Our Road to Exascale:

Particle Accelerator & Laser-Plasma Modeling

EXASCALE COMPUTING PROJECT

Axel Huebl

Lawrence Berkeley National Laboratory

On behalf of the WarpX team (lead: J-L Vay @ LBNL)
LBNL, LLNL, SLAC

+ contributors external to ECP from labs, universities & industry in USA, Europe & Asia

Webinar Series on Best Practices for HPC Software Developers

virtual
March 15th, 2023

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

Outline

- **Intro -** particle accelerator modeling
- **Then (<2016) -** Warp
- **The journey -** from Warp to WarpX (2016-2023)
- **Now (2023+) -** WarpX + ecosystem
- **Conclusion**

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

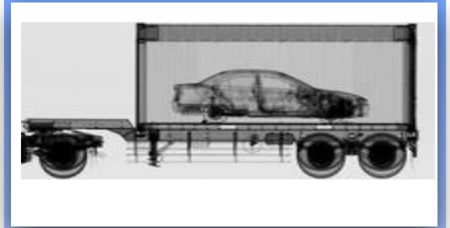

- **Cargo** scanning
- Active interrogation
- **Stockpile stewardship**: materials character-ization, radiography, support of non-proliferation

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

Opportunity for much bigger impact by reducing size and cost.

Modeling: Exploration → Understanding → Design

The Modeling of Particle Accelerators is Very Complex

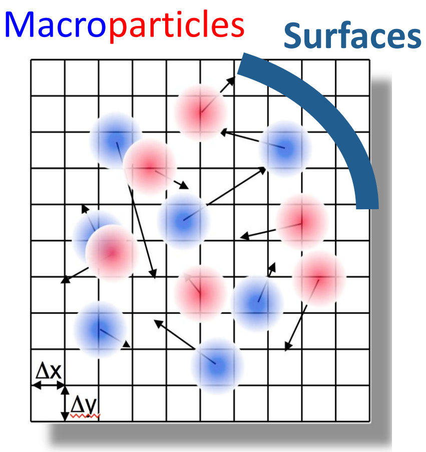

electromagnetic (EM) fields on a grid

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.

Typical computer representations:

- **particles:** macroparticles 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

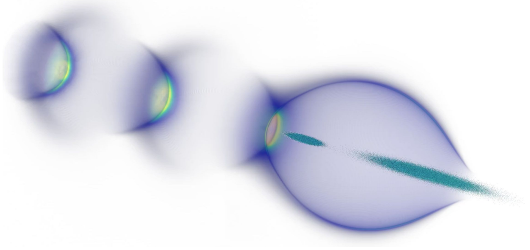

AMP has pioneered algorithms to cut on

- # of meshes: adaptive mesh refinement for beams & plasmas[1]
- # of time steps: Lorentz boosted frame method[2]

[1] J.-L. Vay *et al*, *Phys. Plasmas* **11**, 2928 (2004)

[2] J.-L. Vay, *Phys. Rev. Lett.* **98**, 130405 (2007)

All Accelerators in the World Rely on Modeling
and Increasingly on *High-Performance Computing (HPC)*

HEP Accelerator and Beam Physics Grand Challenges*

#1 beam intensity: Increase beam intensities by orders of magnitude.

⇒ increase # of particles by orders of magnitude.

#2 beam quality: Increase beam phase-space density by orders of
mag

#3 b

⇒ simulate all the particles.

Next generation of accelerators needs next generation of HPC modeling tools!

#4 beam prediction: Develop predictive "virtual particle accelerators"?

⇒ simulate everything: all the particles, conductors, dark currents, many turns, ...

*S. Nagaitsev *et al.*, "Accelerator and Beam Physics: Grand Challenges and Research Opportunities," Snowmass 2021 LOI

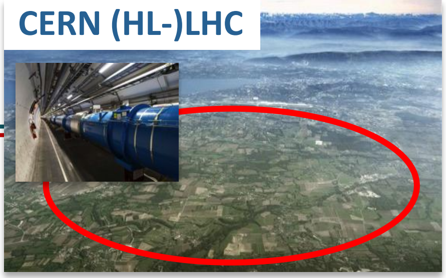

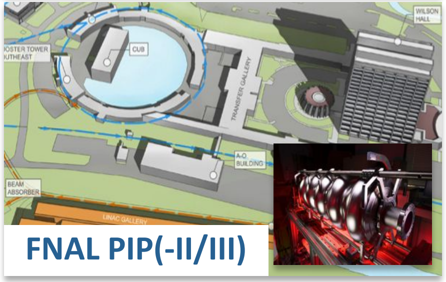

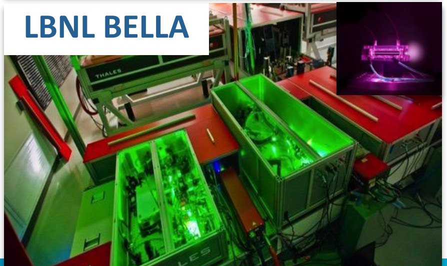

Ultimate goal: *virtual accelerator* with *on-the-fly tunability* of physics & numerics complexity to users

Fast

Great for ensemble runs for ***design studies***

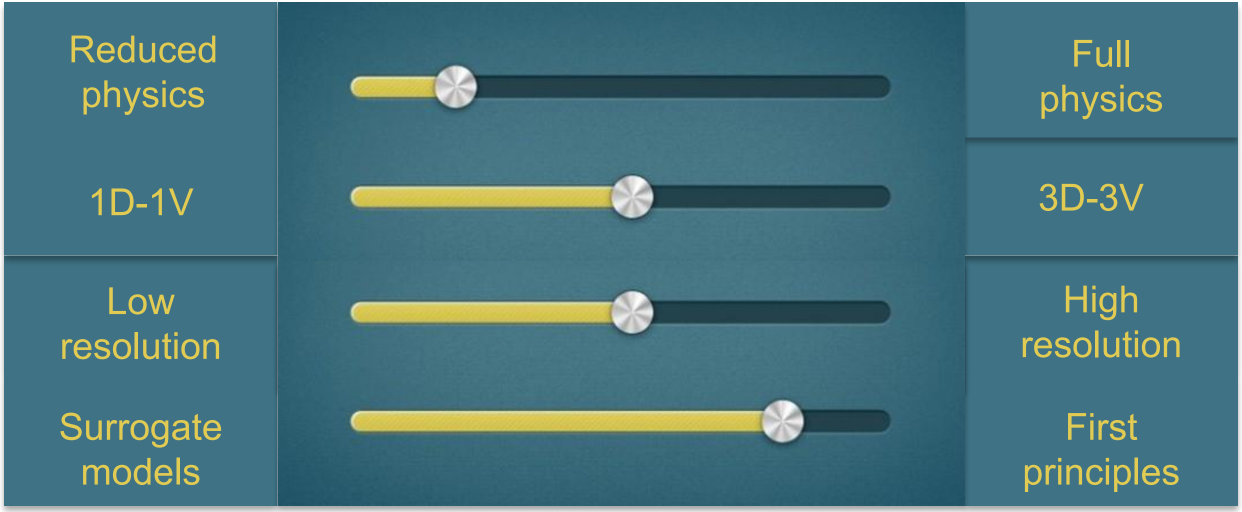

Accurate

Great for detailed runs for ***physics studies***

Mean toward goal
Open software ecosystem with tunable physics & numerics

BLAST
BEAM PLASMA & ACCELERATOR SIMULATION TOOLKIT

Start-to-End Modeling R&D

- advanced models: numerics, AI/ML surrogates
- speed & scalability: team science with computer sci.
- flexibility & reliability: modern software ecosystem

ECP EXASCALE COMPUTING PROJECT

Then:
Warp prior to 2016

Warp: historical roadmap & funding

Fortran + BASIS + PVM

2000 Fortran + Python + MPI

Open source BSD license 2013

2014 NESAP NERSC

2016 ECP EXASCALE COMPUTING PROJECT → WarpX

1989 Warp | ES-PIC → ES- & EM-PIC

A. Friedman

I. Haber

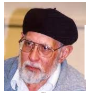

UNIVERSITY OF MARYLAND

1991 D. Grote

1995 S. Lund

2000 J.-L. Vay

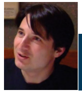

BERKELEY LAB

Other contributors: R. Kishek, R. Cohen, W. Sharp, M. deHoon, S. Chilton, A. Persaud, …

UNIVERSITY OF MARYLAND BERKELEY LAB

Fortran + PVM/MPI

1993 BPIC | EM-PIC

J.-L. Vay

UNIVERSITÉ PARIS SUD BERKELEY LAB

More contributors:

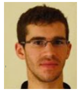

H. Vincenti,
R. Lehe,
M. Thévenet,
M. Kirchen,
P. Lee,
I. Dornmair,
G. Blaclard,
M. Lobet,
H. Kallala,
…

BERKELEY LAB cea PARIS-SACLAY DESY. UNIVERSITÉ PARIS SUD

Warp had become a PIC framework with many applications

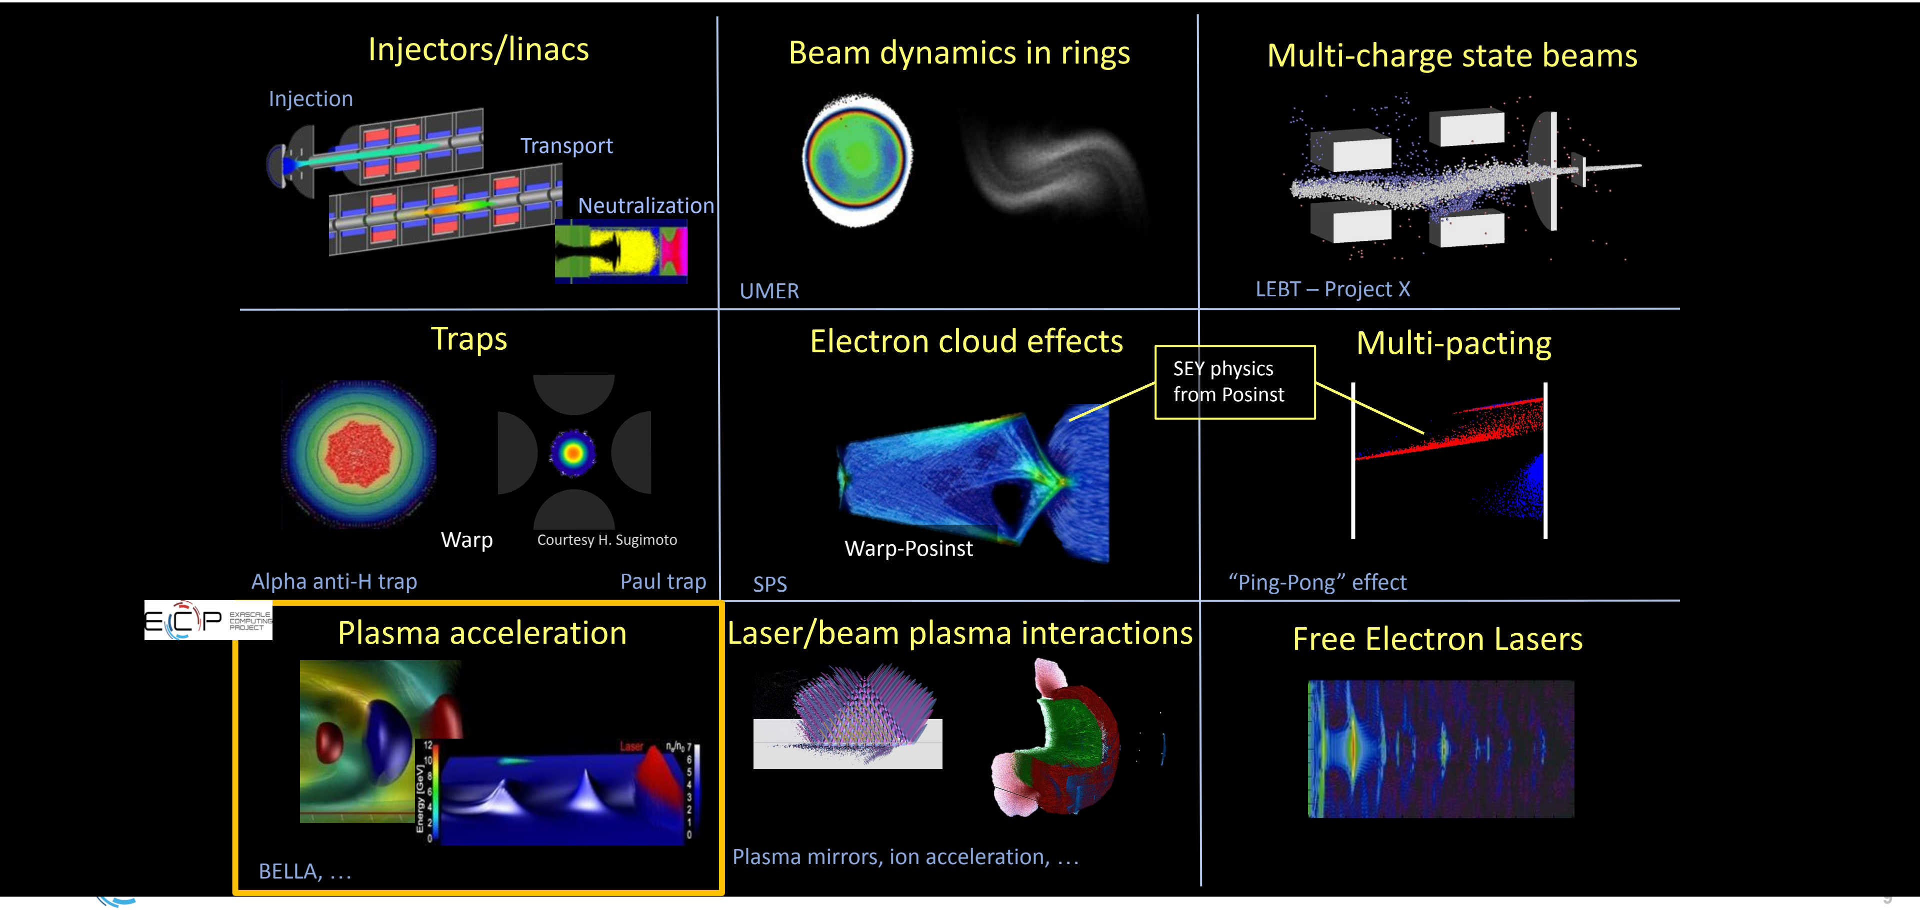

What made Warp unique: algorithmic innovation

Particle-In-Cell with Adaptive Mesh Refinement

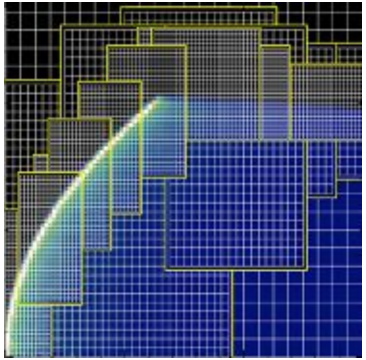

J.-L. Vay, et al, *Comput. Phys. Comm.* **164**, 171 & 297 (2004)

Use of special relativity to slash # time steps by orders of magnitude

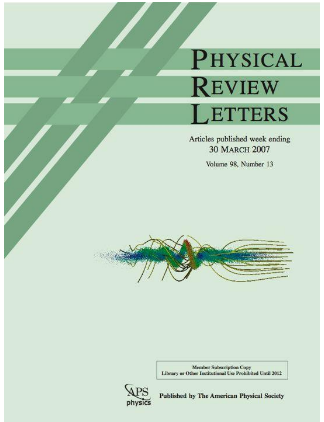

J.-L. Vay, *Phys. Rev. Lett.* **98**, 130405 (2007)

Novel particle pushers

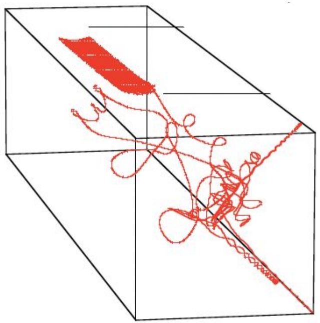

R. Cohen, et al, *Phys. Plasmas* **12**, 056708 (2005)

J.-L. Vay, *Phys. Plasmas* **15**, 056701 (2008)

Control of numerical Cherenkov instability

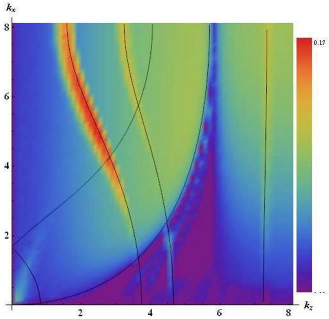

B. Godfrey, et al, *J. Comp. Phys.* **248**, 33 (2013)

R. Lehe, et al, *Phys. Rev. E* **94**, 053305 (2016)

Scalable spectral Maxwell solvers w/ domain decomposition

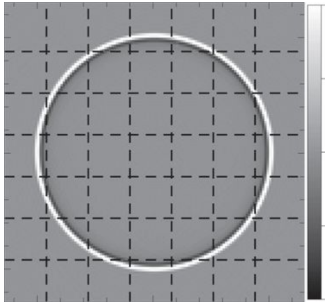

J.-L. Vay, et al, *J. Comp. Phys.* **243**, 260 (2013)

H. Kallala, et al, *Comp. Phys. Comm.* **244**, 25 (2019)

Sample algorithm innovations pioneered in Warp (some adopted in other codes)

EXASCALE COMPUTING PROJECT

Warp: limitations as of 2016

- **Half code in Fortran + half in Python**
  - Python: difficulty scaling on large supercomputers (>1 hour start time at scale)
  - Complicating transition to new hardware: manycore, GPUs, etc.
- **Small core team (2+ physicists w/ science projects) made it difficult to**
  - Transition to new hardware
  - Keep up performance optimization with fast pace of algorithmic innovation
  - Support growing users (& developers) base
  - Maintain documentation, test suite
- **SciDAC funding**
  - Supported a project that included 6+ (independent) Particle-In-Cell codes/frameworks with diverse science targets: dilution

⇒ ECP provided opportunity for focused effort with sufficient critical mass, & more…

The Journey:
from Warp to WarpX (2016-2023)

Power-Limits Seed a *Cambrian Explosion* of Compute Architectures

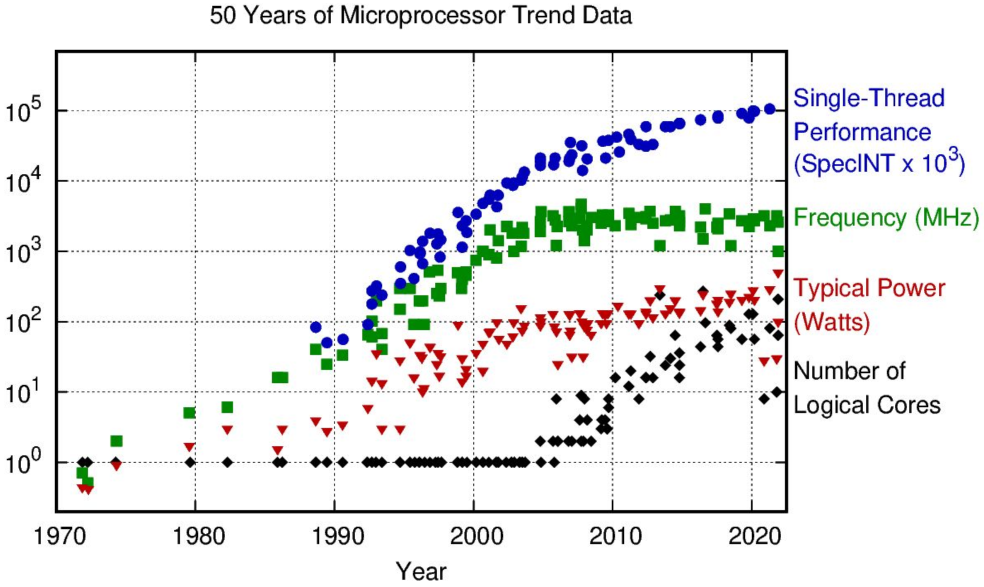

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

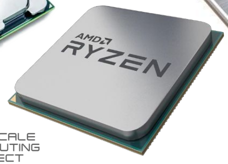

ARM

Power™

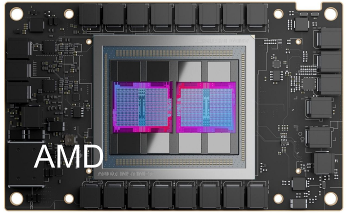

AMD

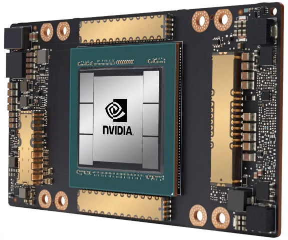

ECP EXASCALE COMPUTING PROJECT

Power-Limits Seed a *Cambrian Explosion* of Compute Architectures

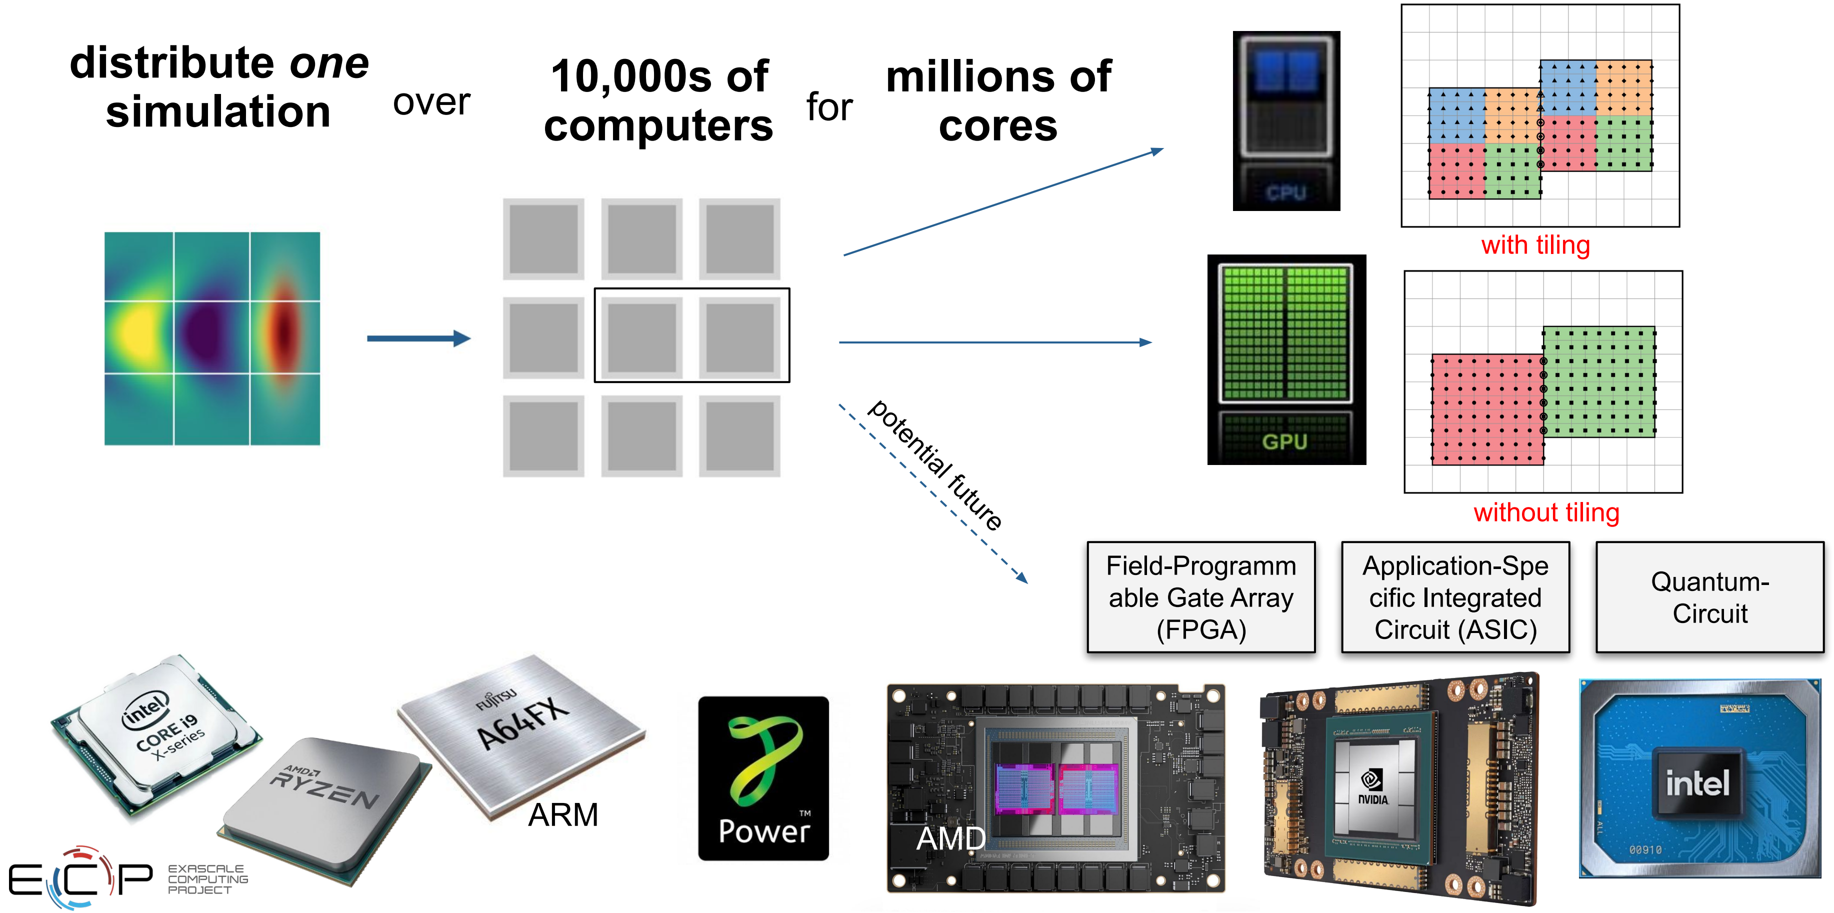

Portable Performance through Exascale Programming Model

AMReX library

- Domain decomposition & MPI **communications:** MR & load balance

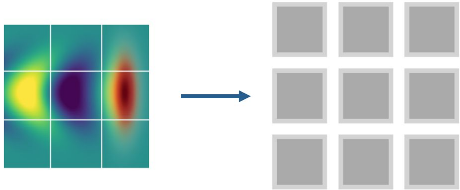

- **Performance-Portability Layer:** GPU/CPU/KNL

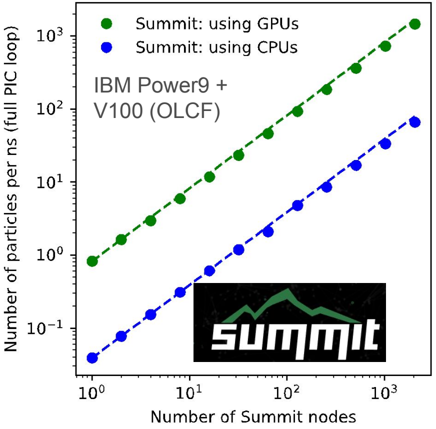

A100 gives additional ~< 2x

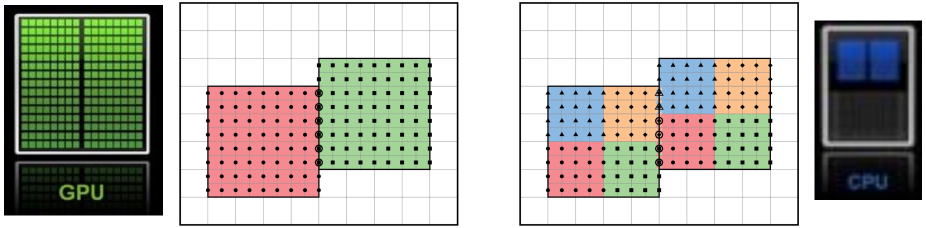

without tiling

with tiling

Data Structures

- **Write the code once**, specialize at ***compile-time***

`ParallelFor(/Scan/Reduce)`

```
amrex::ParallelFor( n_particles,
    [=] AMREX_GPU_DEVICE (long i) {

        UpdatePosition( x[i], y[i], z[i],
                ux[i], uy[i], uz[i], dt );

    });
```

- **Parallel linear solvers** (e.g. multi-grid Poisson solvers)
- **Embedded boundaries**

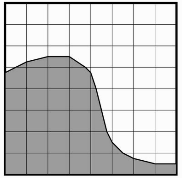

- Runtime parser for user-provided math expressions (incl. GPU)

A. Myers et al., "Porting WarpX to GPU-accelerated platforms," Parallel Computing 108, 102833 (2021)

WarpX is a GPU-Accelerated PIC Code for Exascale

BLAST
BEAM PLASMA & ACCELERATOR SIMULATION TOOLKIT

Available Particle-in-Cell Loops

- electrostatic & electromagnetic (fully kinetic)

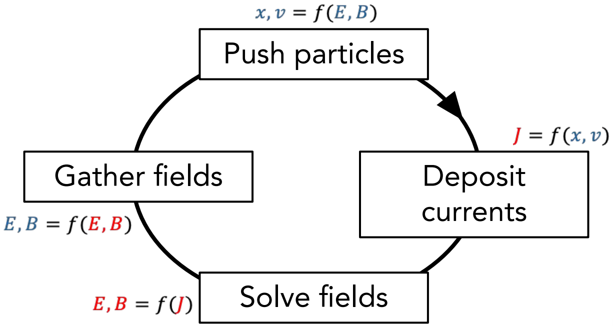

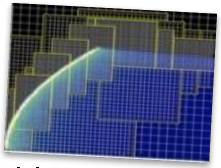

Advanced algorithms

boosted frame, spectral solvers, Galilean frame, embedded boundaries + CAD, MR, ...

Multi-Physics Modules

field ionization of atomic levels, Coulomb collisions, QED processes (e.g. pair creation), macroscopic materials

Geometries

- 1D3V, 2D3V, 3D3V and RZ (quasi-cylindrical)

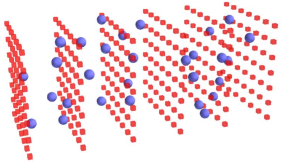

3D Cartesian grid

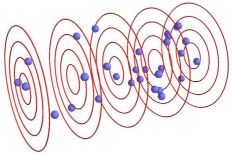

Cylindrical grid (schematic)

Multi-Node parallelization

- MPI: 3D domain decomposition
- dynamic load balancing

On-Node Parallelization

- GPU: CUDA, HIP and SYCL
- CPU: OpenMP

Scalable, Standardized I/O

- PICMI Python interface
- openPMD (HDF5 or ADIOS)
- in situ diagnostics

mac OS

EXASCALE COMPUTING PROJECT

Developed by a multidisciplinary, multi-institution team

EXASCALE COMPUTING PROJECT

U.S. DEPARTMENT OF ENERGY | Office of Science

NNSA National Nuclear Security Administration

Jean-Luc Vay (ECP PI) | Arianna Formenti | Marco Garten | Axel Huebl | Rémi Lehe | Ryan Sandberg | Olga Shapoval | Yinjian Zhao | Edoardo Zoni

ATAP ACCELERATOR TECHNOLOGY & APPLIED PHYSICS DIVISION

BLAST BEAM PLASMA & ACCELERATOR SIMULATION TOOLKIT

BERKELEY LAB

Ann Almgren (ECP coPI) | John Bell | Kevin Gott (NESAP) | Junmin Gu | Revathi Jambunathan | Hannah Klion | Prabhat Kumar | Andrew Myers | Weiqun Zhang

AMCRD AMReX

NERSC NESAP

David Grote (ECP coPI)

\+ a growing list of contributors from labs, universities…

CEA PARIS-SACLAY (France): Henri Vincenti | Luca Fedeli | Thomas Clark | Neïl Zaim | Pierre Bartoli

DESY (Germany): Maxence Thévenet

SLAC

Marc Hogan (ECP coPI) | Lixin Ge | Cho Ng

CERN (Switzerland): Lorenzo Giacomel

…& private sector

MODERN ELECTRON | Intense Computing | AVALANCHE | tae TECHNOLOGIES

Warp to WarpX: the journey

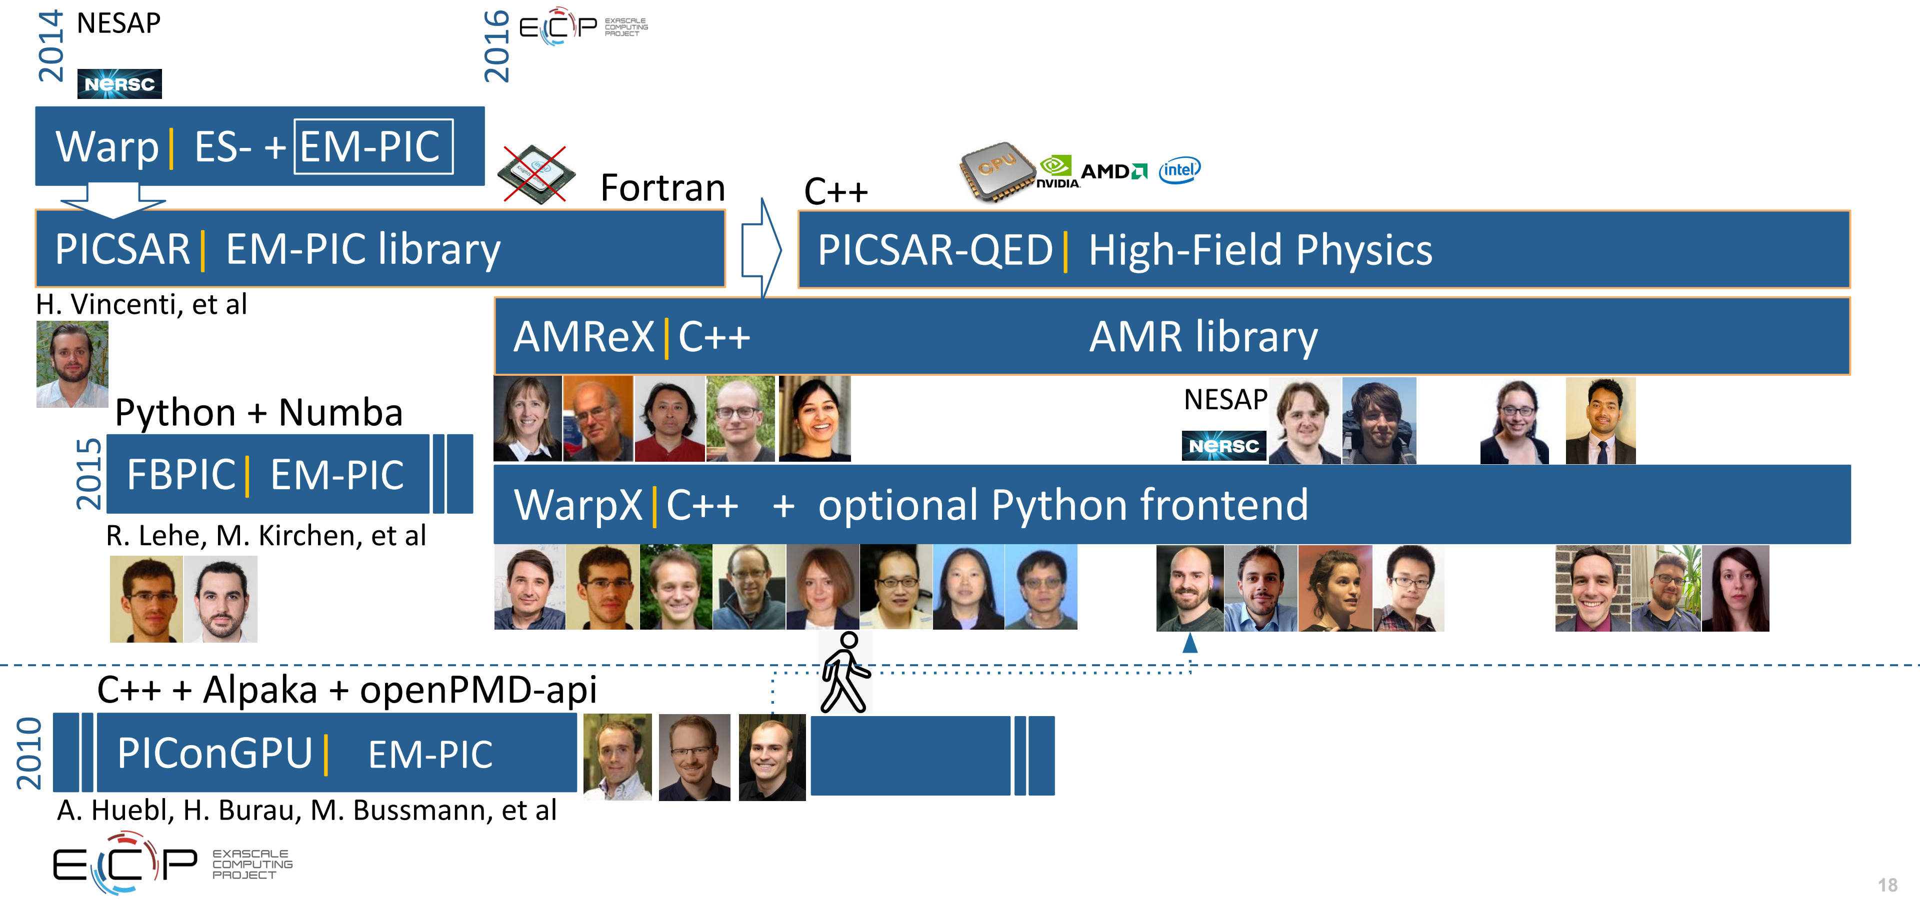

Software Productivity & Sustainability

Sustainable *documentation/knowledge, code* and *maintenance* (e.g., testing)

Online Documentation:
warpx|hipace|impactx.readthedocs.io

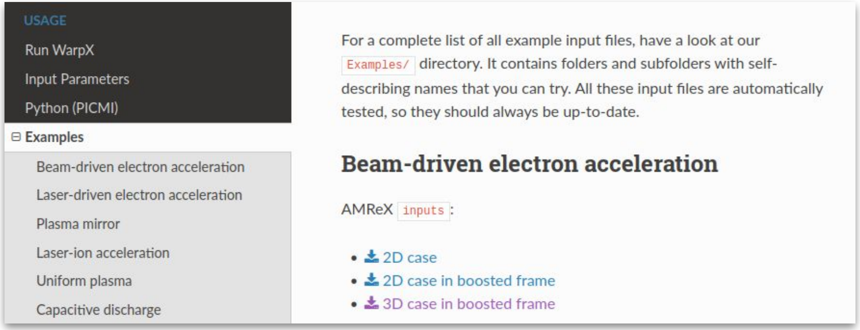

Open-Source Development & Benchmarks:
github.com/ECP-WarpX

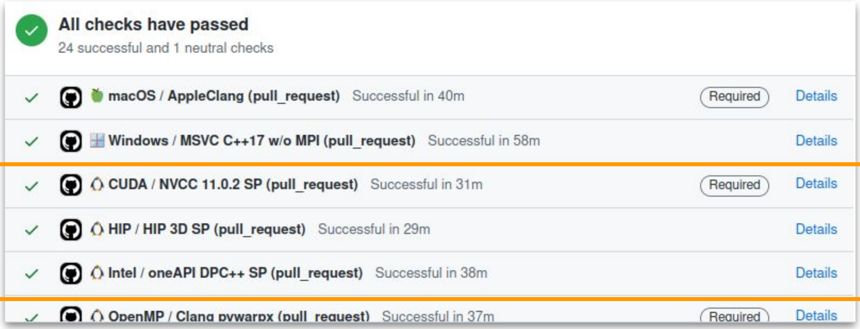

188 physics benchmarks *run on every code change* of WarpX
13 physics benchmarks + 32 tests *for* ImpactX

PSIP processes: onboarding, etc.

Productivity and **S**ustainability **I**mprovement **P**lanning

RateYourProject.org
bssw.io/psip

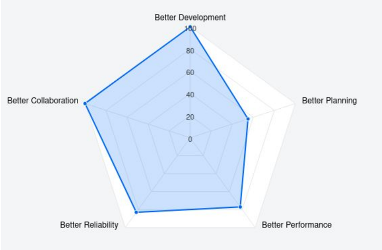

Automated Performance on Targets

- OLCF Ascent: IBM w/ V100
- NERSC GitLab: Cray w/ A100

Continuous Deployment

Rapid and easy installation on any platform:

conda install
-c conda-forge warpx

spack install warpx
spack install py-warpx

python3 -m pip install .

brew tap ecp-warpx/warpx
brew install warpx

cmake -S . -B build
cmake --build build
--target install

E4S

module load warpx
module load py-warpx

WarpX

Azure Pipelines succeeded | nightly packages succeeded | docs passing | spack v23.01 | conda-forge v23.01 | chat on gitter

platforms linux | osx | win | commits since 23.01 10 | supported by ECP | language C++17 | language Python

license BSD-3-Clause-LBNL | DOI (source) 10.5281/zenodo.4571577 | DOI (paper) 10.1016/j.parco.2021.102833

Summary

Jobs

Name
Ubuntu pip from dev w/ OpenMPI
Ubuntu pip from dev
Ubuntu conda
Ubuntu conda w/ OpenMPI
Ubuntu mamba
Ubuntu mamba w/ OpenMPI
Ubuntu spack
Ubuntu spack CUDA
Ubuntu spack w/o MPI
macOS pip from dev w/ OpenMPI
macOS pip from dev w/o OMP
macOS conda
macOS conda w/ OpenMPI
macOS mamba
macOS mamba w/ OpenMPI
macOS spack
macOS spack from dev w/o MPI
Windows pip from dev
Windows conda
Windows mamba

IDEAS productivity

ECP EXASCALE COMPUTING PROJECT

Now (2023+):
WarpX + Ecosystem

WarpX in ECP: Staging of Laser-Driven Plasma Acceleration

Goal: deliver & scientifically use the nation's first exascale systems

- **ExaFLOP:** a quintillion (10^{18}) calculations per second
- ensure ***all*** the necessary pieces are *concurrently* in place

Our DOE science case is in **HEP**, our methods are **ASCR**:

first 3D simulation of a chain of plasma accelerator stages for future colliders

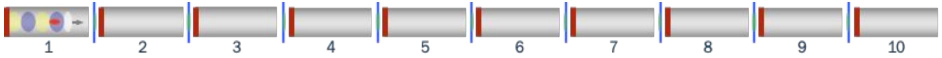

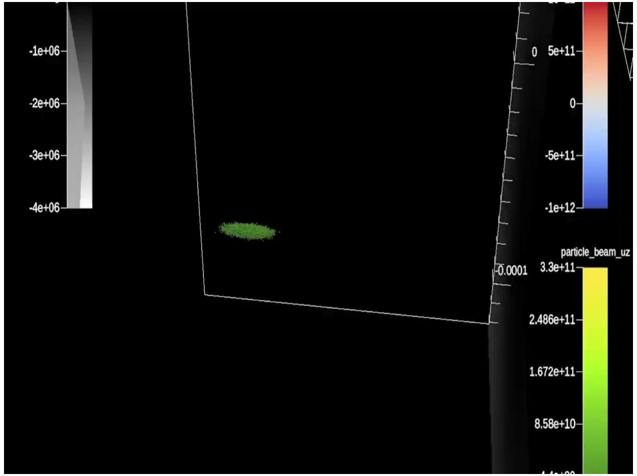

First-of-their-kind platforms: NERSC (Intel, then Nvidia)→Exascale: OLCF (AMD), ALCF (Intel)

J.-L. Vay, A. Huebl et al., ISAV'20 Workshop Keynote (2020) and PoP 28.2, 023105 (2021); L. Fedeli, A. Huebl et al., SC22 (2022)

2022 ACM Gordon Bell Prize: using the First Exascale Supercomputer

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

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

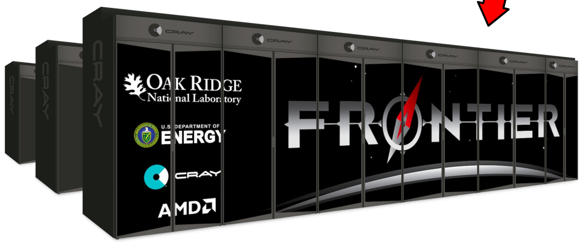

modeling of novel plasma e- beam injection scheme

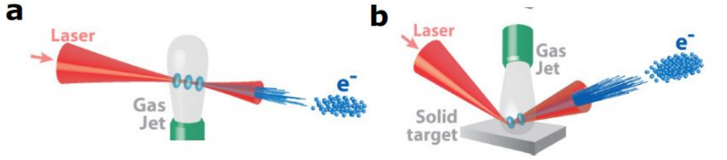

Fig. 1: Sketches showing the focusing of a high-power femtosecond laser (a) into a gas jet (b) onto a hybrid solid-gas target.

CEA PARIS-SACLAY · BERKELEY LAB · CNRS · GENCI · RIKEN · arm · Atos

A success story of a multidisciplinary, multi-institutional team!

2022 ACM Gordon Bell Prize: using the First Exascale Supercomputer

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

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

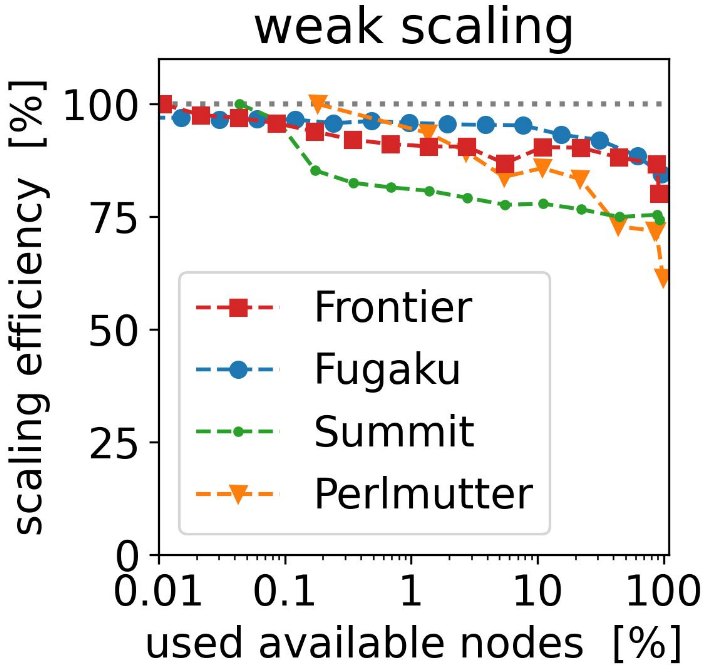

Note: Perlmutter & Frontier are pre-acceptance measurements!

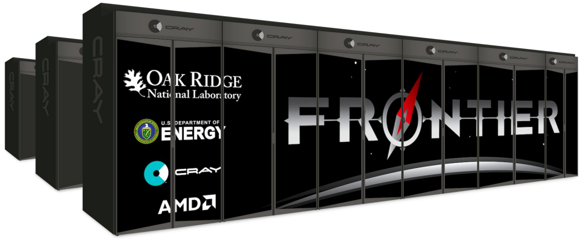

Figure-of-Merit: weighted updates / sec

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

110x

500x

EXASCALE COMPUTING PROJECT

2022 ACM Gordon Bell Prize: using the First Exascale Supercomputer

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

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

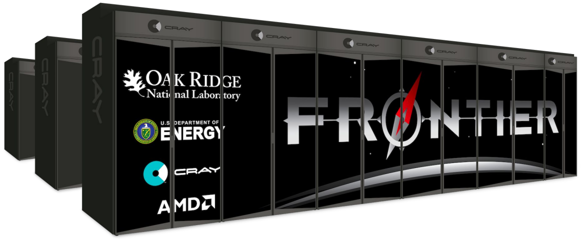

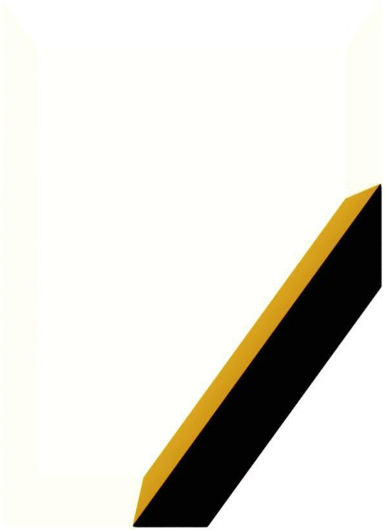

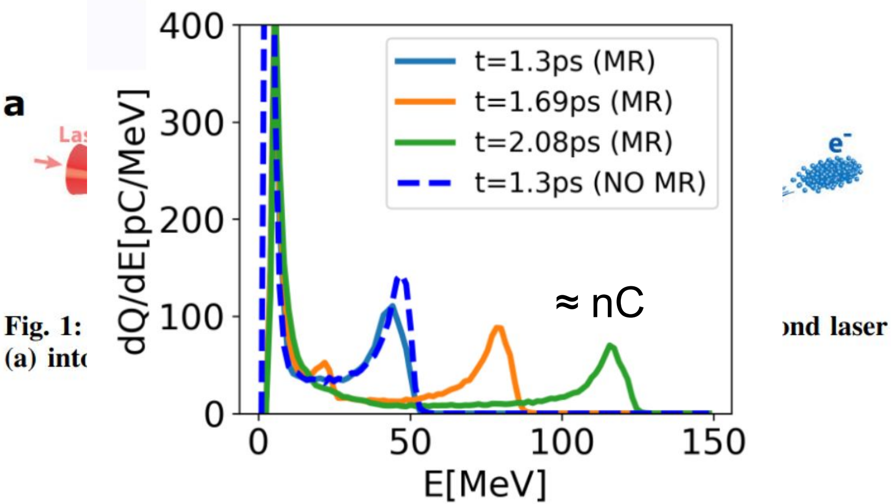

A success story of a multidisciplinary, multi-institutional team!

cea PARIS-SACLAY · BERKELEY LAB · cnrs · GENCI · RIKEN · arm · Atos

2022 ACM Gordon Bell Prize: using the First Exascale Supercomputer

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

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

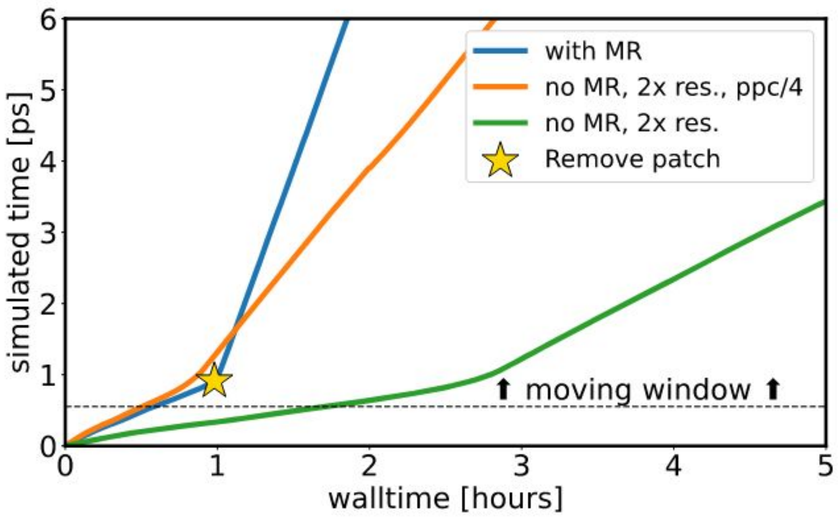

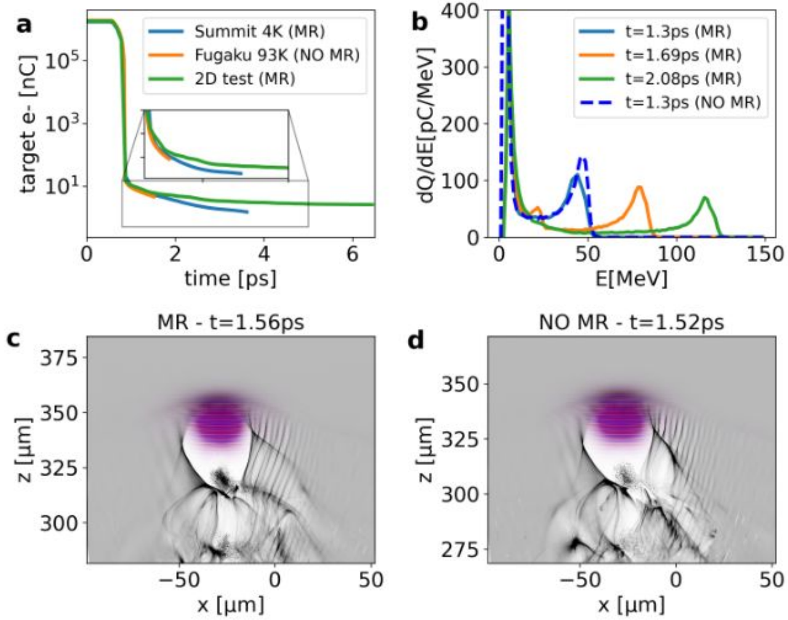

A success story of a multidisciplinary, multi-institutional team!

Is an ExaFlop/s (2022) 1,000x "*faster*" than a PetaFlop/s (2008)?

For the **exact same simulation size**, time-to-solution is *at best down by 20-100x!*

Distribute *one* Simulation over **10,000s of computers**

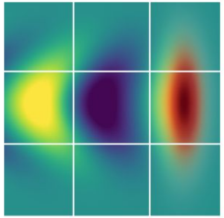

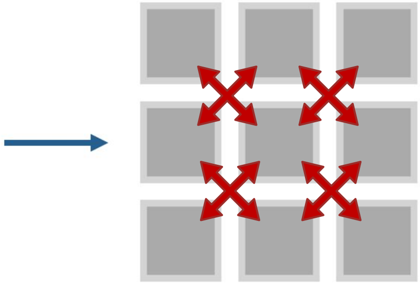

Data Communication

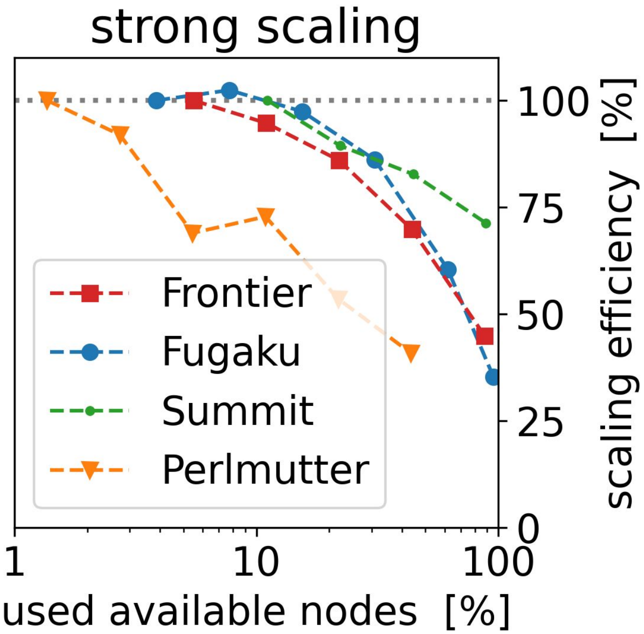

Note: Perlmutter & Frontier are pre-acceptance measurements!

Is an ExaFlop/s (2022) 1,000x "*faster*" than a PetaFlop/s (2008)?

For the **exact same simulation size**, time-to-solution is *at best down by 20-100x!*

Distribute *one* Simulation over **10,000s of computers**

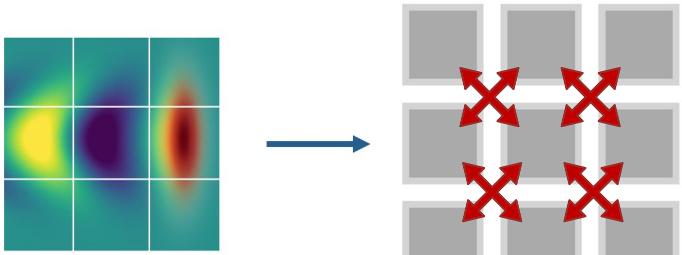

Data Communication

We now have **more parallelism**!
Let's model **more physics**:

- higher grid resolution
- more particles
- resolve ion motion & collisions
- resolve emittance growth from collisions
- 2D → 3D
- long-term stable, advanced solvers
- add high-field effects
- …

HEP Science Drivers Require Exascale

AAC Linear collider design: PBA & LWFA

PBA XFELs design
Extreme field science

High-Field Physics
Advanced ion sources

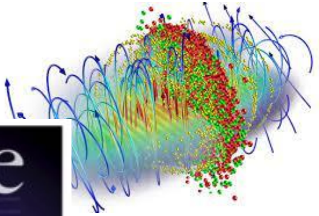

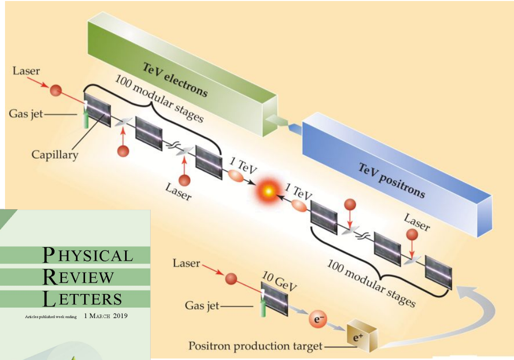

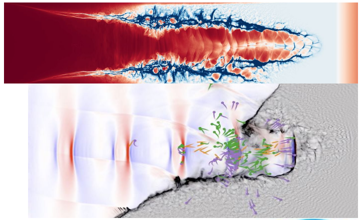

UCLA
+ LLE, Beijing Normal, Tsinghua

BERKELEY LAB
cea PARIS-SACLAY
DESY.

Energy-scalable plasma accelerators

Traveling-wave electron acceleration (TWEAC) simultaneously circumvents the LWFA limitations of diffraction, dephasing and depletion.

- non-rotationally symmetric 3D geometry
- long acceleration lengths

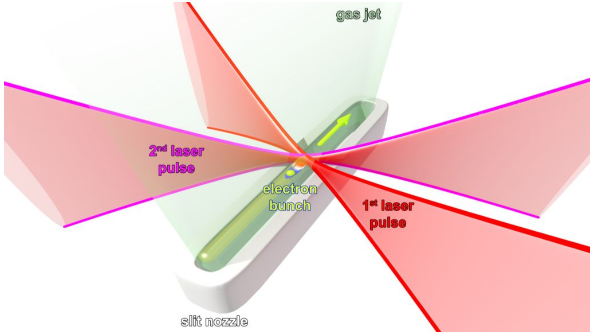

University of Delaware
CASUS Center for Advanced Systems Understanding
HZDR Helmholtz Zentrum Dresden Rossendorf
Oak Ridge National Laboratory Leadership Computing Facility
Georgia Tech
EuroHPC Joint Undertaking

Plasma-PEPSC
Centre of Excellence

Exascale PIC Modeling Benefits Plasma Science Beyond HEP

Light Sources for Basic Energy Science & National Security

LPA beams for light sources

all-optical undulators

Industry

EUV semiconductor lithography

commercial fusion energy research

Fusion-Energy Sciences

BELLA iP2 target chamber

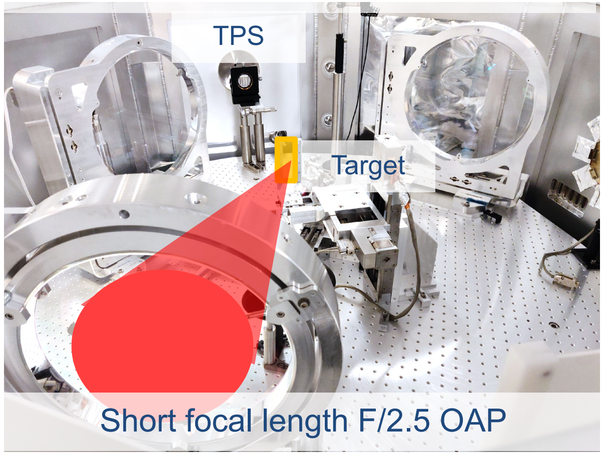

BERKELEY LAB

LaserNetUS

Life Sciences & Medicine

FLASH radio-biology

Astrophysics

origin of strong astrophysical fields & largest energy electrons

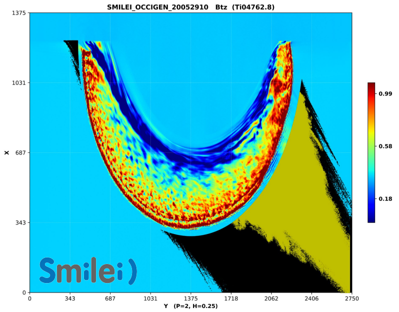

em. radiation, turbulence & electron beams from solar flares

magnetic reconnection

Our ECP Software Stack

BLAST
BEAM PLASMA & ACCELERATOR SIMULATION TOOLKIT

Spack

E4S

Python: Modules, PICMI interface, Workflows

WarpX
full PIC, LPA/LPI

HiPACE++
quasi-static, PWFA

ARTEMIS
microelectronics

ImpactX
accelerator lattice design

Object-Level Python Bindings
extensible, AI/ML

pyAMReX

PICSAR
QED Modules

ABLASTR library: common PIC physics

AMReX
Containers, Communication, Portability, Utilities

Diagnostics
I/O
code coupling

openPMD

ADIOS2

HDF5

ZFP

...

Ascent

VTK-m

FFT
on- or multi-device

Lin. Alg.
BLAS++
LAPACK++

MPI

CUDA, OpenMP, SYCL, HIP

mac OS

Desktop to HPC

ECP EXASCALE COMPUTING PROJECT

Develop, CI/CD **I/O** **Data & Compute** **Physics**

Conclusions

The ECP *multi-year* strategy enabled

- **integration**: dependencies on vendors and ST require iteration
  - continuous improvements through ***multiple release cycles***
  - coherent, feature-complete, algorithmically innovative ***software products***
- **redesigns**: adjustments to technical execution were possible
  - compilers lacked support: decided for a Fortran to C++ migration mid project
  - Python infrastructure overhaul to GPU-capable methods
  - performance optimizations led to redesign of particle data structures
- **risk mitigation**: e.g., replace self-made with better I/O *methods from ECP ST*
  - *Contingency Funding Requests*: efficient process & funds
- **trust & collaboration**: attracted national & internat. contributors
  - trusted us: their FOSS contributions will be maintained
  - enabled us: leveraging of investment in WarpX spin-off/follow-up projects that *contribute back*

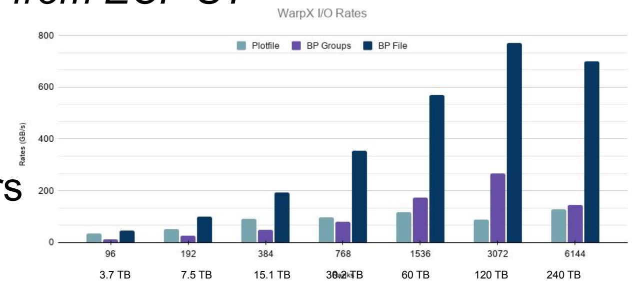

Write strategies: plotfiles → ADIOS BP per rank & step → ADIOS BP w/ append to subfiles

EXASCALE COMPUTING PROJECT

Trust & Collaboration: Ecosystem Release Cycles

Dependencies become *"Team of Teams"*

- vendors, ST, centers, …
- own sub-libs/modules, contributions
- machines & deployed environments

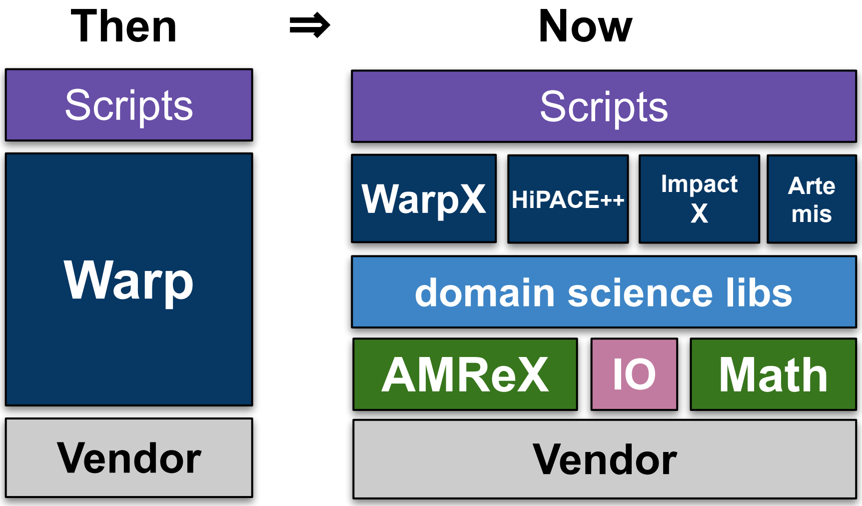

Incentives to integrate at all levels

- part of KPP's
- aimed at *production level*, not just prototypes

Prioritization of vendor bugs & features

- AMD: CMake scripts, Compiler
- Intel: DPC++ compiler
- Nvidia: compiler bugs, warn host mem access

7yrs ECP vs. shorter multi-yr projects:

- **Scalable collaboration:** easy to add partners
- **Multiple iterations** over release cycles with partners
- **Deep SW stack** that is continuously developed together

We are Establishing an Open Community Ecosystem with Standards

Complexities in accelerator modeling **require to work together**

common input → **Particle-In-Cell Modeling Interface** (PICMI Standard) → **Code A** / **Code B** / **...** → **open Particle Mesh Data standard** (openPMD) → common output

We established a **leadership role** in our community

- push to **organize**: enhance cooperation, avoid duplication
- **sustainable** development: code ecosystem

US

We lead the
Consortium for Advanced
Modeling of Particle Accelerators

BERKELEY LAB, SLAC, Fermilab, UCLA

CAMPA DOE HEP GARD - now DOE SciDAC-5

Lawrence Livermore National Laboratory, radiasoft, TECH-X

Internationally

HZDR Helmholtz-Zentrum Dresden-Rossendorf, BERKELEY LAB, CASUS Center for Advanced Systems Understanding, Oak Ridge National Laboratory, Technische Universität Dresden, European XFEL, Perimeter Institute, IOQ Jena Friedrich-Schiller-Universität, Lawrence Livermore National Laboratory, Fermilab, SLAC National Accelerator Laboratory, Maison de la Simulation, DESY, UHH, Cornell University, CBG Max Planck Institute of Molecular Cell Biology and Genetics, UCLA, NVIDIA

SciDAC Scientific Discovery through Advanced Computing

Conclusion

- **ECP enabled**
  - Complete rewrite of efficient code for CPUs+GPUs for focused area of research
  - Ability to combine innovation in algorithms with very efficient implementation
  - Collaboration WarpX+AMReX is the best that many of us have experienced

  ⇒ great outcome: ***500x FOM*** for Laser-Particle Acceleration; ***Gordon Bell Prize '22***

- **Sustainable development over nearly a decade**
  - across teams, labs and technologies: ***one "felt" as a part a community***
  - integrated teams of comput. phys. + applied math. + comput. sc. + software eng.
- **Management Structure**
  - Focused main deliverables, flexibility to correct course, ***freedom to innovate***
  - Felt a good balance between oversight and overhead

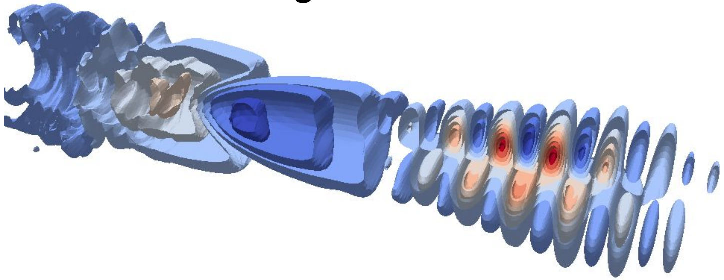

Funding Support

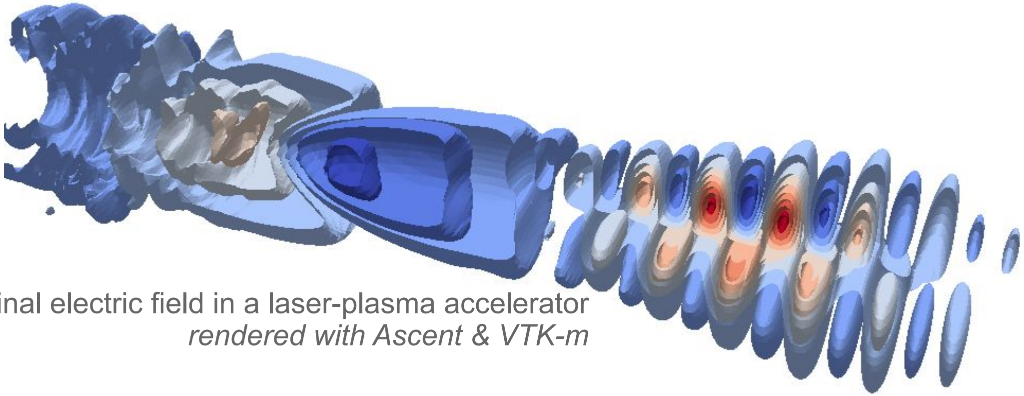

WarpX: longitudinal electric field in a laser-plasma accelerator
rendered with Ascent & VTK-m

This research was supported by the **Exascale Computing Project** (17-SC-20-SC), a collaborative effort of two **U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration)** responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering and early testbed platforms, in support of the nation's exascale computing imperative. This work was also performed in part by the **Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory** under U.S. Department of Energy Contract No. DE-AC02-05CH11231, **Lawrence Livermore National Laboratory** under Contract No. DE-AC52–07NA27344 and **SLAC National Accelerator Laboratory** under Contract No. AC02–76SF00515. This research used resources of the **Oak Ridge Leadership Computing Facility**, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725, the **National Energy Research Scientific Computing Center (NERSC)**, a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231, and the supercomputer Fugaku provided by **RIKEN**.

Backup Slides