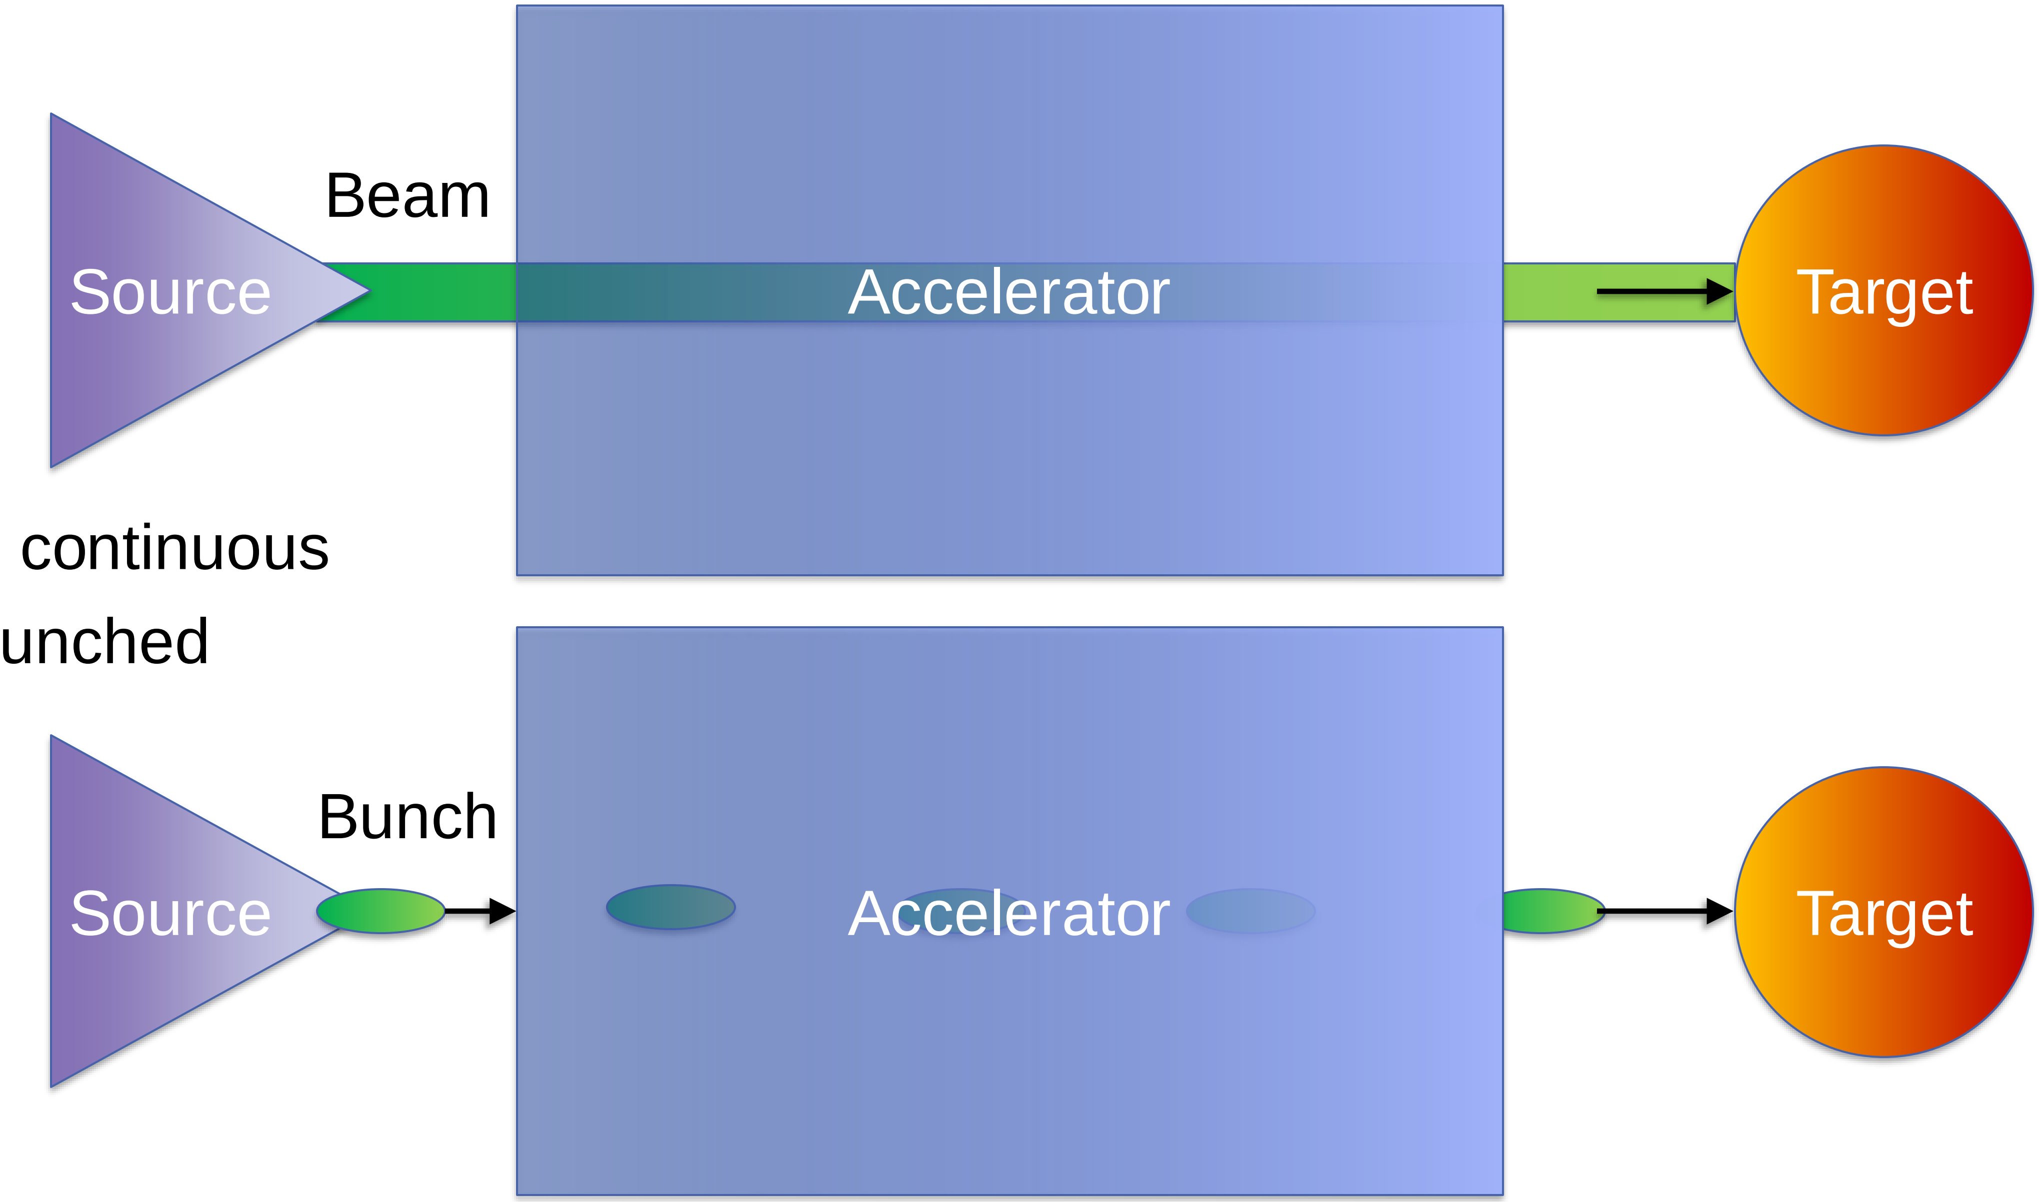

Beam is continuous
or bunched

Particle accelerators: building blocks & typical configurations

Charged particles:

e^-, e^+, p, $\bar{p}$, Au^{n+}, …

Electric fields to accelerate:

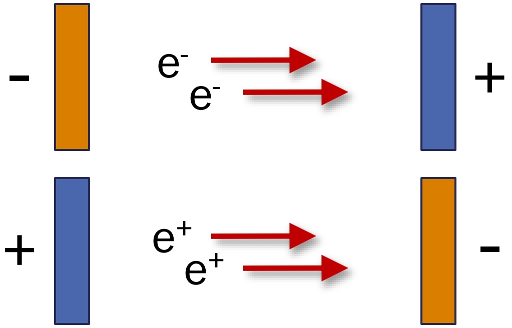

Magnetic fields to bend:

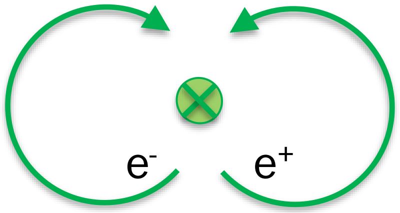

Circular

Cyclotron

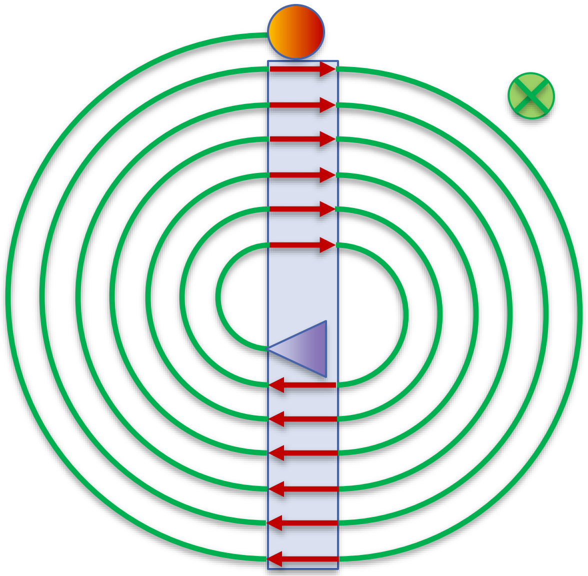

Synchrotron

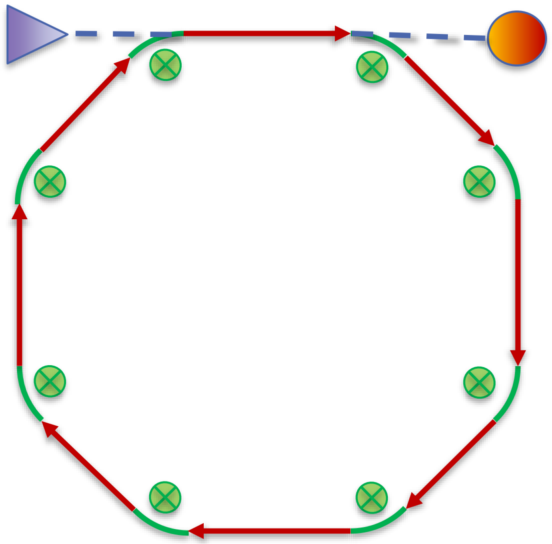

Linear

RF Linac

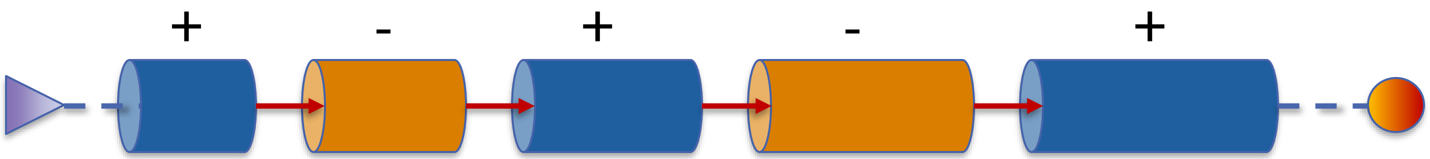

Each has pros/cons. Choice based on needs, size, cost, etc.

Keeping the beam together

- Accelerating one particle at a time is one thing. Accelerating many particles in a beam and keeping them together is another.

- There are two main causes for the beam to expand:
  - Velocity spread (“temperature”)

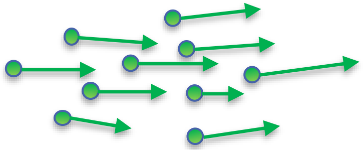

a “cold” beam has no velocity spread

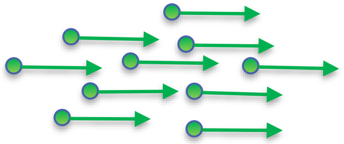

  - Repulsion from particles with same charge

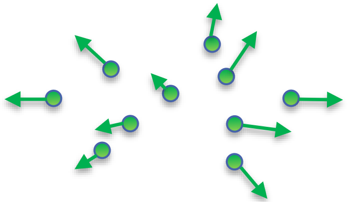

➔ need for periodic confinement

Transversally, it is usually done using Focusing-Defocusing (FODO) quadrupole

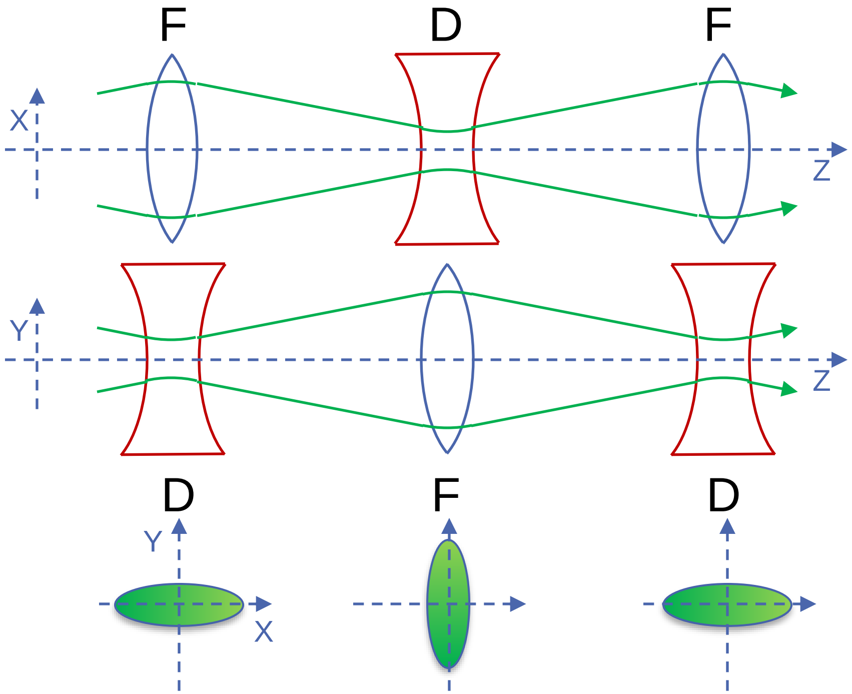

Particle accelerators started small

E. O. Lawrence

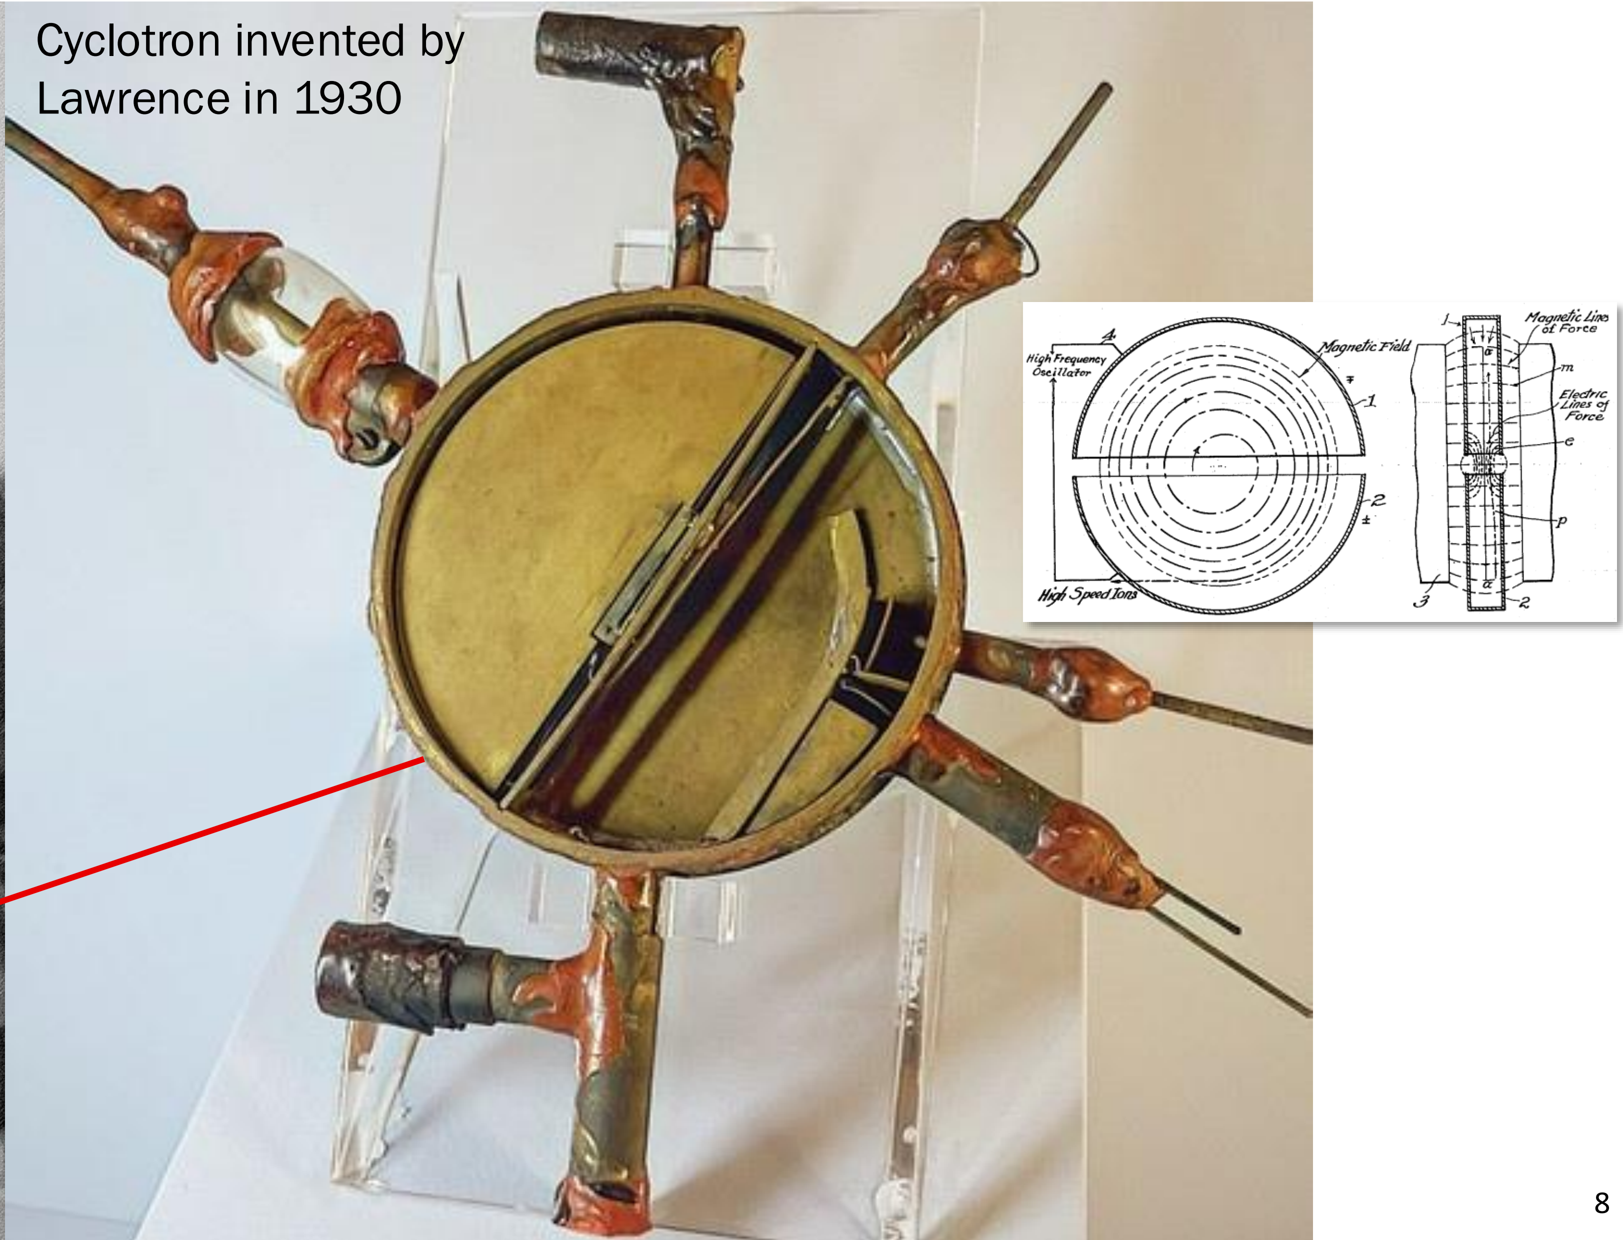

Cyclotron invented by Lawrence in 1930

E. Lawrence founded RadLab at UC Berkeley in 1931

***

where he pioneered “team science” aka “big science”

***

for particle accelerator research & applications

BERKELEY LAB

1938

Luis Alvarez
Nobel 1968

Edwin McMillan
Nobel 1951

Robert
Oppenheimer

Ernest Lawrence
Nobel 1939

Magnet of 60” cyclotron

Particle accelerators can be very large & expensive

Example 1: Large Hadron Collider (circular accelerator)

- Circumference: ~27 kms
- Construction cost: ~$5B
- Consumption: ~200MW

  (total CERN complex)

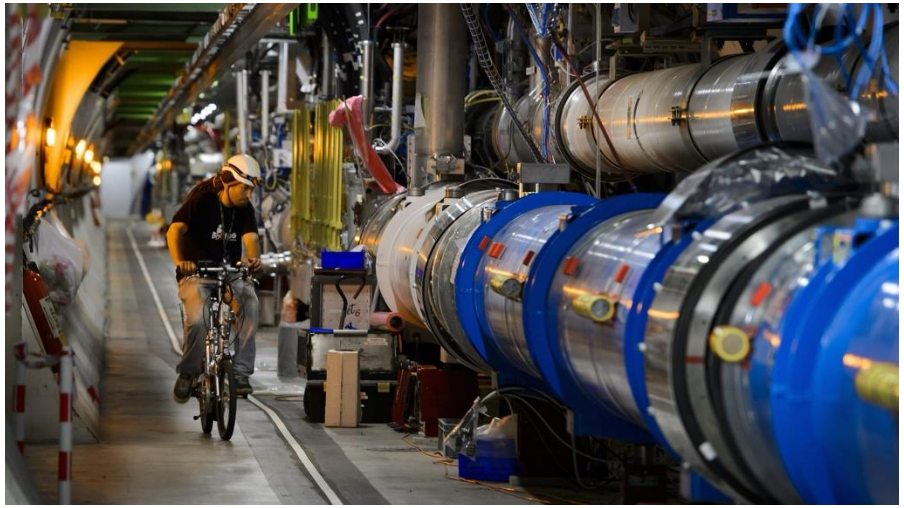

Particle accelerators can be very large & expensive

Example 2: LCLS-II light source (linear accelerator)

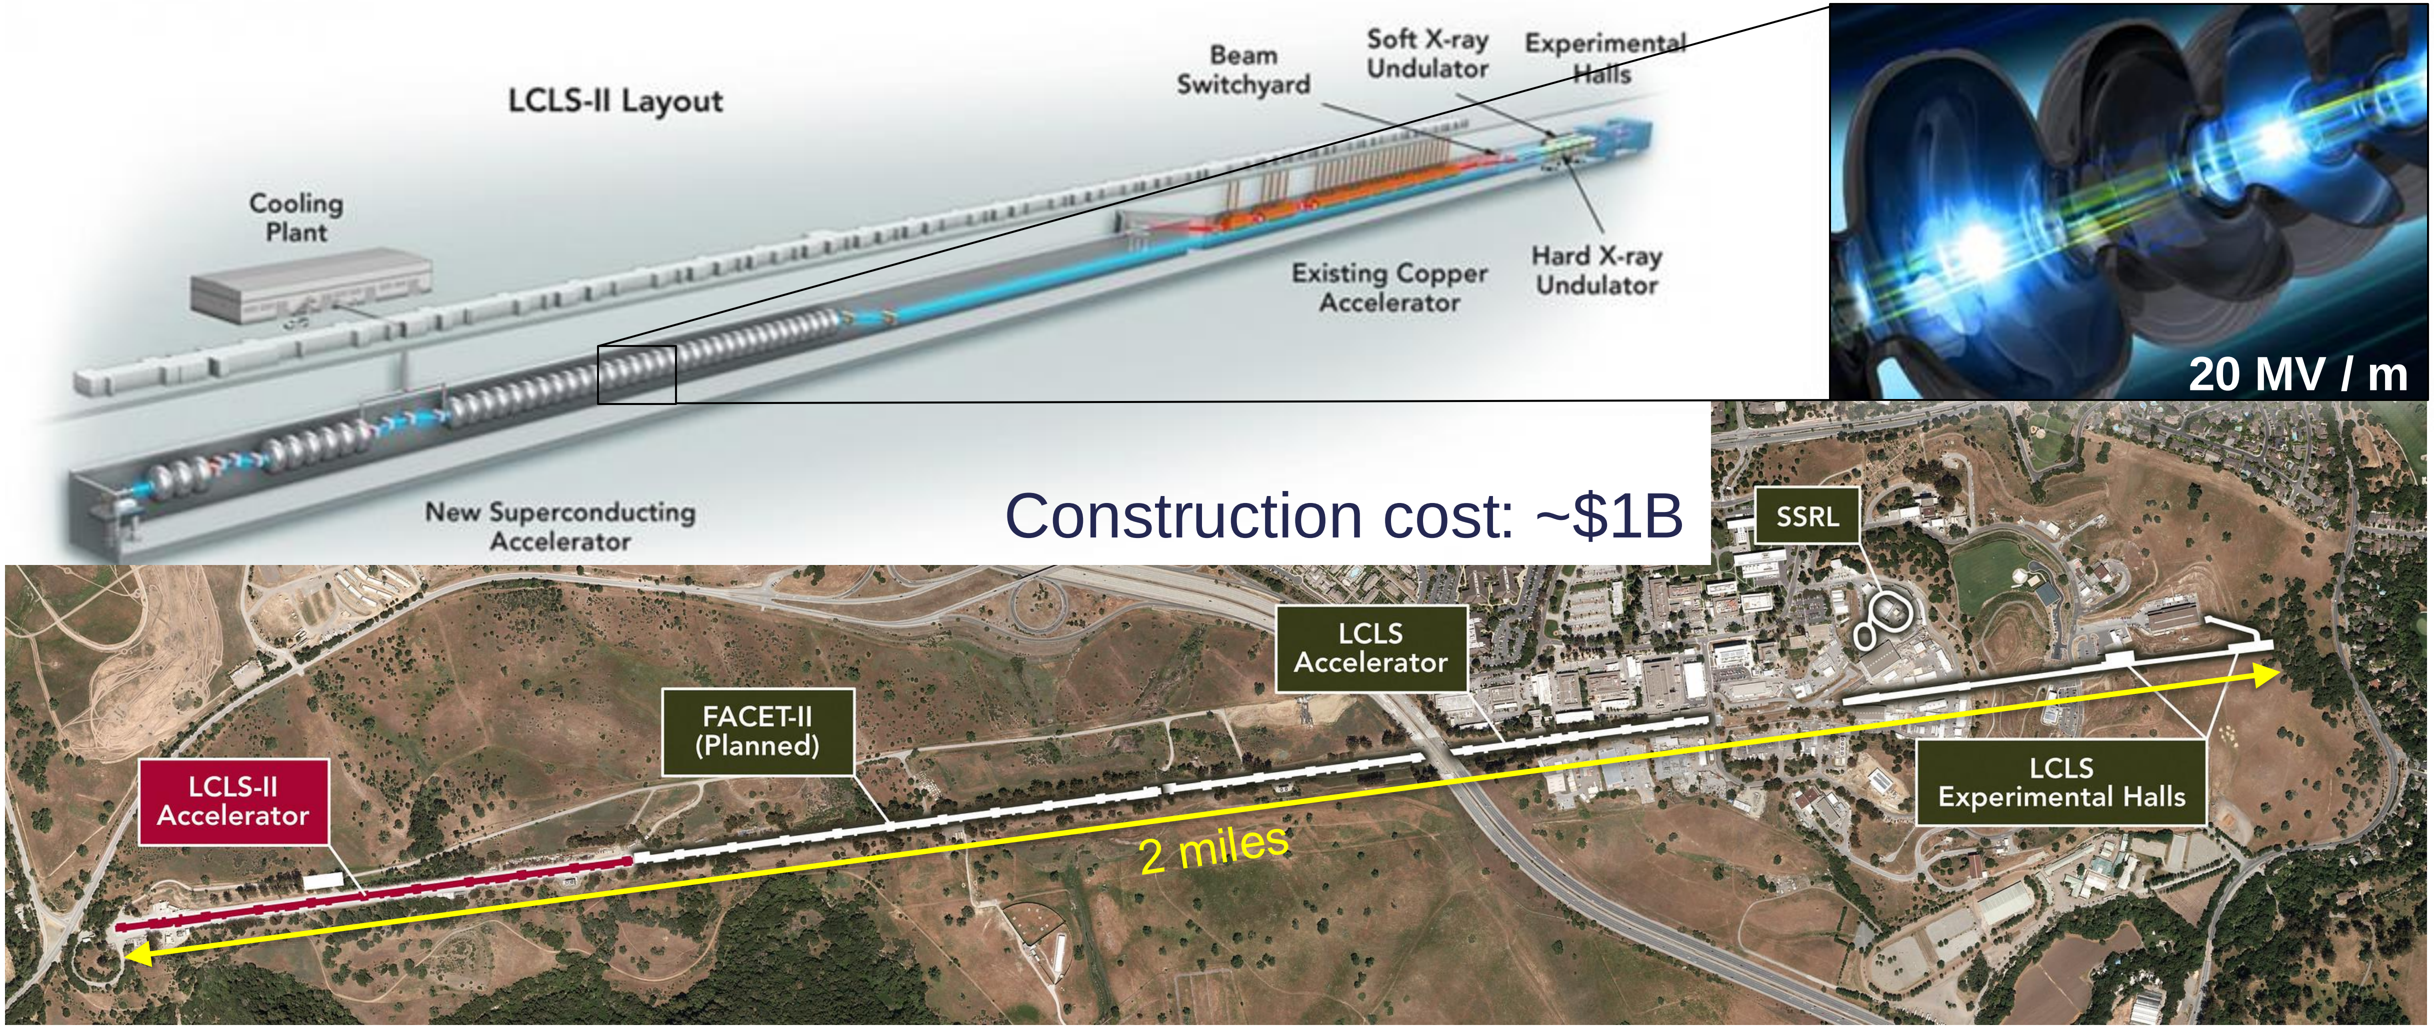

Construction cost: ~$1B

Why are particle accelerators so large?

Particle accelerators typically involve a metallic pipe with vacuum inside.

⇒ breakdown occurs if electric field is too high!

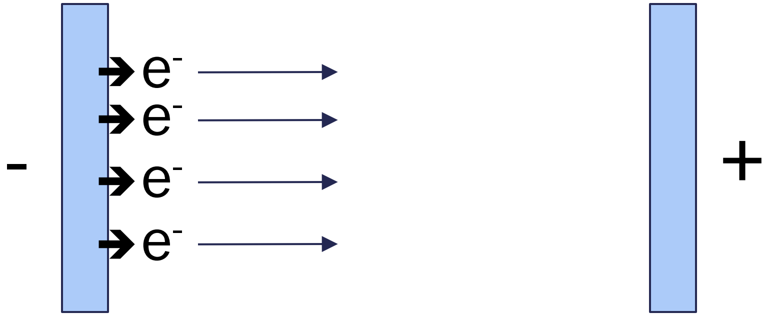

Possible solution to reach higher accelerating fields?

⇒ plasmas.

Plasma as a solution to shrink plasma accelerators

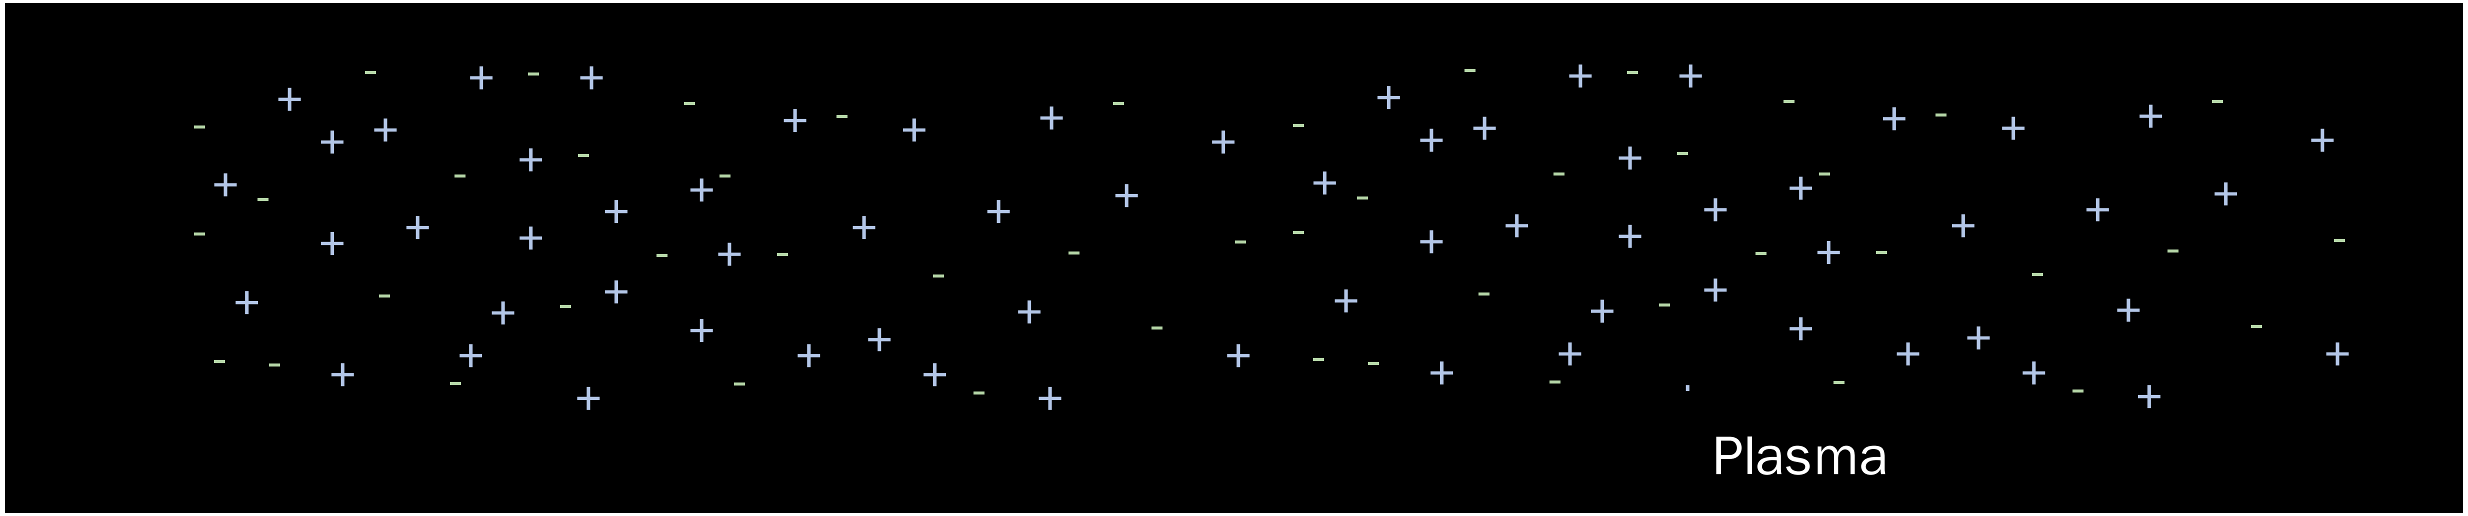

Plasma as a solution to shrink plasma accelerators

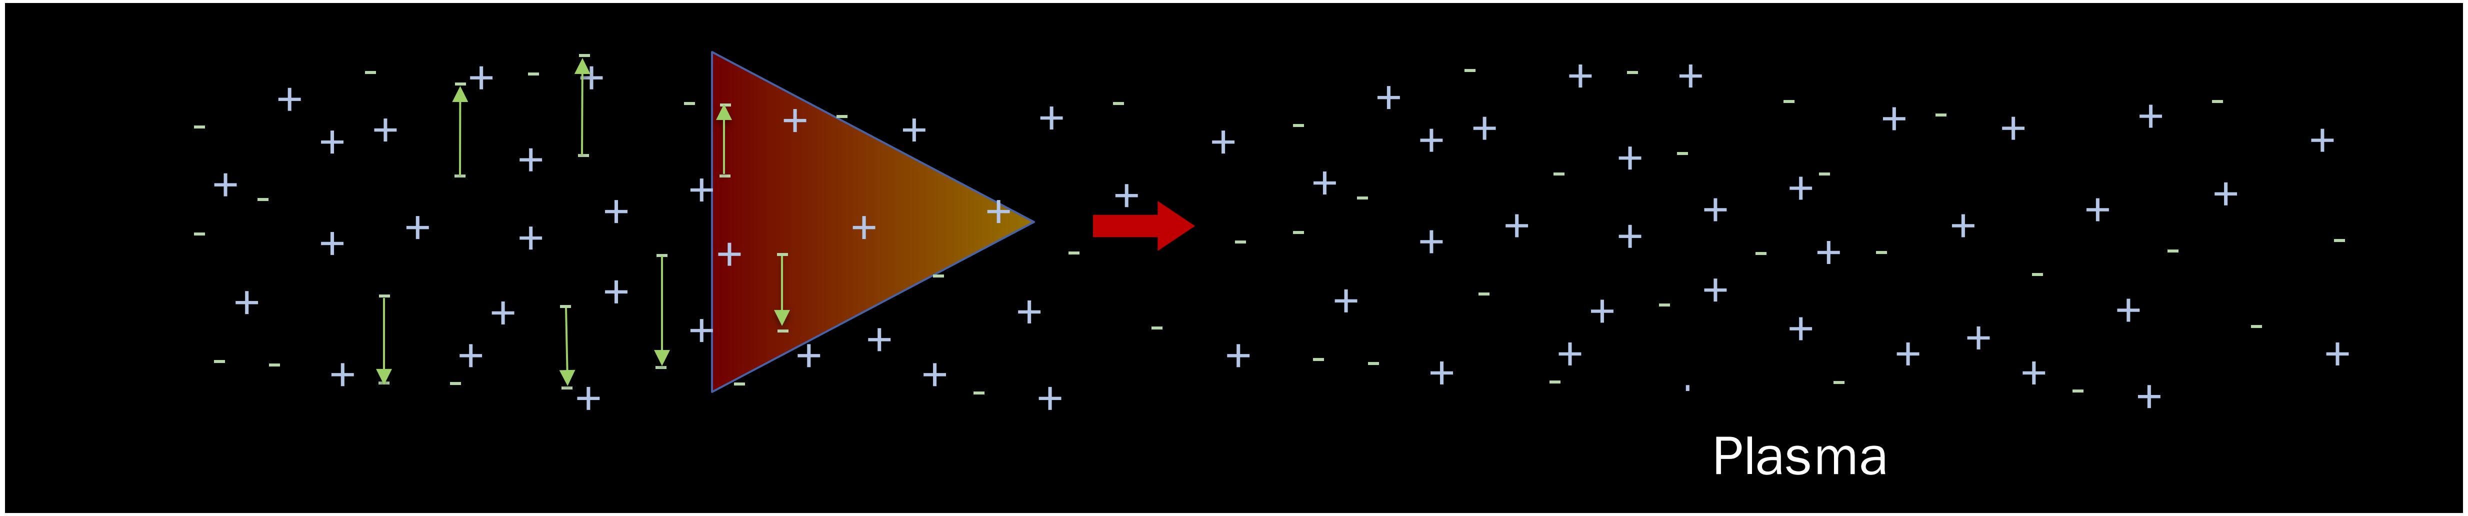

The separation of electrons and protons creates electric fields **orders of magnitude larger** than in conventional particle accelerators

➔ opportunity to accelerate (and guide) charged particle beams **over much shorter distances**

T. Tajima & J. M. Dawson, "Laser Electron Accelerator", *Phys. Rev. Lett.* **43**, 267 (1979)

Particle Accelerators can be Very Large - Can We Shrink Them?

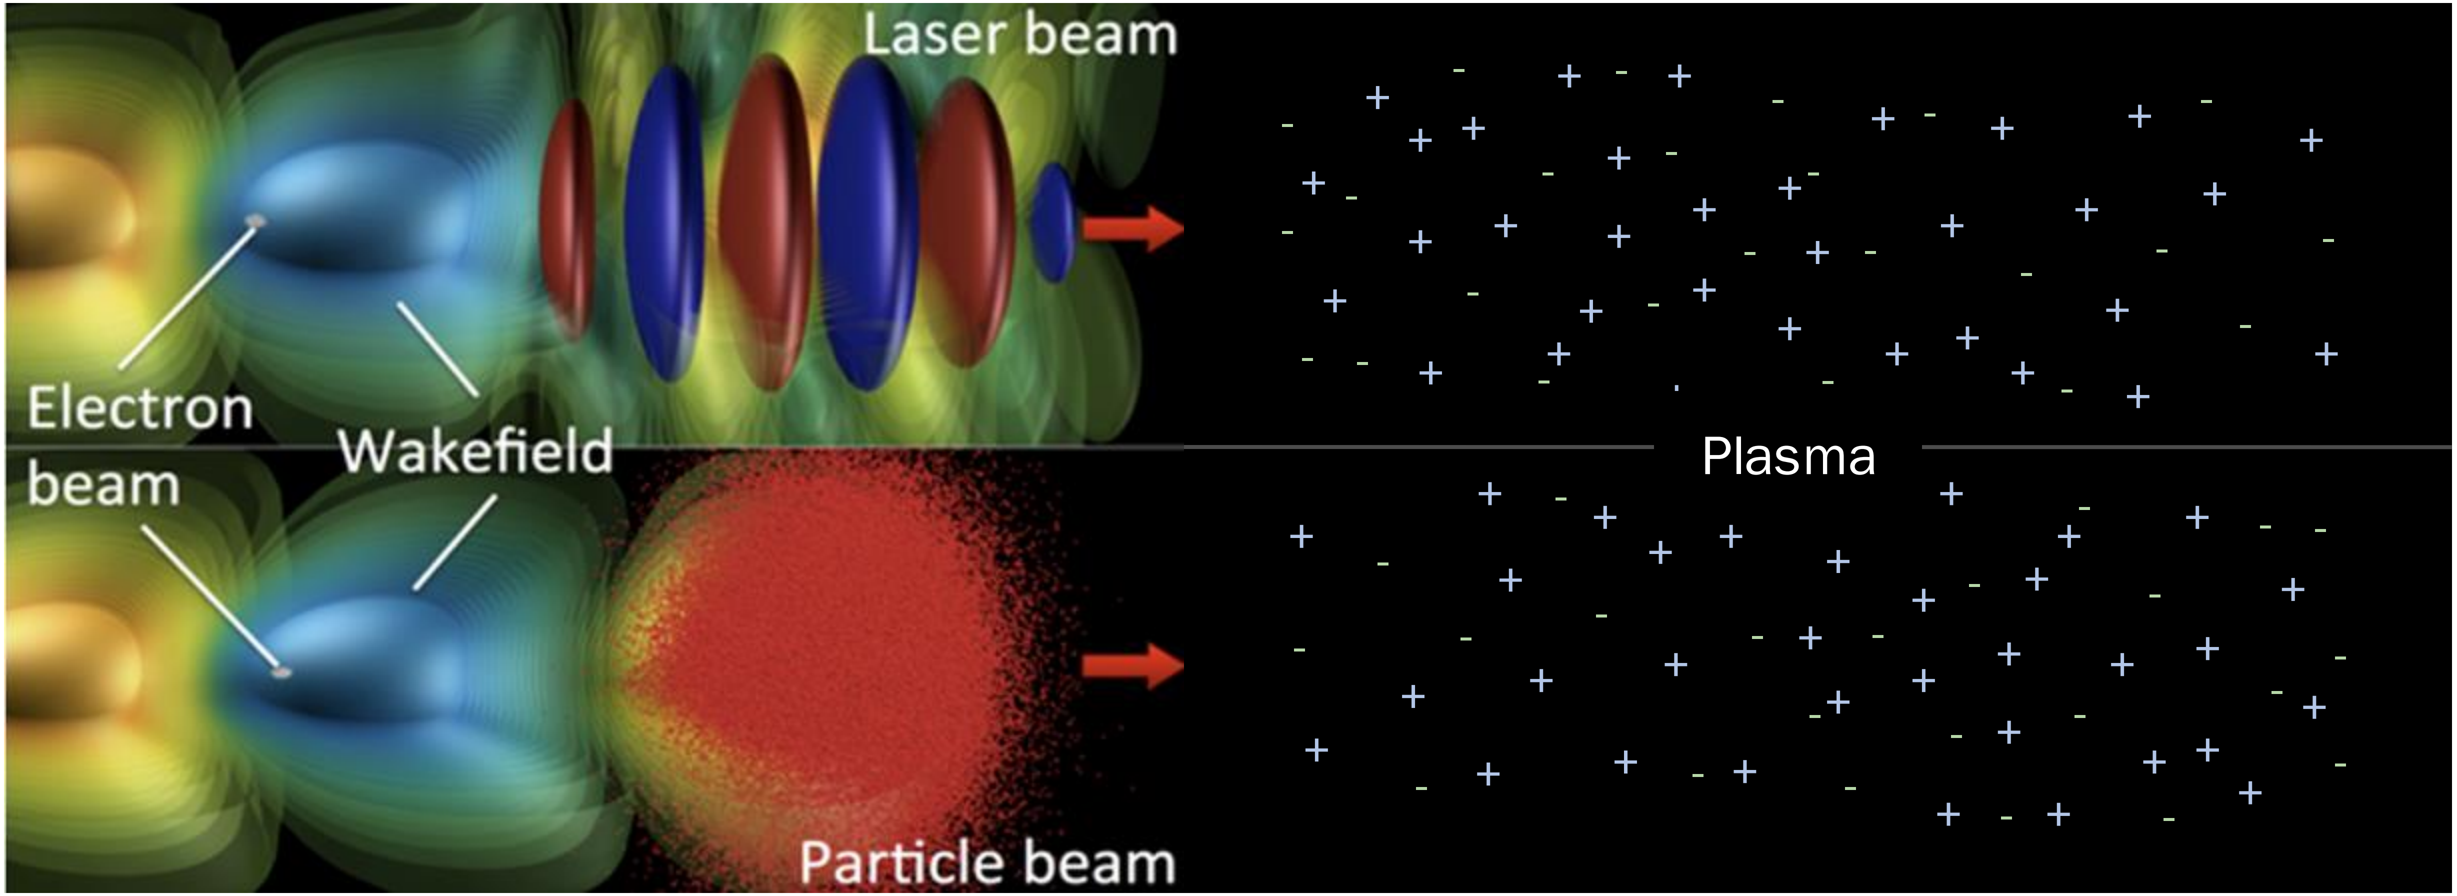

T. Tajima & J. M. Dawson, "Laser Electron Accelerator”, *Phys. Rev. Lett.* **43**, 267 (1979)

There are many designs & ideas for future Higgs factories and 10 TeV colliders

kms<length/circumference<90kms | \$5B<construction cost<\$50B | 100MW<consumption<500MW*

Computer simulations is essential to particle accelerator R&D to minimize size, cost & consumption

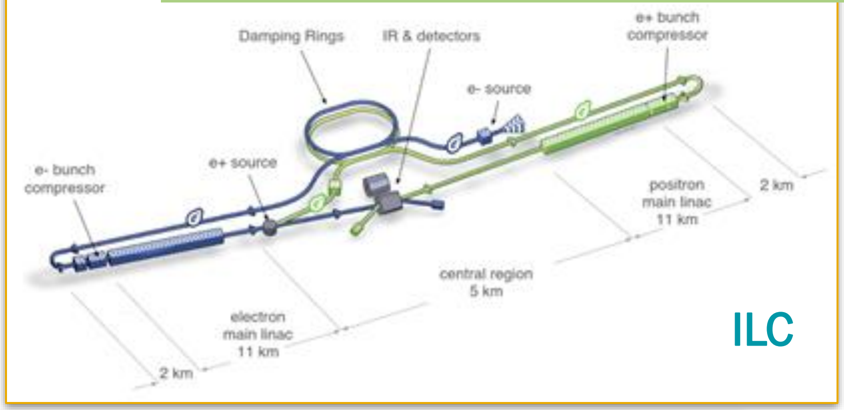

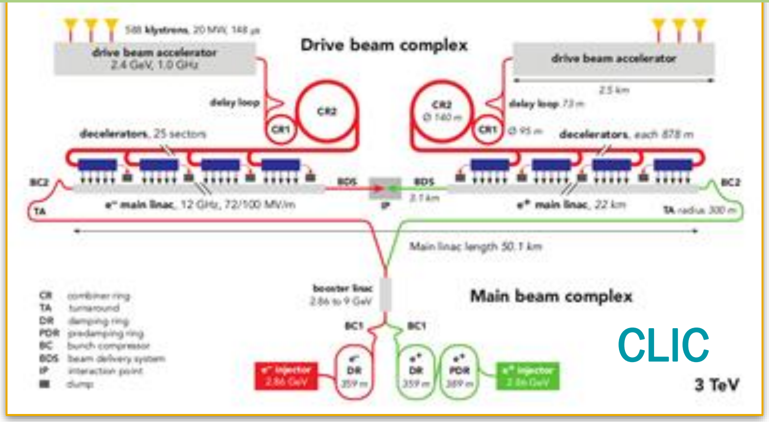

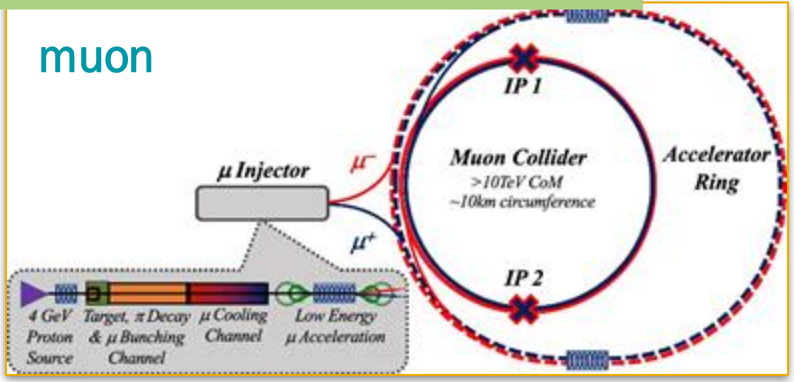

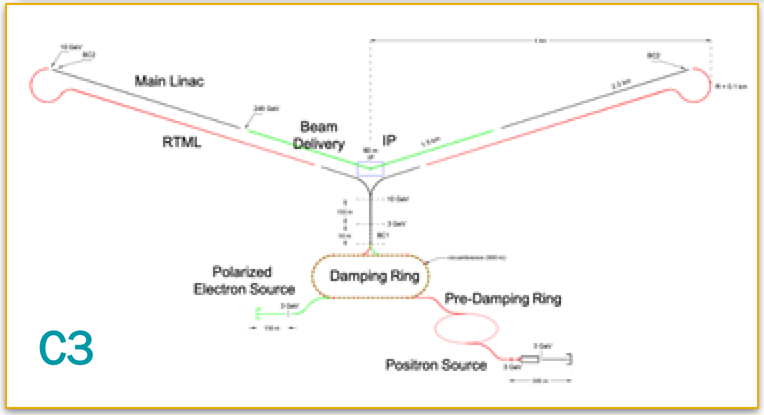

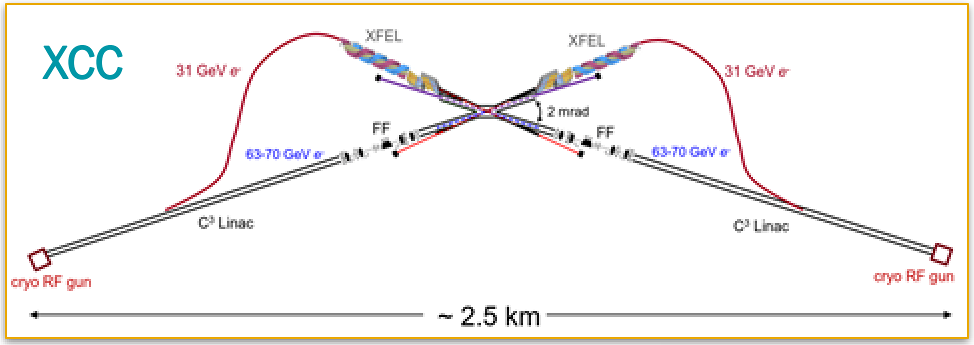

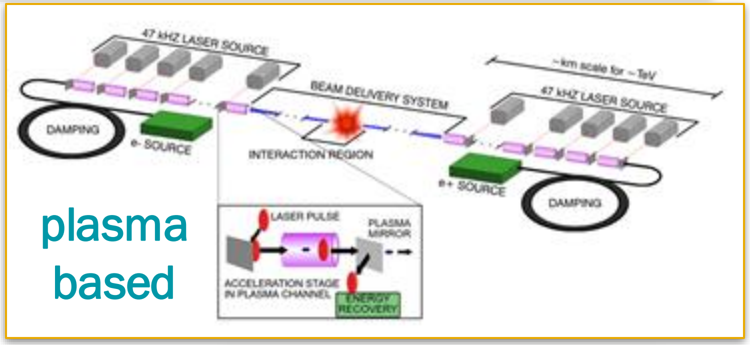

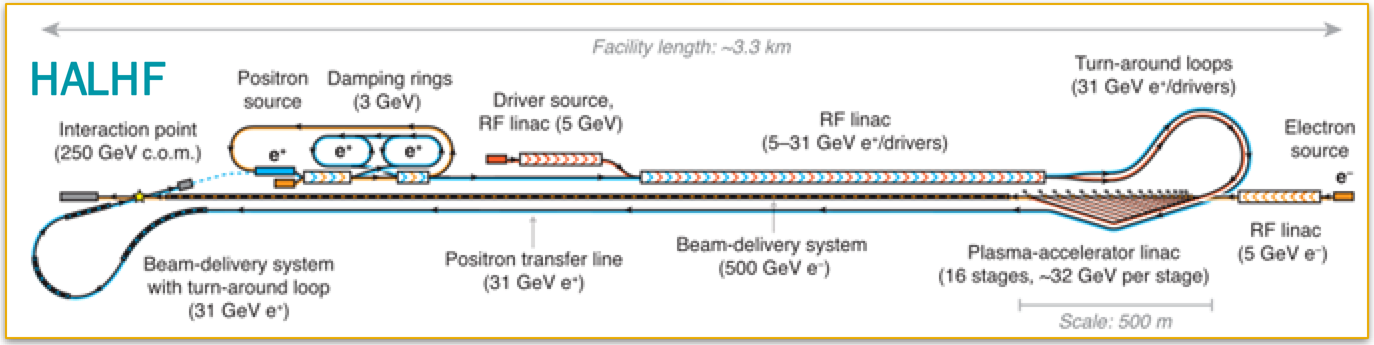

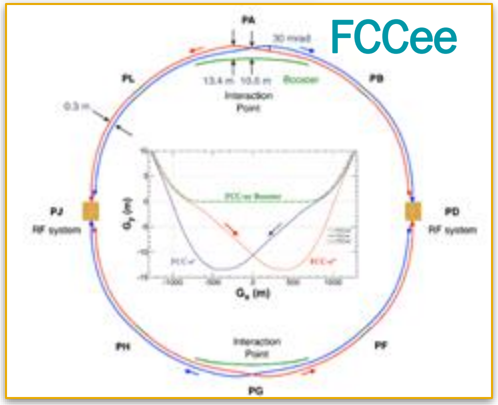

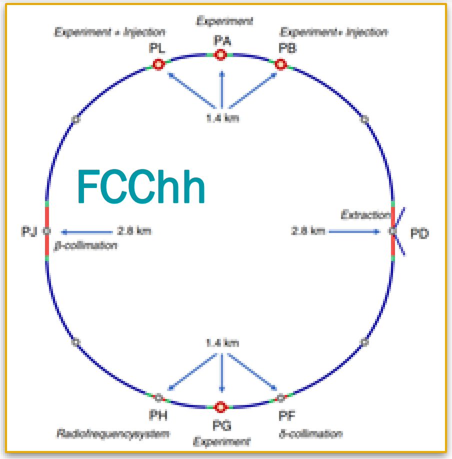

*T. Roser, et al, "On the feasibility of future colliders: report of the Snowmass'21 Implementation Task Force", arXiv 2208.06030 (2023)

Accelerator Modeling is Very Complex

Involves the modeling of the intricate interactions of

- **relativistic particles**: beams, plasmas, halo, stray electrons
- **EM fields**: accelerating/focusing fields, beam self-fields, laser/plasma fields
- **structures**: metals, dielectrics.
- **periodic structures & motion:** resonant coupling, instabilities.

Typical computer representations based on the Particle-In-Cell method:

- **particles:** macro particles representing each 1-10^6 particles
- **fields:** electromagnetic, on a grid
- **structures:** surfaces interacting with grid and macroparticles

Many space- and time scales to cover:

- from **µm** (e.g., plasma structures, e^--surface interactions) **to km** (e.g., LHC)
- from **ns** (beam passing one element) **to seconds or more** (beam lifetime)

⇒ needs best algorithms on largest & fastest computers

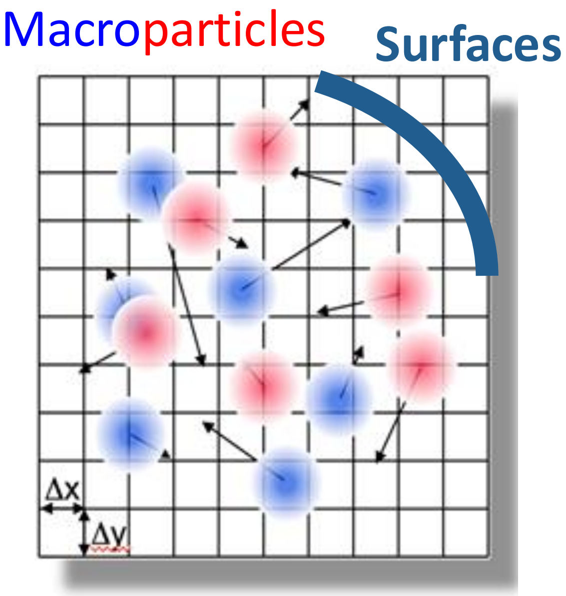

electromagnetic (EM) fields on a grid

Modeling of Particle Accelerators @ NERSC: Examples

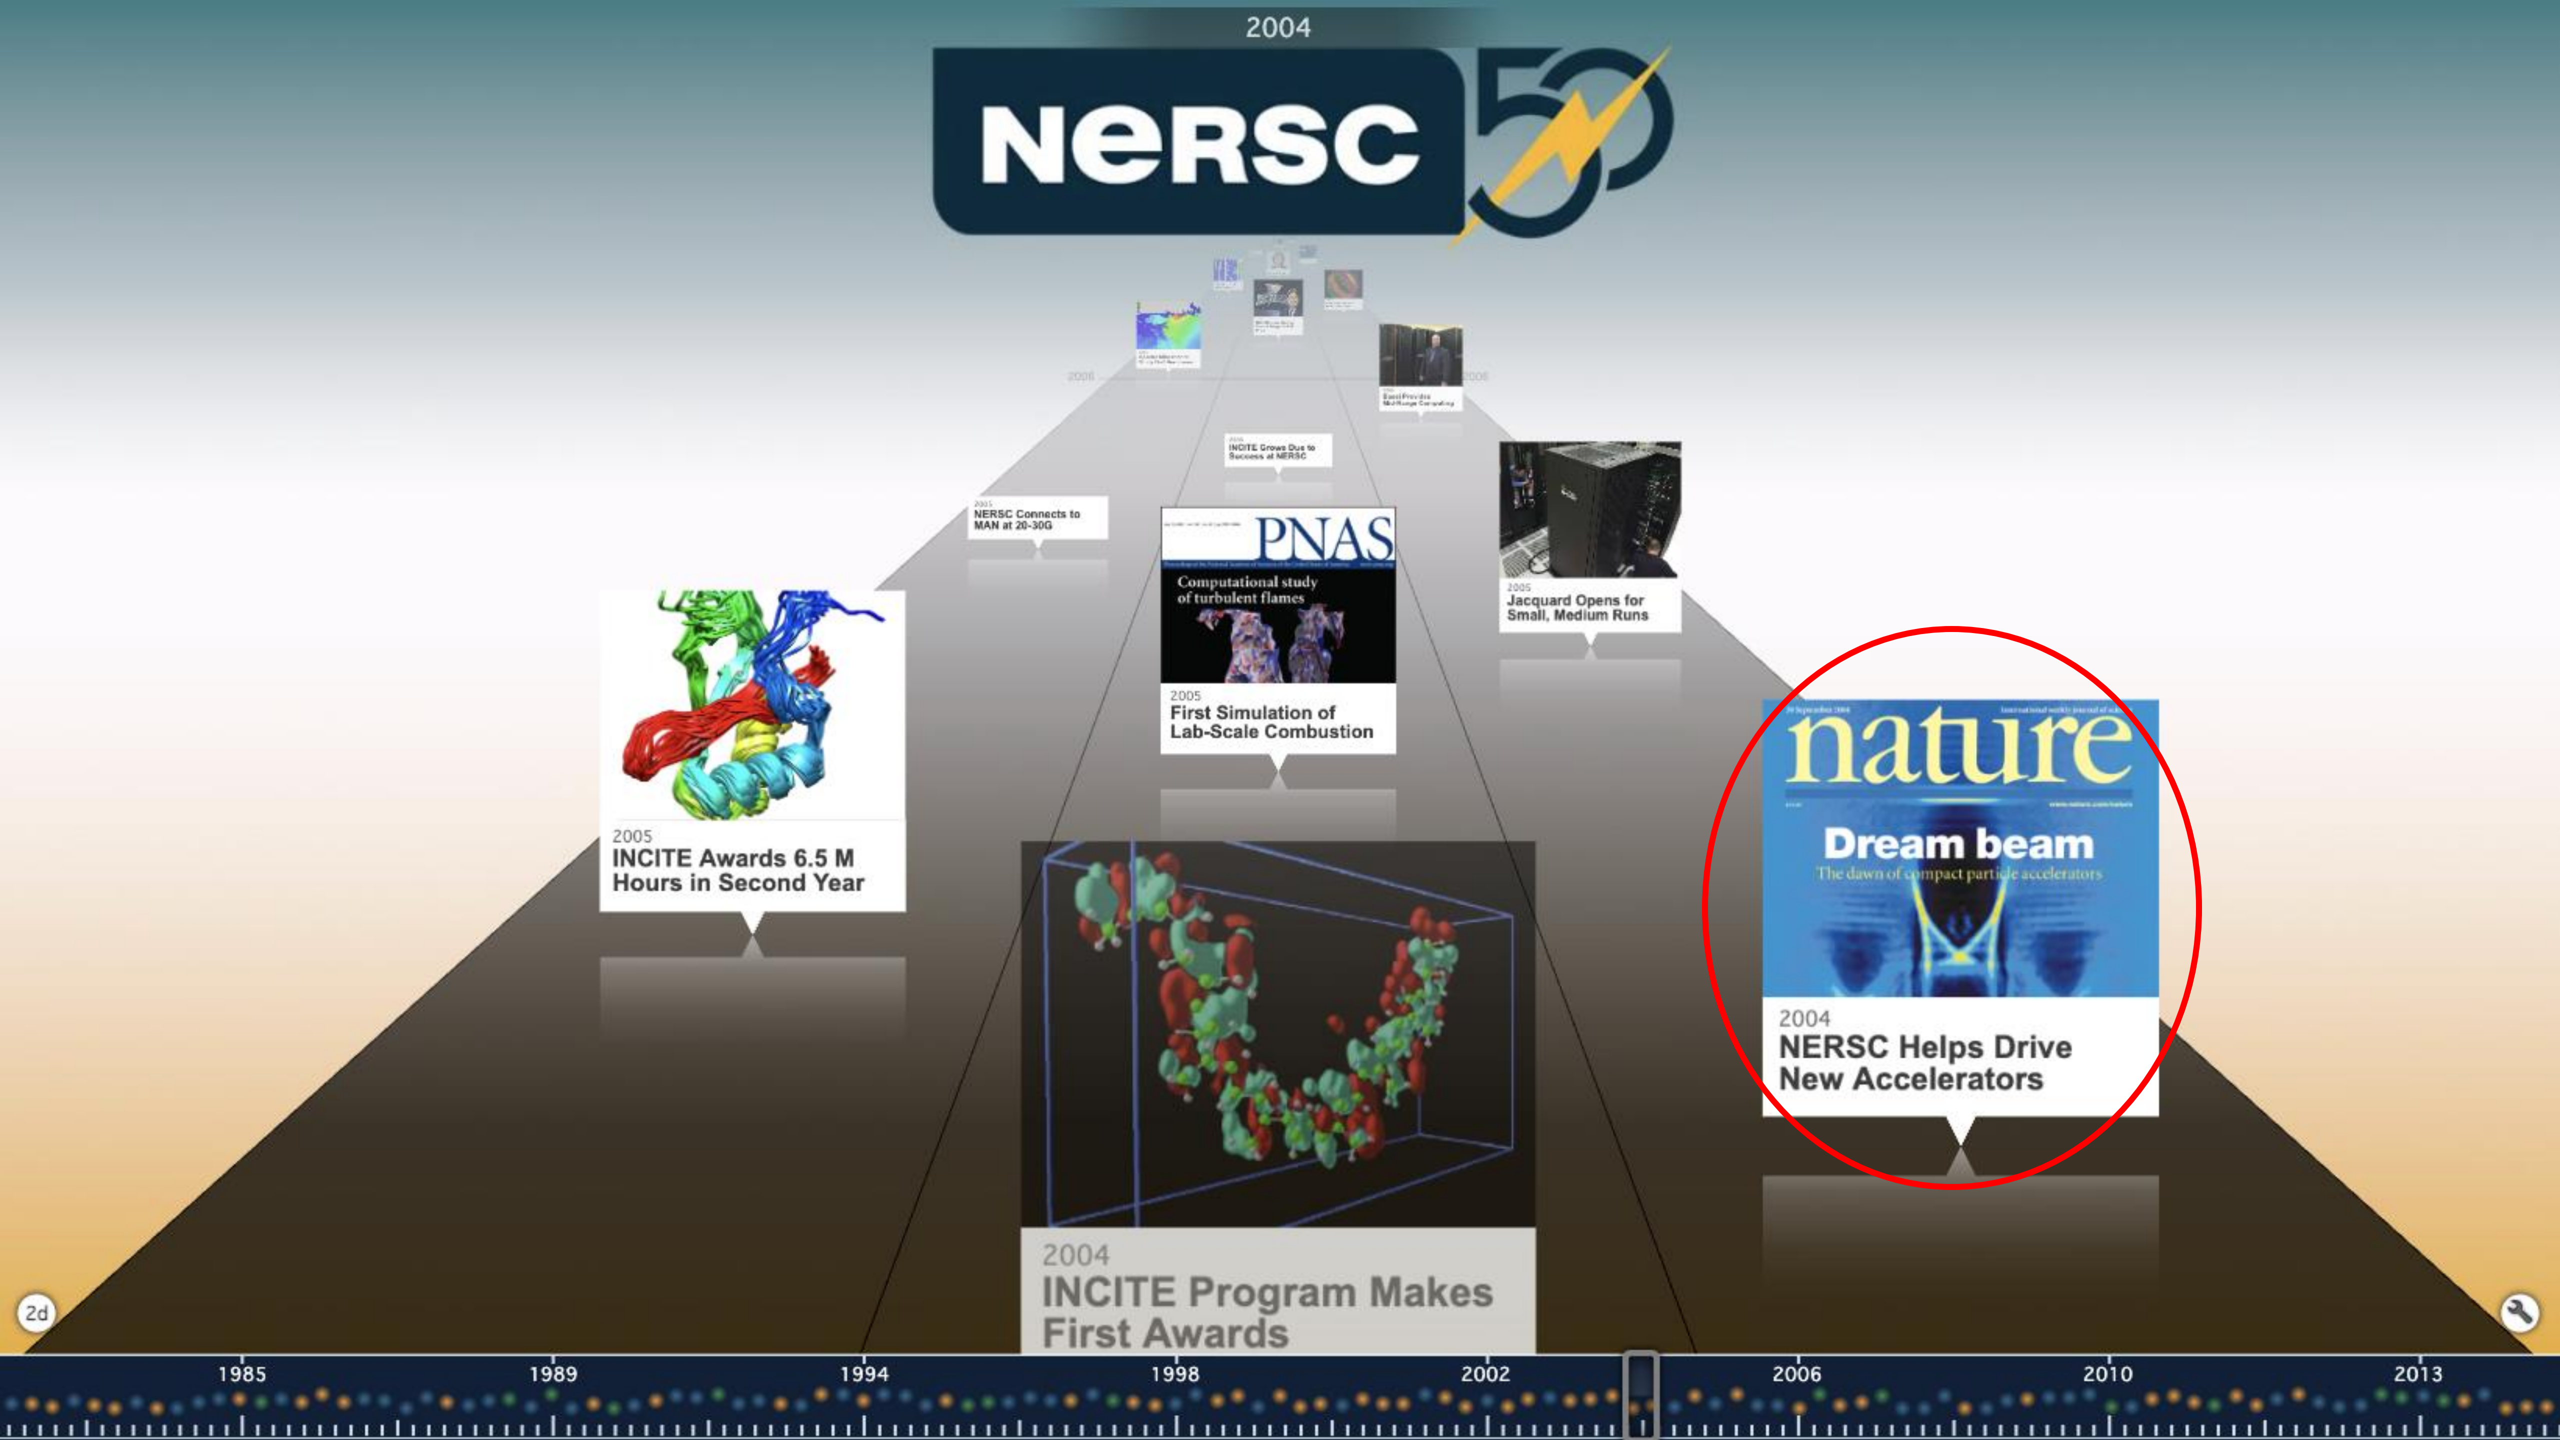

2004

2005
NERSC Connects to MAN at 20-30G

INCITE Grows Due to Success at NERSC

Bassi Provides Mid-Range Computing

2005
INCITE Awards 6.5 M Hours in Second Year

2005
First Simulation of Lab-Scale Combustion

2005
Jacquard Opens for Small, Medium Runs

2004
INCITE Program Makes First Awards

2004
NERSC Helps Drive New Accelerators

NERSC simulations unveiled the physics of narrow energy spread in high-gradient plasma-based particle accelerators transforming the field

Simulations provided key support to 2004 experimental observation at BELLA
of beams with particles near a single energy – critical to applications –

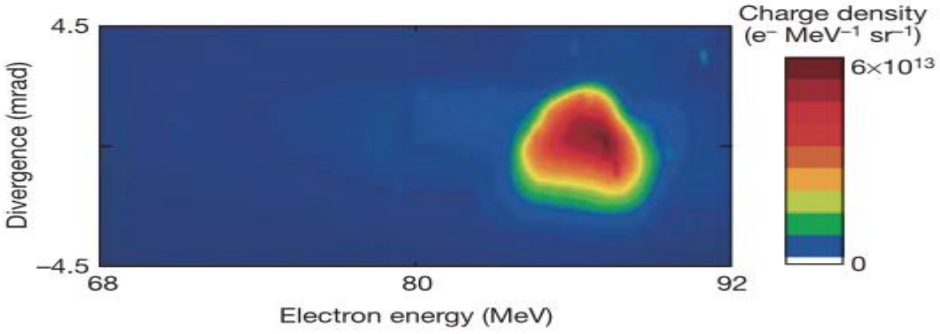

- BELLA (LBNL)
- Code: Vorpal
- NERSC computer: Seaborg
- Qualitative picture: 16-64 cores
- Accurate 3D trapping (INCITE): 1024 cores

Simulating the Inaccessible, Discovering the Unknown

SciDAC REVIEW

Number 13 | Scientific Discovery through Advanced Computing | Summer 2009

Advanced Computing

Simulations Drive Next-Generation Technology

"Supercomputers are the only tools that can be used to reconcile problems we find through our virtual approach."
Dr. Dimitri Kusnezov, p7
Director of the Office of Research and Development for National Security Science and Technology

An intense laser excites a wave in electron density in a plasma

30 September 2004

International weekly journal of science

nature

www.nature.com/nature

Dream beam

The dawn of compact particle accelerators

Disease control Europe plays catch-up

Protein folding Escape from the ribosome

The Earth's hum Sounds of air and sea

Human ancestry One from all and all from one

technology feature RNA interference

Geddes et al, *Nature* **431**, 540 (2004); *SciDAC Review* **13**, 13 (2009)

2009: Simulations uncover default of fabrication

2009: Simulations + UQ uncovered reasons for acceleration cavity measurements that were off specs & observed Beam Breakup (BBU) instability in operation
➔ the 1 meter cavity was 8mm shorter than designed

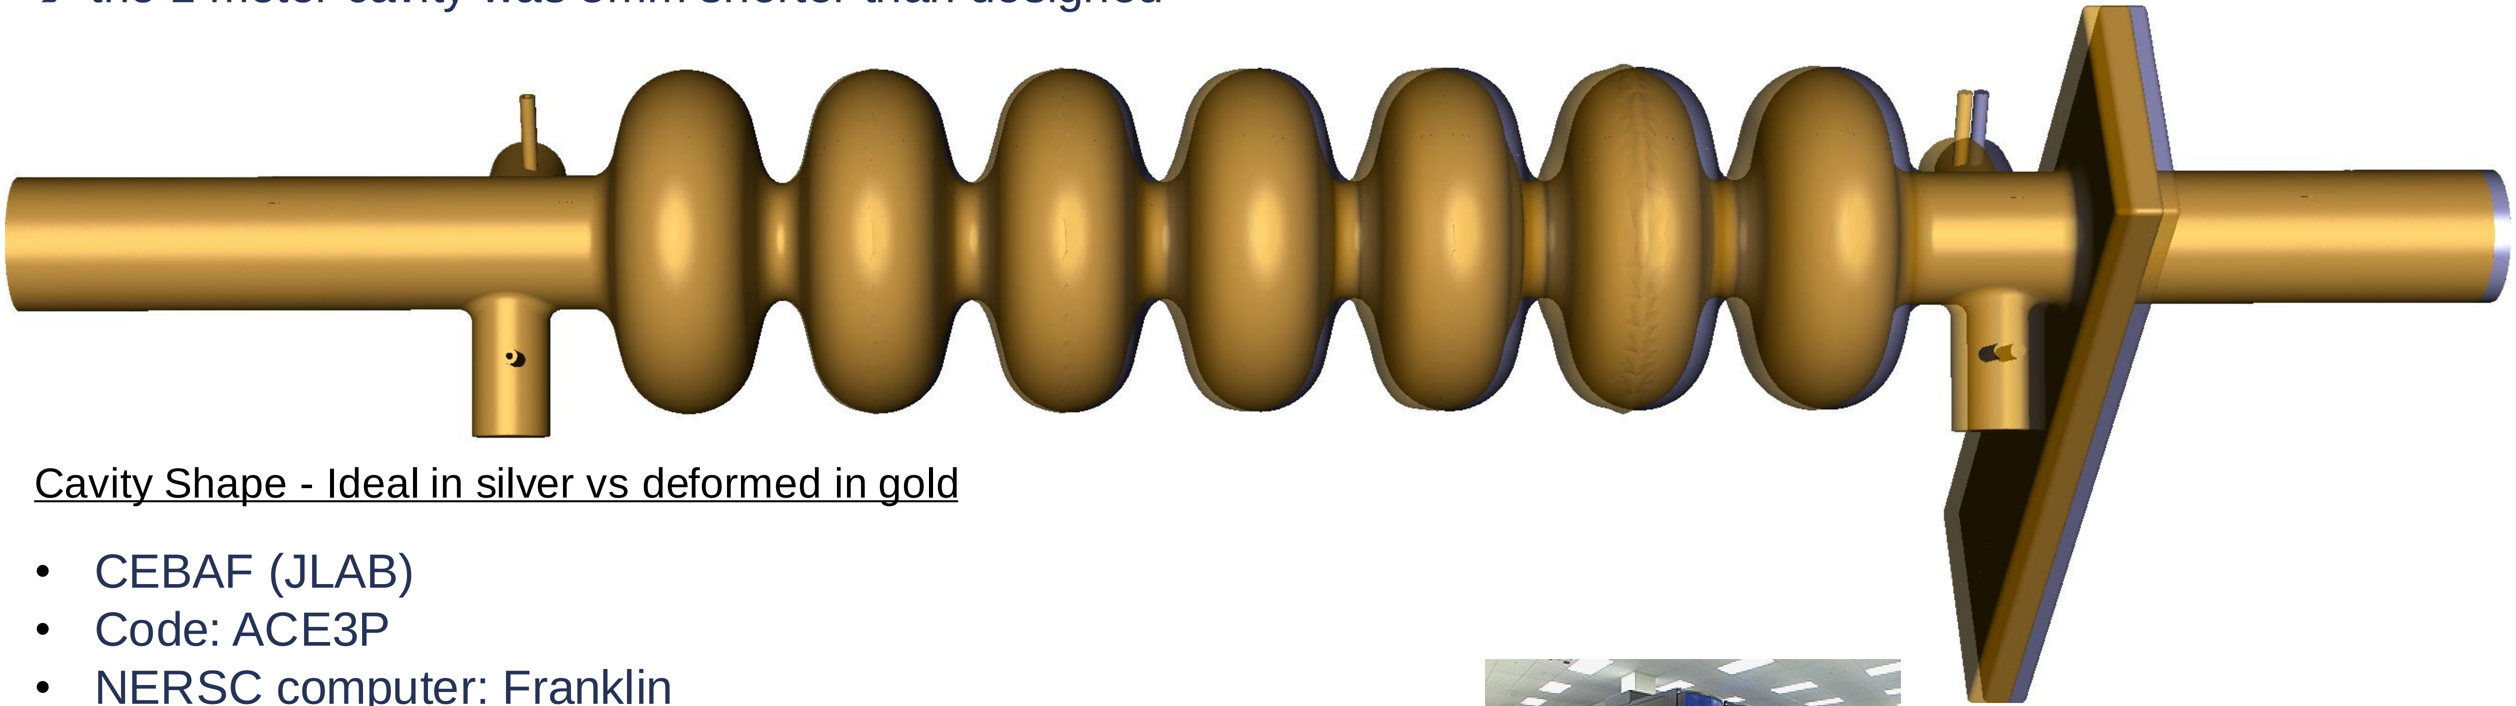

Cavity Shape - Ideal in silver vs deformed in gold

- CEBAF (JLAB)
- Code: ACE3P
- NERSC computer: Franklin
- Individual runs used 256 cores
- 37k CPU hours needed to obtain solve the inverse problem from measured data & recover the deformed cavity

2010: Simulation of electron cloud (e-cloud) effects in CERN SPS

Electrons can interact with beam & beam pipe:

➔Multiply ➔ enter a resonant “headtail” instability

➔ spoils the beam quality

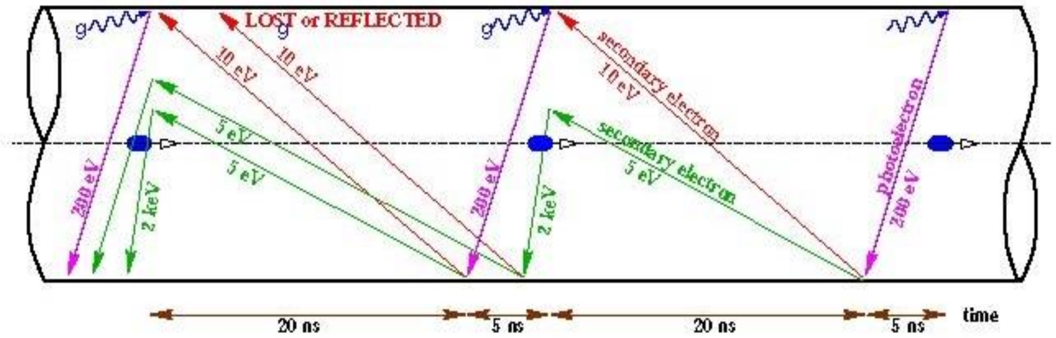

2010:
first fully self-consistent coupled (Warp+Posinst) simulations of e-cloud build w/ beam interaction

Turn 1

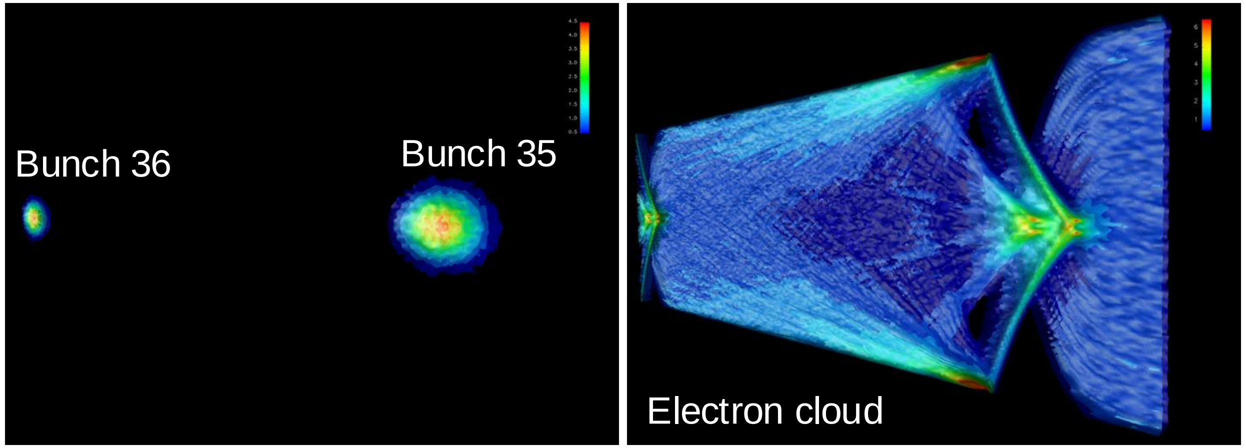

Turn 500

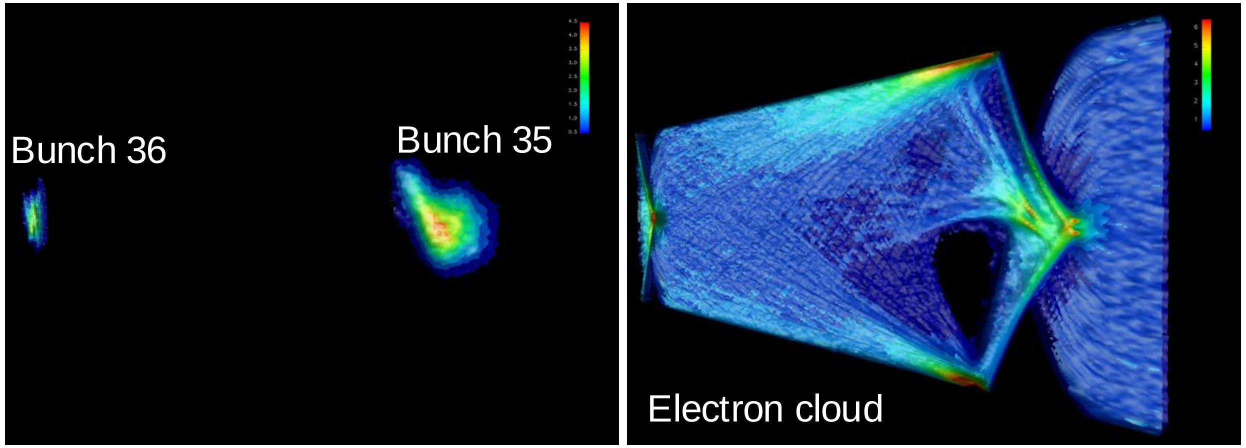

- SPS (CERN); circumference=6.9km
- 3 batches of 72 bunches/batch (=216 bunches)
- 1000 turns (w/ 10 stations/turn)
- Code: Warp; # cores: 9,600
- NERSC computer: Franklin

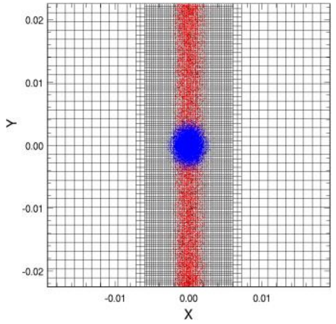

Mesh refinement enabled simulations in reasonable time (~10h)

SciDAC
Scientific Discovery through Advanced Computing

BERKELEY LAB
LAWRENCE BERKELEY NATIONAL LABORATORY

2013: simulations started to use the physical number of particles
2017: application to modeling of LCLS with excellent agreement w/ expt

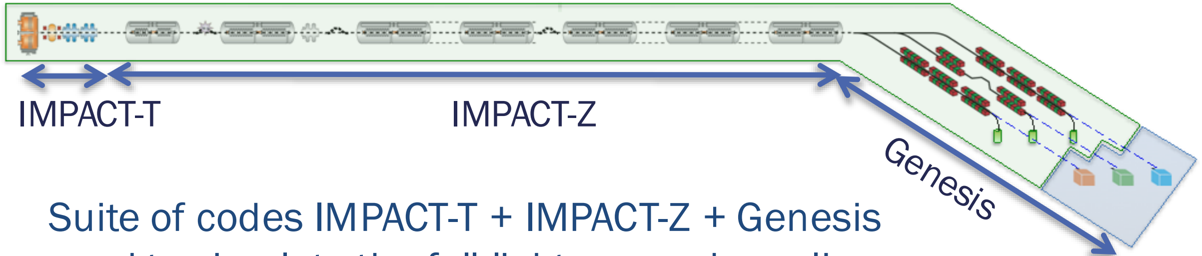

Suite of codes IMPACT-T + IMPACT-Z + Genesis used to simulate the full light source beamline

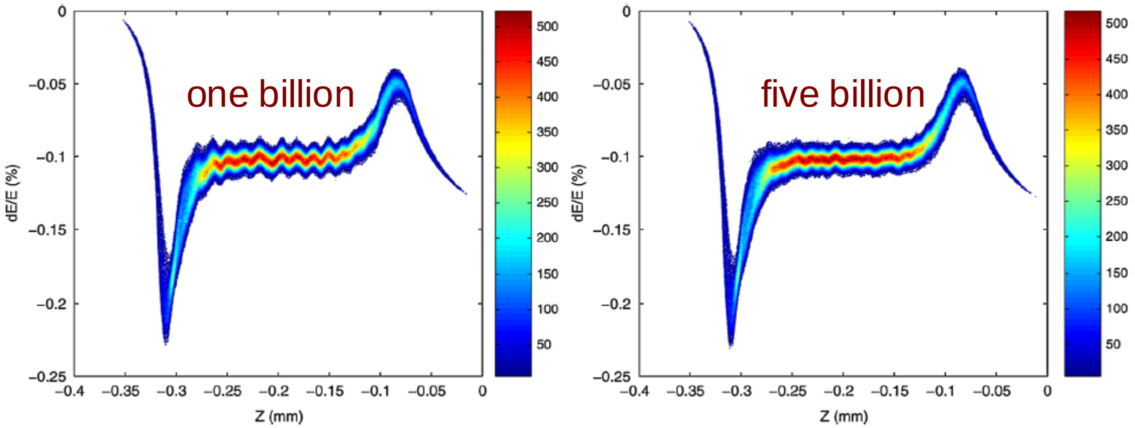

Using (for the first time) the real number of electrons matters to get the shot noise right!

J. Qiang *et al.*, *Phys. Rev. Accel. Beams* **17**, 030701 (2013)

- LCLS (SLAC);
- # cores: 2,048 for 14h
- NERSC computers: Hopper

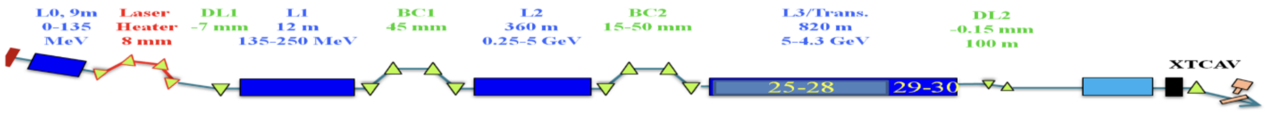

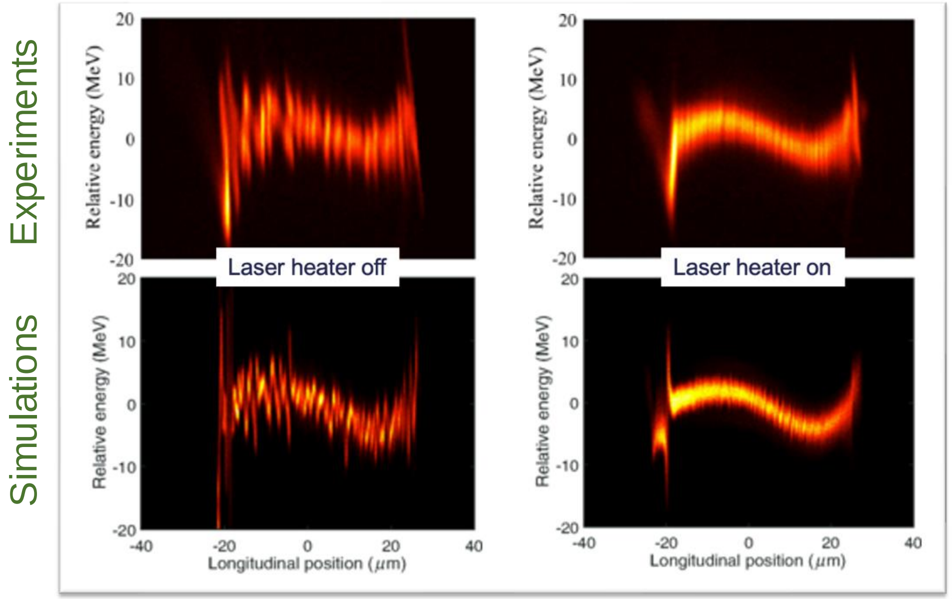

Start-to-end, one-to-one modeling reproduces microbunching in the LCLS X-ray FEL.

J. Qiang *et al.*, *Phys. Rev. Accel. Beams* **20**, 054402 (2017).

- LCLS (SLAC);
- # cores: 2,048 for 6h
- NERSC computers: Edison

BERKELEY LAB
LAWRENCE BERKELEY NATIONAL LABORATORY

SLAC

2016: Modeling of (unwanted) dark currents in LCLS-II cryomodule

2016

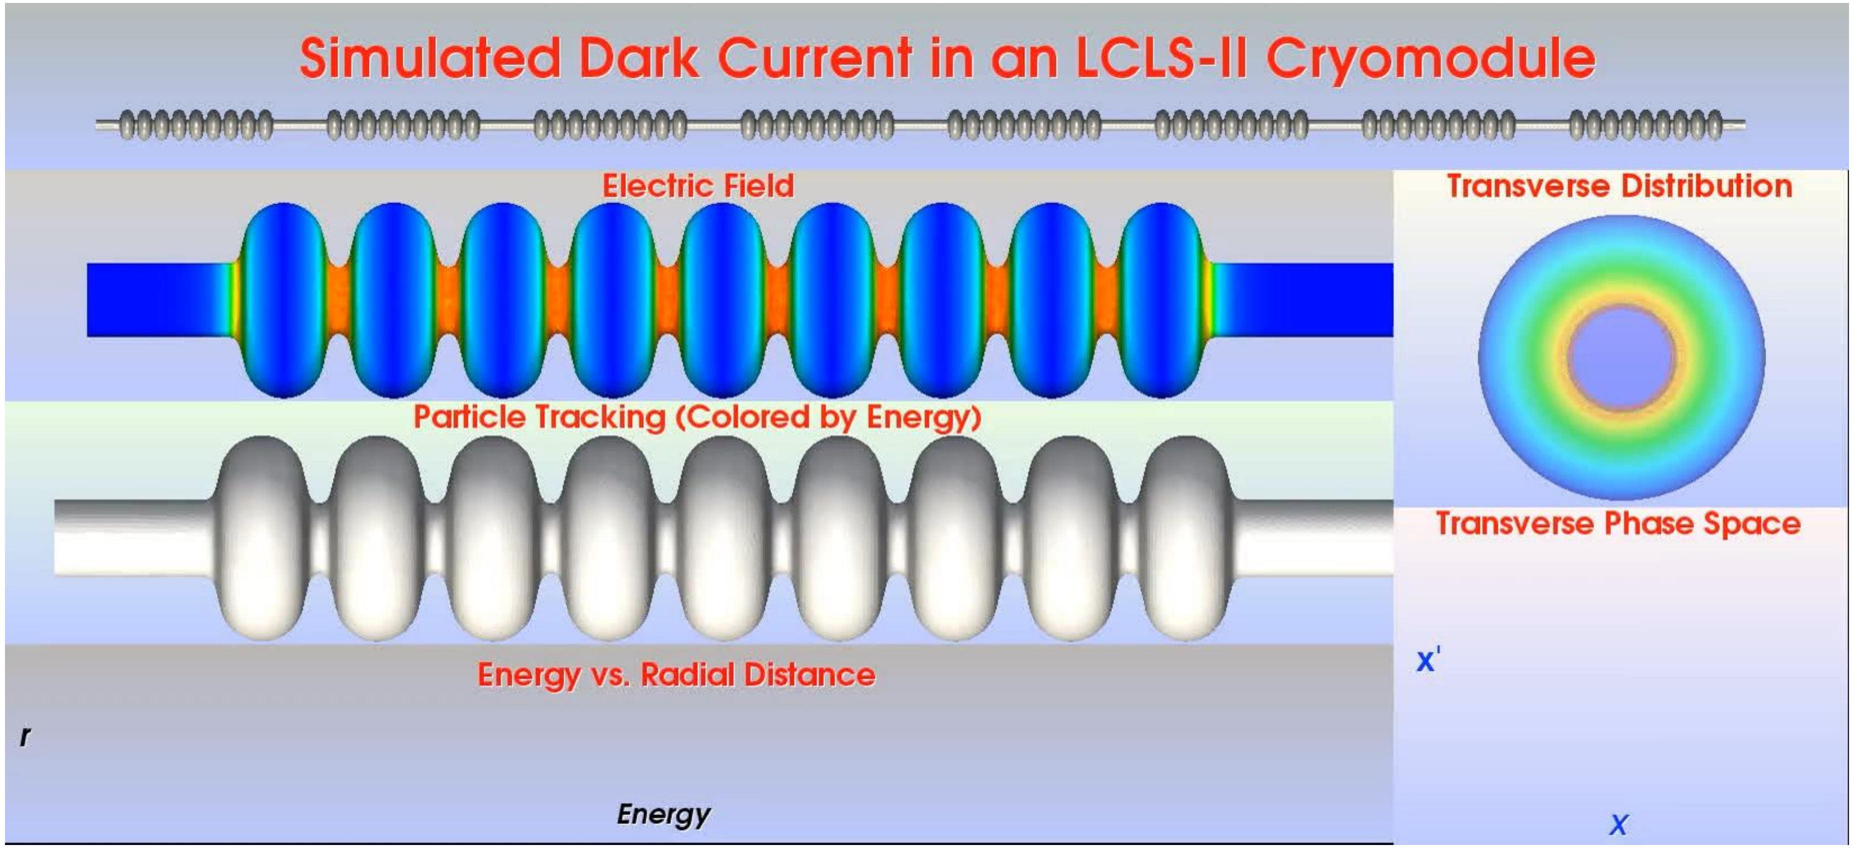

- LCLS-II (SLAC)
- Code: ACE3P (Track3P)
- NERSC computer: Edison
- Individual runs used 240 cores (10 nodes)

SciDAC
Scientific Discovery through Advanced Computing

SLAC

2007-2017:~2 orders of magnitude speedup in cryomodule modeling

Higher-order modes in TESLA cryomodules

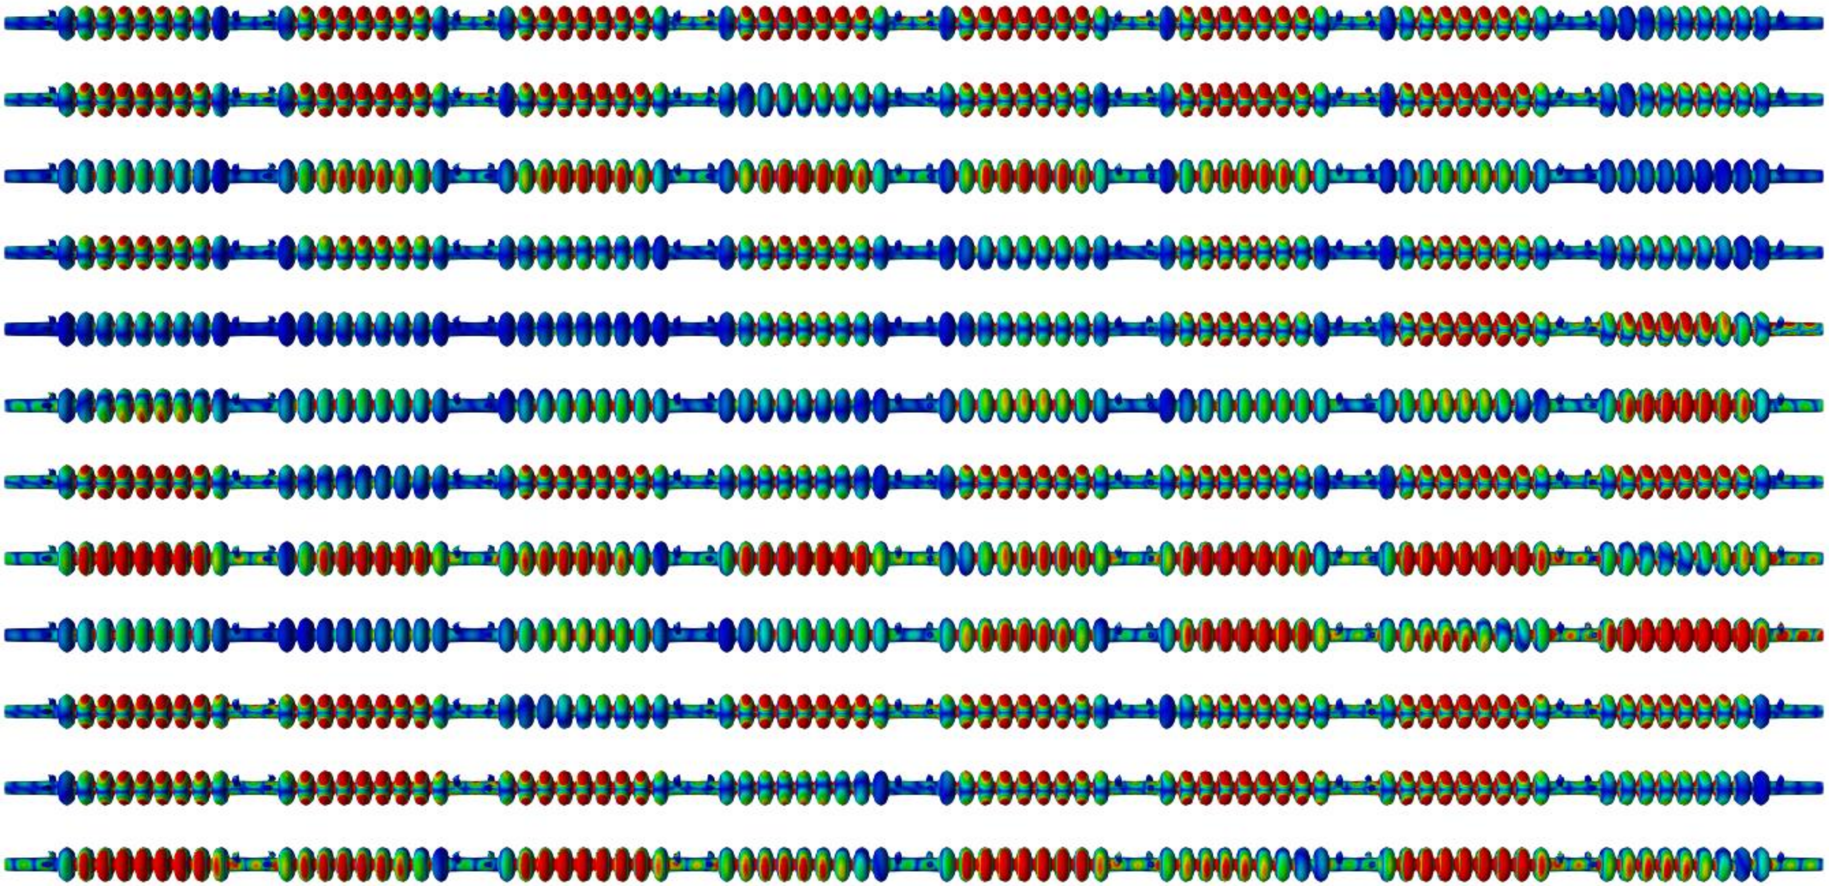

- *1 hour per mode* using 1500 cores on Seaborg in 2007
- *< 1 minute per mode* using 960 cores on Edison in 2017
- Speedup from advances in hardware and algorithms

Helping the transition to Exascale with NESAP

Getting Started in NESAP

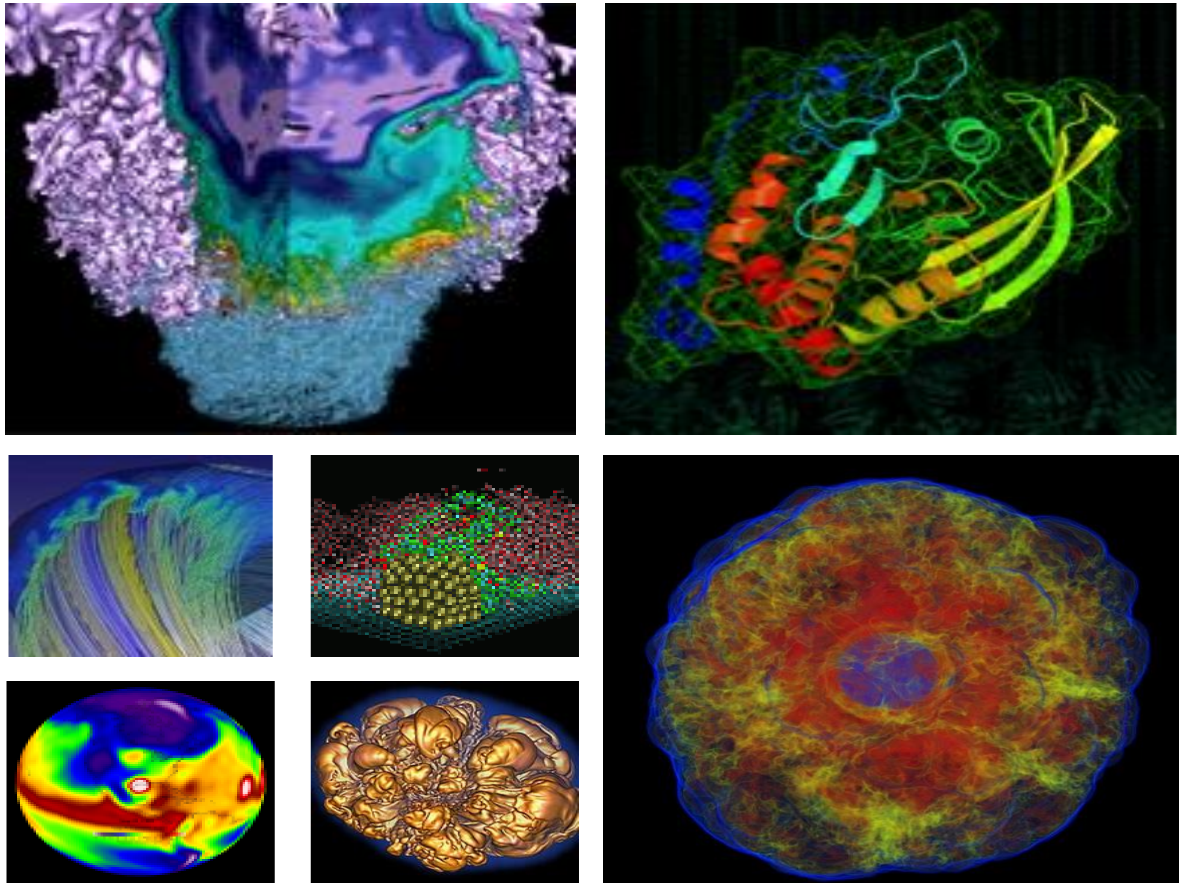

NERSC App Readiness Team

NESAP PI Briefing
September 15, 2014

20 NESAP Codes

ASCR (2)

Almgren (LBNL) – **BoxLib AMR Framework** used in combustion, astrophysics

Trebotich (LBNL) – **Chombo-crunch** for subsurface flow

BES (5)

Kent (ORNL) – **Quantum Espresso**
Deslippe (NERSC) – **BerkeleyGW**
Chelikowsky (UT) – **PARSEC** for excited state materials
Bylaska (PNNL) – **NWChem**
Newman (LBNL) – **EMGeo** for geophysical modeling of Earth

BER (5)

Smith (ORNL) – **Gromacs** Molecular Dynamics
Yelick (LBNL) – **Meraculous** genomics
Ringler (LANL) – **MPAS-O** global ocean modeling
Johansen (LBNL) – **ACME** global climate
Dennis (NCAR) – **CESM**

HEP (3)

Vay (LBNL) – **WARP & IMPACT-** accelerator modeling
Toussaint (U Arizona) – **MILC** Lattice QCD
Habib (ANL) – **HACC** for *n*-Body cosmology

NP (3)

Maris (U. Iowa) – **MFDn** *ab initio* nuclear structure
Joo (JLAB) – **Chroma** Lattice QCD
Christ/Karsch (Columbia/BNL) – **DWF/HISQ** Lattice QCD

FES (2)

Jardin (PPPL) – **M3D** continuum plasma physics
Chang (PPPL) – **XGC1** PIC plasma

NESAP Postdocs

Taylor Barnes
Quantum ESPRESSO

Brian Friesen
Boxlib

Andrey Ovsyannikov
Chombo-Crunch

Mathieu Lobet
WARP

Tuomas Koskela
XGC1

Tareq Malas
EMGeo

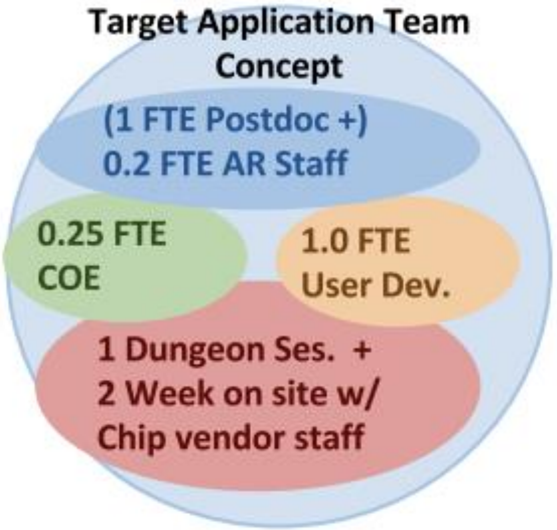

Target Application Team Concept

(1 FTE Postdoc +)
0.2 FTE AR Staff

0.25 FTE
COE

1.0 FTE
User Dev.

1 Dungeon Ses. +
2 Week on site w/
Chip vendor staff

2016: Performance of optimized vs original code

demonstrated 3-8X speedup for Warp

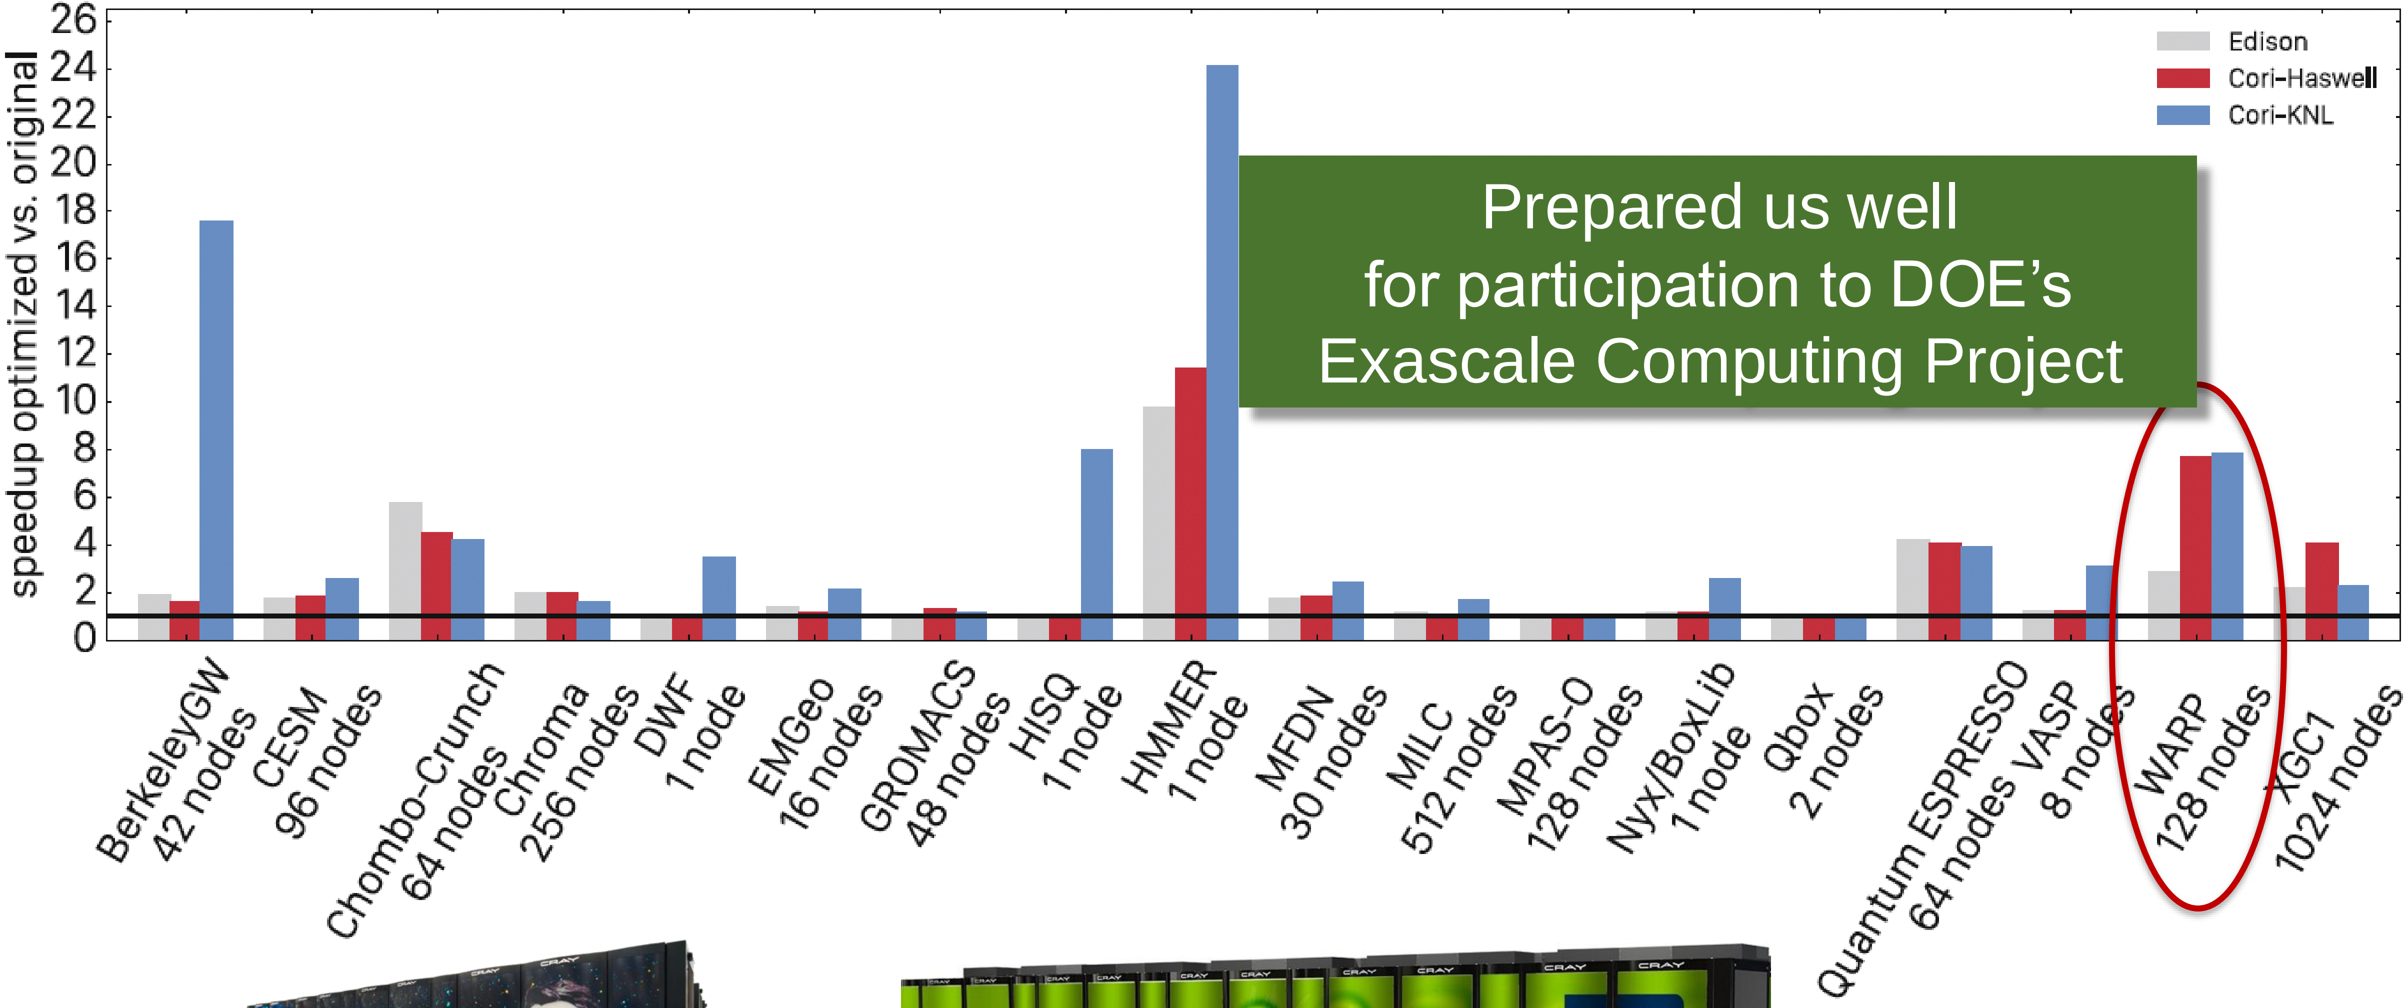

Prepared us well for participation to DOE's Exascale Computing Project

U.S. DOE Exascale Computing Initiative (ECI) – 2016-2023

WarpX among 21 applications selected to cover broad range of science

Exascale Computing Initiative

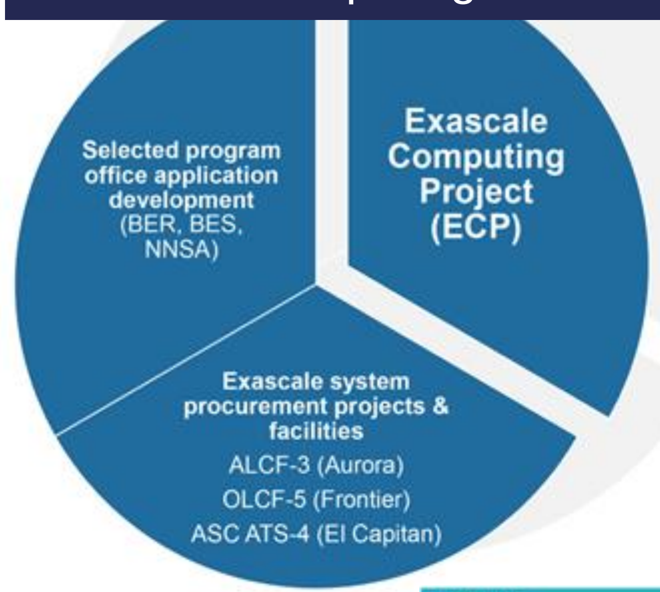

Hardware

5x Summit=1EF

Software

Applications

50x performance

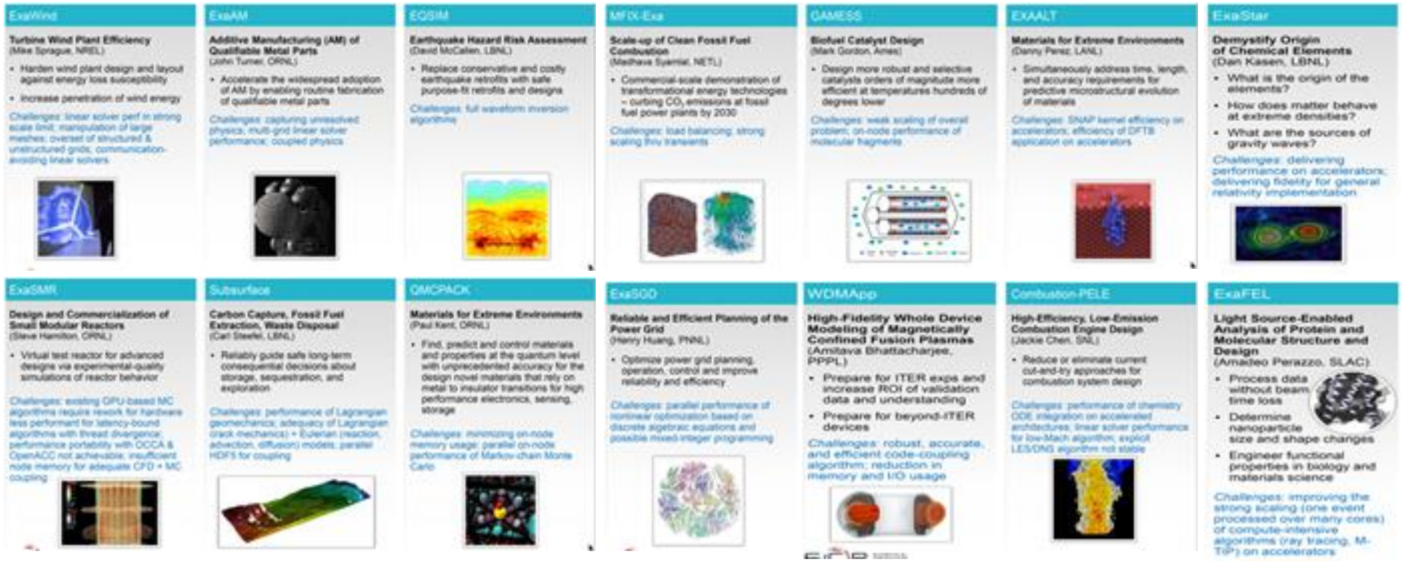

ExaWind
Turbine Wind Plant Efficiency
(Mike Sprague, NREL)
- Harden wind plant design and layout against energy loss susceptibility
- Increase penetration of wind energy

Challenges: linear solver perf in strong scale limit; manipulation of large meshes; overset of structured & unstructured grids; communication-avoiding linear solvers

ExaAM
Additive Manufacturing (AM) of Qualifiable Metal Parts
(John Turner, ORNL)
- Accelerate the widespread adoption of AM by enabling routine fabrication of qualifiable metal parts

Challenges: capturing unresolved physics; multi-grid linear solver performance; coupled physics

EQSIM
Earthquake Hazard Risk Assessment
(David McCallen, LBNL)
- Replace conservative and costly earthquake retrofits with safe purpose-fit retrofits and designs

Challenges: full waveform inversion algorithms

MFIX-Exa
Scale-up of Clean Fossil Fuel Combustion
(Madhava Syamlal, NETL)
- Commercial-scale demonstration of transformational energy technologies – curbing CO_2 emissions at fossil fuel power plants by 2030

Challenges: load balancing; strong scaling thru transients

GAMESS
Biofuel Catalyst Design
(Mark Gordon, Ames)
- Design more robust and selective catalysts orders of magnitude more efficient at temperatures hundreds of degrees lower

Challenges: weak scaling of overall problem; on-node performance of molecular fragments

EXAALT
Materials for Extreme Environments
(Danny Perez, LANL)
- Simultaneously address time, length, and accuracy requirements for predictive microstructural evolution of materials

Challenges: SNAP kernel efficiency on accelerators; efficiency of DFTB application on accelerators

ExaStar
Demystify Origin of Chemical Elements
(Dan Kasen, LBNL)
- What is the origin of the elements?
- How does matter behave at extreme densities?
- What are the sources of gravity waves?

Challenges: delivering performance on accelerators; delivering fidelity for general relativity implementation

ExaSMR
Design and Commercialization of Small Modular Reactors
(Steve Hamilton, ORNL)
- Virtual test reactor for advanced designs via experimental-quality simulations of reactor behavior

Challenges: existing GPU-based MC algorithms require rework for hardware less performant for latency-bound algorithms with thread divergence; performance portability with OCCA & OpenACC not achievable; insufficient node memory for adequate CFD + MC coupling

Subsurface
Carbon Capture, Fossil Fuel Extraction, Waste Disposal
(Carl Steefel, LBNL)
- Reliably guide safe long-term consequential decisions about storage, sequestration, and exploration

Challenges: performance of Lagrangian geomechanics; adequacy of Lagrangian crack mechanics) + Eulerian (reaction, advection, diffusion) models; parallel HDF5 for coupling

QMCPACK
Materials for Extreme Environments
(Paul Kent, ORNL)
- Find, predict and control materials and properties at the quantum level with unprecedented accuracy for the design novel materials that rely on metal to insulator transitions for high performance electronics, sensing, storage

Challenges: minimizing on-node memory usage; parallel on-node performance of Markov-chain Monte Carlo

ExaSGD
Reliable and Efficient Planning of the Power Grid
(Henry Huang, PNNL)
- Optimize power grid planning, operation, control and improve reliability and efficiency

Challenges: parallel performance of nonlinear optimization based on discrete algebraic equations and possible mixed-integer programming

WDMApp
High-Fidelity Whole Device Modeling of Magnetically Confined Fusion Plasmas
(Amitava Bhattacharjee, PPPL)
- Prepare for ITER exps and increase ROI of validation data and understanding
- Prepare for beyond-ITER devices

Challenges: robust, accurate, and efficient code-coupling algorithm; reduction in memory and I/O usage

Combustion-PELE
High-Efficiency, Low-Emission Combustion Engine Design
(Jackie Chen, SNL)
- Reduce or eliminate current cut-and-try approaches for combustion system design

Challenges: performance of chemistry ODE integration on accelerated architectures; linear solver performance for low-Mach algorithm; explicit LES/DNS algorithm not stable

ExaFEL
Light Source-Enabled Analysis of Protein and Molecular Structure and Design
(Amadeo Perazzo, SLAC)
- Process data without beam time loss
- Determine nanoparticle size and shape changes
- Engineer functional properties in biology and materials science

Challenges: improving the strong scaling (one event processed over many cores) of compute-intensive algorithms (ray tracing, M-TIP) on accelerators

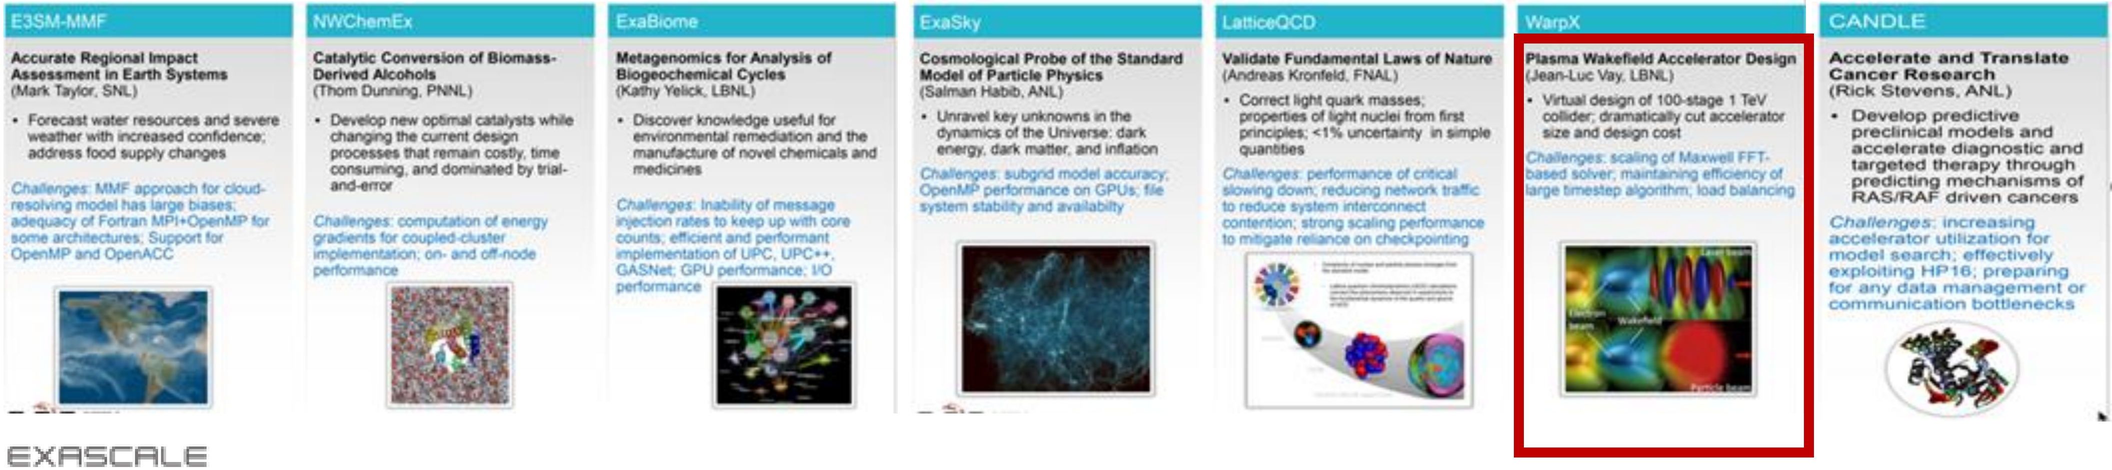

E3SM-MMF
Accurate Regional Impact Assessment in Earth Systems
(Mark Taylor, SNL)
- Forecast water resources and severe weather with increased confidence; address food supply changes

Challenges: MMF approach for cloud-resolving model has large biases; adequacy of Fortran MPI+OpenMP for some architectures; Support for OpenMP and OpenACC

NWChemEx
Catalytic Conversion of Biomass-Derived Alcohols
(Thom Dunning, PNNL)
- Develop new optimal catalysts while changing the current design processes that remain costly, time consuming, and dominated by trial-and-error

Challenges: computation of energy gradients for coupled-cluster implementation; on- and off-node performance

ExaBiome
Metagenomics for Analysis of Biogeochemical Cycles
(Kathy Yelick, LBNL)
- Discover knowledge useful for environmental remediation and the manufacture of novel chemicals and medicines

Challenges: Inability of message injection rates to keep up with core counts; efficient and performant implementation of UPC, UPC++, GASNet; GPU performance; I/O performance

ExaSky
Cosmological Probe of the Standard Model of Particle Physics
(Salman Habib, ANL)
- Unravel key unknowns in the dynamics of the Universe: dark energy, dark matter, and inflation

Challenges: subgrid model accuracy; OpenMP performance on GPUs; file system stability and availabity

LatticeQCD
Validate Fundamental Laws of Nature
(Andreas Kronfeld, FNAL)
- Correct light quark masses; properties of light nuclei from first principles; <1% uncertainty in simple quantities

Challenges: performance of critical slowing down; reducing network traffic to reduce system interconnect contention; strong scaling performance to mitigate reliance on checkpointing

WarpX
Plasma Wakefield Accelerator Design
(Jean-Luc Vay, LBNL)
- Virtual design of 100-stage 1 TeV collider; dramatically cut accelerator size and design cost

Challenges: scaling of Maxwell FFT-based solver; maintaining efficiency of large timestep algorithm; load balancing

CANDLE
Accelerate and Translate Cancer Research
(Rick Stevens, ANL)
- Develop predictive preclinical models and accelerate diagnostic and targeted therapy through predicting mechanisms of RAS/RAF driven cancers

Challenges: increasing accelerator utilization for model search; effectively exploiting HP16; preparing for any data management or communication bottlenecks

EXASCALE COMPUTING PROJECT

Years of collaboration with NERSC + NESAP were key to selection

Challenge of the ECP WarpX Project

Go from modeling 1 or 2 plasma stages to tens of stages in 3D from first principles for plasma collider R&D and design

Laser driven

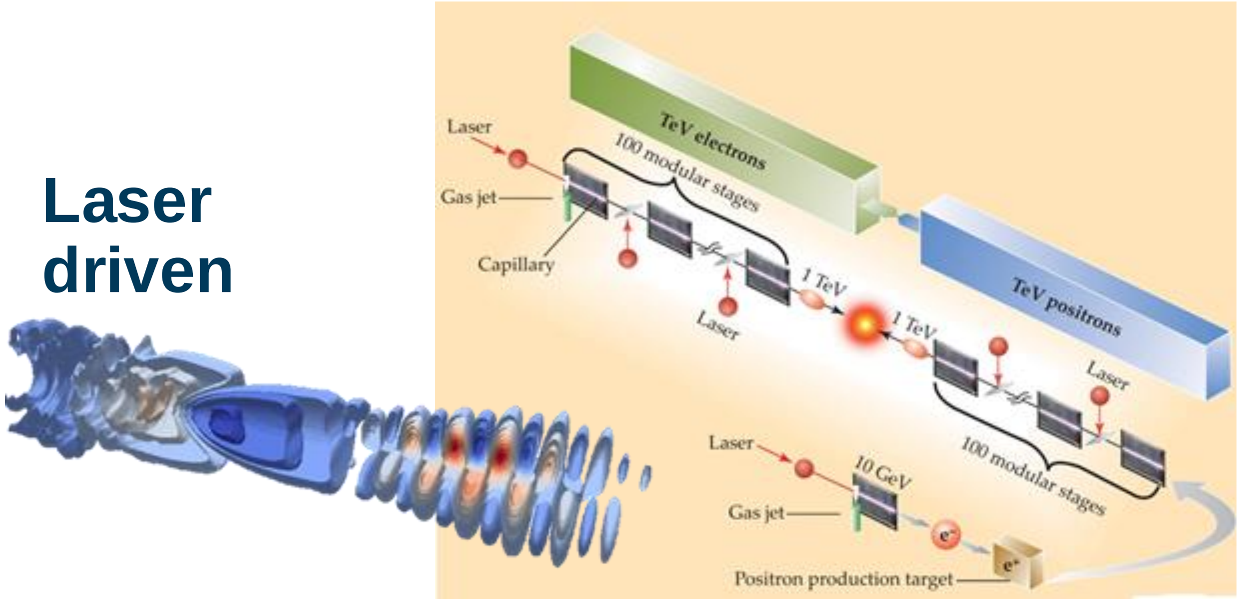

Particle driven

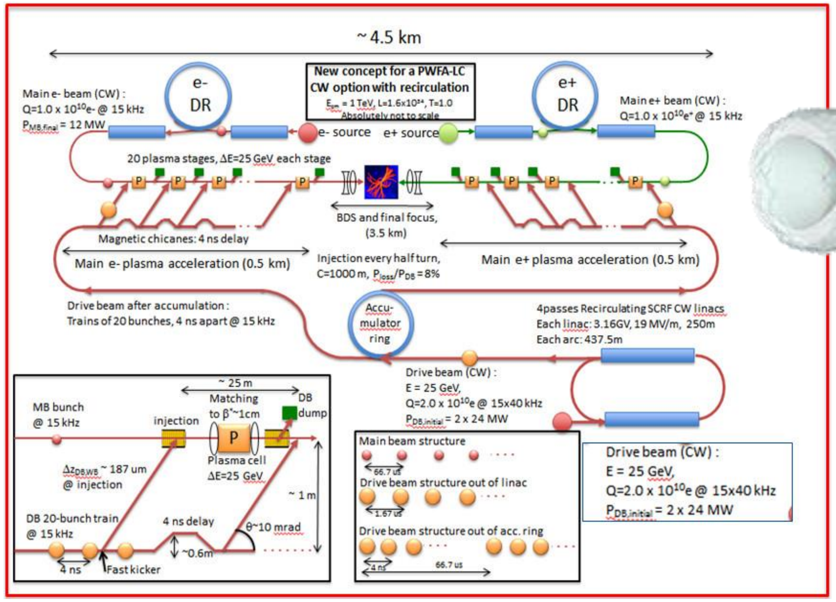

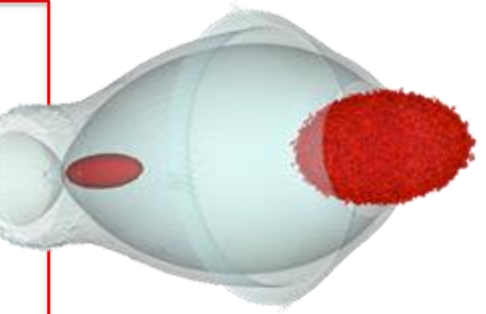

➔ advanced algorithms on fastest/largest supercomputers

Key partnership with AMReX

3 labs collaboration

Welcome to NESAP for Perlmutter

NERSC
March 28/29, 2019

(liaison)
Kevin Gott
NERSC (staff)
WarpX

2019-2021

Michael Rowan
NERSC (postdoc)
WarpX

2023-2024

Muhammad Haseeb
NERSC (postdoc)
WarpX

NESAP led to significant WarpX FOM increase on Perlmutter

also increase efficiency of dynamic load balancing & binary collisions kernel

FOM ~ # particles/runtime

Date	Machine	Nodes	FOM
3/19	Cori (Warp)	6625	2.2e10
3/19	Edison	4 694	6.7e10
3/19	Cori	6 625	1.0e11
7/21	Perlmutter	960	1.1e12

16x

50x

3.8x speedup w/ adaptive load balancing

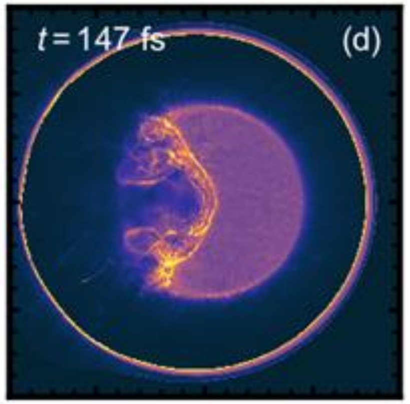

particles/cell

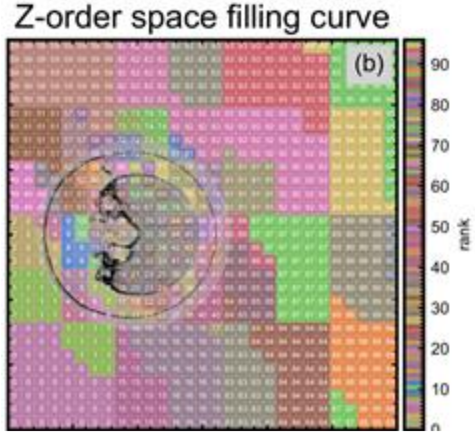

Domain decomposition

4x speedup of Coulomb collision module

New kernel:

```
MultiParticleContainer::doCollisions()
Avg. per step = 0.1375590809 s
Total Time                        : 14.58777054
```

WarpX dev branch:

```
MultiParticleContainer::doCollisions()
Avg. per step = 0.6979252147
Total Time                        : 70.43095279
```

Applied to other binary collisions modules in WarpX (nuclear fusion, DSMC, etc.)

ECP WarpX team benefitted from NESAP to win 22 Gordon Bell Prize

April-July 2022: WarpX on **world's largest HPCs**

L. Fedeli, A. Huebl et al., *Gordon Bell Prize Winner* in SC'22, 2022

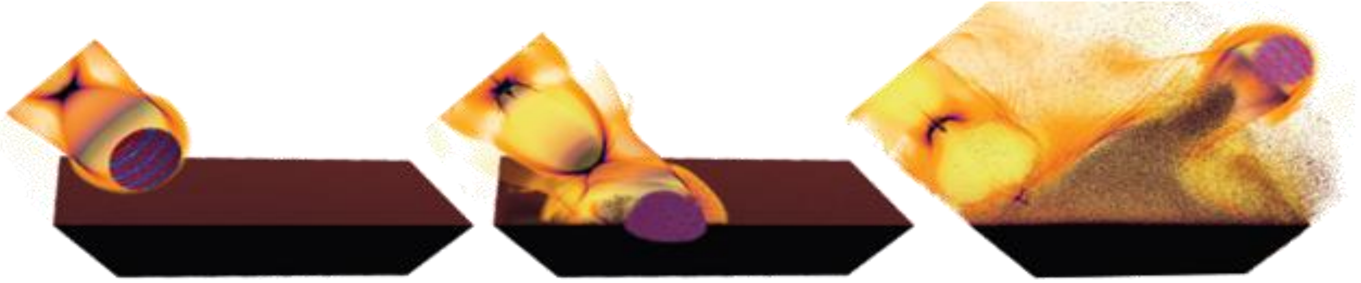

Modeling of novel plasma e- beam injection scheme

- 3 levels of parallelism, scalable & portable
- adaptive mesh refinement
- efficient dynamic loal balancing

PLASMA SIMULATION CODE WINS 2022 ACM GORDON BELL PRIZE

Figure-of-Merit: weighted updates / sec

Date	Code	Machine	N_c/Node	Nodes	FOM
3/19	Warp	Cori	0.4e7	6 625	2.2e10
3/19	WarpX	Cori	0.4e7	6 625	1.0e11
6/19	WarpX	Summit	2.8e7	1 000	7.8e11
9/19	WarpX	Summit	2.3e7	2 560	6.8e11
1/20	WarpX	Summit	2.3e7	2 560	1.0e12
2/20	WarpX	Summit	2.5e7	4 263	1.2e12
6/20	WarpX	Summit	2.0e7	4 263	1.4e12
7/20	WarpX	Summit	2.0e8	4 263	2.5e12
3/21	WarpX	Summit	2.0e8	4 263	2.9e12
6/21	WarpX	Summit	2.0e8	4 263	2.7e12
7/21	WarpX	Perlmutter	2.7e8	960	1.1e12
12/21	WarpX	Summit	2.0e8	4 263	3.3e12
4/22	WarpX	Perlmutter	4.0e8	928	1.0e12
4/22	WarpX	Perlmutter†	4.0e8	928	1.4e12
4/22	WarpX	Summit	2.0e8	4 263	3.4e12
4/22	WarpX	Fugaku†	3.1e6	98 304	8.1e12
6/22	WarpX	Perlmutter	4.4e8	1 088	1.0e12
7/22	WarpX	Fugaku	3.1e6	98 304	2.2e12
7/22	WarpX	Fugaku†	3.1e6	152 064	9.3e12
7/22	WarpX	Frontier	8.1e8	8 576	1.1e13

U.S. DEPARTMENT OF ENERGY | Office of Science

ECP EXASCALE COMPUTING PROJECT

NERSC NESAP

Looking forward: more HPC
+ superfacilities (aka IRI)
+ workflows

Simulations of Advanced Laser-Plasma Ion Acceleration Mechanisms

3D WarpX simulations supported **commissioning of BELLA iP2 beamline** in 2022

Directed Coulomb Explosion (DCE)

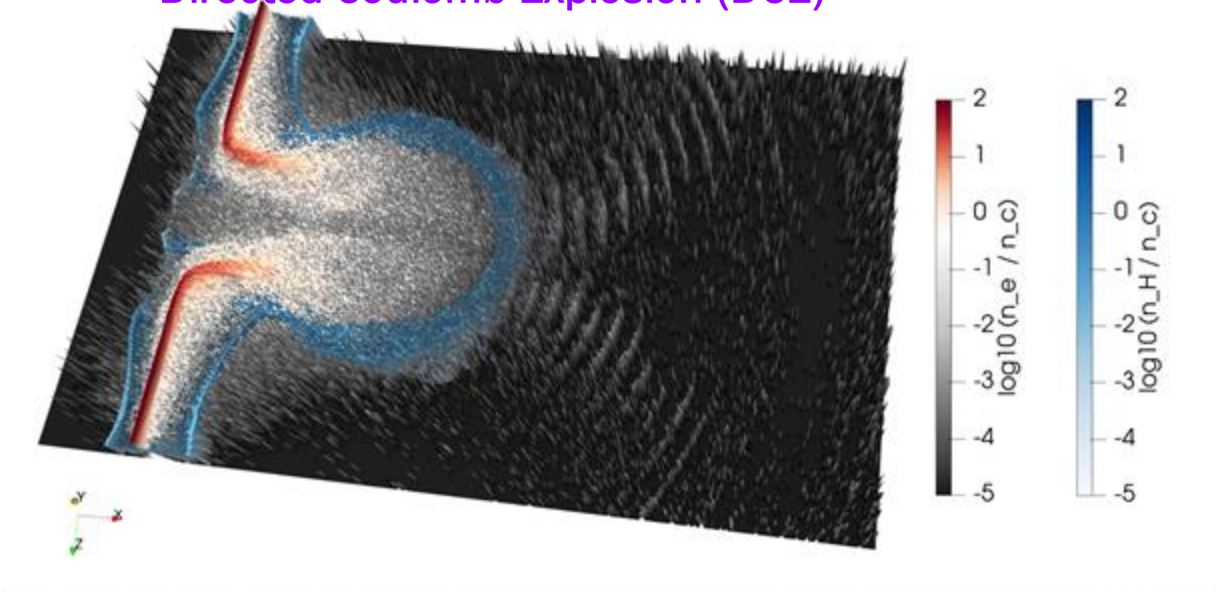

https://newscenter.lbl.gov/2022/12/01/laser-upgrade-research-possibilities/

Magnetic Vortex Acceleration (MVA)

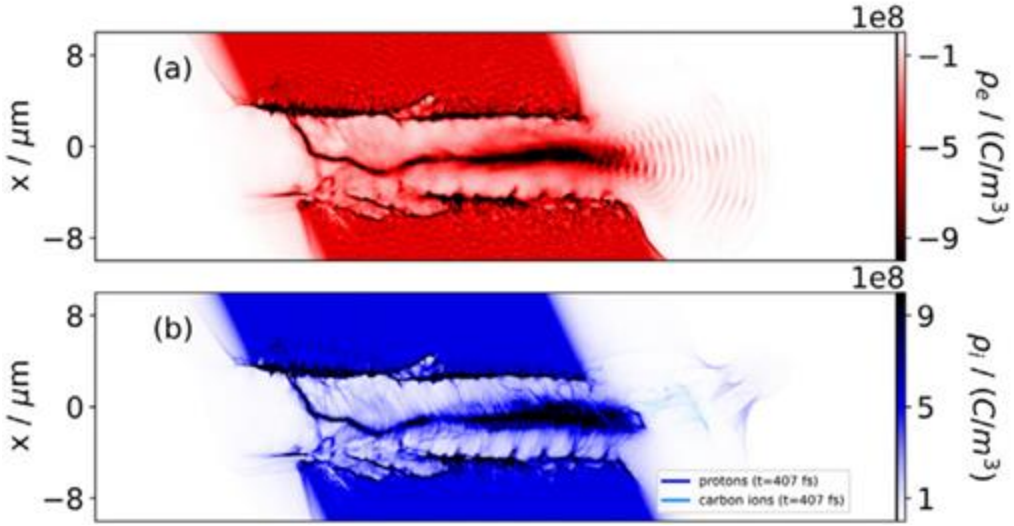

S. Hakimi *et al.*, Phys. Plasmas **29**, 083102 (2022)

- BELLA iP2 (LBNL)
- Code: WarpX
- NERSC computer: Perlmutter
- # nodes: 100s

Simulations demonstrating Boosting of Intense Ion Beam Energies using New Concept with Hollow-Channel Laser-Plasma Stages

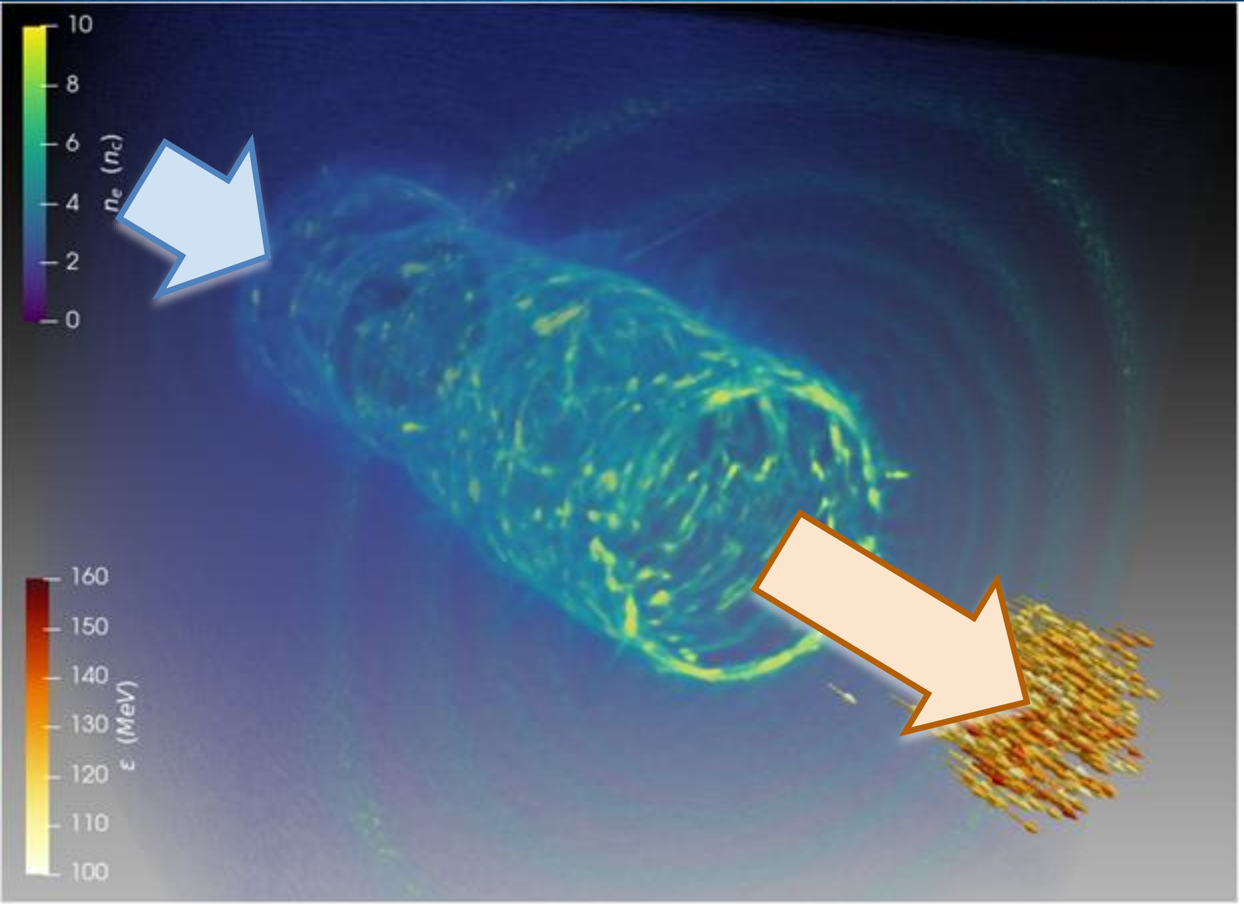

☑ Key beam quality parameters are conserved

Charge, energy spread, and emittance are conserved well

☑ State-of-the-art PW laser facility parameters are sufficient

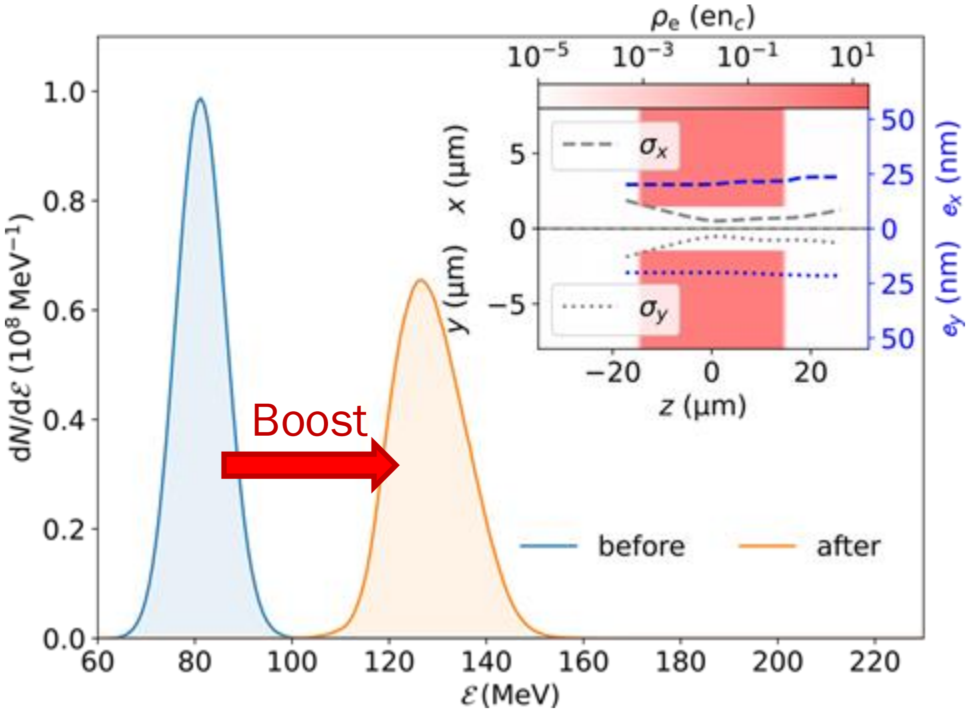

- Code: WarpX
- NERSC computer: Perlmutter
- # nodes: 768 for 8.5h

M. Garten *et al.*, Phys. Rev. Research **6**, 033148 (2024)

Exploiting High-Quality HPC Data for ML-Boosted Collider Design

Start-to-end modeling of chain of plasma accelerator stages for colliders can be very expensive with PIC.

Under some conditions (low beam charge, repetition of similar stages), ML surrogate can be trained & replace PIC.

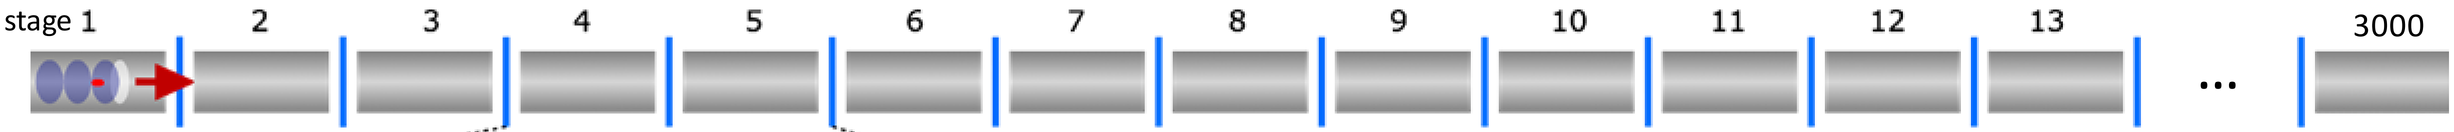

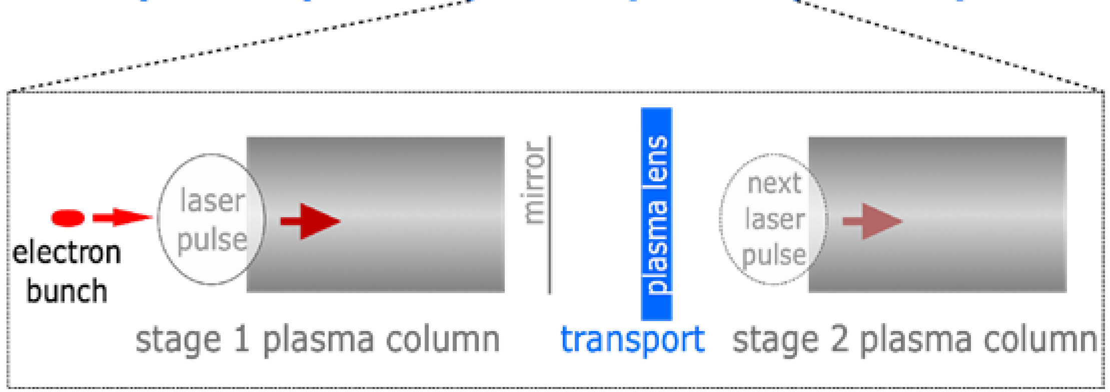

tightly-coupled LPA-*neural networks* inside **ImpactX**

WarpX start-to-end simulation
256 GPUs
1 simulation / 5.1 hours

⇩

ImpactX with WarpX-trained NNs
1 GPU
2-4 simulations / sec

LPA + Transport Optimization
with 1000s of evaluations

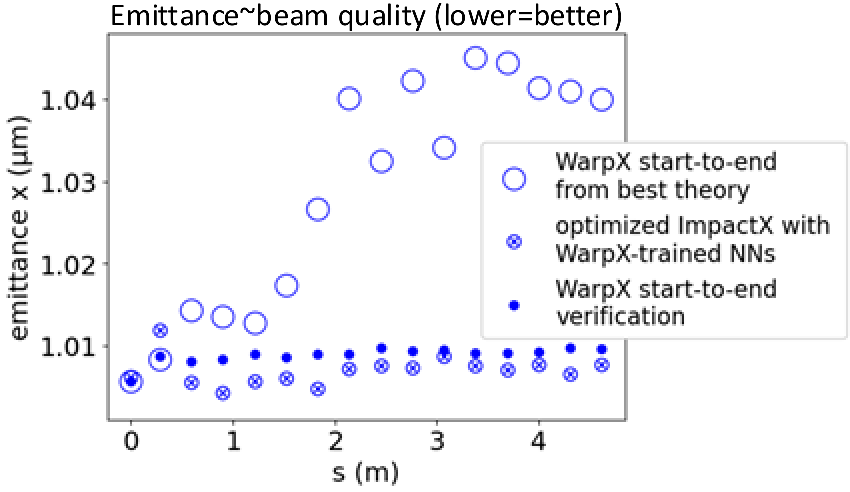

RT Sandberg et al., IPAC23, DOI:10.18429/JACoW-IPAC2023-WEPA101 (2023)
RT Sandberg et al., ***PASC24 Best Paper*** (2024)

LDRD

LDRD in collaboration with NERSC to build Superfacility prototype

Remi Lehe (PI)

Bella iP2 Experiment

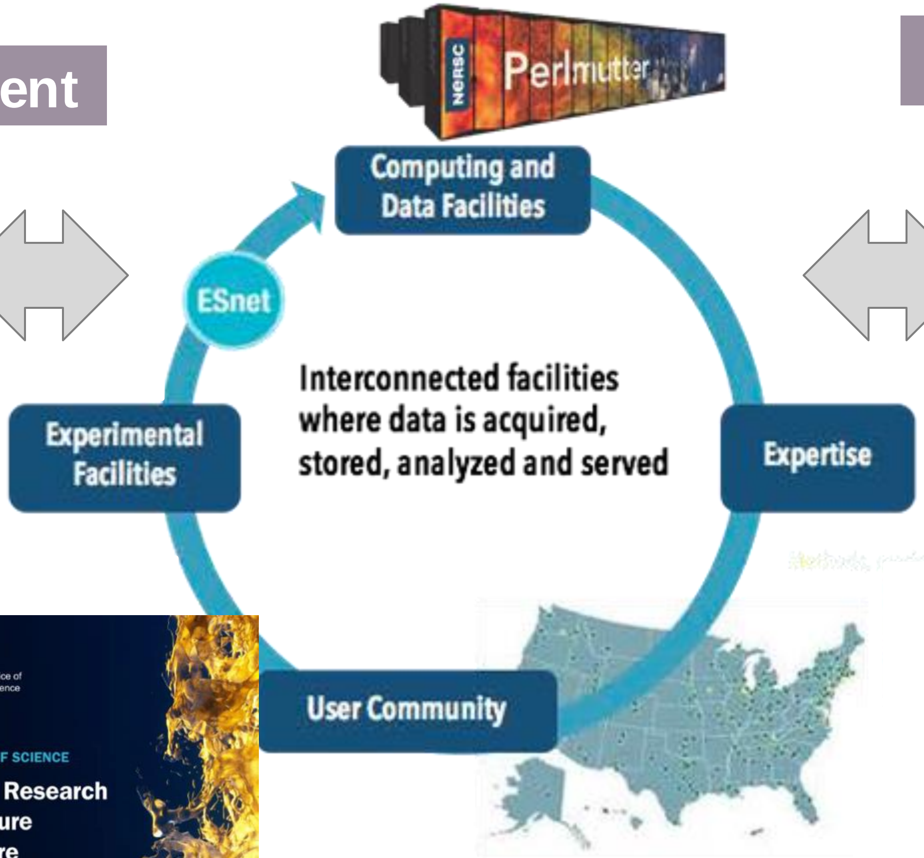

Computing and Data Facilities

ESnet

Interconnected facilities where data is acquired, stored, analyzed and served

Experimental Facilities

Expertise

User Community

Exascale code WarpX

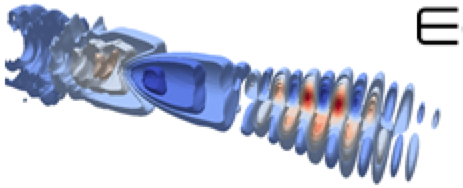

2022
Superfacility Project Report
Lawrence Berkeley National Laboratory
Computing Sciences

THE DOE OFFICE OF SCIENCE
Integrated Research Infrastructure Architecture Blueprint Activity
FINAL REPORT
2023

NERSC Superfacility concept and tools

First application for fusion, with applicability to accelerator R&D.

New NESAP project on Integrated Plasma Simulation Workflows

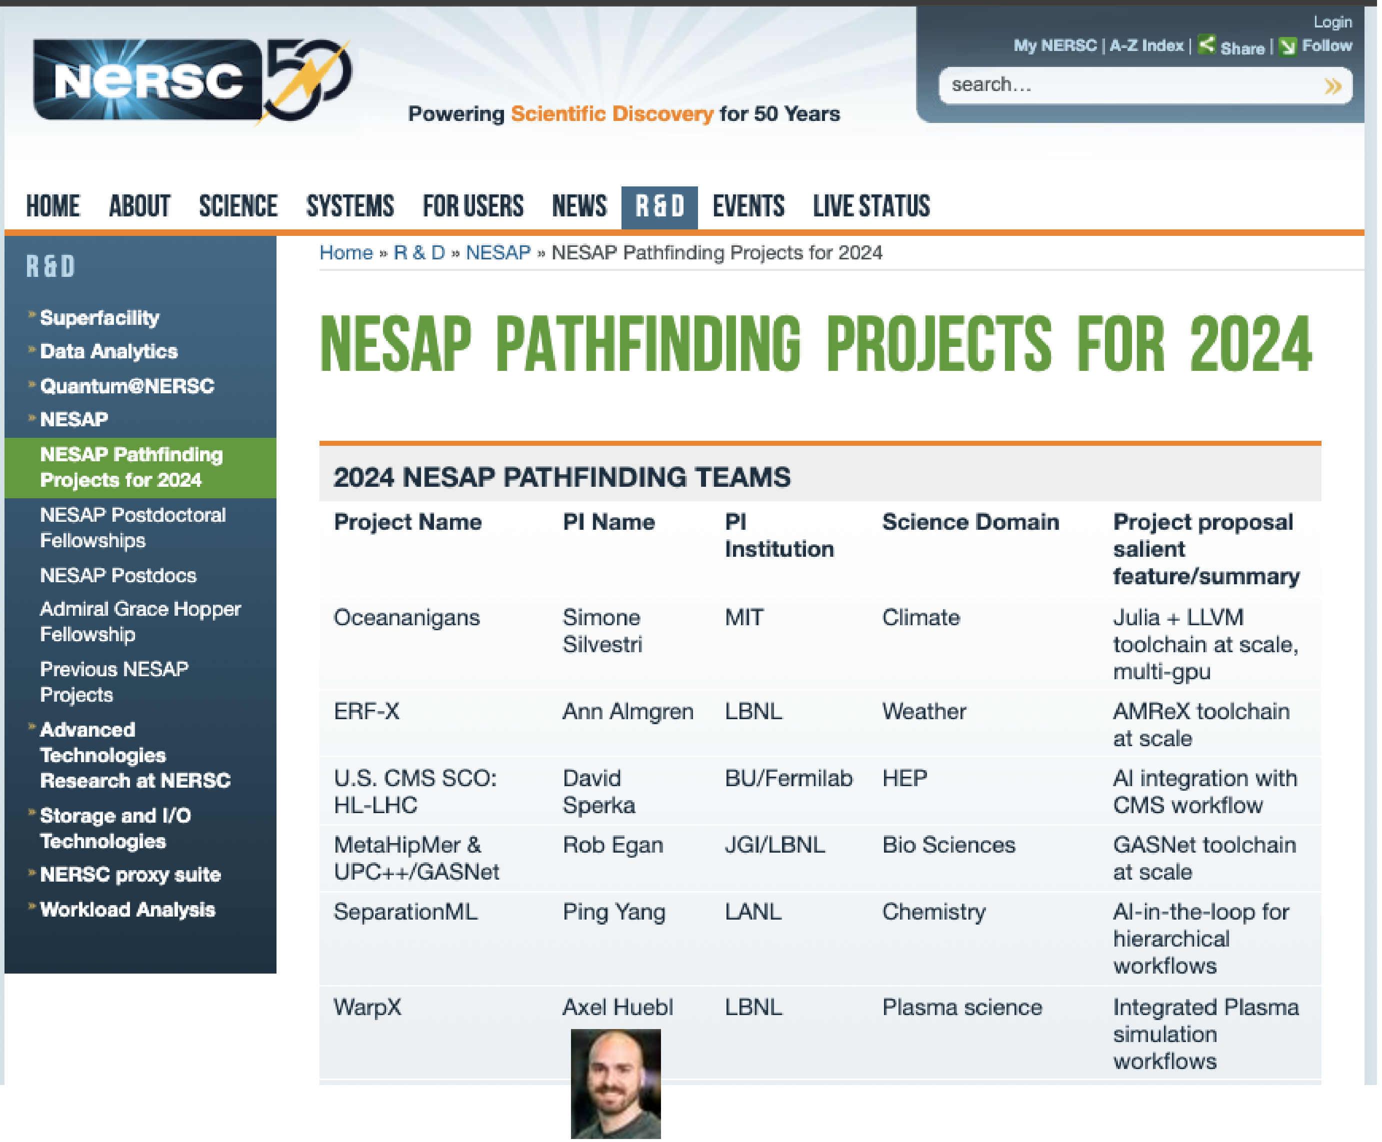

Login
My NERSC | A-Z Index | Share | Follow
search…

Powering Scientific Discovery for 50 Years

HOME ABOUT SCIENCE SYSTEMS FOR USERS NEWS R&D EVENTS LIVE STATUS

R&D
- **Superfacility**
- **Data Analytics**
- **Quantum@NERSC**
- **NESAP**
  - **NESAP Pathfinding Projects for 2024**
  - NESAP Postdoctoral Fellowships
  - NESAP Postdocs
  - Admiral Grace Hopper Fellowship
  - Previous NESAP Projects
- **Advanced Technologies Research at NERSC**
- **Storage and I/O Technologies**
- **NERSC proxy suite**
- **Workload Analysis**

Home » R & D » NESAP » NESAP Pathfinding Projects for 2024

NESAP PATHFINDING PROJECTS FOR 2024

2024 NESAP PATHFINDING TEAMS

Project Name	PI Name	PI Institution	Science Domain	Project proposal salient feature/summary
Oceananigans	Simone Silvestri	MIT	Climate	Julia + LLVM toolchain at scale, multi-gpu
ERF-X	Ann Almgren	LBNL	Weather	AMReX toolchain at scale
U.S. CMS SCO: HL-LHC	David Sperka	BU/Fermilab	HEP	AI integration with CMS workflow
MetaHipMer & UPC++/GASNet	Rob Egan	JGI/LBNL	Bio Sciences	GASNet toolchain at scale
SeparationML	Ping Yang	LANL	Chemistry	AI-in-the-loop for hierarchical workflows
WarpX	Axel Huebl	LBNL	Plasma science	Integrated Plasma simulation workflows

NESAP Pathfinding Projects

NESAP pathfinding projects are one-year projects intended to prepare for and better use advanced workflow capabilities such as hardware acceleration, reconfigurable storage, advanced scheduling, and integration with edge services, as well as alignment with DOE's IRI program.

Urjoshi Sinha
HPC Engineer
Data & AI Group
NERSC

Project goals

- **Performance goals**
  - *Scalability* (solvers, ML hybrids, load balancing)
  - *GPU performance* (kernels)
- **Advanced Capability goals**
  - *Containerization &* automated *performance regression testing*
  - *Jupyter-centric simulation lifecycle*

BERKELEY LAB
LAWRENCE BERKELEY NATIONAL LABORATORY

A quick note on my own experience with NERSC

I joined LBNL shortly after NERSC dedication following move from LLNL

1998
Conversion to HPSS Completed

1998
One T3E Is Better than Two

1998
ESnet Conducts Differentiated Services Demo

1997
Simon Writes About Reinventing NERSC

1997
NERSC First to Checkpoint/Restart MPP System

1996
NERSC Dedicated at Berkeley Lab after Move from LLNL

CLICK TO FIND OUT MORE
1997
Second Cray T3E-900 Arrives

1998
Conversion to HPSS Completed

1998
One T3E Is Better than Two

1998
ESnet Conducts Differentiated Services Demo

NERSC
National Energy Research Scientific Computing Center

1997
Simon Writes About Reinventing NERSC

1997
NERSC First to Checkpoint/Restart MPP System

1996, Nov. 4
Jean-Luc Vay starts postdoc @LBNL

1996
NERSC Dedicated at Berkeley Lab after Move from LLNL

CLICK TO FIND OUT MORE

1997
Second Cray T3E-900 Arrives

started using NERSC shortly after…

…needing support…

NERSC has always provided great support!

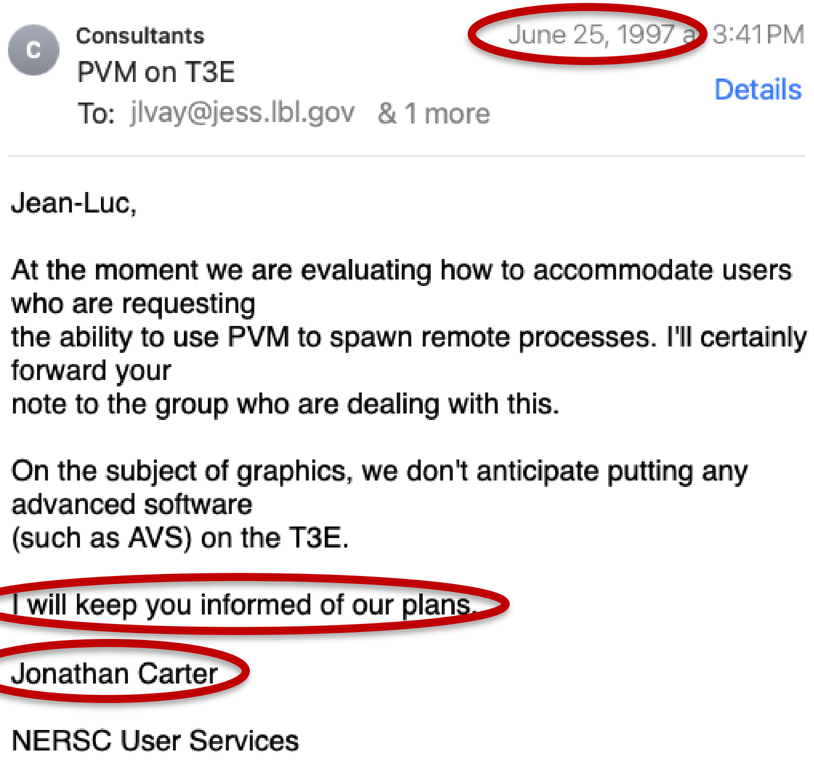

Consultants June 25, 1997 a 3:41PM

PVM on T3E

To: jlvay@jess.lbl.gov & 1 more

Jean-Luc,

At the moment we are evaluating how to accommodate users
who are requesting
the ability to use PVM to spawn remote processes. I'll certainly
forward your
note to the group who are dealing with this.

On the subject of graphics, we don't anticipate putting any
advanced software
(such as AVS) on the T3E.

I will keep you informed of our plans.

Jonathan Carter

NERSC User Services

Jonathan Carter

Jonathan Carter, Ph.D.
Associate Lab Director for Computing Sciences
jtcarter@lbl.gov
Phone: +1 510 486 7514
Lawrence Berkeley National Laboratory

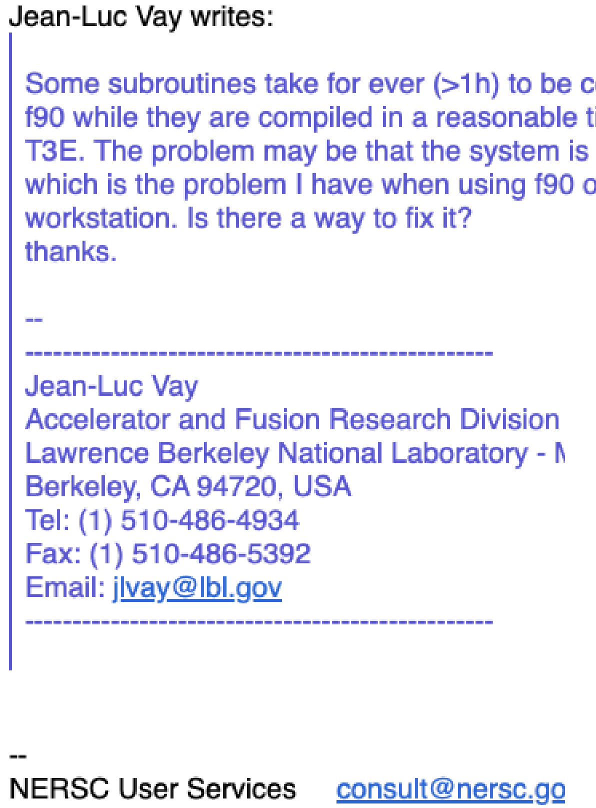

Jean-Luc Vay writes:

> Some subroutines take for ever (>1h) to be compiled on Killeen with f90 while they are compiled in a reasonable time (<10min) on the T3E. The problem may be that the system is swaping with the hard-disk, which is the problem I have when using f90 on the choga.lbl.gov SUN workstation. Is there a way to fix it?
> thanks.
>
> --
> --
> Jean-Luc Vay
> Accelerator and Fusion Research Division
> Lawrence Berkeley National Laboratory -
> Berkeley, CA 94720, USA
> Tel: (1) 510-486-4934
> Fax: (1) 510-486-5392
> Email: jlvay@lbl.gov
> --

--
NERSC User Services consult@nersc.go

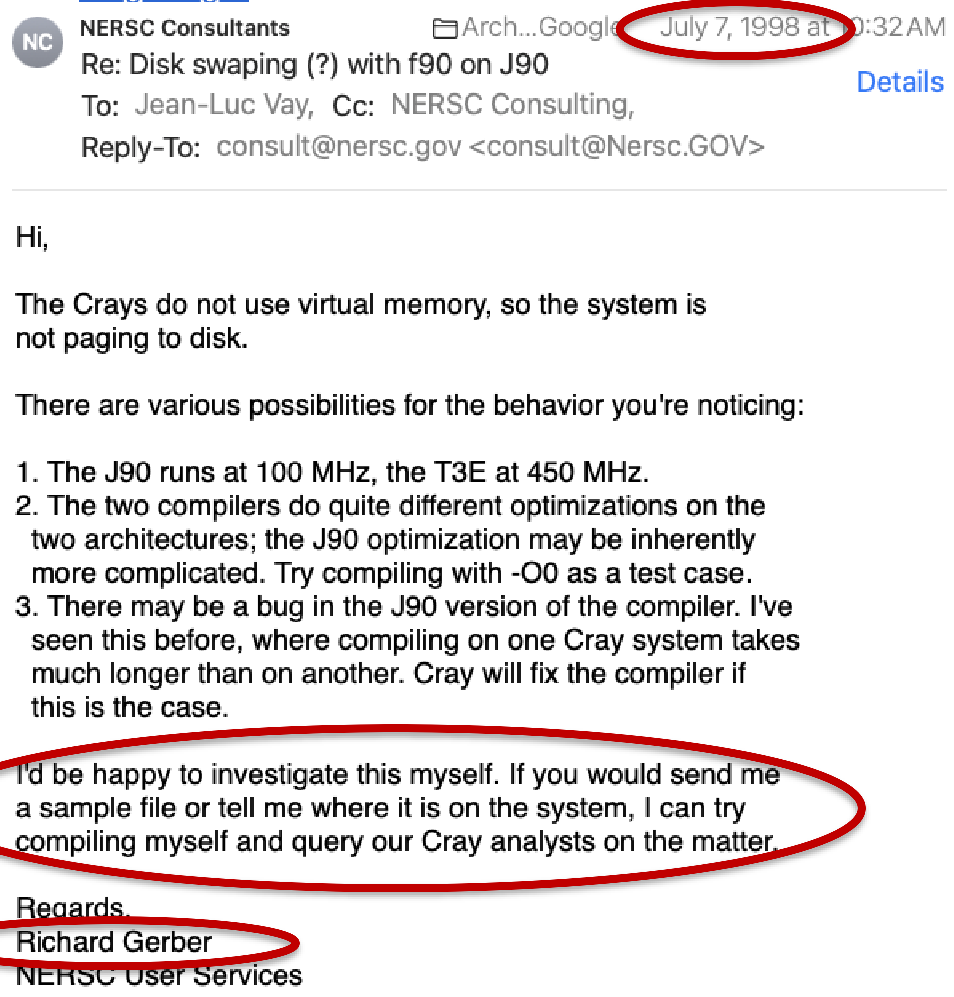

NERSC Consultants Arch...Googl July 7, 1998 at 0:32 AM

Re: Disk swaping (?) with f90 on J90

To: Jean-Luc Vay, Cc: NERSC Consulting,

Reply-To: consult@nersc.gov <consult@Nersc.GOV>

Hi,

The Crays do not use virtual memory, so the system is
not paging to disk.

There are various possibilities for the behavior you're noticing:

1. The J90 runs at 100 MHz, the T3E at 450 MHz.
2. The two compilers do quite different optimizations on the
  two architectures; the J90 optimization may be inherently
  more complicated. Try compiling with -O0 as a test case.
3. There may be a bug in the J90 version of the compiler. I've
  seen this before, where compiling on one Cray system takes
  much longer than on another. Cray will fix the compiler if
  this is the case.

I'd be happy to investigate this myself. If you would send me
a sample file or tell me where it is on the system, I can try
compiling myself and query our Cray analysts on the matter.

Regards,
Richard Gerber
NERSC User Services

RICHARD GERBER

Richard A. Gerber, Ph.D.
NERSC Senior Science Advisor
RAGerber@lbl.gov
Phone: +1 (510) 486-6820
1 Cyclotron Road
M/S 59R3103
Berkeley, CA 94720 us

Thanks to NERSC for all the amazing computing & support

Looking forward to future work @ & w/ NERSC

Thank you all for your attention