Modeling of Advanced-Concepts of Particle Accelerators at the Exascale

EXASCALE COMPUTING PROJECT

Jean-Luc Vay, 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

CSE 6230 – HPC Tools and Applications
GATech

Guest Lecture
virtual

April 13, 2023

National Nuclear Security Administration

BERKELEY LAB

ATAP
ACCELERATOR TECHNOLOGY & APPLIED PHYSICS DIVISION

U.S. DEPARTMENT OF ENERGY | Office of Science

Outline

- **Who we are**
- **Intro -** particle accelerator modeling
- **Exascale project WarpX**
  - Overview & advanced algorithms
  - Preparing WarpX for the world's largest supercomputers
  - Open Science in HPC
  - 2022 Gordon Bell Prize
  - Machine Learning
- **Outlook**

Who we are

Accelerator Modeling Program (AMP)
in Accelerator Technology & Applied Physics Division (ATAP)
in Physical Sciences Area (PSA)
@ Lawrence Berkeley National Laboratory (LBNL, aka Berkeley Lab)

Jean-Luc Vay
Senior Scientist
- Head of AMP
- PI of ECP project WarpX
- PI of SciDAC Collaboration on Advanced Modeling of Particle Accelerators (CAMPA)
- >30 years experience in Particle-in-Cell codes development & application

Axel Huebl
Research Scientist
- Lead software architect of WarpX & BLAST
- Lead developer of PIConGPU, first Particle-in-Cell code at scale on Titan (OLCF)
- >14 years of experience in SWE of HPC, FOSS, data sci.; laser-plasma modeling

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

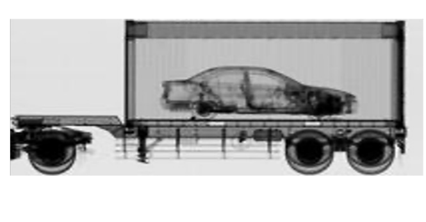

- **Cargo** scanning
- Active interrogation
- **Stockpile stewardship**: materials character-ization, radiography, support of non-proliferation

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

There are many types of particle accelerators: cyclotrons

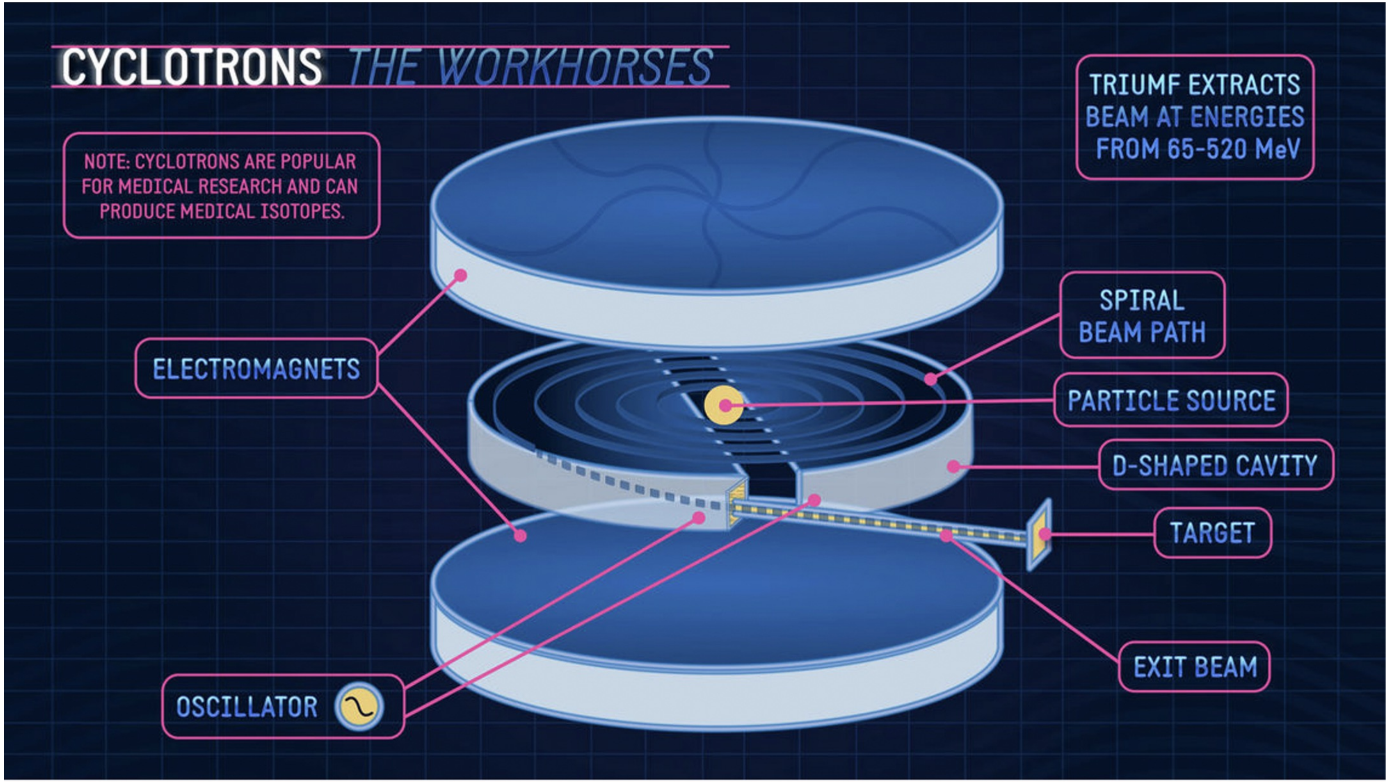

Artwork by Sandbox Studio, Chicago with Jill Preston

“A primer on particle accelerators”, Signe Brewster, Symmetry Magazine 07/12/2016

There are many types of particle accelerators: synchrotrons

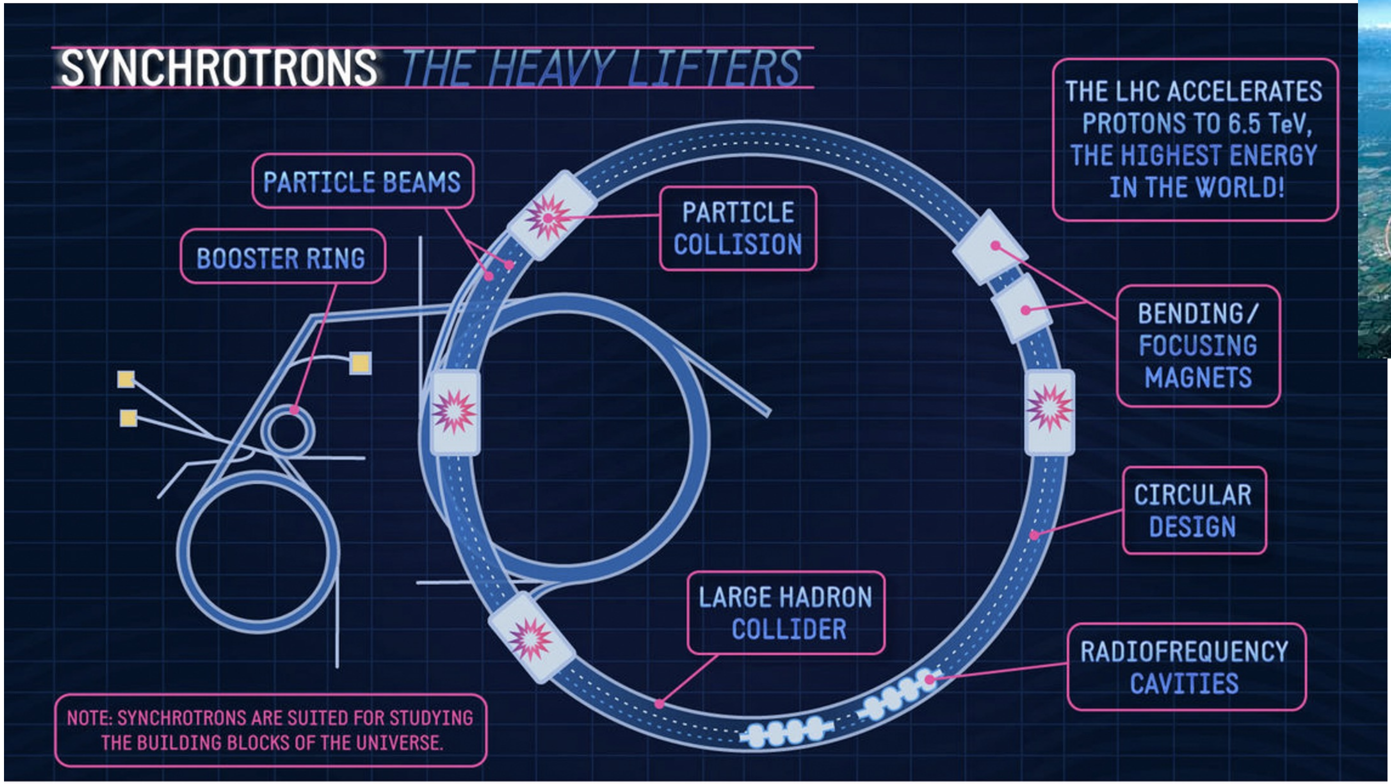

Artwork by Sandbox Studio, Chicago with Jill Preston

"A primer on particle accelerators", Signe Brewster, Symmetry Magazine 07/12/2016

EXASCALE COMPUTING PROJECT

There are many types of particle accelerators: linacs

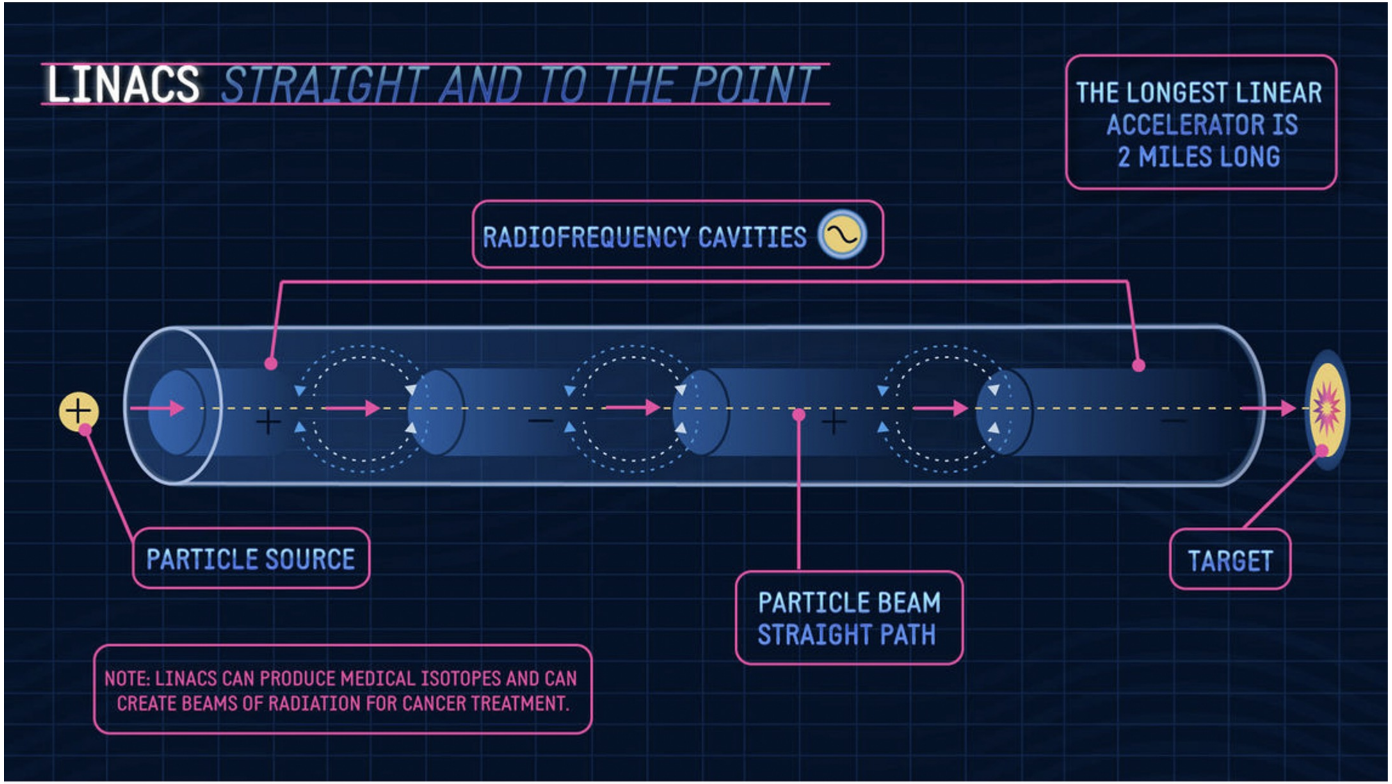

Artwork by Sandbox Studio, Chicago with Jill Preston

"A primer on particle accelerators", Signe Brewster, Symmetry Magazine 07/12/2016

ECP EXASCALE COMPUTING PROJECT

All these types of "conventional" accelerators involve a metallic pipe with vacuum inside.

⇒ breakdown occurs if electric field is too high!

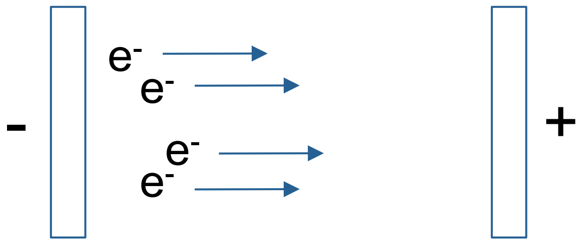

Possible solution
to reach higher accelerating fields?
⇒ plasmas.

In plasma accelerators, a driver beam displaces plasma electrons

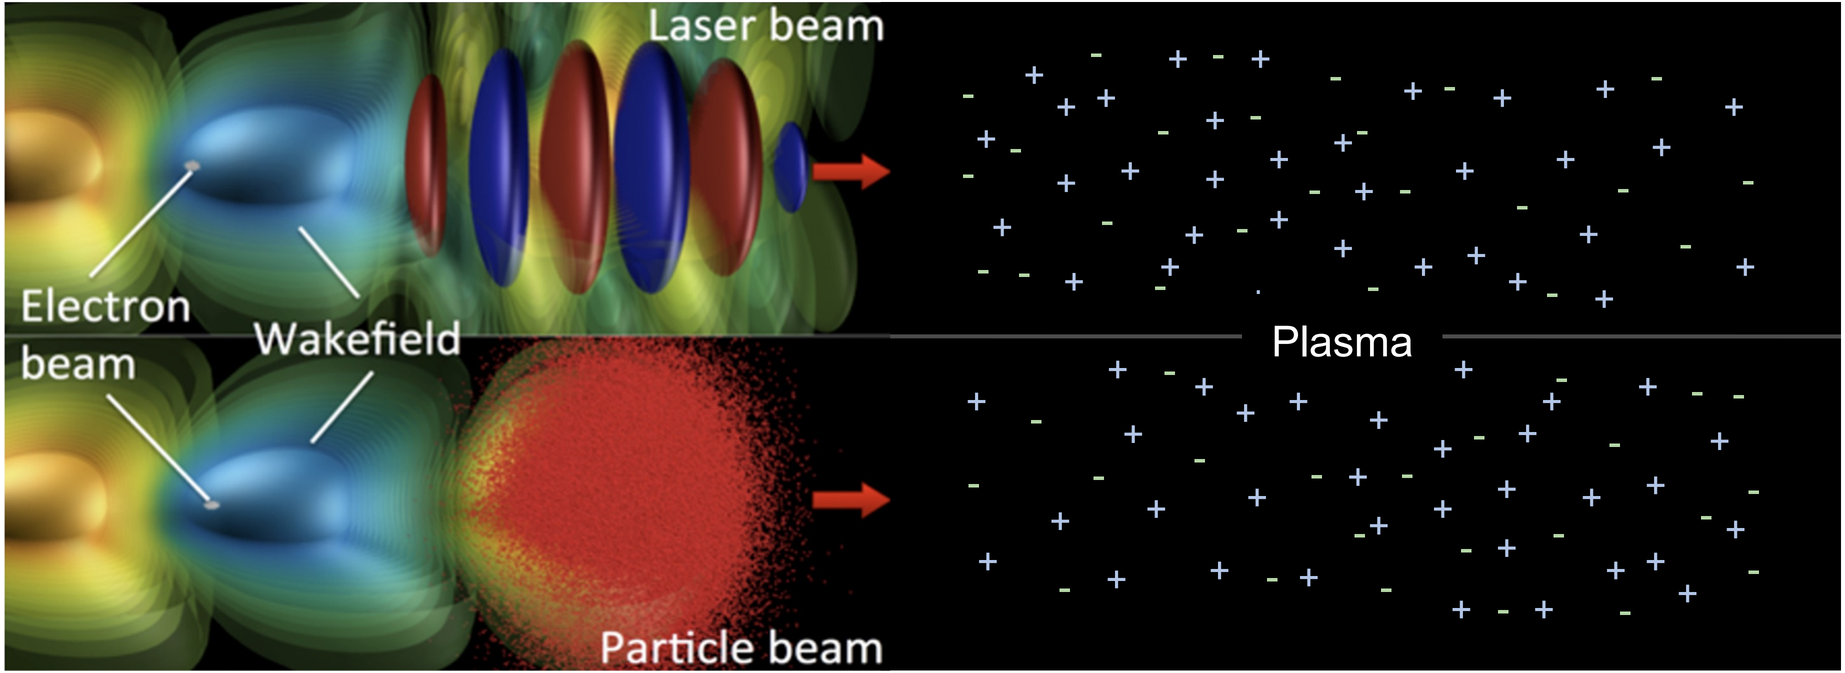

Most used modeling approach is based on the Particle-In-Cell method

Lagrangian macro-particles

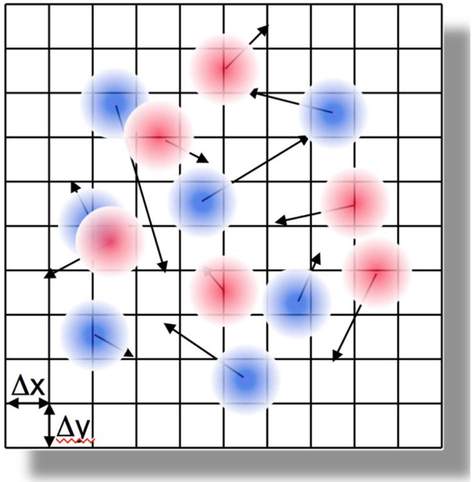

Eulerian fields on grids
(usually Cartesian)

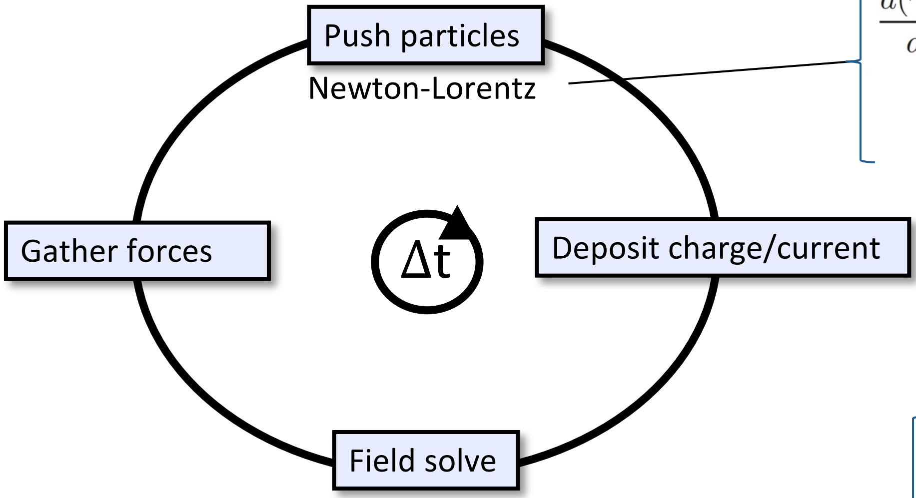

$$\frac{d(\gamma \mathbf{v})}{dt} = \frac{q}{m}(\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

$$\frac{d\mathbf{x}}{dt} = \mathbf{v}$$

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E},$$

$$\frac{\partial \mathbf{E}}{\partial t} = \nabla \times \mathbf{B} - \mathbf{J}$$

Plasma accelerators are challenging to model

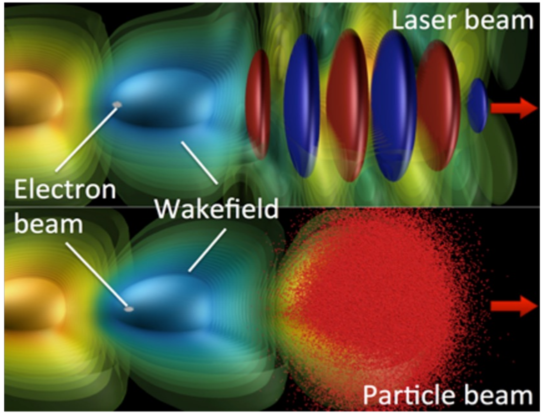

Short driver+wake propagates through long plasma

Many time steps.

E.g., for a 10 GeV LPA scale stage:
~1µm wavelength laser propagates into **~1m** plasma
millions of time steps needed (even with moving window)

Very small features:

Many grid cells.

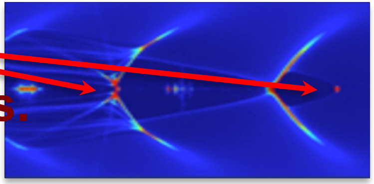

Simulations (in 2D) can take days for 1 stage (at insufficient resolution for collider beam quality).

For multi-TeV collider, need for ×10s-1000s stages ×10s-1000s (high res. 3D) ×10s (ensembles)!

Need for advanced algorithms and supercomputing!

The Exascale Computing Project

WarpX

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

Exascale Computing Initiative

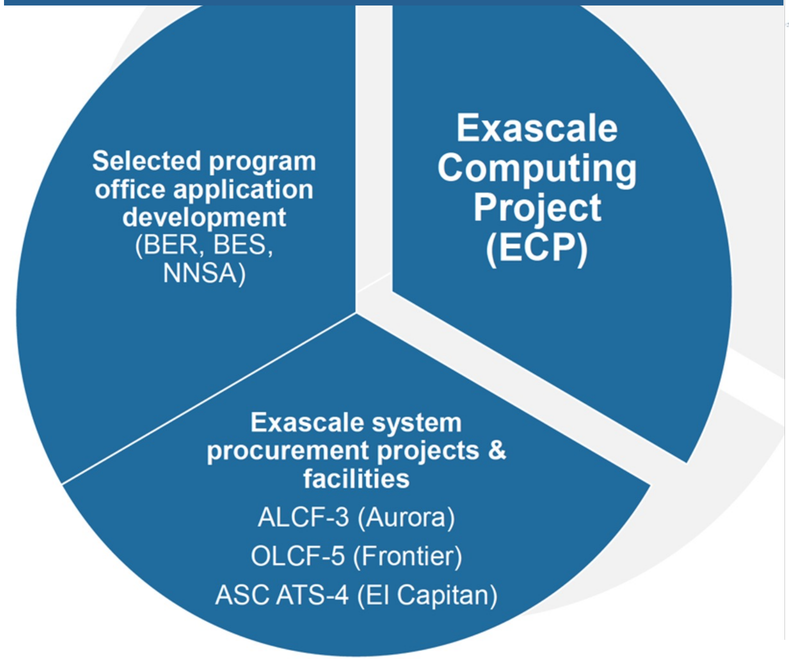

Hardware — Delivers 5x performance of 200 petaflop Summit (ORNL)

Software — Software stack for broad spectrum of apps & workloads

Applications — Wide range of apps that deliver high-fidelity solutions faster to more complex problems

D. Kothe – April 30, 2019

WarpX among 21 applications selected to cover broad range of science

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

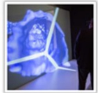

ExaAM

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

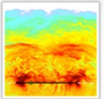

MFIX-Exa

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

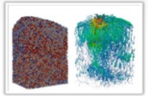

GAMESS

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

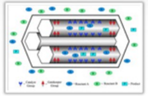

EXAALT

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

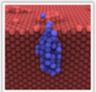

ExaStar

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

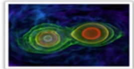

ExaSMR

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

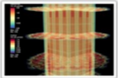

Subsurface

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

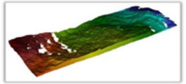

QMCPACK

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

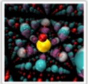

ExaSGD

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

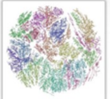

WDMApp

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

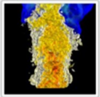

ExaFEL

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

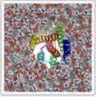

ExaBiome

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 availabilty

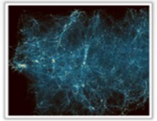

LatticeQCD

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

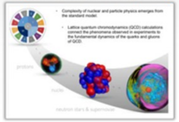

WarpX

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

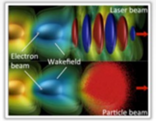

CANDLE

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

WarpX is the new version of previous code Warp

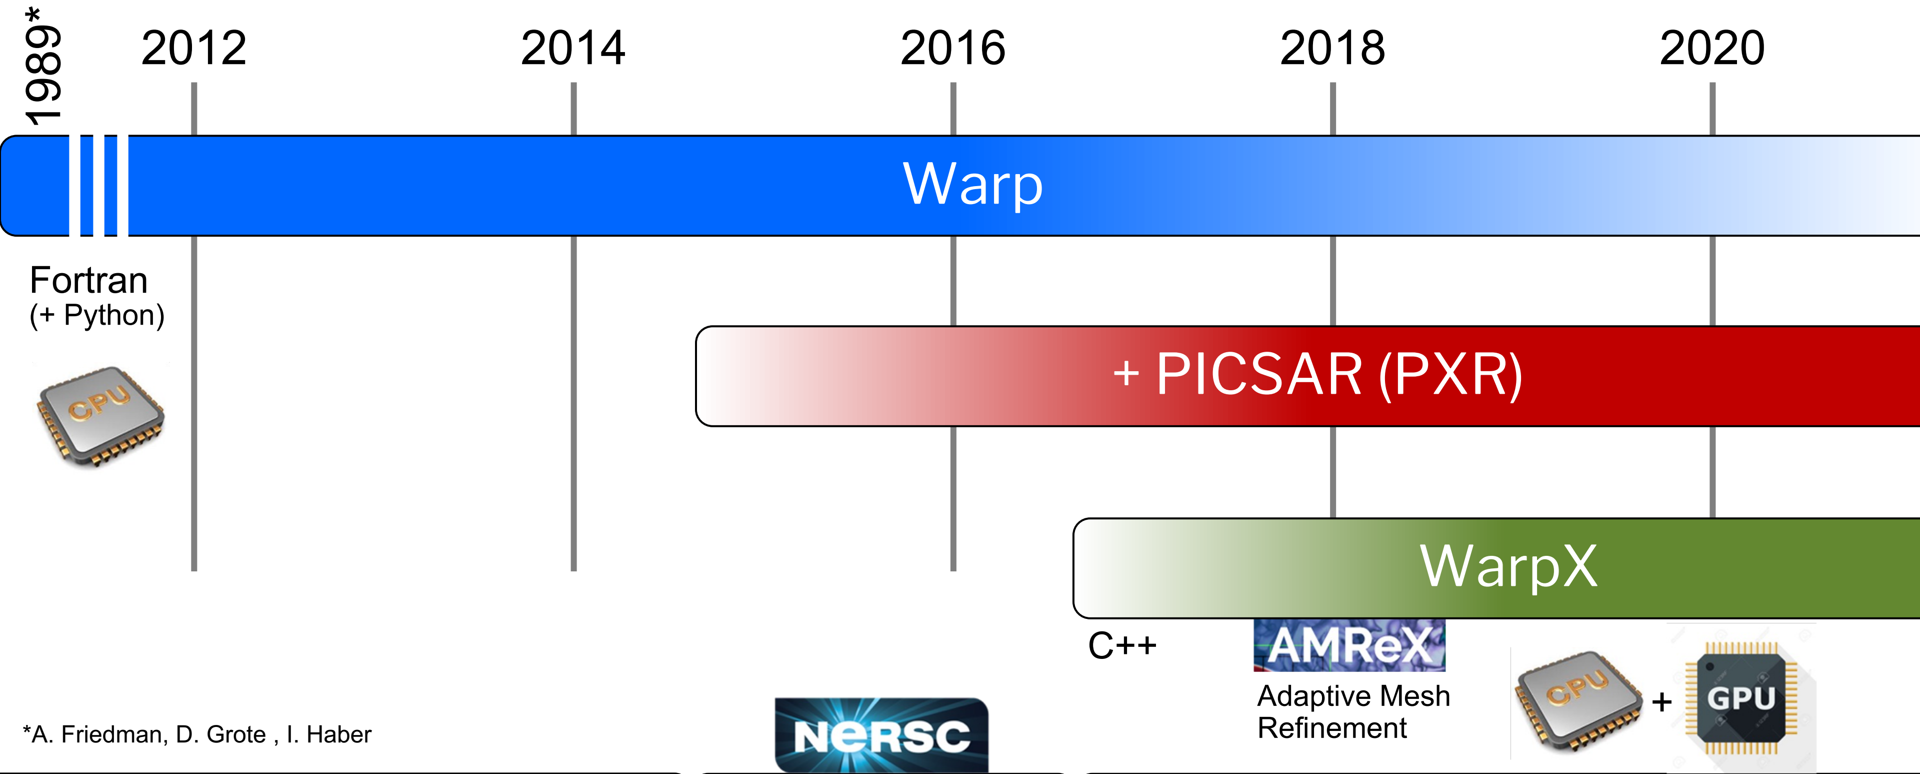

*A. Friedman, D. Grote , I. Haber

U.S. DEPARTMENT OF ENERGY | Office of Science

SciDAC Scientific Discovery through Advanced Computing

NESAP

ECP EXASCALE COMPUTING PROJECT

Timeline for Exascale project WarpX

From initial code coupling to ensemble of chains of multi GeV-scale stages

- **2030 -** Collider designs
- **2023** Modeling of chain of 50-100 multi-GeV-scale plasma accelerator stages
- **2022** Modeling chain of 20-50 multi-GeV-scale plasma accelerator stages
- **2021** Modeling chain of 10-20 multi-GeV-scale plasma accelerator stages
- **2020** Modeling chain of 5-10 multi-GeV-scale plasma accelerator stages
- **2019** Modeling chain of 3 multi-GeV-scale plasma accelerator stages
- **2018** Modeling a single multi-GeV-scale plasma accelerator stage
- **2017** Modeling of single plasma accelerator stage with static mesh refinement
- **2016** Initial coupling of AMReX and Warp/PICSAR

Frontier (ORNL)
Aurora (ANL)
>1 ExaFlops

All C++
Port to GPUs.

EXASCALE COMPUTING PROJECT

WarpX: conceived & developed by a multidisciplinary, multi-institution team

EXASCALE COMPUTING PROJECT

U.S. DEPARTMENT OF ENERGY | Office of Science

NNSA National Nuclear Security Administration

Advanced algorithms are needed in addition to supercomputing

Lower # time steps

- optimal Lorentz boosted frame

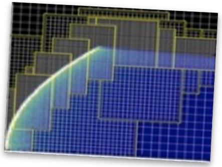

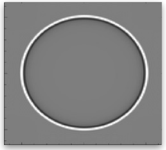

Higher accuracy

- (Adaptive) Mesh Refinement
- Spectral (FFT-based) Maxwell solvers (PSATD)

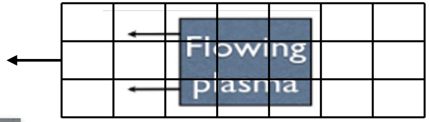

Higher stability

- Galilean PSATD suppresses Numerical Cherenkov Instability (NCI)

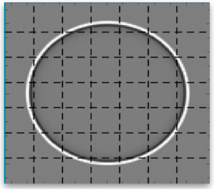

Higher scalability

- PSATD: FFT on local subdomains

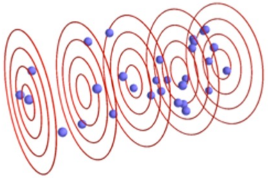

Lower dimensionality, reduced physics

- Axisymmetric solver with azimuthal Fourier decomposition
- Envelope laser solvers

Use Lorentz boosted frame of reference cuts simulation times drastically

Use Lorentz boosted frame instead of Lab frame*

time steps α L/λ

time steps α L’/λ′

Speedup = $(L'/\lambda')/(L/\lambda) = \gamma^2(1+\beta)$

For γ=100, speedup > 10,000!

Price to pay:

- Physics looks different in boosted frame and lab frame ⇒ need to transform between boost & lab frame.
- Potential numerical instabilities (numerical Cherenkov) ⇒ solve Maxwell in Galilean frame (see below).

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

Time

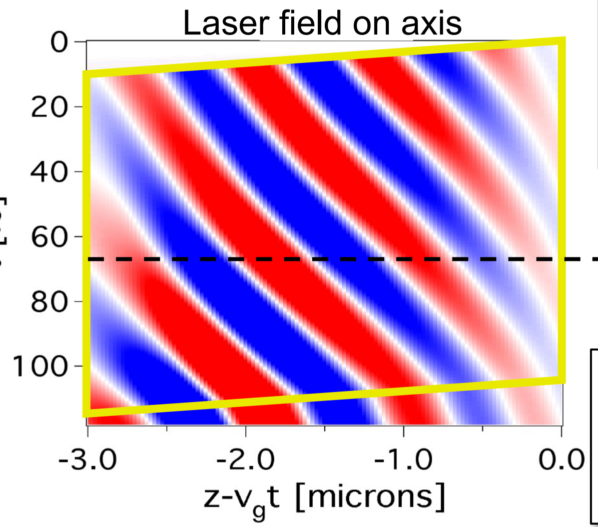

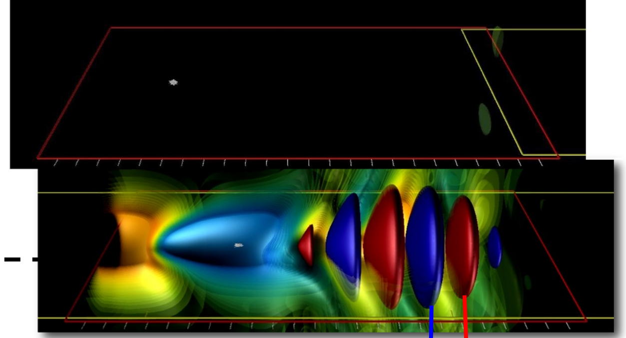

Lab frame

Hyperbolic rotation from Lorentz Transformation converts laser…

…*spatial oscillations* into *time beating*

Time

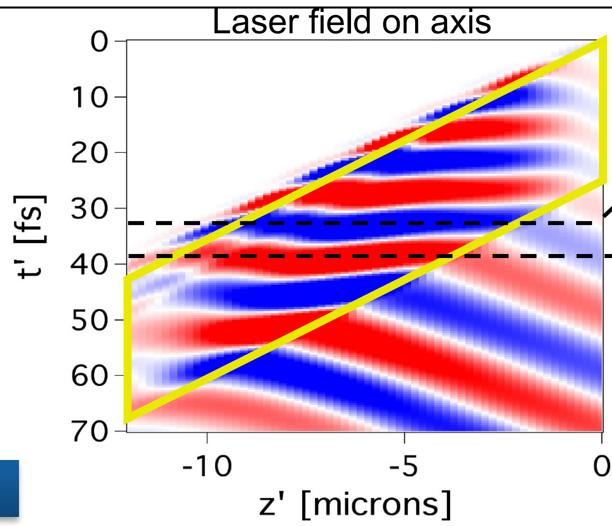

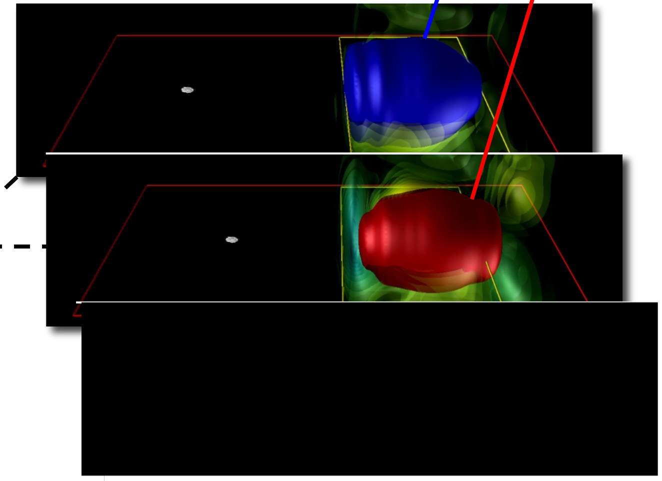

Wake frame

Care needed to ensure frame-independent initial conditions

Initial conditions known in lab frame:

1. Lorentz transform to boosted frame.
2. Perform injection of particle & laser beams through a moving plane.

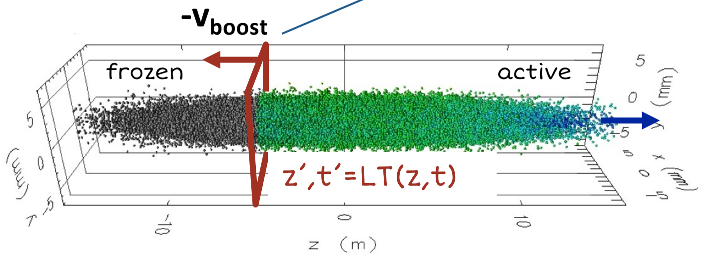

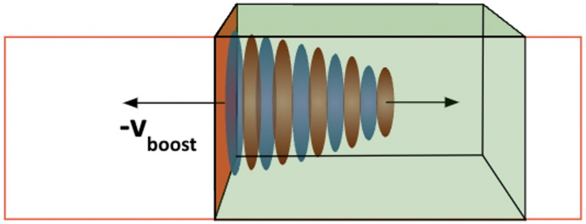

Also need to reconstruct output data in lab frame.

Reconstruction of output data from boosted frame to laboratory frame

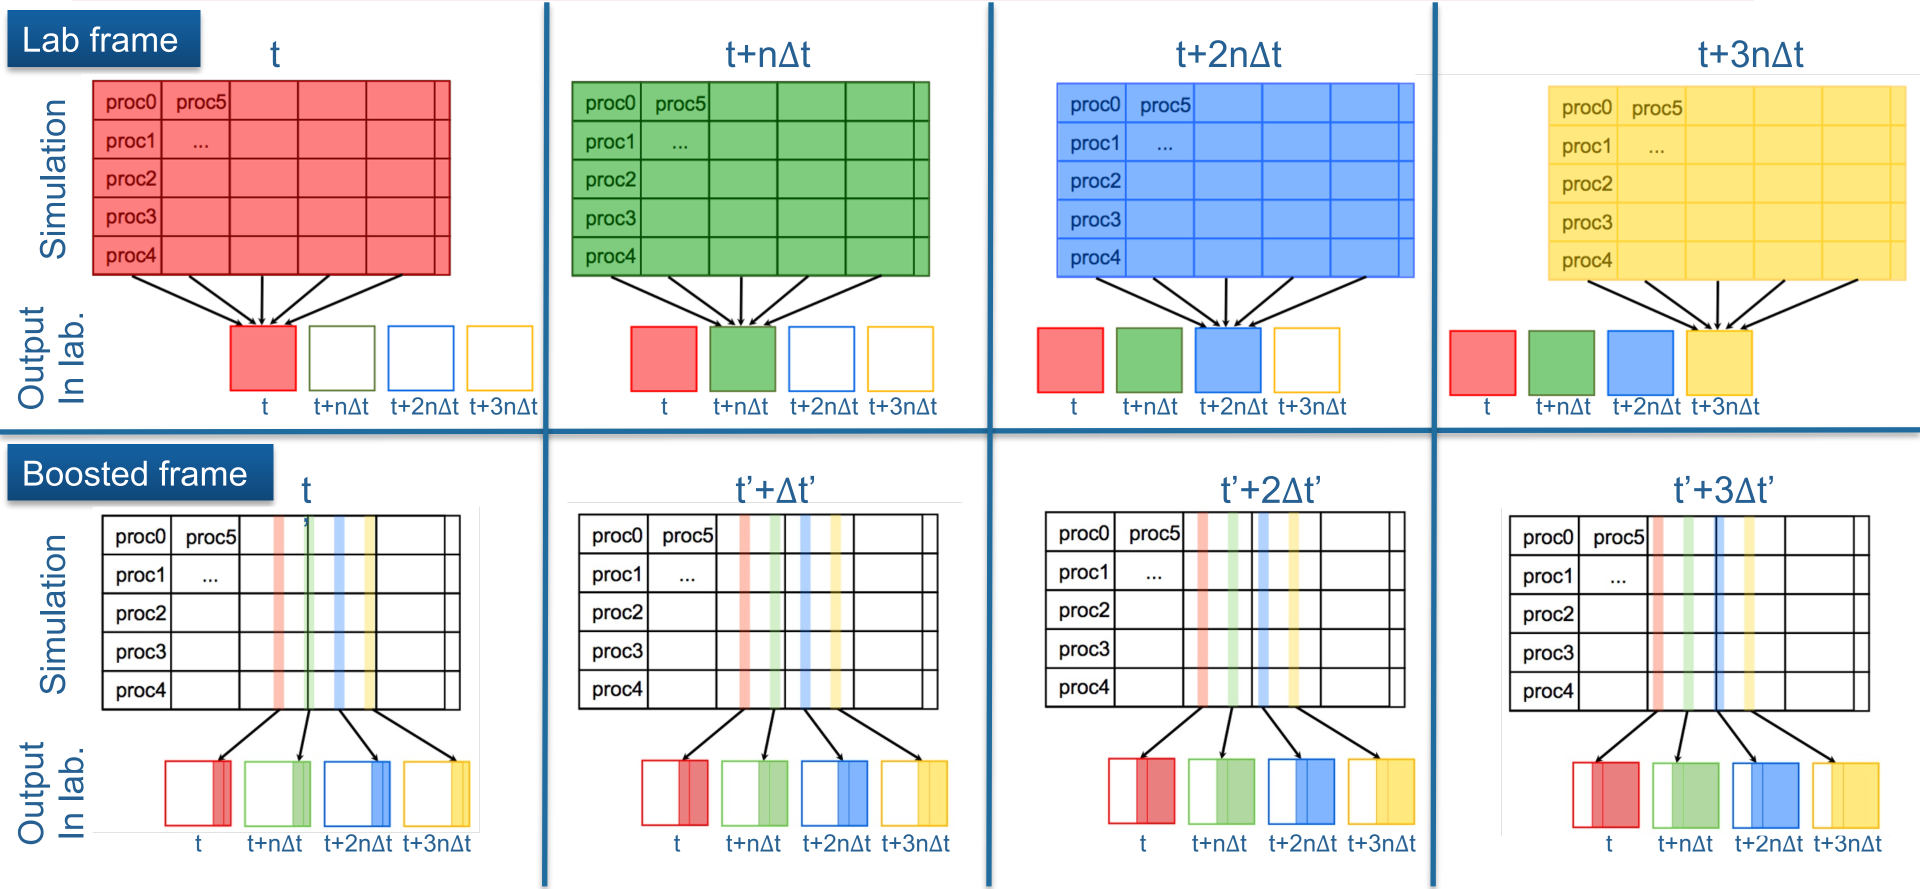

EXASCALE COMPUTING PROJECT

Advanced algorithms are needed in addition to supercomputing

Lower # time steps

- optimal Lorentz boosted frame

Higher accuracy

- (Adaptive) Mesh Refinement

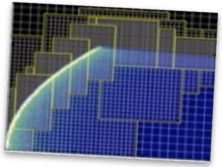

Mesh refinement requires special care

Jump of resolution can induce various side effects.

Need to avoid spurious:

1. self-forces[1]
2. wave reflections[2]
3. Numerical dispersion mismatch[2]
4. Numerical transition radiation

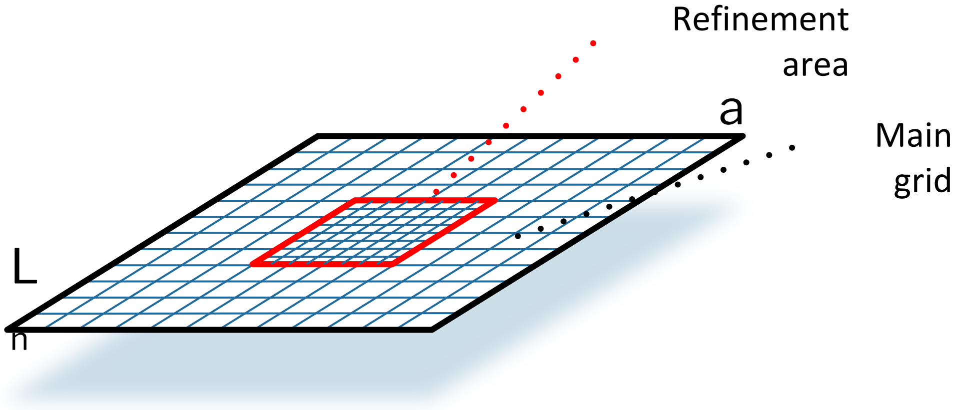

[1]J.-L. Vay, P. Colella, P. McCorquodale, B. Van Straalen, A. Friedman, D. P. Grote, *Laser & Particle Beams* **20,** 569 (2002)
[2]J.-L. Vay, J.-C. Adam, A. Héron, *Computer Physics Comm.* **164**, 171-177 (2004).

Hence mesh refinement requires special algorithm

Need to avoid spurious:

1. self-forces
2. wave reflections
3. dispersion mismatch

⇒

1. buffer regions[1]
2. multiple grids with PMLs around patches[2]
3. pseudo-spectral solvers[3]

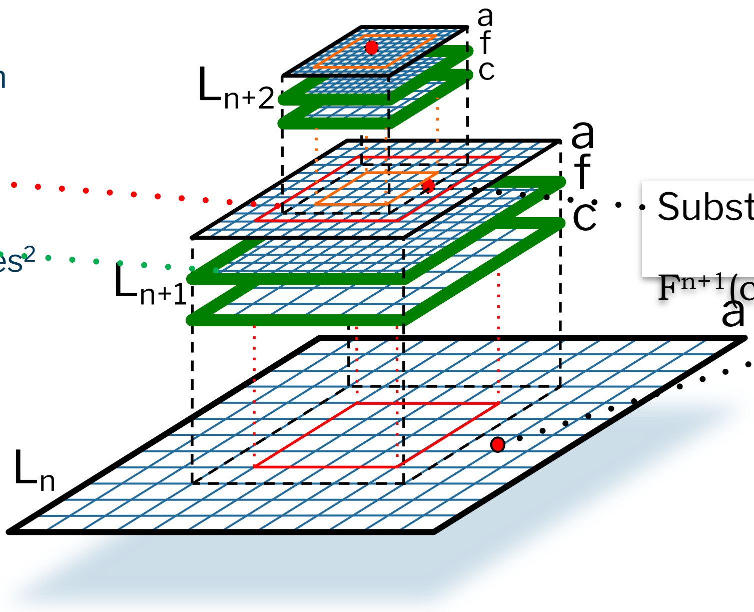

(green bar)	absorbing layer (PML)
a	auxilliary
f	fine
c	coarse

Substitution:
$F^{n+1}(a) = \mathbf{I}[F^n(a)-F^{n+1}(c)]+F^{n+1}(r)$

Main grid: $F_n(a)$

[1]J.-L. Vay, D. P. Grote, R. H. Cohen, & A. Friedman, *Computational Science & Discovery* **5**, 014019 (2012).
[2]J.-L. Vay, J.-C. Adam, A. Héron, *Computer Physics Comm.* **164**, 171-177 (2004).
[3]J.-L. Vay, I. Haber, B. B. Godfrey, *J. Comput. Phys.* **243**, 260 (2013).

Advanced algorithms are needed in addition to supercomputing

Lower # time steps

- optimal Lorentz boosted frame

Higher accuracy

- (Adaptive) Mesh Refinement
- Spectral (FFT-based) Maxwell solvers (PSATD)

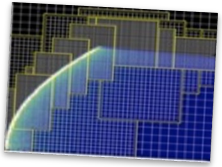

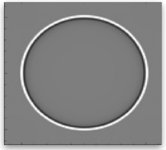

Arbitrary-order Maxwell solver offers flexibility in accuracy

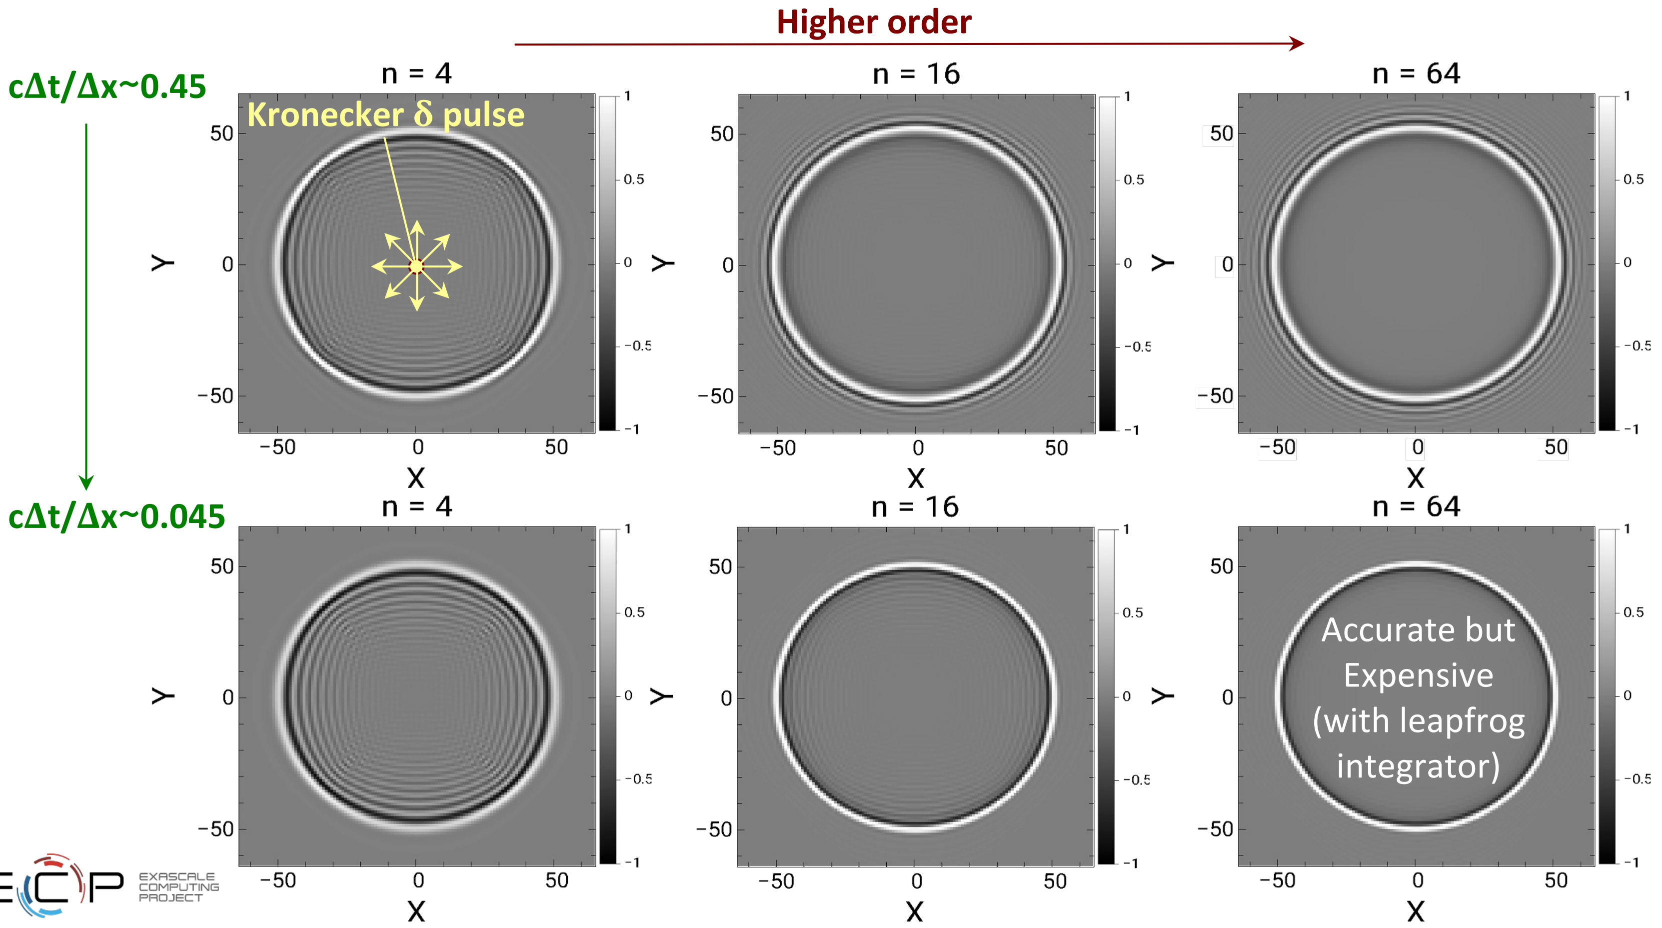

EXASCALE COMPUTING PROJECT

Analytical integration in Fourier space offers infinite order

Pseudo-Spectral Analytical Time-Domain[1] (PSATD)

$$B_z^{n+1} = F^{-1}\left(C F\left(B_z^n\right)\right) + F^{-1}\left(iSk_y F\left(E_x\right)\right) - F^{-1}\left(iSk_x F\left(E_y\right)\right)$$

with $C = \cos\left(kc\Delta t\right)$; $S = \sin\left(kc\Delta t\right)$; $k = \sqrt{k_x^2 + k_y^2}$

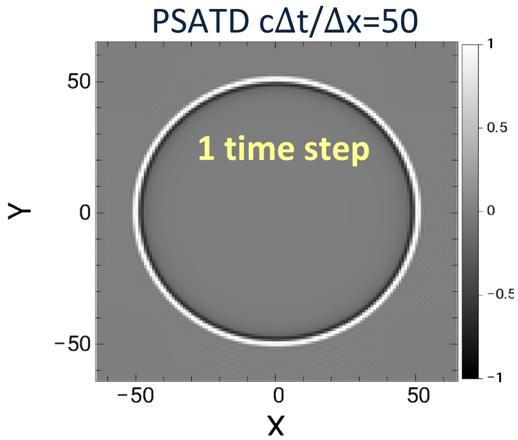

Easy to implement arbitrary-order *n* with PSATD ($k=k^{\infty}$ → k^n).

[1]I. Haber,R. Lee,H. Klein & J. Boris,*Proc. Sixth Conf. on Num. Sim. Plasma,* Berkeley, CA, 46-48 (1973)

EXASCALE COMPUTING PROJECT

Advanced algorithms are needed in addition to supercomputing

Lower # time steps

- optimal Lorentz boosted frame

Higher accuracy

- (Adaptive) Mesh Refinement
- Spectral (FFT-based) Maxwell solvers (PSATD)

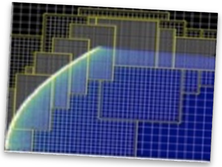

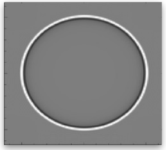

Higher stability

- Galilean PSATD suppresses Numerical Cherenkov Instability (NCI)

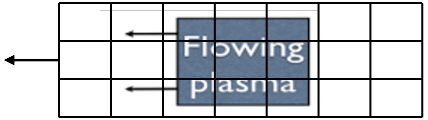

Relativistic plasma PIC subject to numerical Cherenkov instability (NCI)

B. B. Godfrey, “Numerical Cherenkov instabilities in electromagnetic particle codes”, *J. Comput. Phys.* **15** (1974)

Numerical dispersion leads to crossing of EM field and plasma modes -> instability.

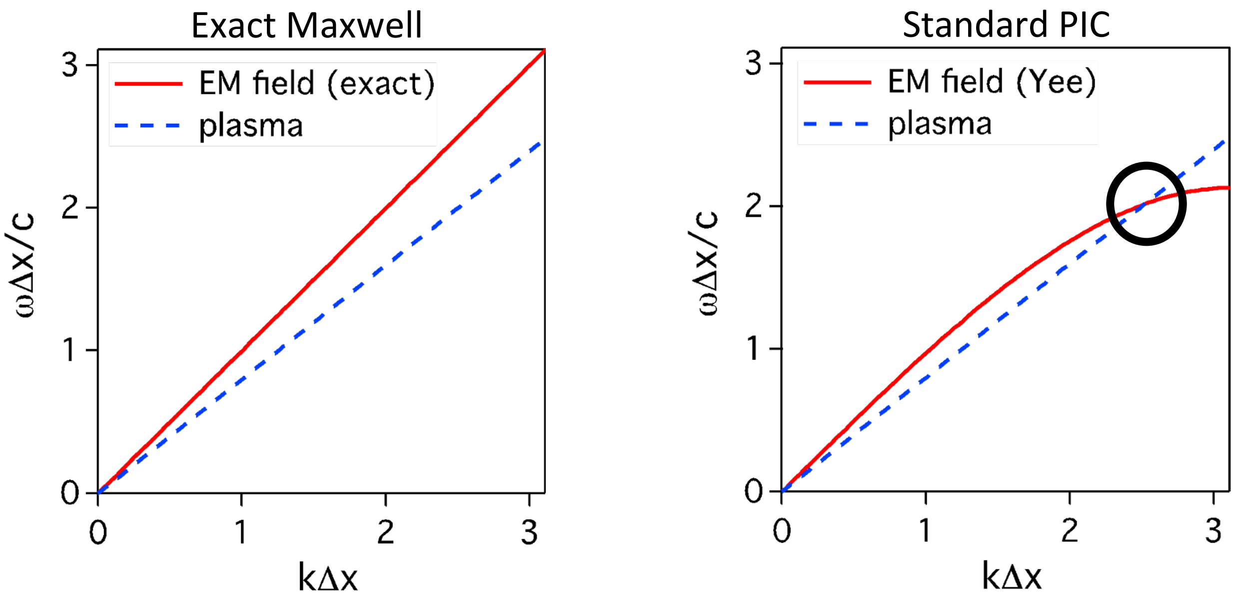

Situation slightly more complex in 2D & 3D

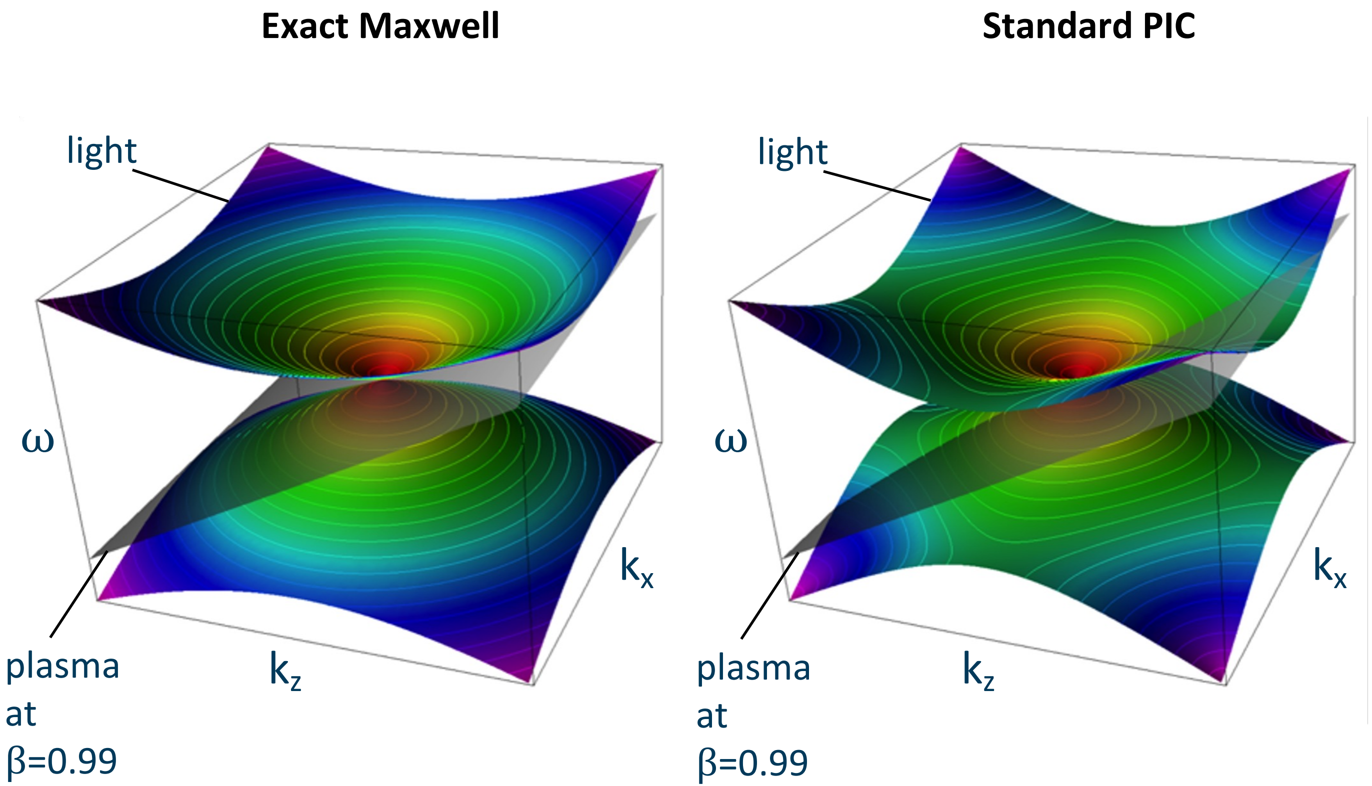

Aliases lead to more crossings in 2D & 3D

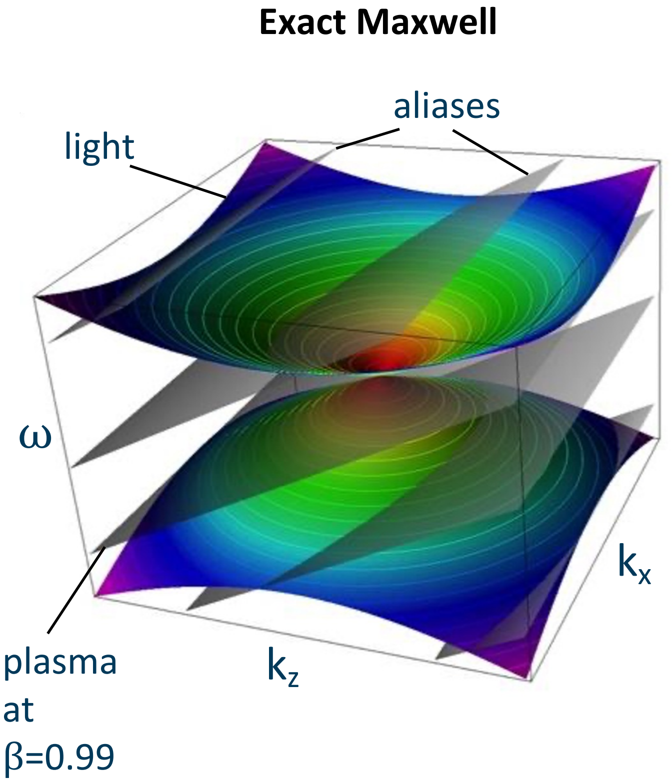

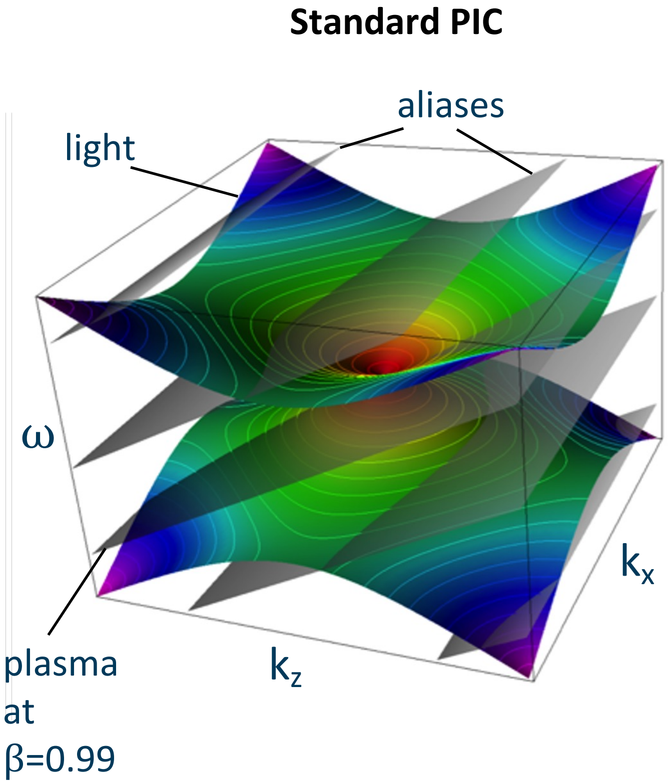

Elegant solution: use PSATD for time integration in Galilean frame

Standard PSATD PIC

Plasma moves through fixed grid.

Galilean PSATD PIC

Grid follows the relativistic plasma.

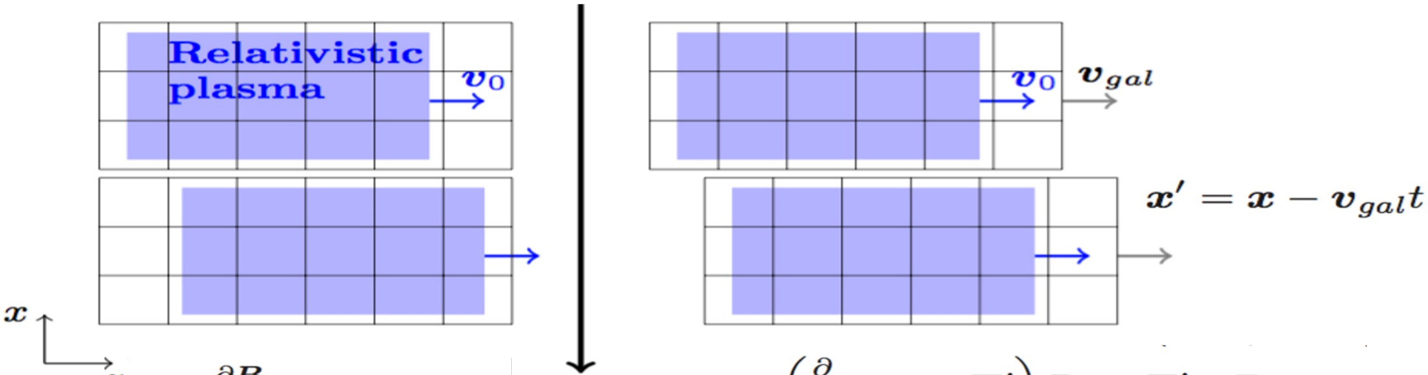

Time

$$\frac{\partial B}{\partial t} = -\nabla \times E$$

$$\frac{1}{c^2}\frac{\partial E}{\partial t} = \nabla \times B - \mu_0 j$$

$$\left(\frac{\partial}{\partial t} - v_{gal} \cdot \nabla'\right) B = -\nabla' \times E$$

$$\frac{1}{c^2}\left(\frac{\partial}{\partial t} - v_{gal} \cdot \nabla'\right) E = \nabla' \times B - \mu_0 j$$

+ integrate analytically, assuming $j(x,t)$ $j(x',t)$ is constant over one timestep.

Original idea by Manuel Kirchen (U. Hamburg)

Concept and applications: *Kirchen et al., Phys. Plasmas 23, 100704 (2016)*

Derivation of the algorithm by Rémi Lehe (Berkeley Lab):

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

EXASCALE COMPUTING PROJECT

Advanced algorithms are needed in addition to supercomputing

Lower # time steps

- optimal Lorentz boosted frame

Higher accuracy

- (Adaptive) Mesh Refinement
- Spectral (FFT-based) Maxwell solvers (PSATD)

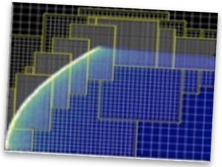

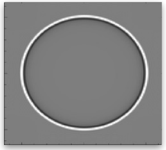

Higher stability

- Galilean PSATD suppresses Numerical Cherenkov Instability (NCI)

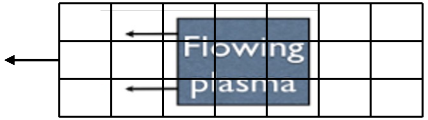

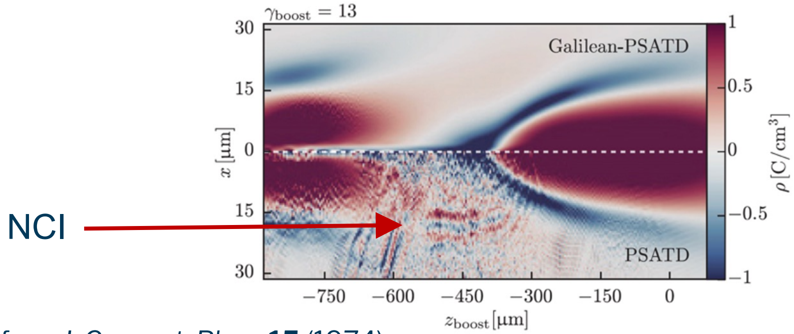

PSATD enables analytical integration of Maxwell in Galilean frame following the plasma[2,3].

[1]B. B. Godfrey, *J. Comput. Phys.* **15** (1974)
[2]R. Lehe, M. Kirchen, B. B. Godfrey, A. R. Maier, J.-L. Vay, *Phys. Rev. E* **94**, 053305 (2016).
[3]M. Kirchen, R. Lehe, B. B. Godfrey, I. Dornmair, S. Jalas, K. Peters, J.-L. Vay J.-L., A. R. Maier, *Phys. Plasmas* **23**, 100704 (2016).

Advanced algorithms are needed in addition to supercomputing

Lower # time steps

- optimal Lorentz boosted frame

Higher accuracy

- (Adaptive) Mesh Refinement
- Spectral (FFT-based) Maxwell solvers (PSATD)

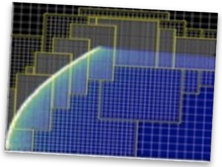

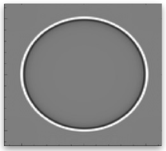

Higher stability

- Galilean PSATD suppresses Numerical Cherenkov Instability (NCI)

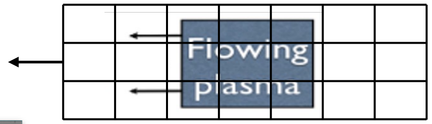

Higher scalability

- PSATD: FFT on local subdomains

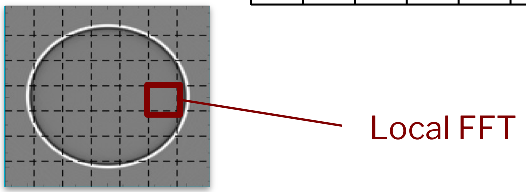

Finite speed of light/arbitray order ⇒ local FFTs ⇒ spectral accuracy + FDTD scaling!

J.-L. Vay, I. Haber, B. B. Godfrey, *J. Comput. Phys.* **243**, 260-268 (2013).
H. Vincenti, J.-L. Vay, *Comput. Phys. Comm.* **200**, 147 (2016).
Jalas, S. and Dornmair, I. and Lehe, R. and Vincenti, H. and Vay, J.-L. and Kirchen, M. and Maier, A. R., *Phys. Plasmas* **24**, 033115 (2017).
M. Kirchen, R. Lehe, S. Jalas, O. Shapoval, J.-L. Vay, and A. R. Maier, *Phys. Rev. E* **102**, 013202 (2020).

Advanced algorithms are needed in addition to supercomputing

Lower # time steps

- optimal Lorentz boosted frame

Higher accuracy

- (Adaptive) Mesh Refinement
- Spectral (FFT-based) Maxwell solvers (PSATD)

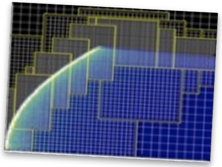

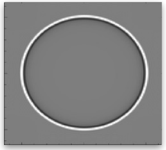

Higher stability

- Galilean PSATD suppresses Numerical Cherenkov Instability (NCI)

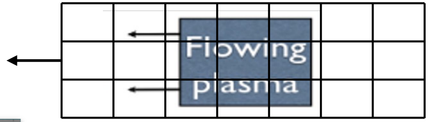

Higher scalability

- PSATD: FFT on local subdomains

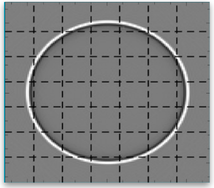

Lower dimensionality, reduced physics

- Axisymmetric solver with azimuthal Fourier decomposition
- Envelope laser solvers

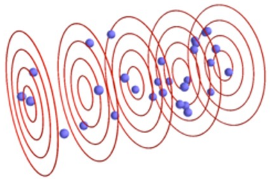

Exascale

Preparing WarpX for the World's Largest Supercomputers

WarpX is a GPU-Accelerated PIC Code for Exascale

BLAST
BEAM PLASMA & ACCELERATOR SIMULATION TOOLKIT

Available Particle-in-Cell Loops

- electrostatic & electromagnetic (fully kinetic)

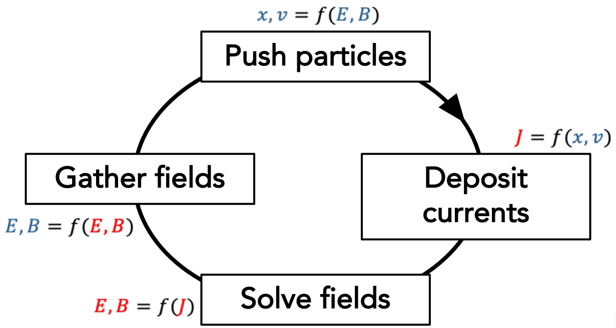

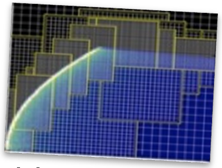

Advanced algorithms

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)

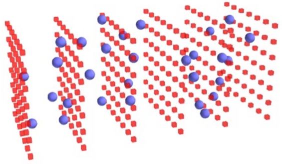

3D Cartesian grid

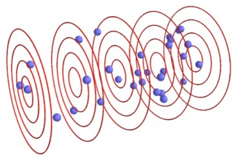

Cylindrical grid (schematic)

Multi-Node parallelization

- MPI: 3D domain decomposition
- dynamic load balancing

On-Node Parallelization

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

mac OS

Scalable, Standardized I/O

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

ECP EXASCALE COMPUTING PROJECT

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

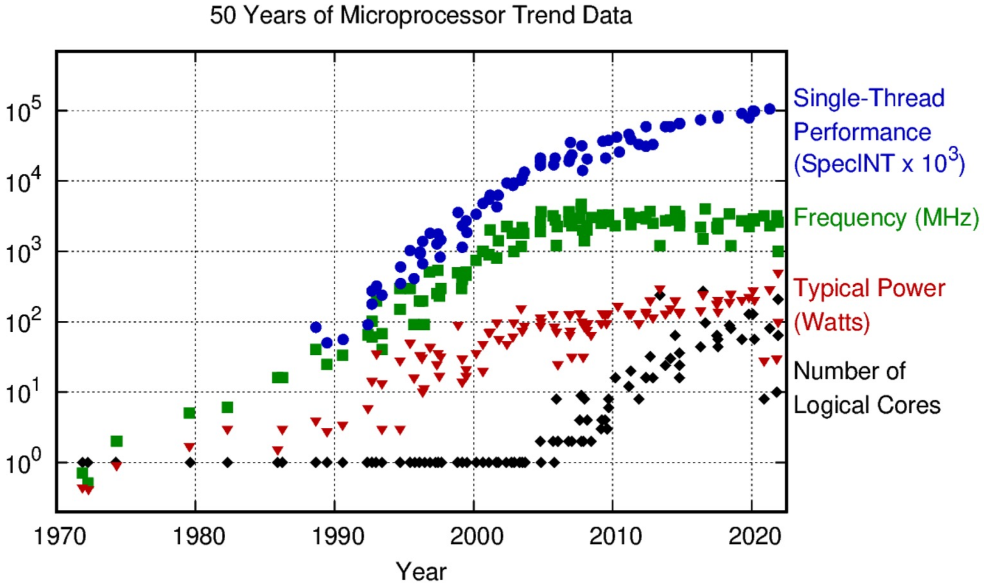

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

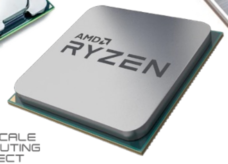

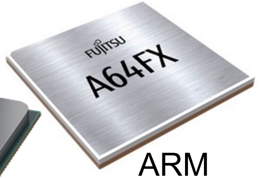

ARM

Power

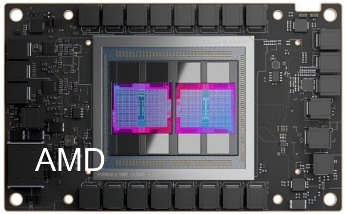

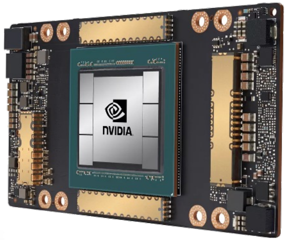

EXASCALE COMPUTING PROJECT

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

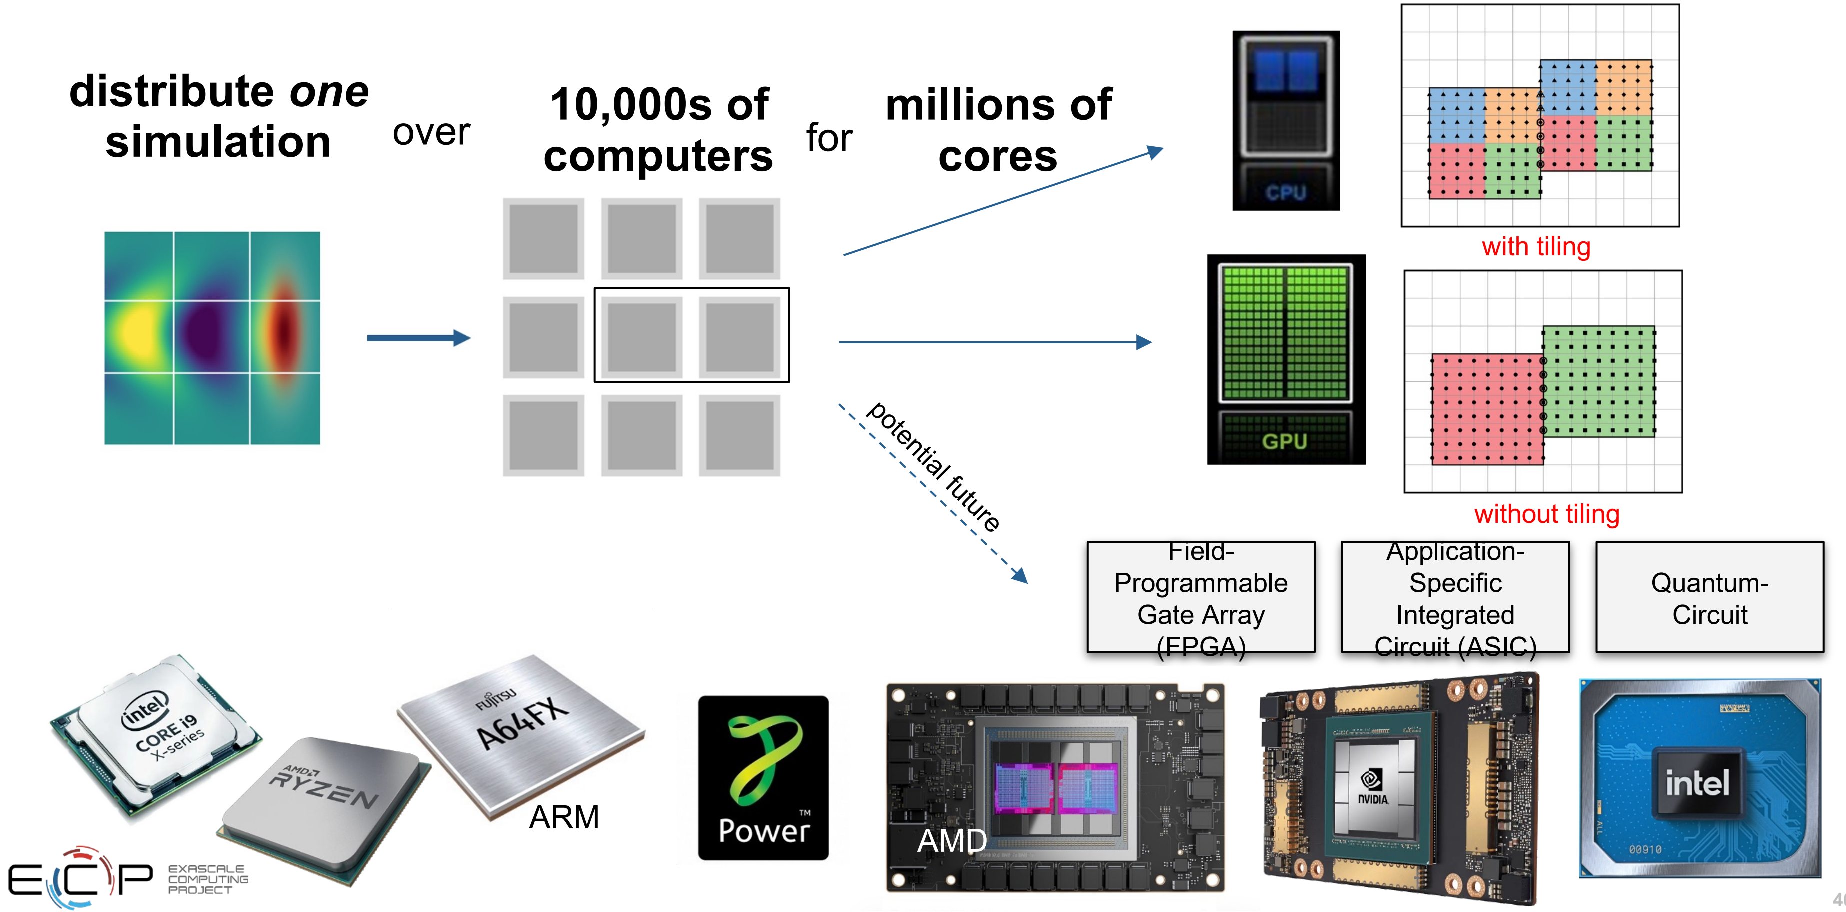

Software Stacks, Standardization & Reuse Opportunities

Applications

- Scripting & Language Bindings (PICMI Standard)
- **Applications & Physics Modules**

Libraries (openPMD)

PIC Algorithms

Communication

Performance Portability

- **I/O** | **Math** | **Containers and Algorithms**
- In-Node Acceleration | Message-Passing

H.C. Edwards, C.R. Trott et al, JPDC (2014); B. Worpitz, MA (2015); E. Zenker, A. Huebl et al., IPDPSW (2016)
E. Zenker, A. Huebl et al., IWOPH (2017); A. Matthes, A. Huebl et al., P3MA (2017); W. Zhang et al, JOSS (2019)
S. Slattery, S.T. Reeve et al., JOSS (2022)

Programming Models

OpenACC (Directives for Accelerators) — OpenMP — AMD ROCm / HIP — NVIDIA CUDA — SYCL

Hardware

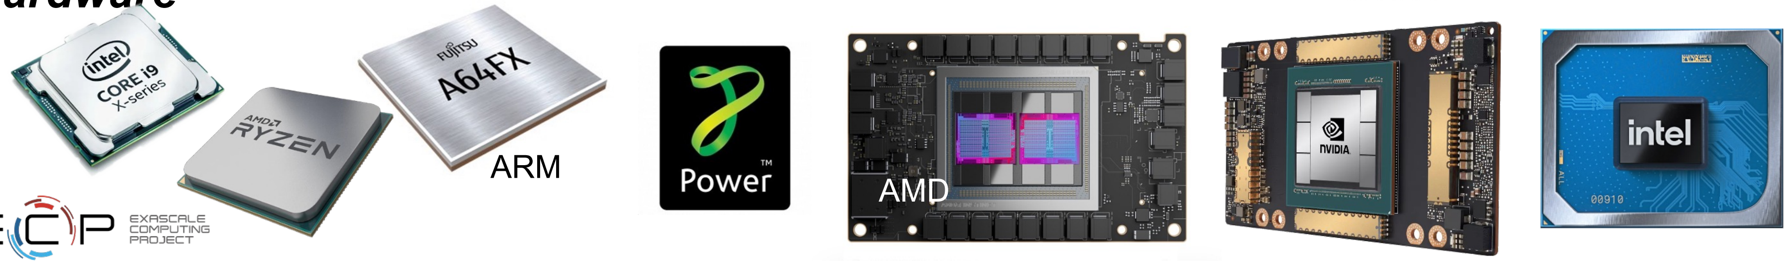

EXASCALE COMPUTING PROJECT

Portable Performance through Exascale Programming Model

AMReX library

Performance-Portability Layer in C++17 GPU/CPU/KNL

- **algorithms** and
- **data structures** for block-structured mesh-refinement: fields & particles

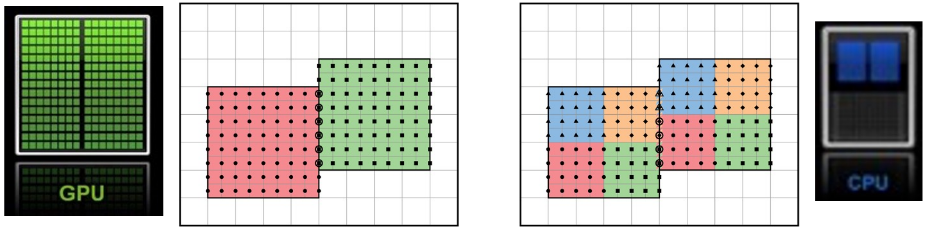

without tiling

with tiling

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 );

    });
```

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

Portable Performance through Exascale Programming Model

AMReX library

Domain decomposition & MPI
Communications:

- domain decomposition
- boundary updates, particle moves, load balancing

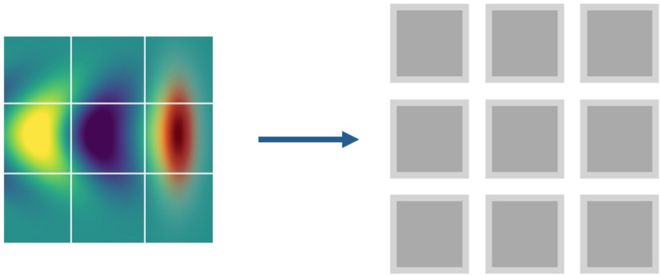

Parallel linear solvers
e.g., multi-grid Poisson solvers

Embedde

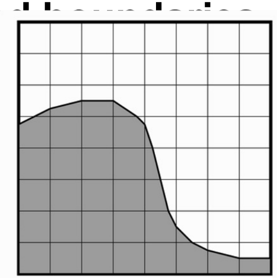

Runtime parser for user-provided math expressions (incl. GPU)

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

GPU Computing at Scale Requires Advanced Load Balancing

Application Challenges

- Plasma Mirrors & Laser-Ion Acceleration: moving fron
- Laser Wakefield Accelerator: Injected Beam Particles

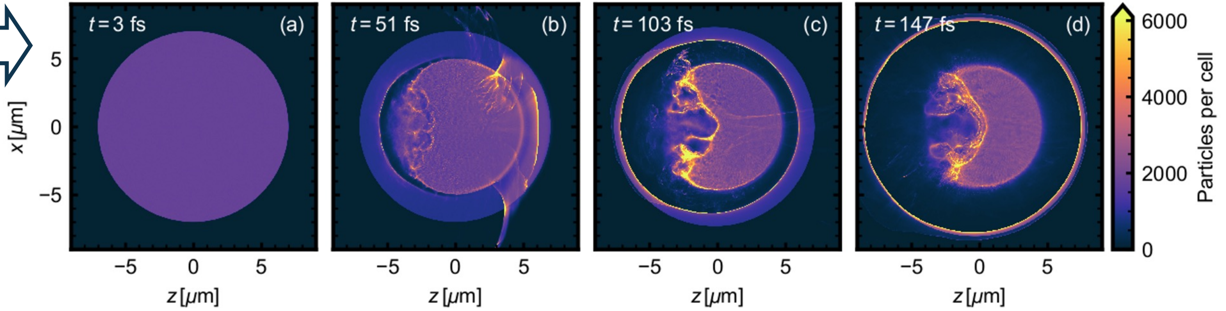

domain decomposition example:

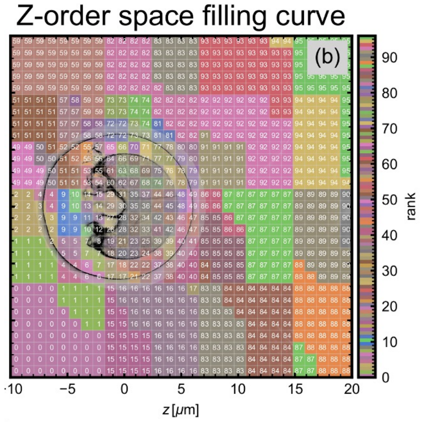

In Situ Cost Analysis

- basis for distribution functions
- realistic cost: kernel timing

Result: 3.8x speedup!

- production-quality, easy-to-use
- larger simulation: mitigate local memory spikes

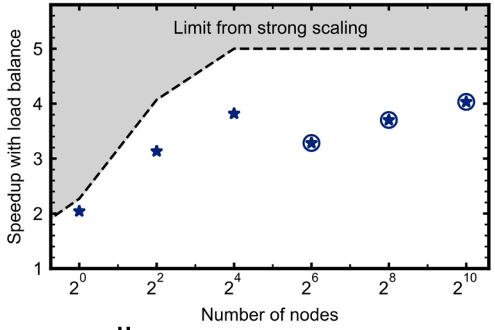

M. Rowan, A. Huebl, K. Gott, R. Lehe, M. Thévenet, J. Deslippe, J.-L. Vay, "In-Situ Assessment of Device-Side Compute Work for Dynamic Load Balancing in a GPU-Accelerated PIC Code," PASC21, DOI:10.1145/3468267.3470614 (2021)

EXASCALE COMPUTING PROJECT

Open Science

In HPC, we collaborate across domains and work with many specialists

Common Data Challenges in HPC

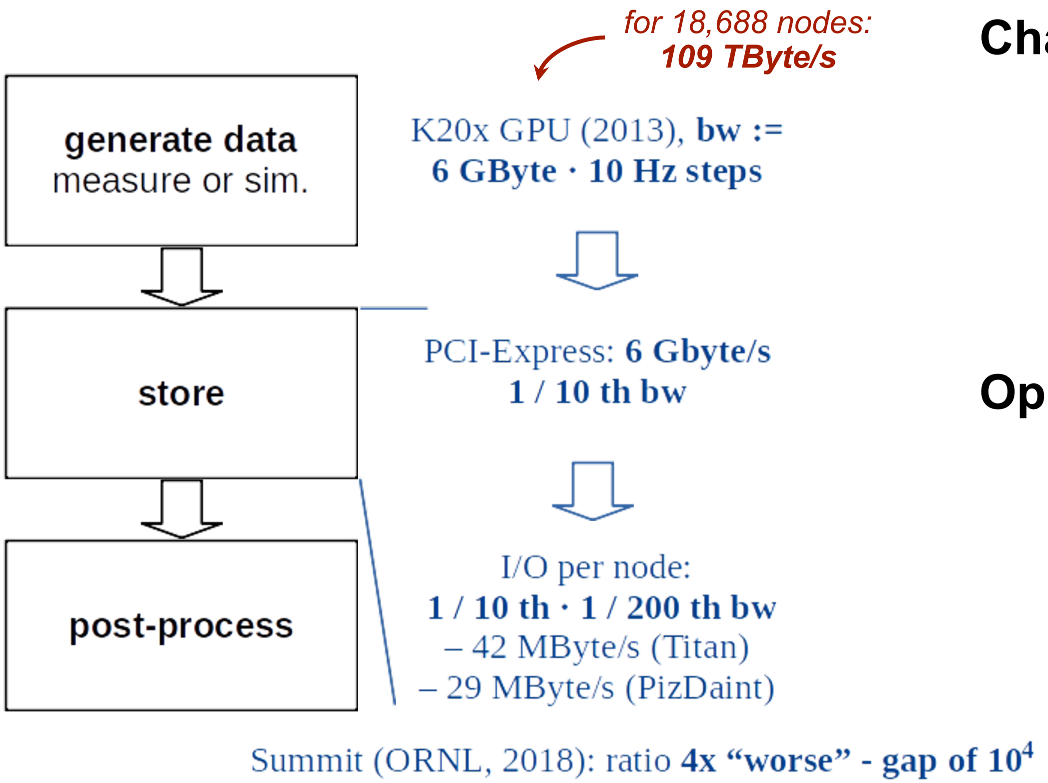

A Huebl et al., ISC 2017, DOI:10.1007/978-3-319-67630-2_2 (2017)

Challenges

- **3 orders of magnitude gap** between producing devices and storage
- "store & analyze everything" is *unaffordable*

Opportunities

- analysis tasks have varying **fidelity** needs
- many common tasks can be done **in situ**
- manual steps: limit the sampling of raw data to **setup phase**

Data Processing & Reduction Examples

Binning of a **spectrogram**
Fitting of an **ellipsoid**

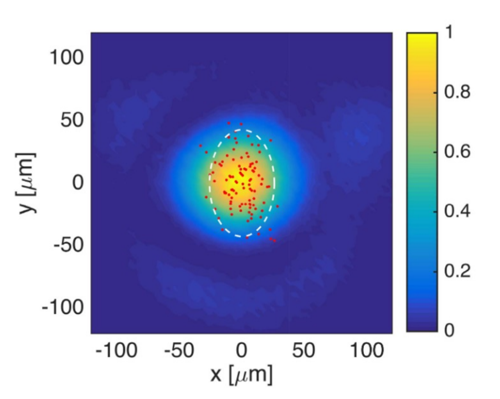

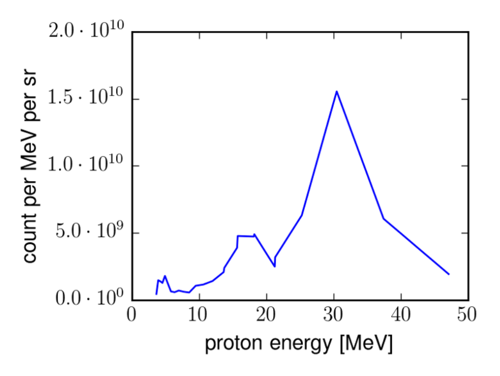

Compression (lossless/lossy)

Ray-casting 3D data,
training a neural network, etc.

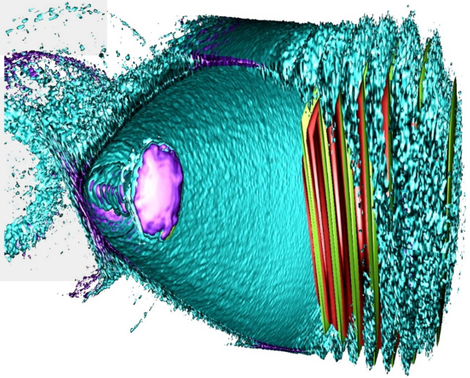

A Matthes, A Huebl et al., ISC 2017, DOI:10.14529/jsfi160403 (2017);
A Huebl et al., ISC 2017, DOI:10.1007/978-3-319-67630-2_2 (2017);

K Nakamura et al., IEEE J. Quantum Electron, DOI:10.1109/JQE.2017.2708601 (2019)

Avoid Backlog: Design Criteria for Data Reduction Pipelines

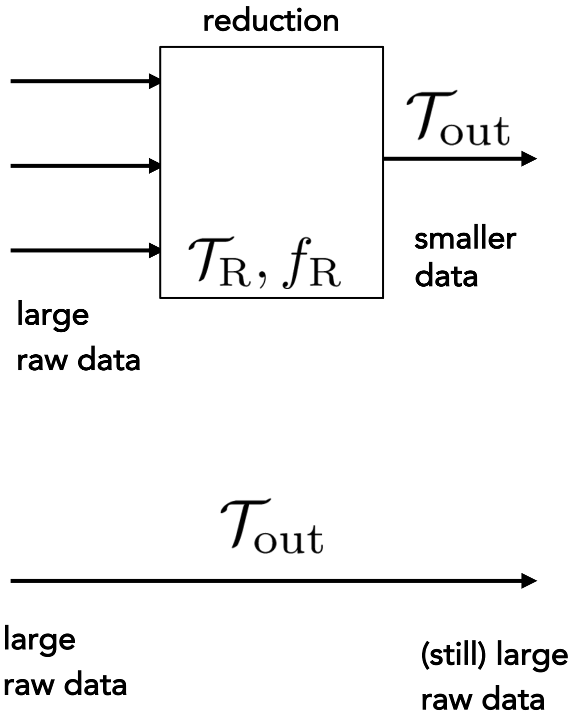

$$\frac{\mathcal{T}_R \times (1 - f_R)}{1 - \mathcal{T}_R} > \mathcal{T}_{out}$$

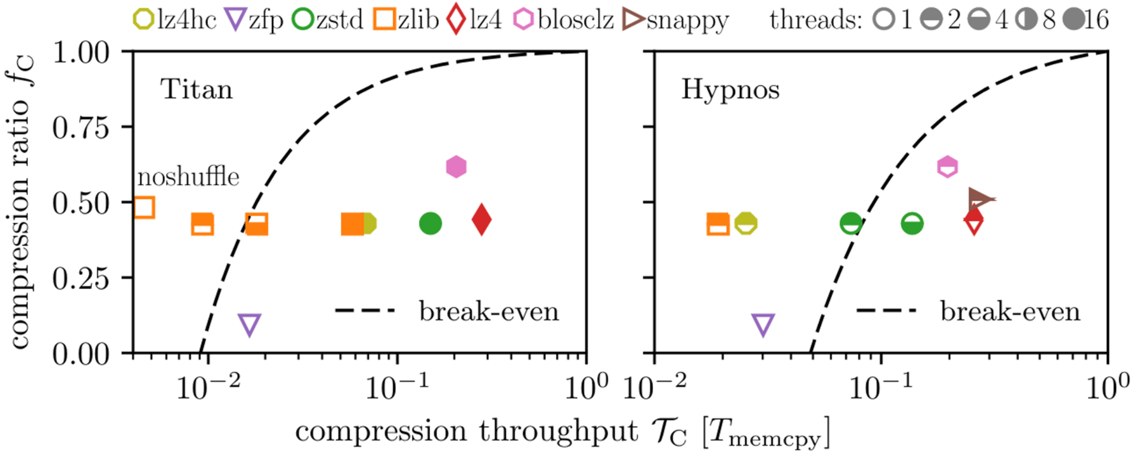

Result: **trade compute for throughput**, >100 GByte perceived application throughput and **280 GByte/s** peak parallel filesystem throughput

A Huebl et al., "On the Scalability of Data Reduction Techniques in Current and Upcoming HPC Systems from an Application Perspective," ISC High Performance Workshops, DOI:10.1007/978-3-319-67630-2_2 (2017)

Reduce Particle I/O with Novel In Situ Visualization

Particle *Adaptive Sampling*

- **emphasis** on "uncommon" properties
- inverse sampling to incidence of a property

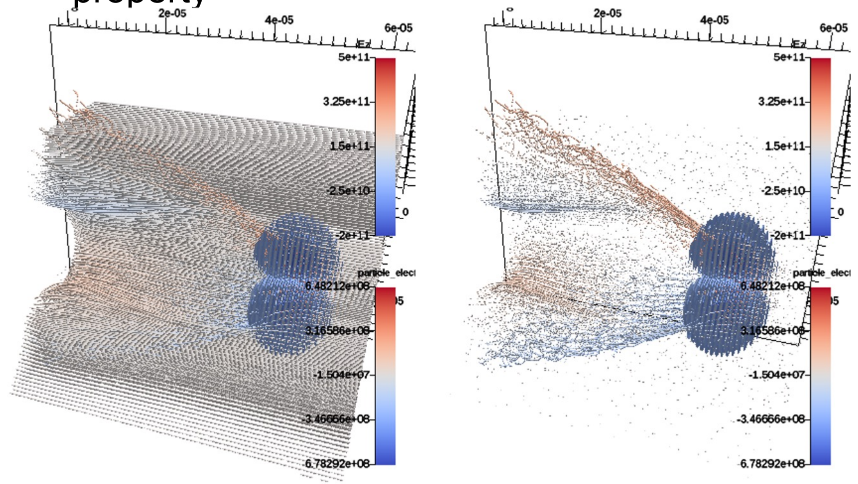

A. Biswas et al., "In Situ Data-Driven Adaptive Sampling for Large-scale Simulation Data Summarization," ISAV18 @SC18 (2018)

Physics-Informed Flow Tracelines

- traditional flow vis. depends only on *local field values*
- plasma particles:
  - **inert**: track *relativistic momentum* on a traceline
  - **Lorentz-Force**: 6 fields (electromag.), leap-frog
- chance to **significantly reduce particle I/O** in real-life workflows through savings on **temporal fidelity**

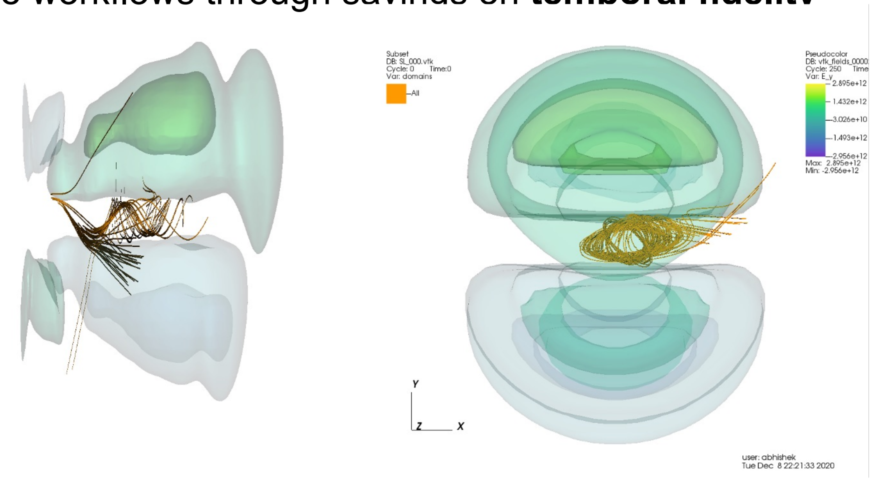

Biswas, Larsen, Lo

Yenpure, Childs, Pugmire

openPMD: Share Data *and* Cutting-Edge Optimizations

Application Challenges

- R&D in: scalable techniques, data layouts, libraries
- scientific data analysis & sharing

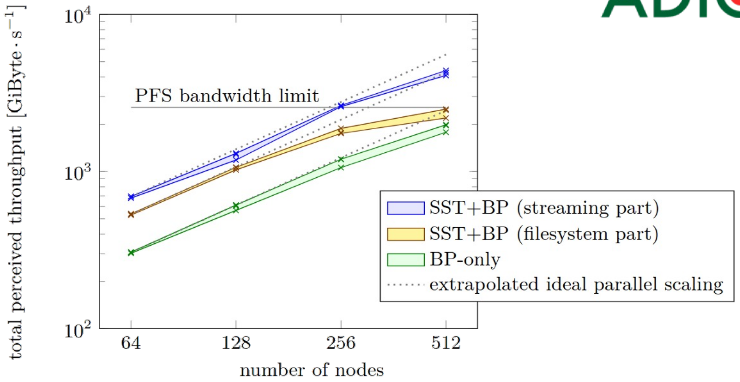

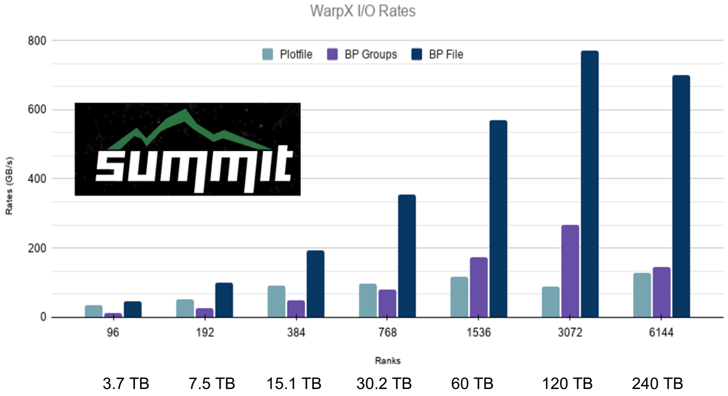

Write: plotfiles → ADIOS BP per rank & step → append to files

Streaming Data Pipelines: arXiv:2107.06108
by F Poeschel, A Huebl et al., SMC21 (2021)

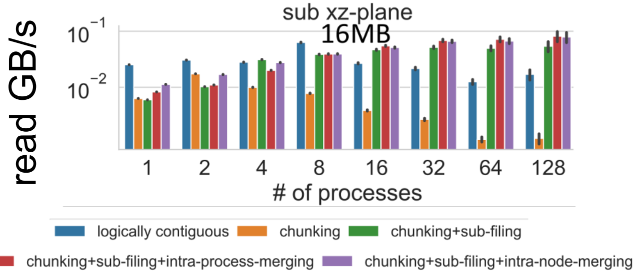

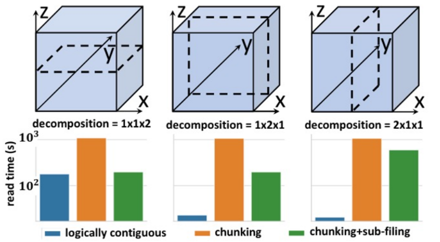

Impact of decomposition schemes when reading

Online Data Layout Reorganization:
DOI:10.1109/TPDS.2021.3100784
by L Wan, A Huebl et al., TPDS (2021)

We are Establishing an Open Community Ecosystem with Standards

In accelerator modeling, we use **specialized codes** for different science questions.
Code usage and **data exchange** *must* become easier to be productive.

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

class **picmistandard.PICMI_Simulation**(*solver=None, time_step_size=None, max_steps=None, max_time=None, verbose=None, particle_shape='linear', gamma_boost=None, cpu_split=None, load_balancing=None, **kw*) [source]

Creates a Simulation object

Parameters

- **solver** (*object*) – An instance of one of the PICMI field solvers ; see Field solvers This is the field solver to be used in the simulation
- **time_step_size** (*float*) – Absolute time step size of the simulation [s] (needed if the CFL is not specified elsewhere)
- **max_steps** (*int*) – Maximum number of time steps (Specify either this, or *max_time*, or use the *step* function directly)
- **max_time** (*float*) – Maximum physical time to run the simulation [s] (Specify either this, or *max_steps*, or use the *step* function directly)
- **verbose** (*int*) – Verbosity flag (A larger integer results in more verbose output.)

- **markup** / schema for arbitrary hierarchical data formats
  ADIOS HDF {JSON} XML
- scientifically **self-describing**
- basis for **open data workflows**

Organizing Scientific Data Records

Example in the Open Particle Mesh Data Standard

- **electric field** $\vec{E}(\vec{r})$

  / … / **meshes** / **E** /

  x **y** **z**

- **temperature** $T(\vec{r})$

  / … / **meshes** /

  T

- **electron position** $\vec{r}$

  / … / **particles** / **electrons** / **position** /

  x **y** **z**

group

data

attributes

extensions to the **base standard** define additional, compatible, **domain-specific** conventions

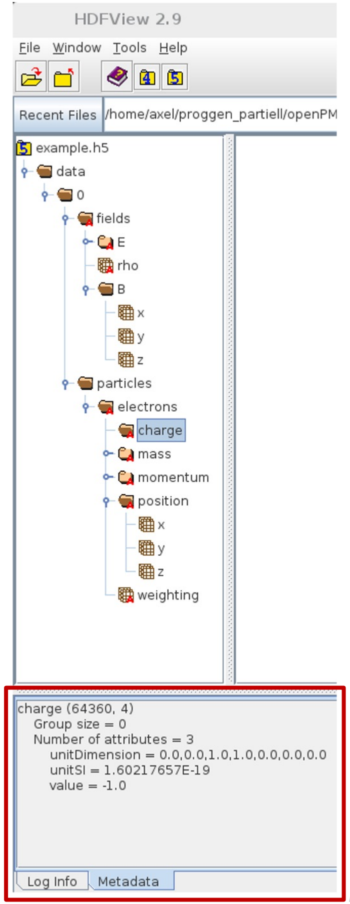

We Develop Openly with the Community

BLAST
BEAM PLASMA & ACCELERATOR SIMULATION TOOLKIT

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

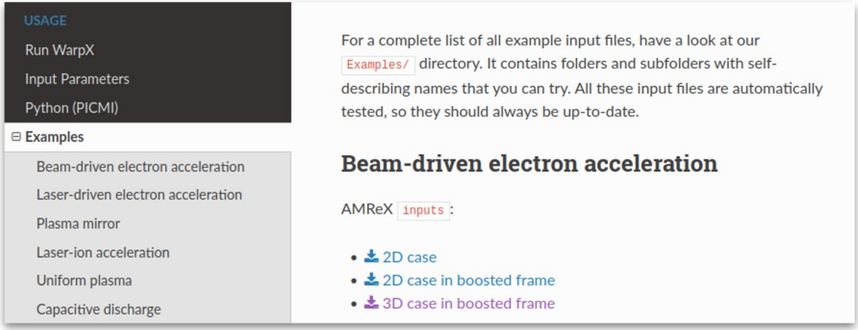

Open-Source Development & Benchmarks:
github.com/ECP-WarpX

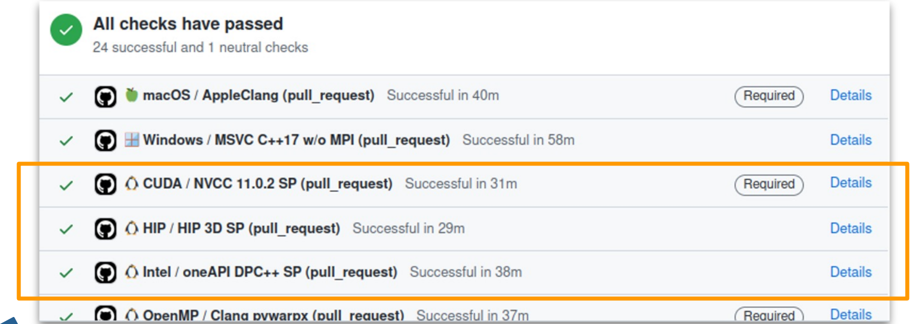

188 physics benchmarks *run on every code change* of WarpX
13 physics benchmarks + 32 tests *for* ImpactX

Rapid and easy installation on any platform:

conda install
-c conda-forge warpx

spack install warpx
spack install py-warpx

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

python3 -m pip install .

brew tap ecp-warpx/warpx
brew install warpx

module load warpx
module load py-warpx

Our Science Case in the 2022 ACM Gordon Bell Prize

Context: radiation therapy techniques for medical treatments

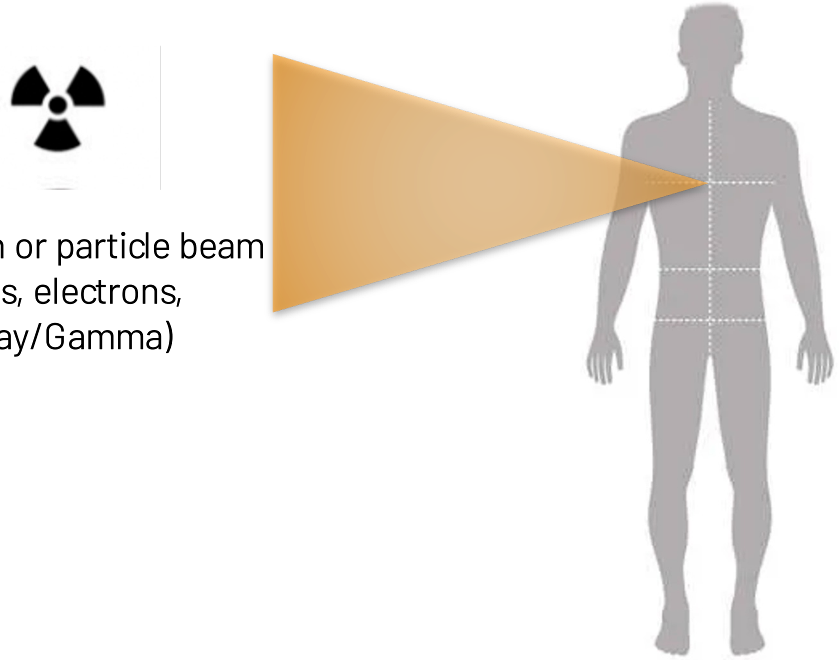

Radiation or particle beam
(ions, electrons,
X-ray/Gamma)

Contributes to 40% of curative treatments for cancer

R. Baskar, *Int. J. Med. Sci.*, 2012

Target tumor cells while sparing healthy cells

Towards a revolution in medical treatments: ultra-high dose rate radiotherapy (FLASH)

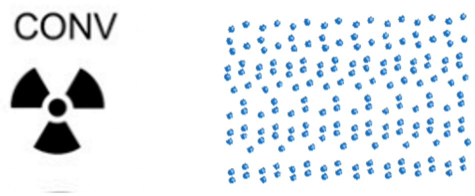

>8s <0.1Gy/s

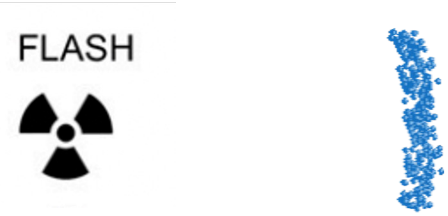

<200ms >40 Gy/s

FLASH-RT significantly reduces radiation toxicity to healthy tissues

Favaudon *et al*, ***Science. Trans. Med.***, 2014

Bourhis *et al*, ***Radiotherapy and Oncology***, 2019

Towards a revolution in medical treatments: ultra-high dose rate radiotherapy (FLASH)

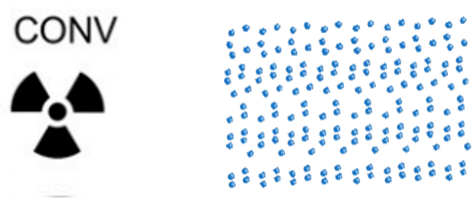

>8s <0.1Gy/s

<200ms >40 Gy/s

FLASH-RT significantly reduces radiation toxicity to healthy tissues

The challenge with FLASH-RT

We need fundamentally new type of particle accelerator technology

- **Ultra-short** - understand & optimize FLASH effect
- **Ultra-compact** - democratize access to treatments

Laser-based sources are very promising candidates for FLASH and beyond....

10^{14}W to 10^{16}W

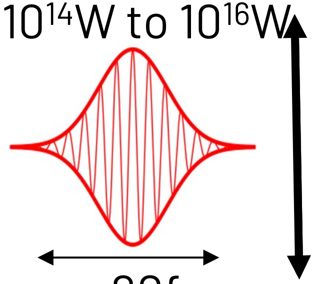

τ=20fs

(1fs=10^{-15}s)

High power CPA laser

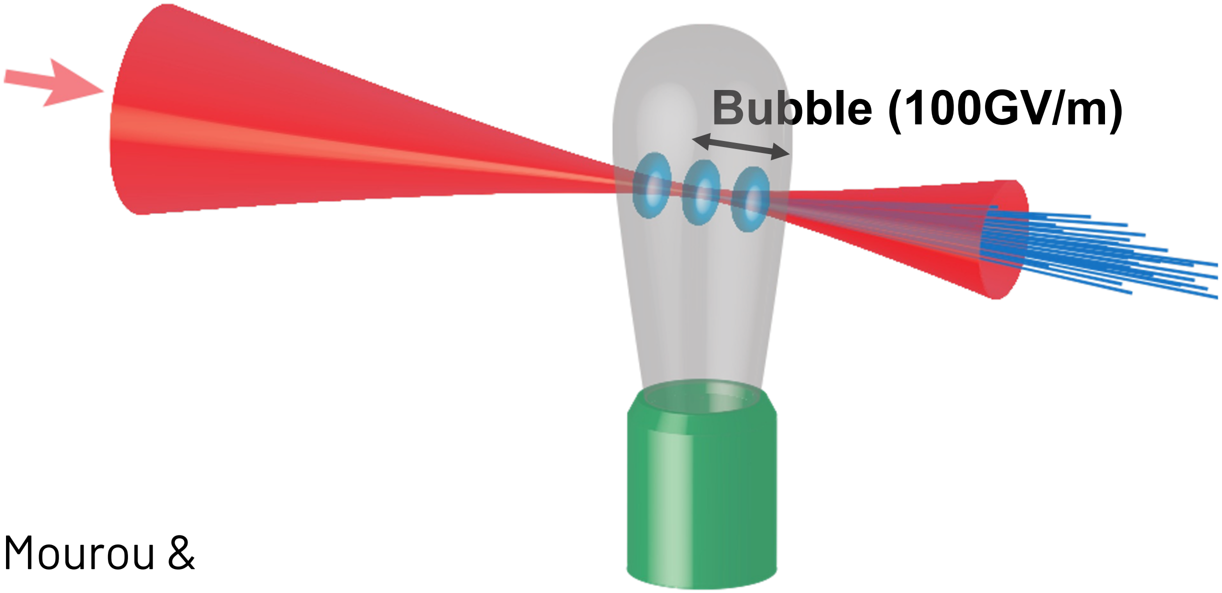

G. Mourou &
D. Strickland
(2018)

Laser-based sources are very promising candidates for FLASH and beyond....

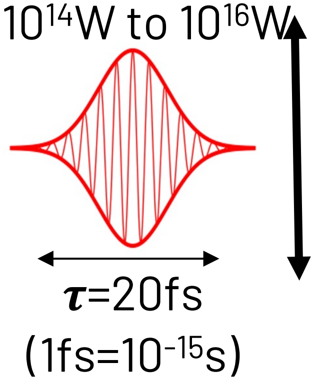

High power CPA laser

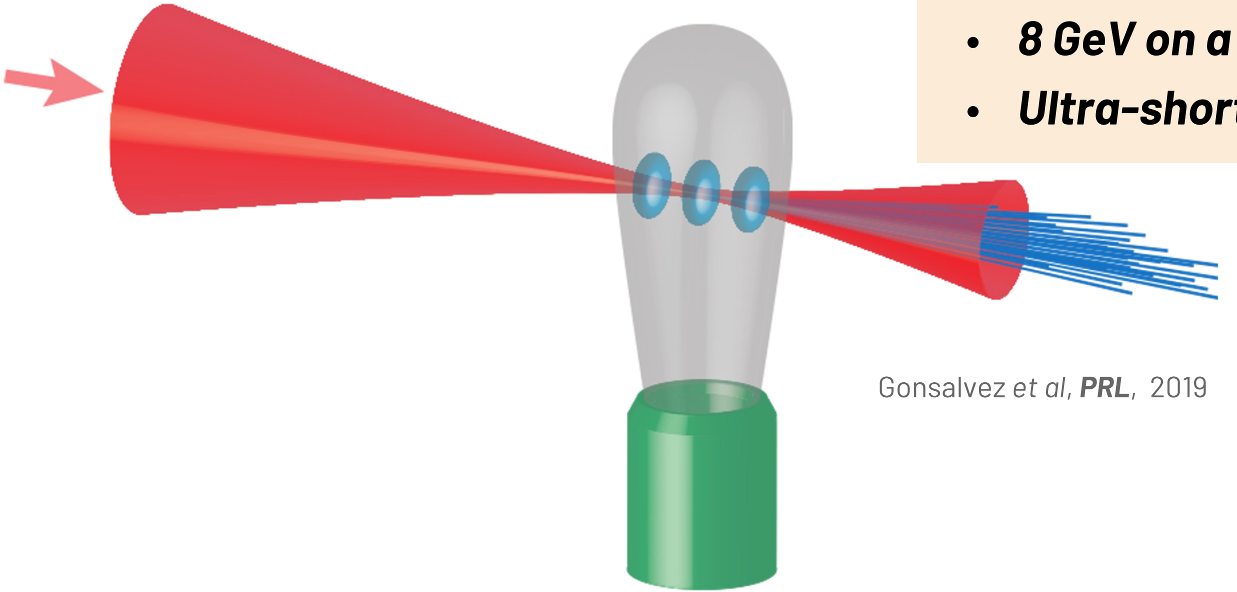

Ultra-compact accelerators

- ***8 GeV on a cm-scale***
- ***Ultra-short (<10 fs)***

Gonsalvez *et al*, ***PRL***, 2019

… but, we need to solve a major limitation of these accelerators

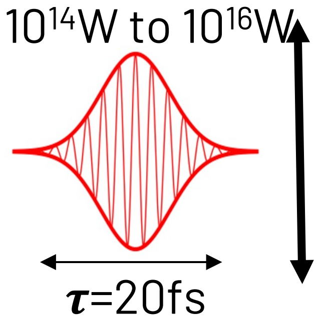

(1fs=10^{-15}s)

High power CPA laser

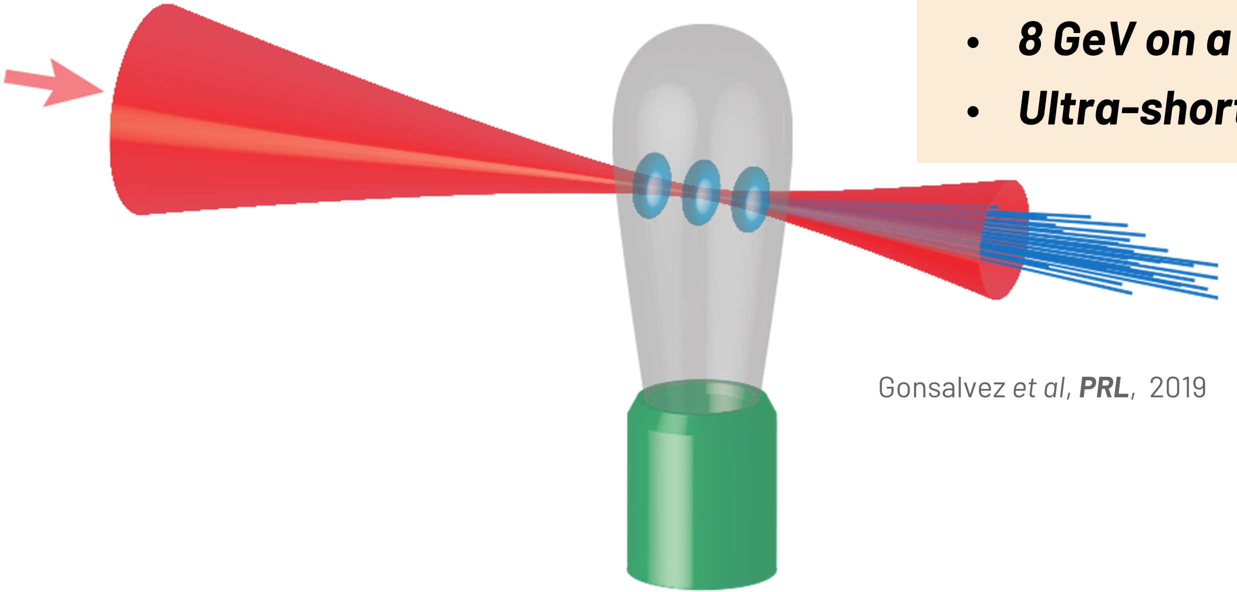

Ultra-compact accelerators

- ***8 GeV on a cm-scale***
- ***Ultra-short (<10 fs)***

Gonsalvez *et al*, ***PRL***, 2019

Major limitation: charge too low at high energy (tens of pC/bunch)

… but, we need to solve a major limitation of these accelerators

10^{14}W to 10^{16}W

τ=20fs ($1fs=10^{-15}s$)

High power CPA laser (λ=0.8µm)

Ultra-compact accelerators

- *10GeV on a cm-scale*
- *Ultra-short (<10 fs)*

The physical challenge

How to level up the charge up to a nC/bunch?

Major limitation: charge too low at high energy (tens of pC/bunch)

A new concept: the hybrid solid-gas target

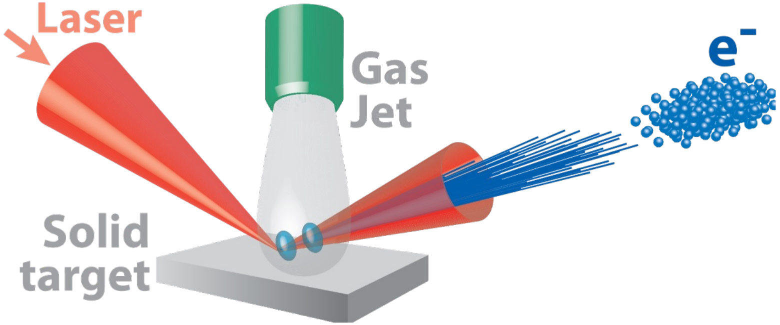

A. Leblanc, L. Fedeli, A. Huebl, J.-L. Vay ... and H. Vincenti, *in preparation*

A new concept: the hybrid solid-gas target

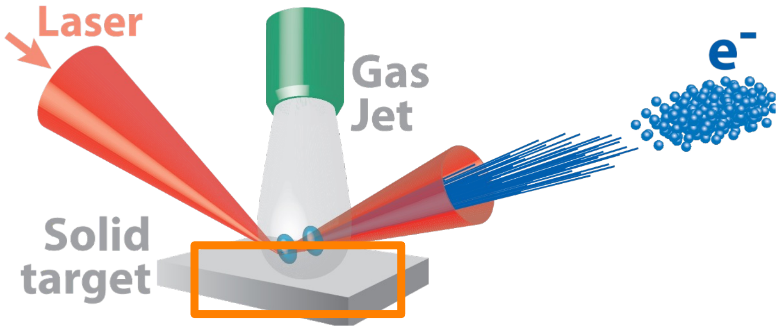

>nC/bunch at high energy!

High resolution (26 nm)

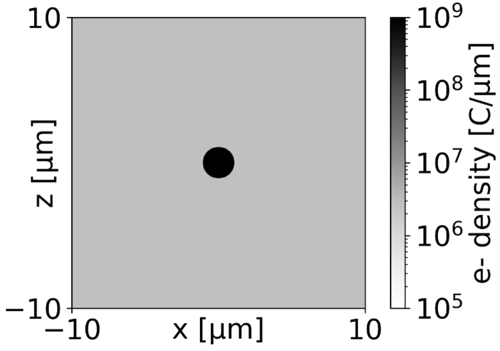

Low resolution (39 nm)

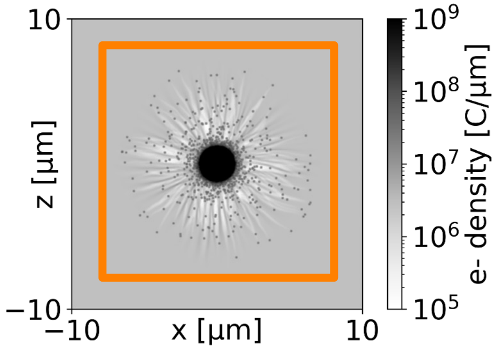

A. Leblanc, L. Fedeli, A. Huebl, J.-L. Vay ... and H. Vincenti, *in preparation*

2022 ACM Gordon Bell Prize: using the First Exascale Supercomputer

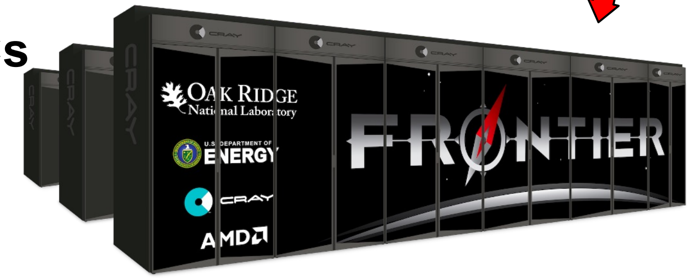

April-July 2022: WarpX on **world's largest HPCs**

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

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

2022 ACM Gordon Bell Prize: using the First Exascale Supercomputer

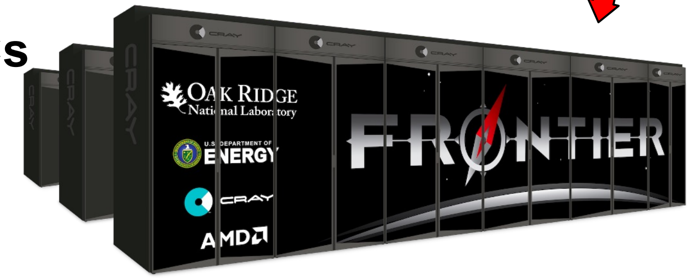

April-July 2022: WarpX on **world's largest HPCs**

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

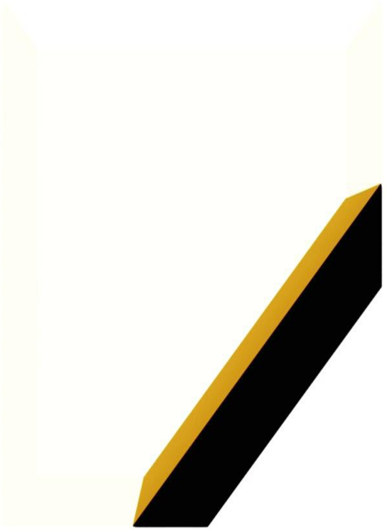

A success story of a multidisciplinary, multi-institutional team!

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

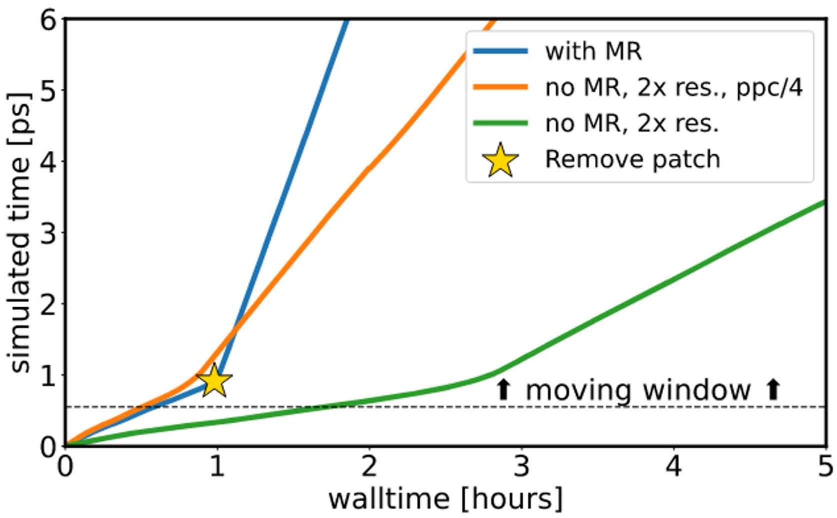

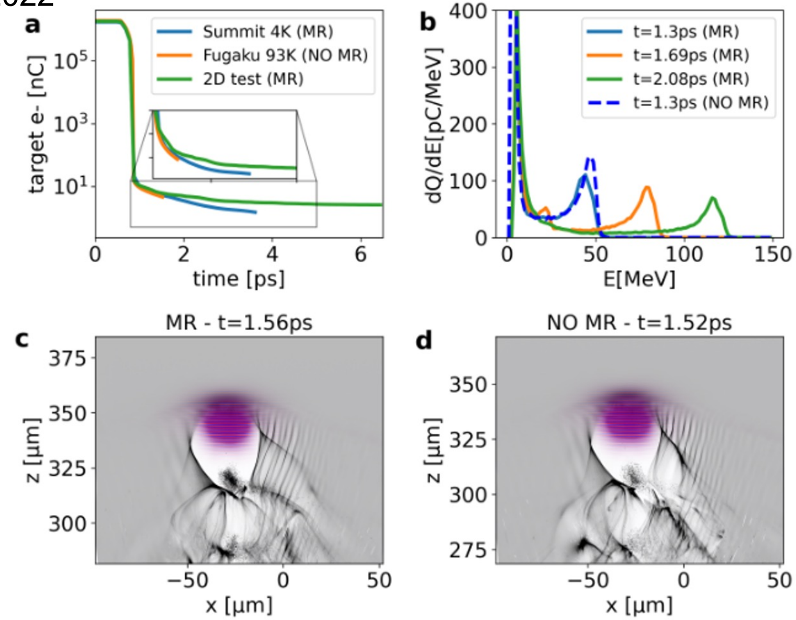

A success story of a multidisciplinary, multi-institutional team!

2D slices of our 3D simulations highlight the acceleration process

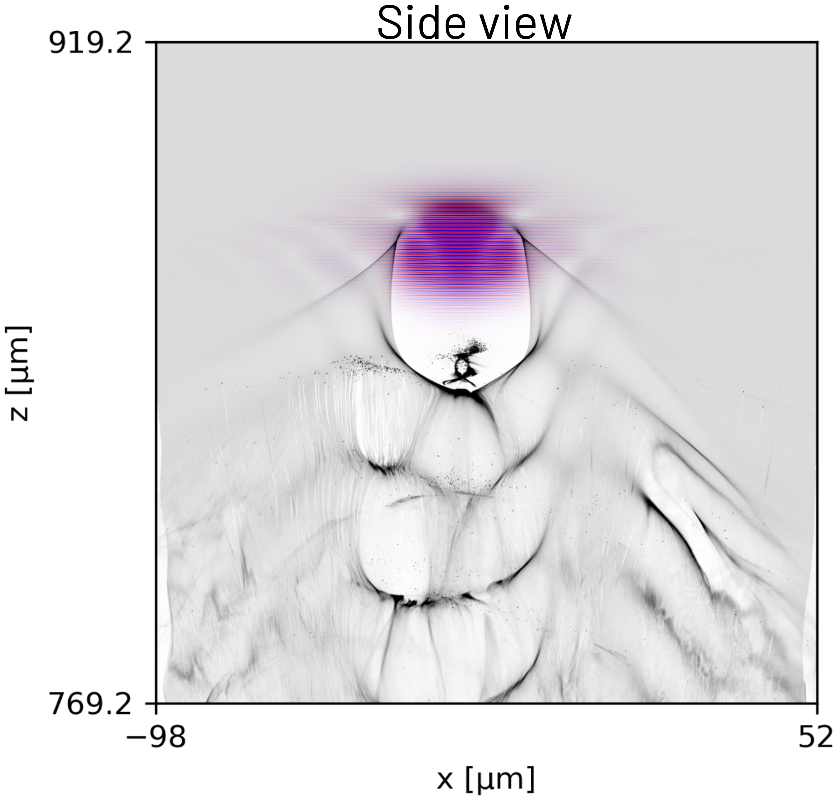

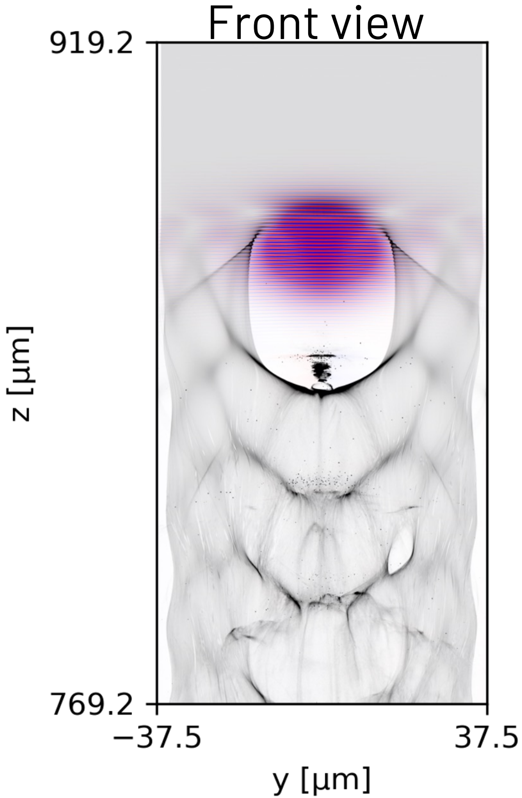

← 3D simulation on 4096 Summit nodes

2D slices of our 3D simulations highlight the acceleration process

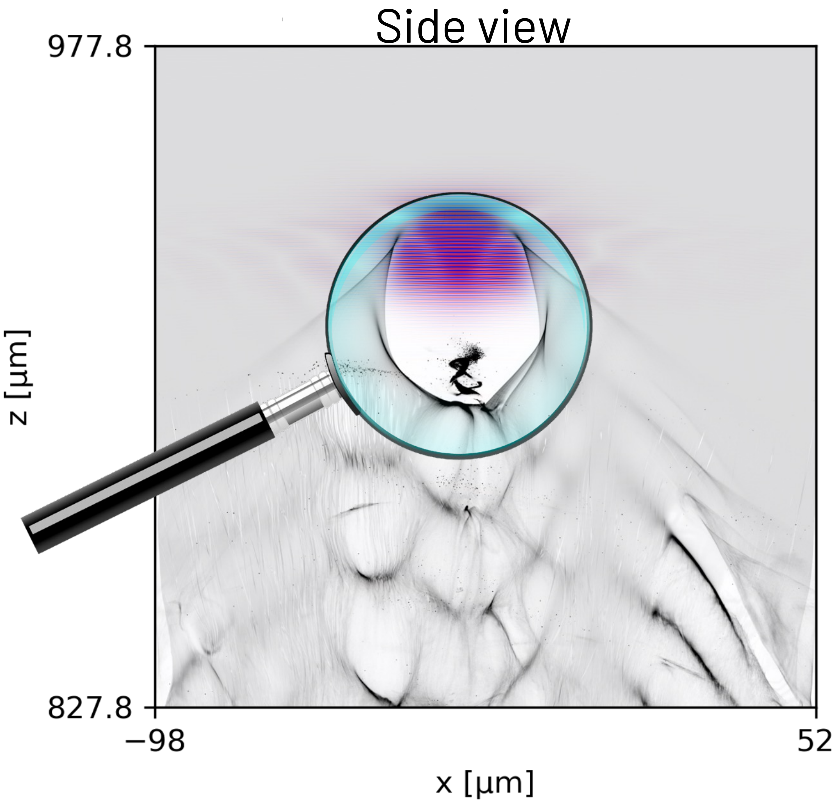

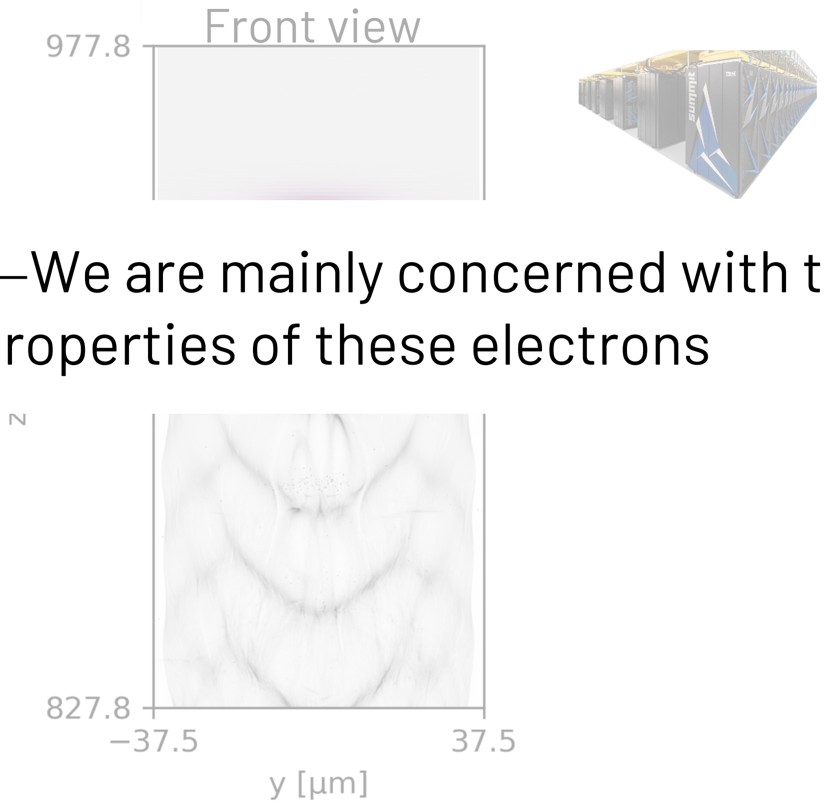

←We are mainly concerned with the properties of these electrons

Our simulations shows that we can accelerate a substantial amount of charge with high quality

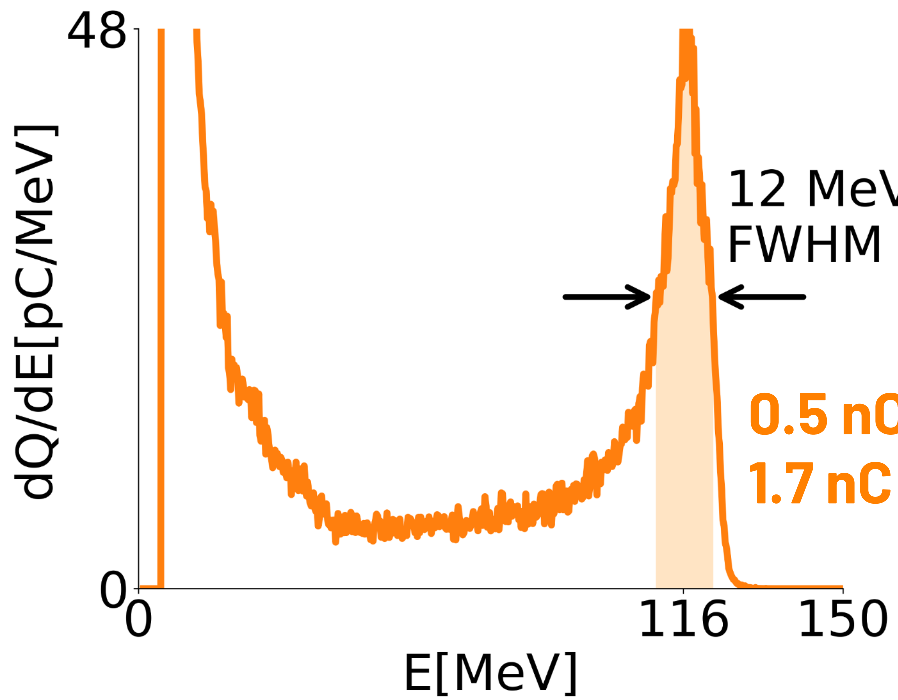

0.5 nC (peak)
1.7 nC (total)

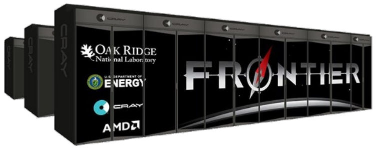

for a PW laser

Production runs on Frontier, Fugaku and Summit

Exascale simulations informed the design of the first experimental validation of our concept

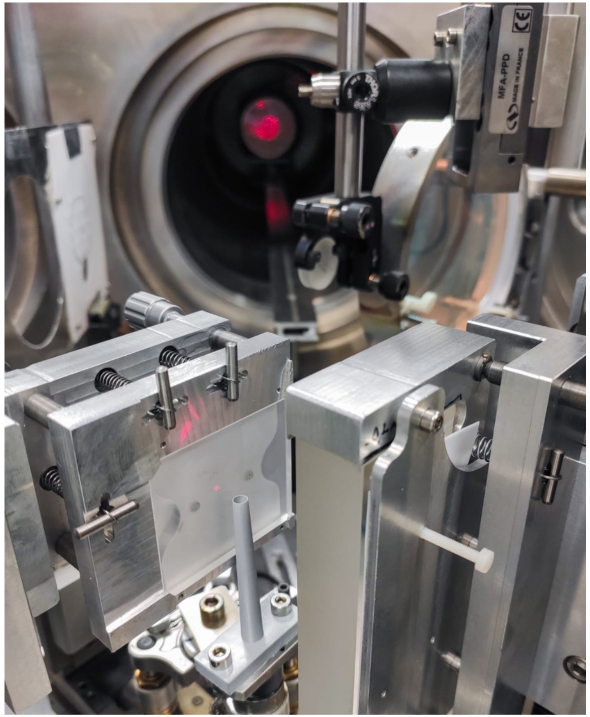

2022

First proof-of principle experiments in March and October 2022 at LOA (France)

loa

Electron energy spectrum

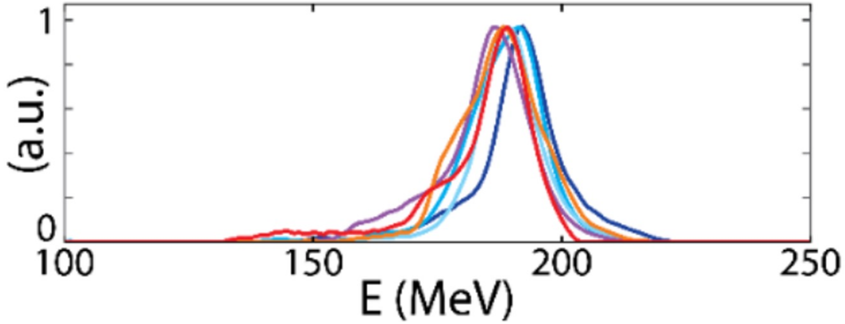

4X increase of accelerated charge with respect to conventional techniques for the same laser energy!

A. Leblanc, L. Fedeli, A. Huebl, J.-L. Vay ... and H. Vincenti, *in preparation*

Sustained Flop/s

Note: Frontier & Perlmutter are pre-acceptance machine results

WarpX can now do science cases 500x larger than pre-ECP.

DP PFlop/s	HPCG	
3.38	223%	Perlmutter A100
11.79	435%	Summit V100
5.31 SP: 17.3	35% x3.3	Fugaku A64FX
43.45	310%	Frontier MI250X


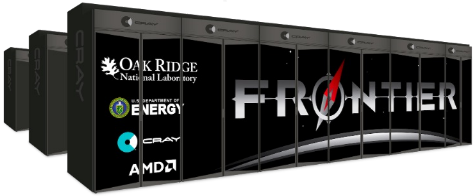

WarpX is now 500x More Performant than its Baseline on Cori

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

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

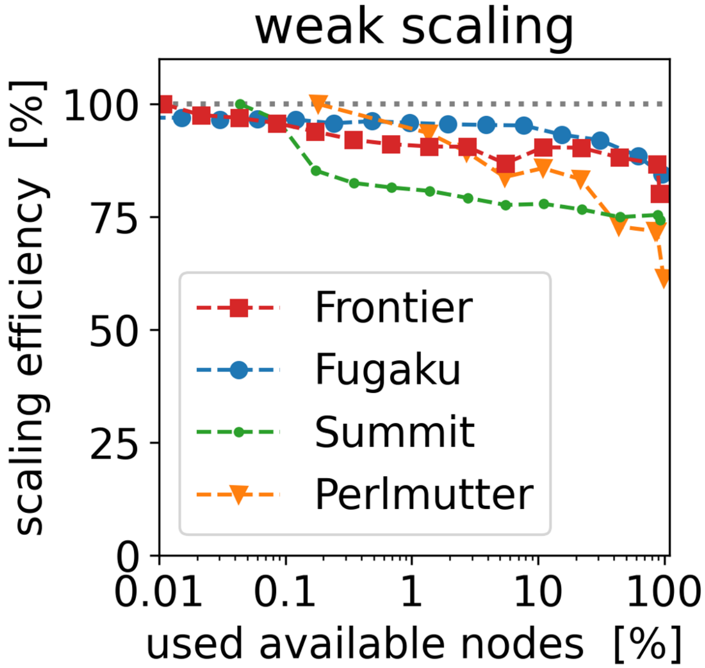

Figure-of-Merit: weighted updates / sec

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

Note: Perlmutter & Frontier are pre-acceptance measurements!

EXASCALE COMPUTING PROJECT

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!*

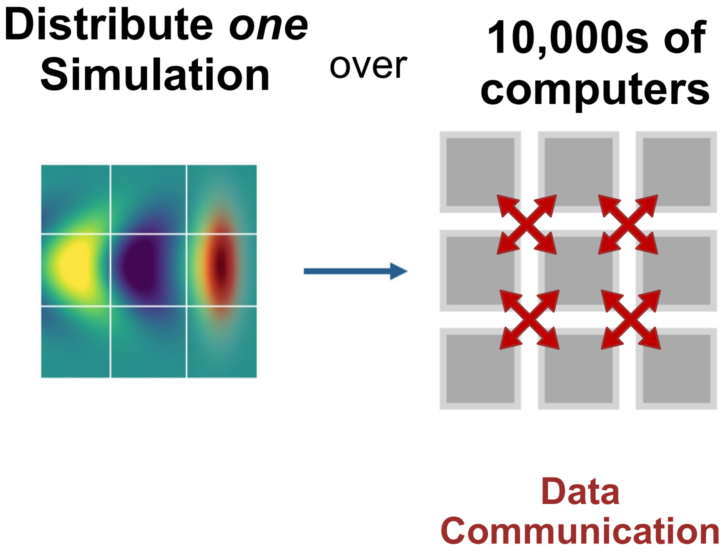

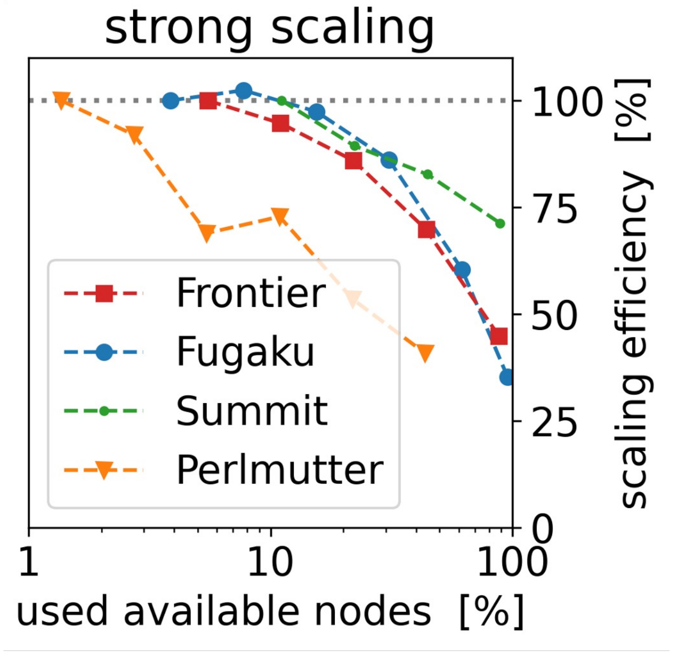

Note: Perlmutter & Frontier are pre-acceptance measurements!

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**

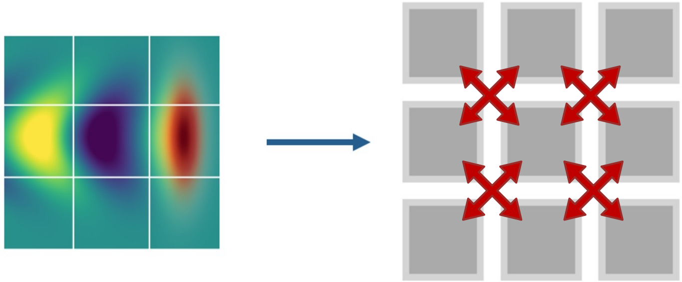

Data Communication

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
- …

The Growing Role of Machine Learning

ML-Guided Optimization: Automate Scans & Design Workflows

Design Optimization:

- ML **finds optima rapidly**, e.g. *Gaussian Processes*, *Bayesian Optimization*
- **Multi-Fidelity** (think: multi-resolution): Learn trends from fast simulations and add precision with large costly sims

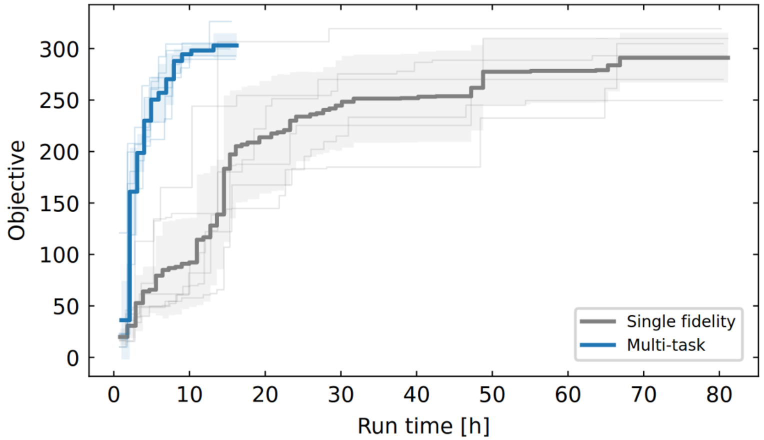

Strongly-correlated case:

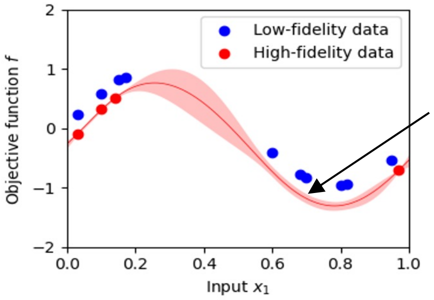

Low uncertainty, despite the absence of high-fidelity data

Un-correlated case:

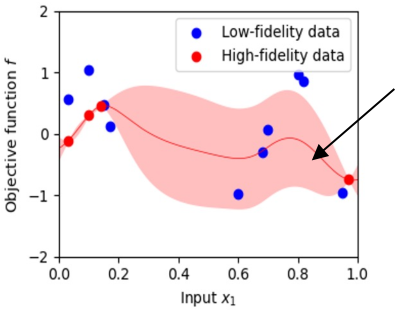

High uncertainty; low-fidelity data is ignored

Bonilla et al., NIPS, (2007); R. Lehe et al., APS DPP (2022); Á. Ferran Pousa et al., IPAC22 (2022) & subm. PRAB (2023)

AI/ML Surrogates: Fast, Advanced Accelerator Elements

Model Speed: for accelerator elements

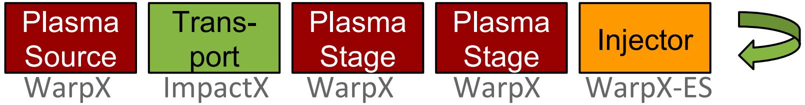

For a well-defined parameter range, we want to replace the **expensive** plasma simulation section with pre-trained surrogate models.

Fast surrogates: Data-driven modeling is
a potential middle ground between
- analytical modeling and
- full-fidelity simulations

for ***beamline design & operations***.

Model Choice:

for complex, nonlinear, many-body systems

pick two* of the following

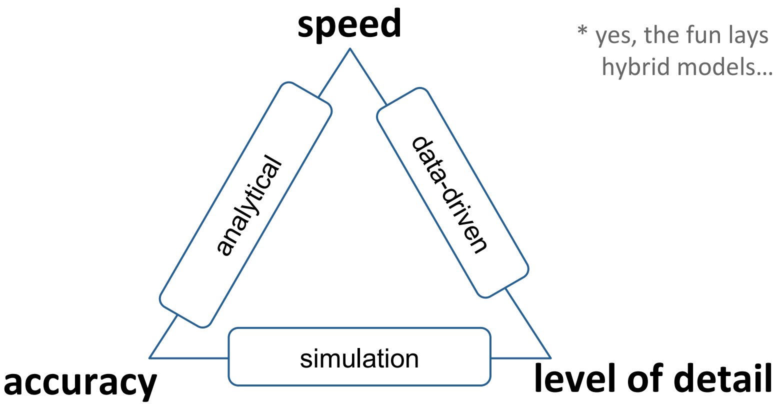

* yes, the fun lays in hybrid models… :)

LDRD

A Huebl, R Sandberg, R Lehe, CE Mitchell et al.

Outlook

Outlook

- **Game changing R&D is on the way to develop smaller & greener, and yet more powerful, particle accelerators**
  - Large impact if successful: in science, health, industry, security, …
  - HPC is key to enable these progress

- **High-Performance Computing** offers exciting paths at the intersection of
  - Computer & Data Science
  - Applied Mathematics
  - Computational Domain Science

Free after the Olympic motto: *Faster, Larger, Smarter - Together*

- **Exciting R&D Challenges** in performance tuning for cutting-edge compute hardware, scalability, novel algorithms & implementations, AI/ML, computer-experiment interplay, and more.

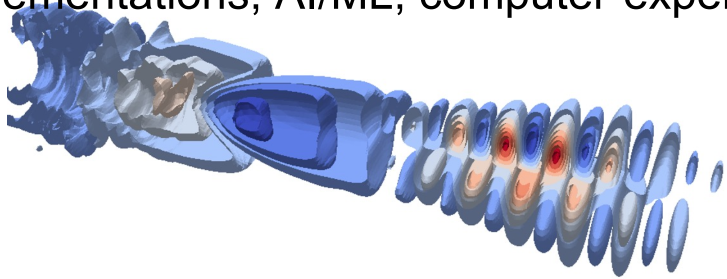

Funding Support

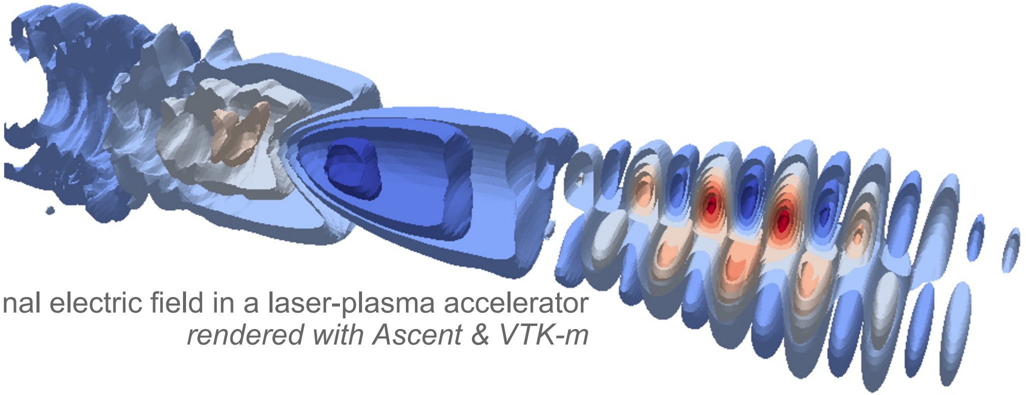

WarpX: longitudinal electric field in a laser-plasma accelerator
rendered with Ascent & VTK-m

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

Our C++17 & Python Software Stack

BLAST
BEAM PLASMA & ACCELERATOR SIMULATION TOOLKIT

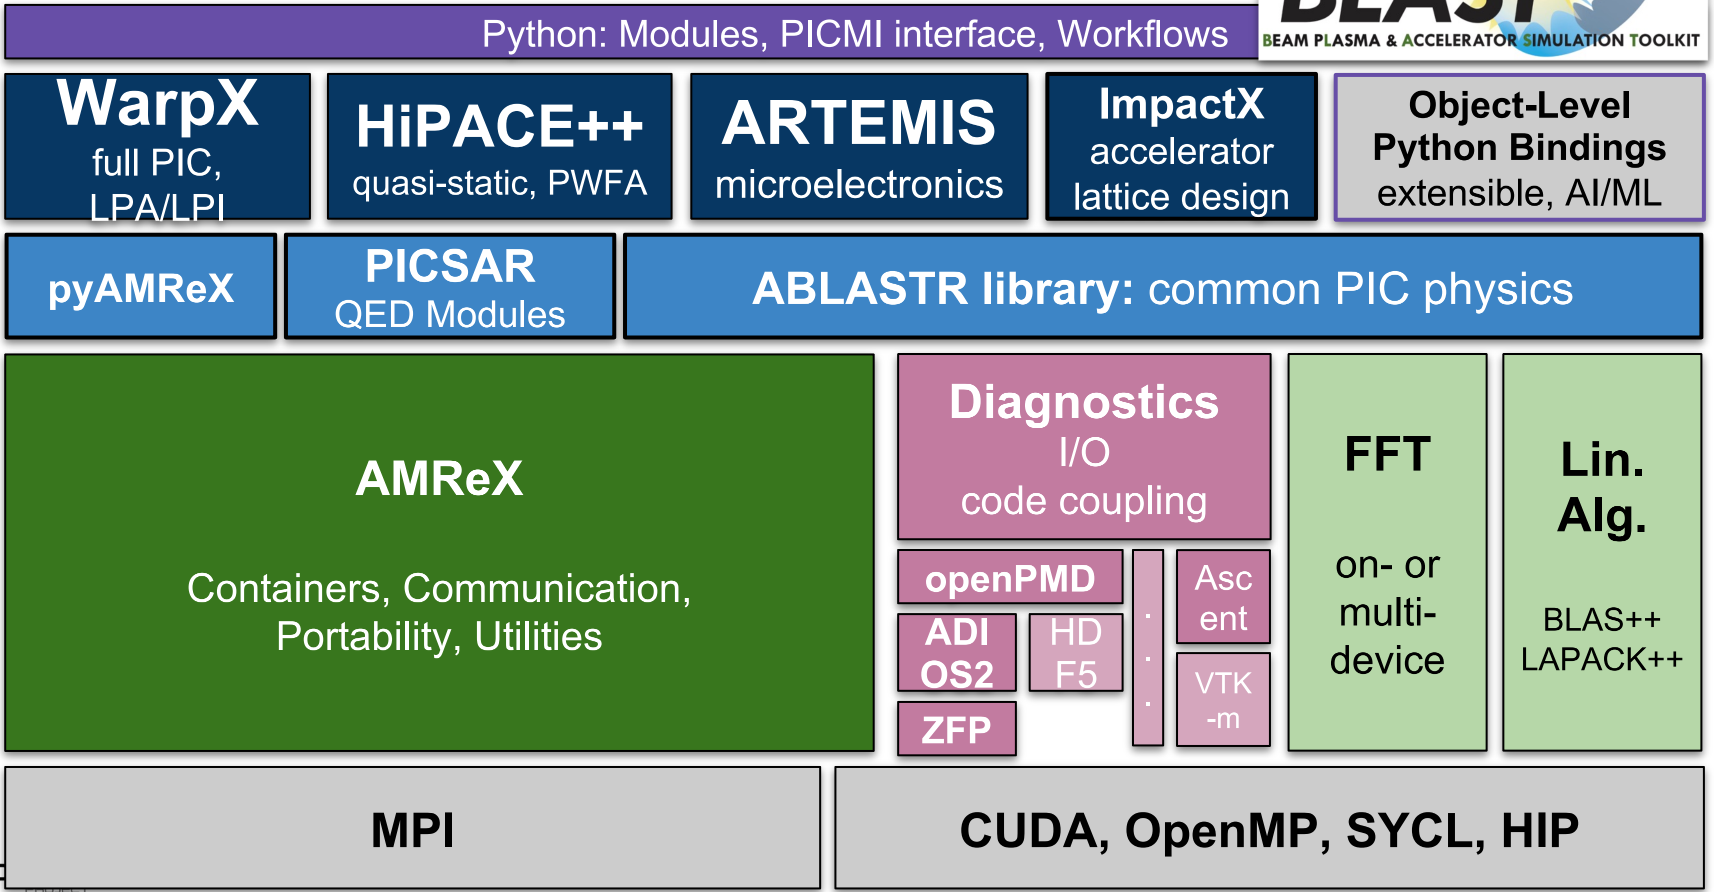

Desktop to HPC