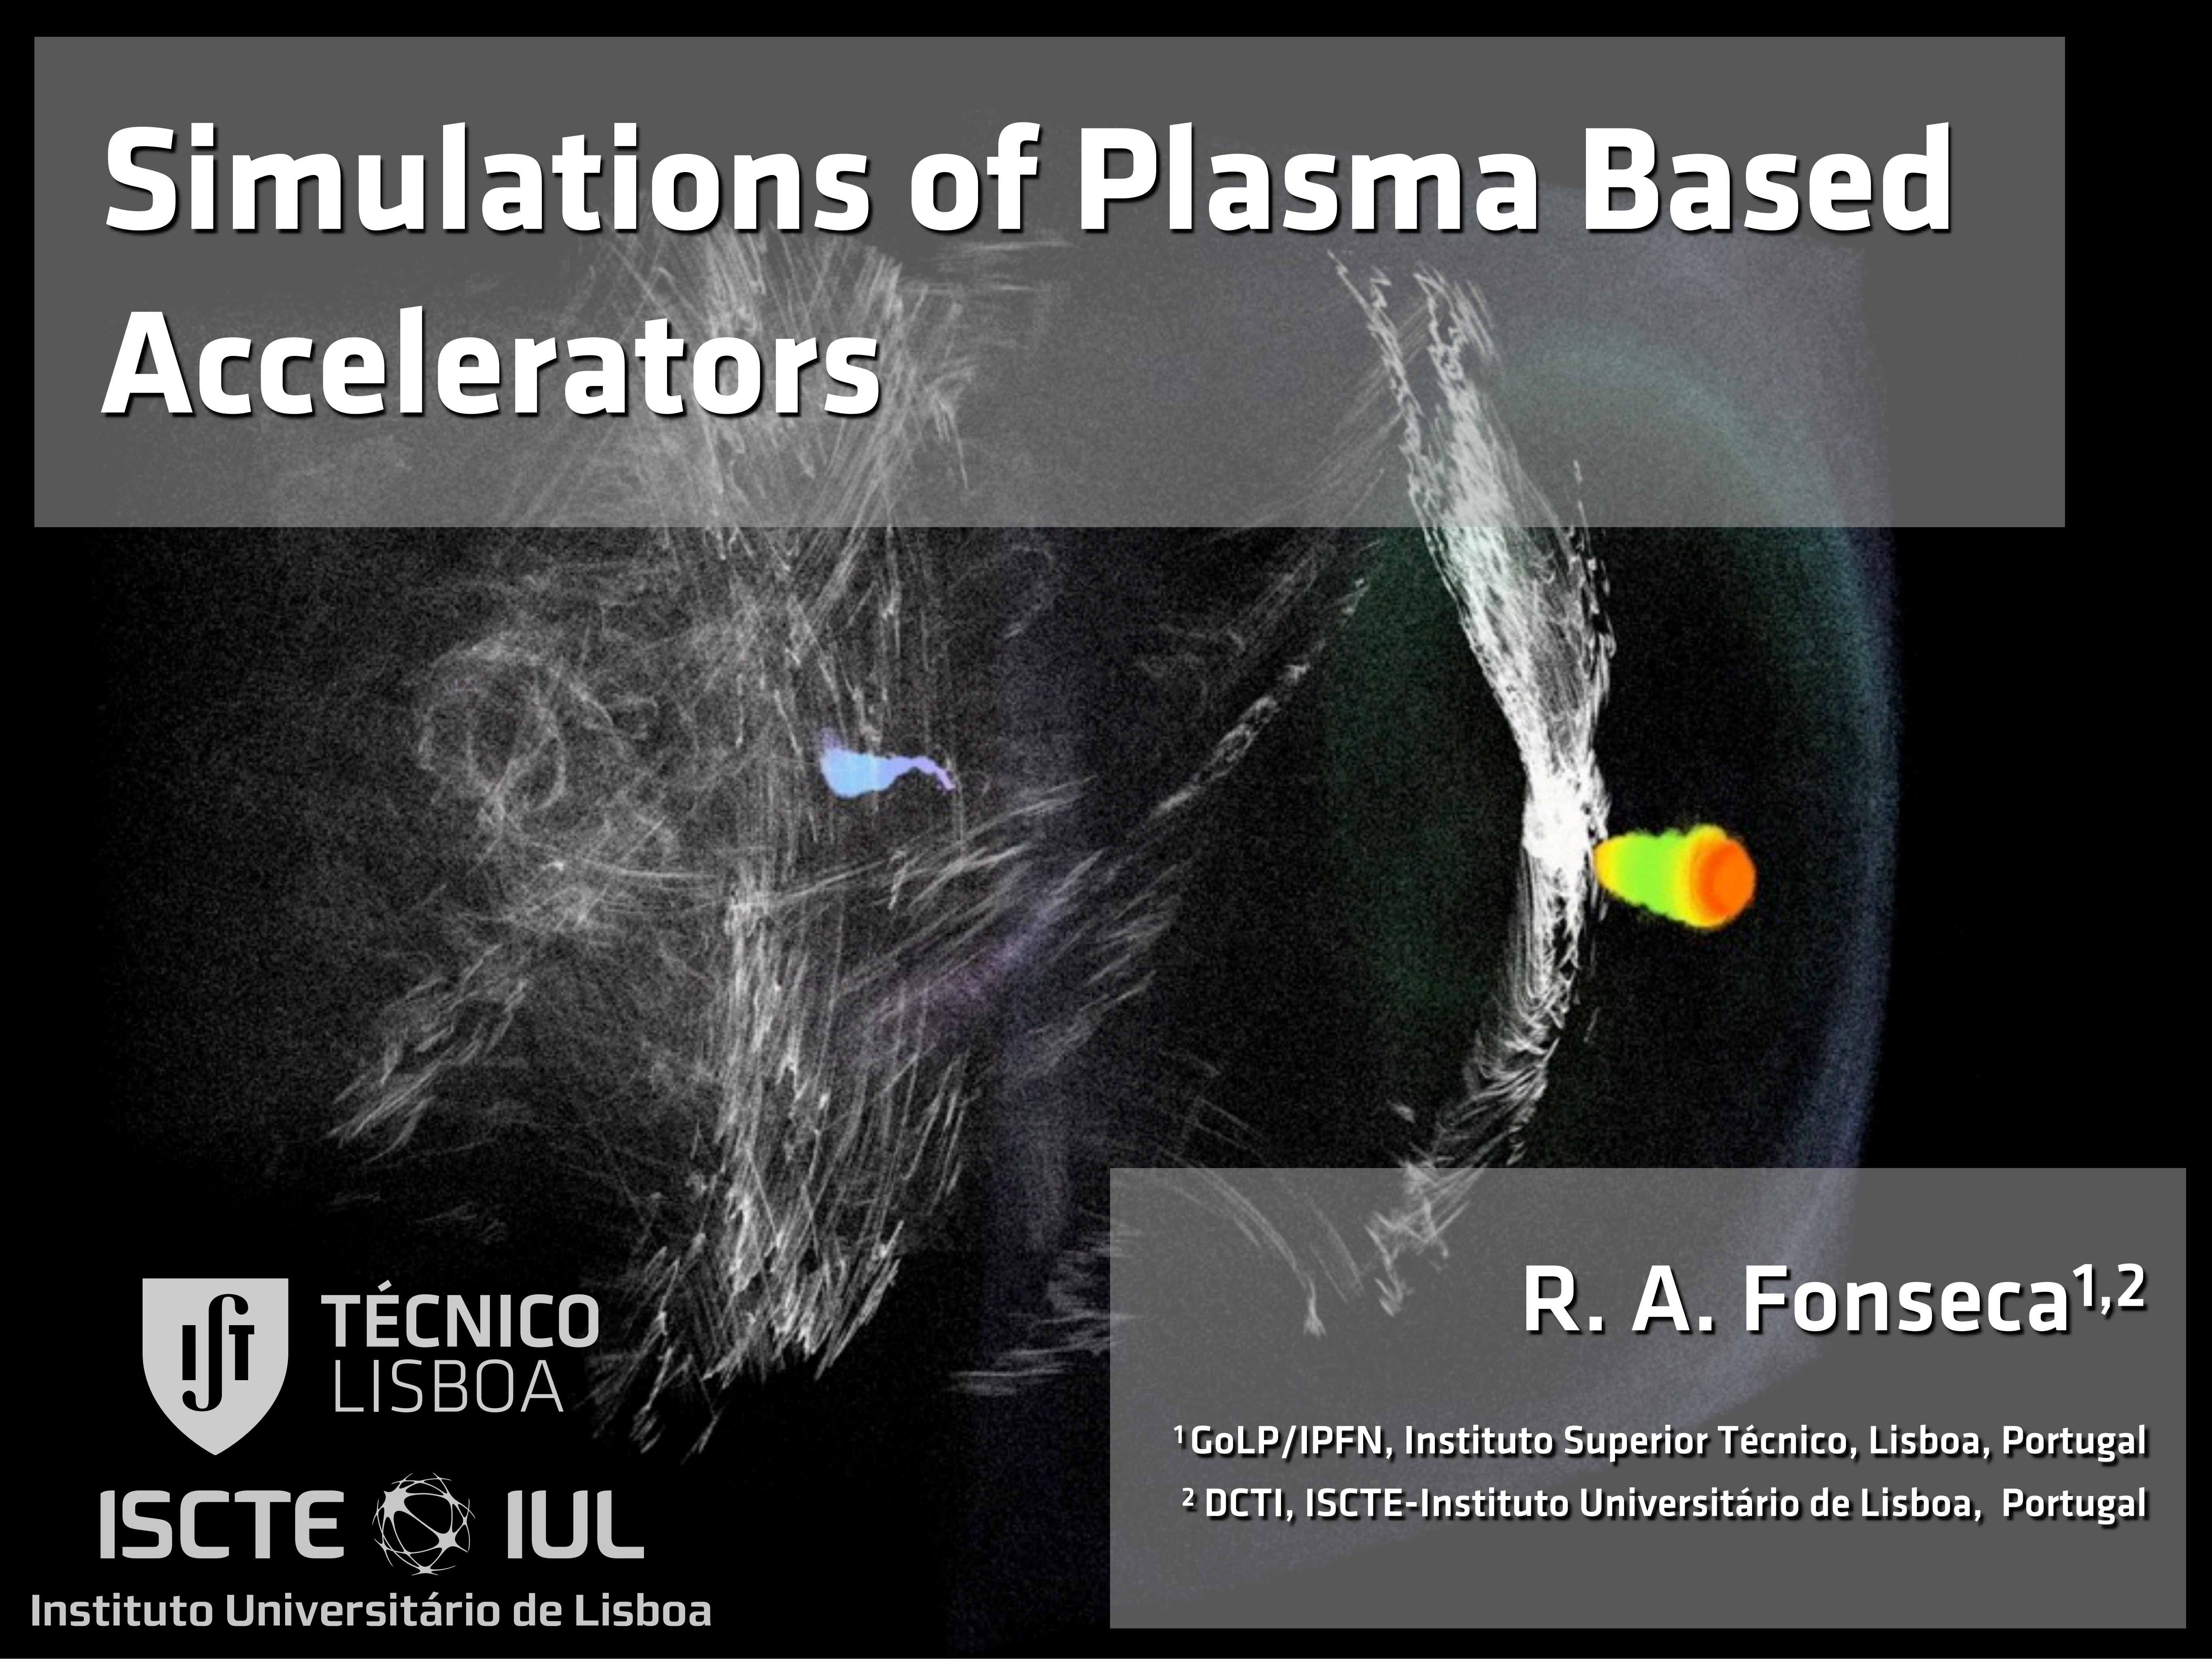

Simulations of Plasma Based Accelerators

R. A. Fonseca[1,2]

[1] GoLP/IPFN, Instituto Superior Técnico, Lisboa, Portugal

[2] DCTI, ISCTE-Instituto Universitário de Lisboa, Portugal

TÉCNICO LISBOA

ISCTE IUL

Instituto Universitário de Lisboa

Acknowledgements

- IST
  - P. Alves, J. L. Martins, T. Grismayer, J. Vieira, V.B. Pathak, P. Ratinho, A. Stockem, M. Vranic, L. O. Silva
- UCLA
  - A. Davidson, J. Tonge, F. Tsung, V.K.Decyk, W. B. Mori, C. Joshi
- Lawrence Livermore National Laboratory
  - F. Fiúza
- Simulation results
  - epp / IST Clusters (IST), Dawson/Hoffman Clusters (UCLA), Jugene/Juqueen (FZ Jülich), Jaguar (ORNL), Supermuc (LRZ), BlueWaters (NCSA)

Computing Power Evolution

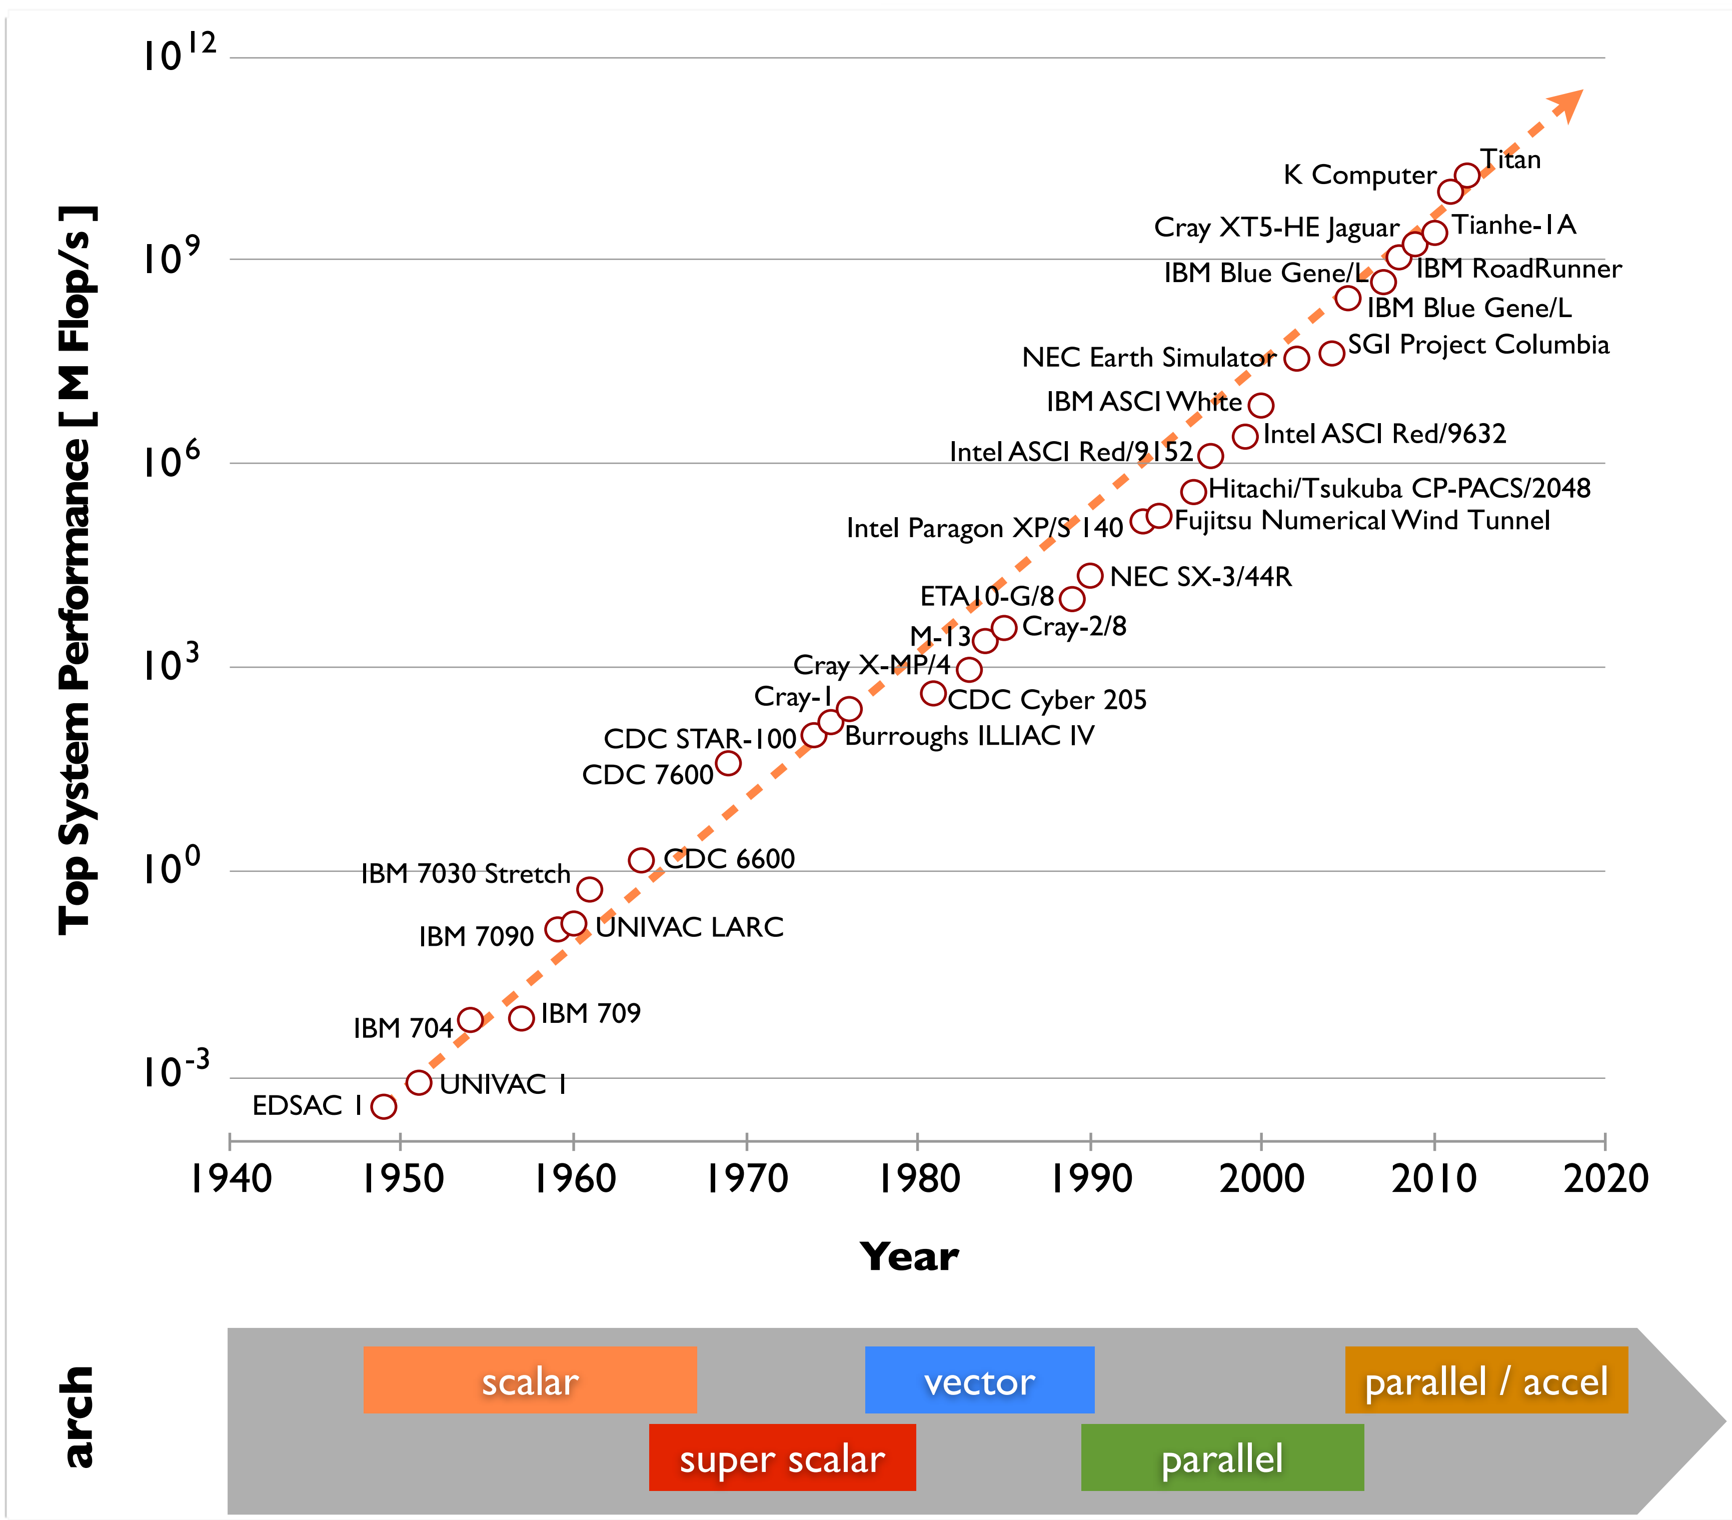

High Performance Computing Power Evolution (data from multiple sources)

Titan Supercomputer @ ORNL

Titan - Cray XK7

18688 Compute Nodes

1× 16 core AMD 6274 @ 2.2 GHz

1 × Nvidia Tesla K20 GPU

R_{peak} 27.1 PFlop/s

R_{max} 17.6 Pflop/s

Outline

- Modeling Plasma Based Accelerators
  - The OSIRIS framework
  - Algorithm
  - Computational Requirements
- Extending the algorithm
  - Lorentz boosted frames
  - Ponderomotive guiding center
  - High density hybrid
  - LWFA in cylindrical geometry
- Scaling to Petascale Systems
  - Multi-level parallelism
  - Multi-dimensional dynamic load balance
- Plasma Acceleration on Tier-0 systems
  - Scalings @ Sequoia & BlueWaters
  - 20 m long beam-plasma interaction @ SuperMUC
  - High-res 1 GeV LWFA @ Jaguar
  - Full scale target LWFA modeling @ SuperMUC
- Overview

Outline

- Modeling Plasma Based Accelerators
  - The OSIRIS framework
  - Algorithm
  - Computational Requirements
- Extending the algorithm
  - Lorentz boosted frames
  - Ponderomotive guiding center
  - High density hybrid
  - LWFA in cylindrical geometry
- Scaling to Petascale Systems
  - Multi-level parallelism
  - Multi-dimensional dynamic load balance
- Plasma Acceleration on Tier-0 systems
  - Scalings @ Sequoia & BlueWaters
  - 20 m long beam-plasma interaction @ SuperMUC
  - High-res 1 GeV LWFA @ Jaguar
  - Full scale target LWFA modeling @ SuperMUC
- Overview

Modeling Plasma Based Accelerators

PIC algorithm

$$\frac{d\mathbf{u}}{dt} = \frac{q}{m}\left(\mathbf{E} + \frac{1}{\gamma c}\mathbf{u}\times\mathbf{B}\right)$$

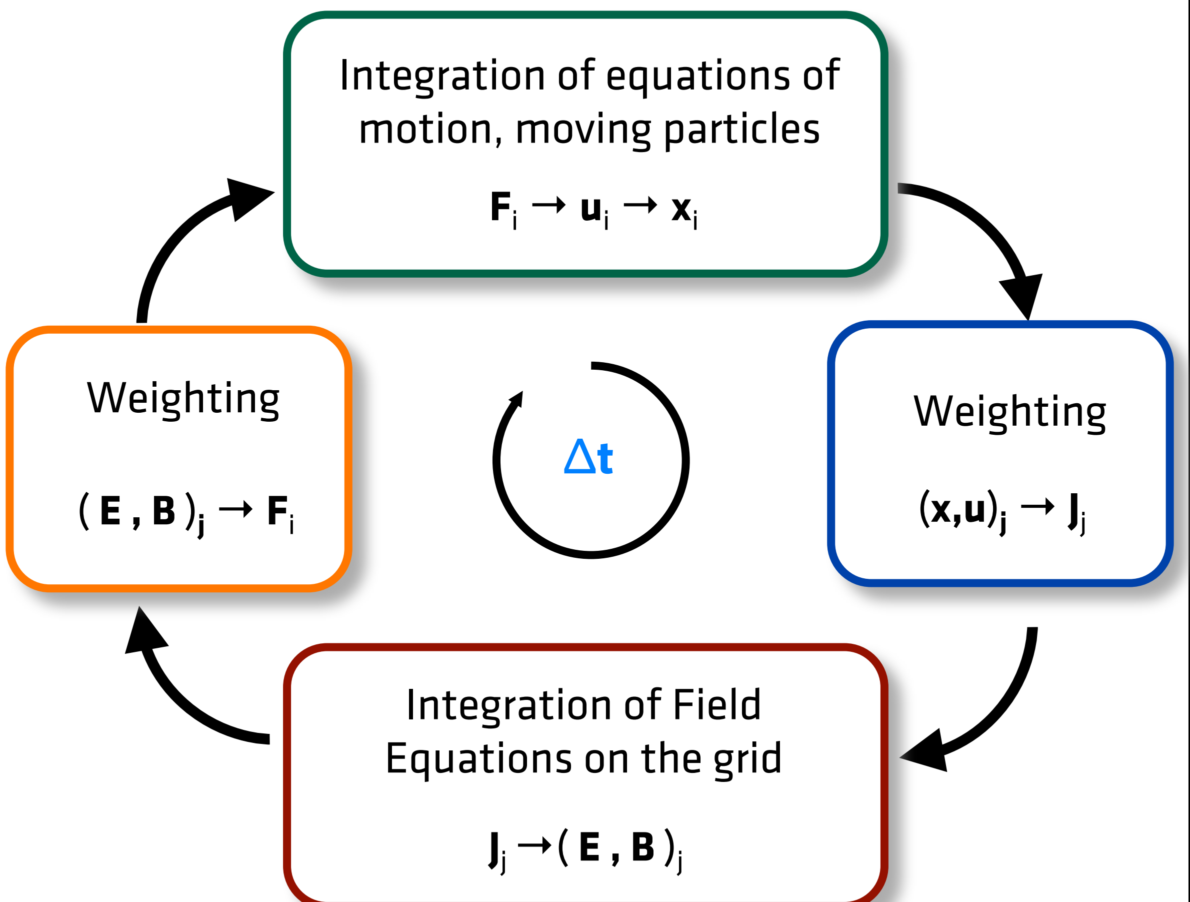

$$\frac{\partial \mathbf{E}}{\partial t} = c\vec{\nabla}\times\mathbf{B} - 4\pi\mathbf{j}$$

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

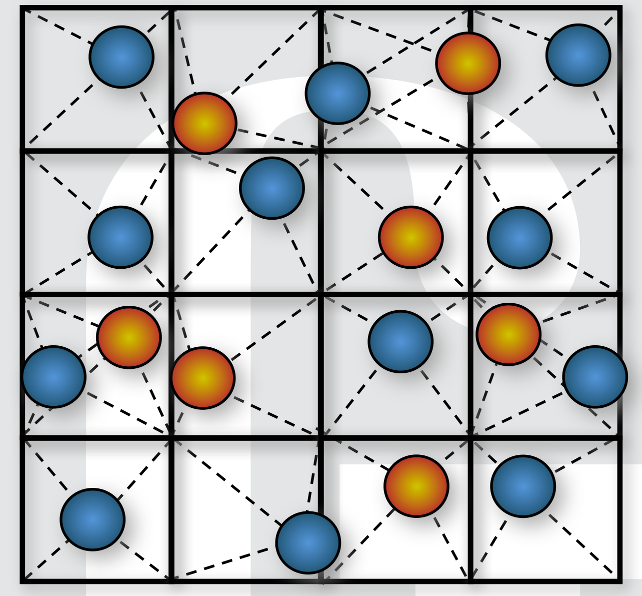

- **Complex, strongly nonlinear interaction**
  - **Kinetic modeling**
- **Modeling particle interactions directly very difficult**
  - **Calculations ~ N_p^2**
- **Collective interactions dominate**
  - **Use a particle-mesh algorithm**
- **Highly collisional plasmas**
  - **Additional MC binary collision module**
- **Ultra relativistic intensities**
  - **Radiation cooling, QED effects**

osiris 2.0

osiris framework

- Massivelly Parallel, Fully Relativistic Particle-in-Cell (PIC) Code
- Visualization and Data Analysis Infrastructure
- Developed by the osiris.consortium
  ⇒ UCLA + IST

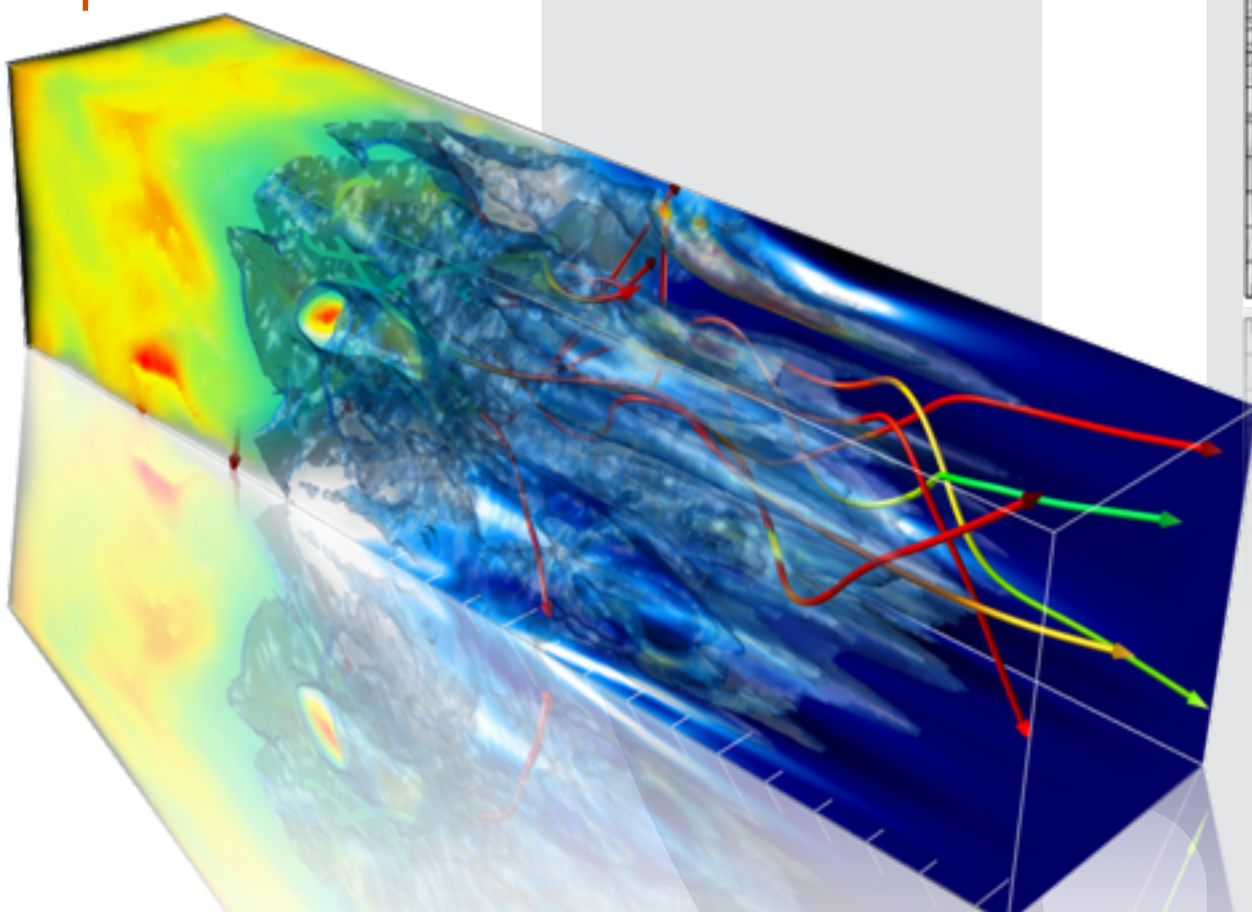

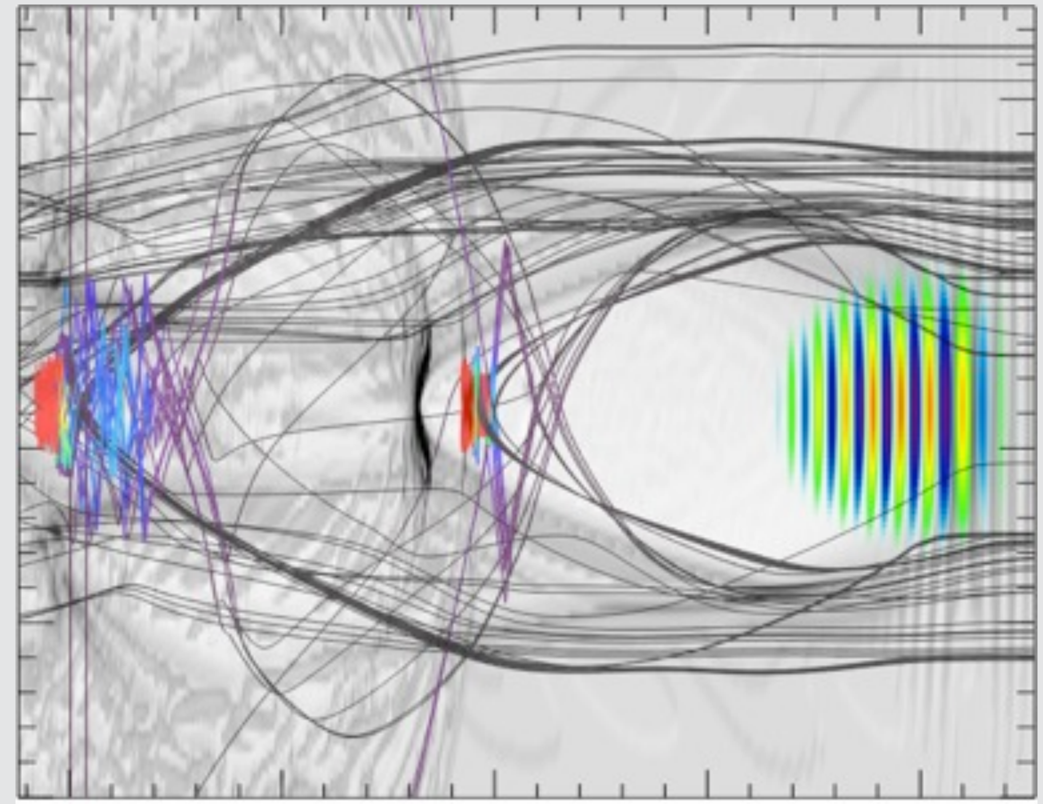

Ricardo Fonseca: ricardo.fonseca@ist.utl.pt
Frank Tsung: tsung@physics.ucla.edu

http://cfp.ist.utl.pt/golp/epp/
http://exodus.physics.ucla.edu/

code features

- Scalability to ~ 1.6 M cores
- SIMD hardware optimized
- Parallel I/O
- Dynamic Load Balancing
- High order splines
- PGC & HD-Hybrid algorithm
- Tight integration with visXD

Modeling Laser Wakefield Acceleration

- Tool of choice
  - 3D EM-PIC algorithm
- Computational Requirements
  - ~ 10^9 grid cells
  - ~ 10^{10} particles
  - Iterations ~ 10^6 - 10^7
  - Memory ~ 1 - 10 TB
  - Operations ~ 10^{18} - 10^{19}
- Petascale Computing
  - ~ 4 cent. / core hour
  - Efficient use of Tier-0 systems

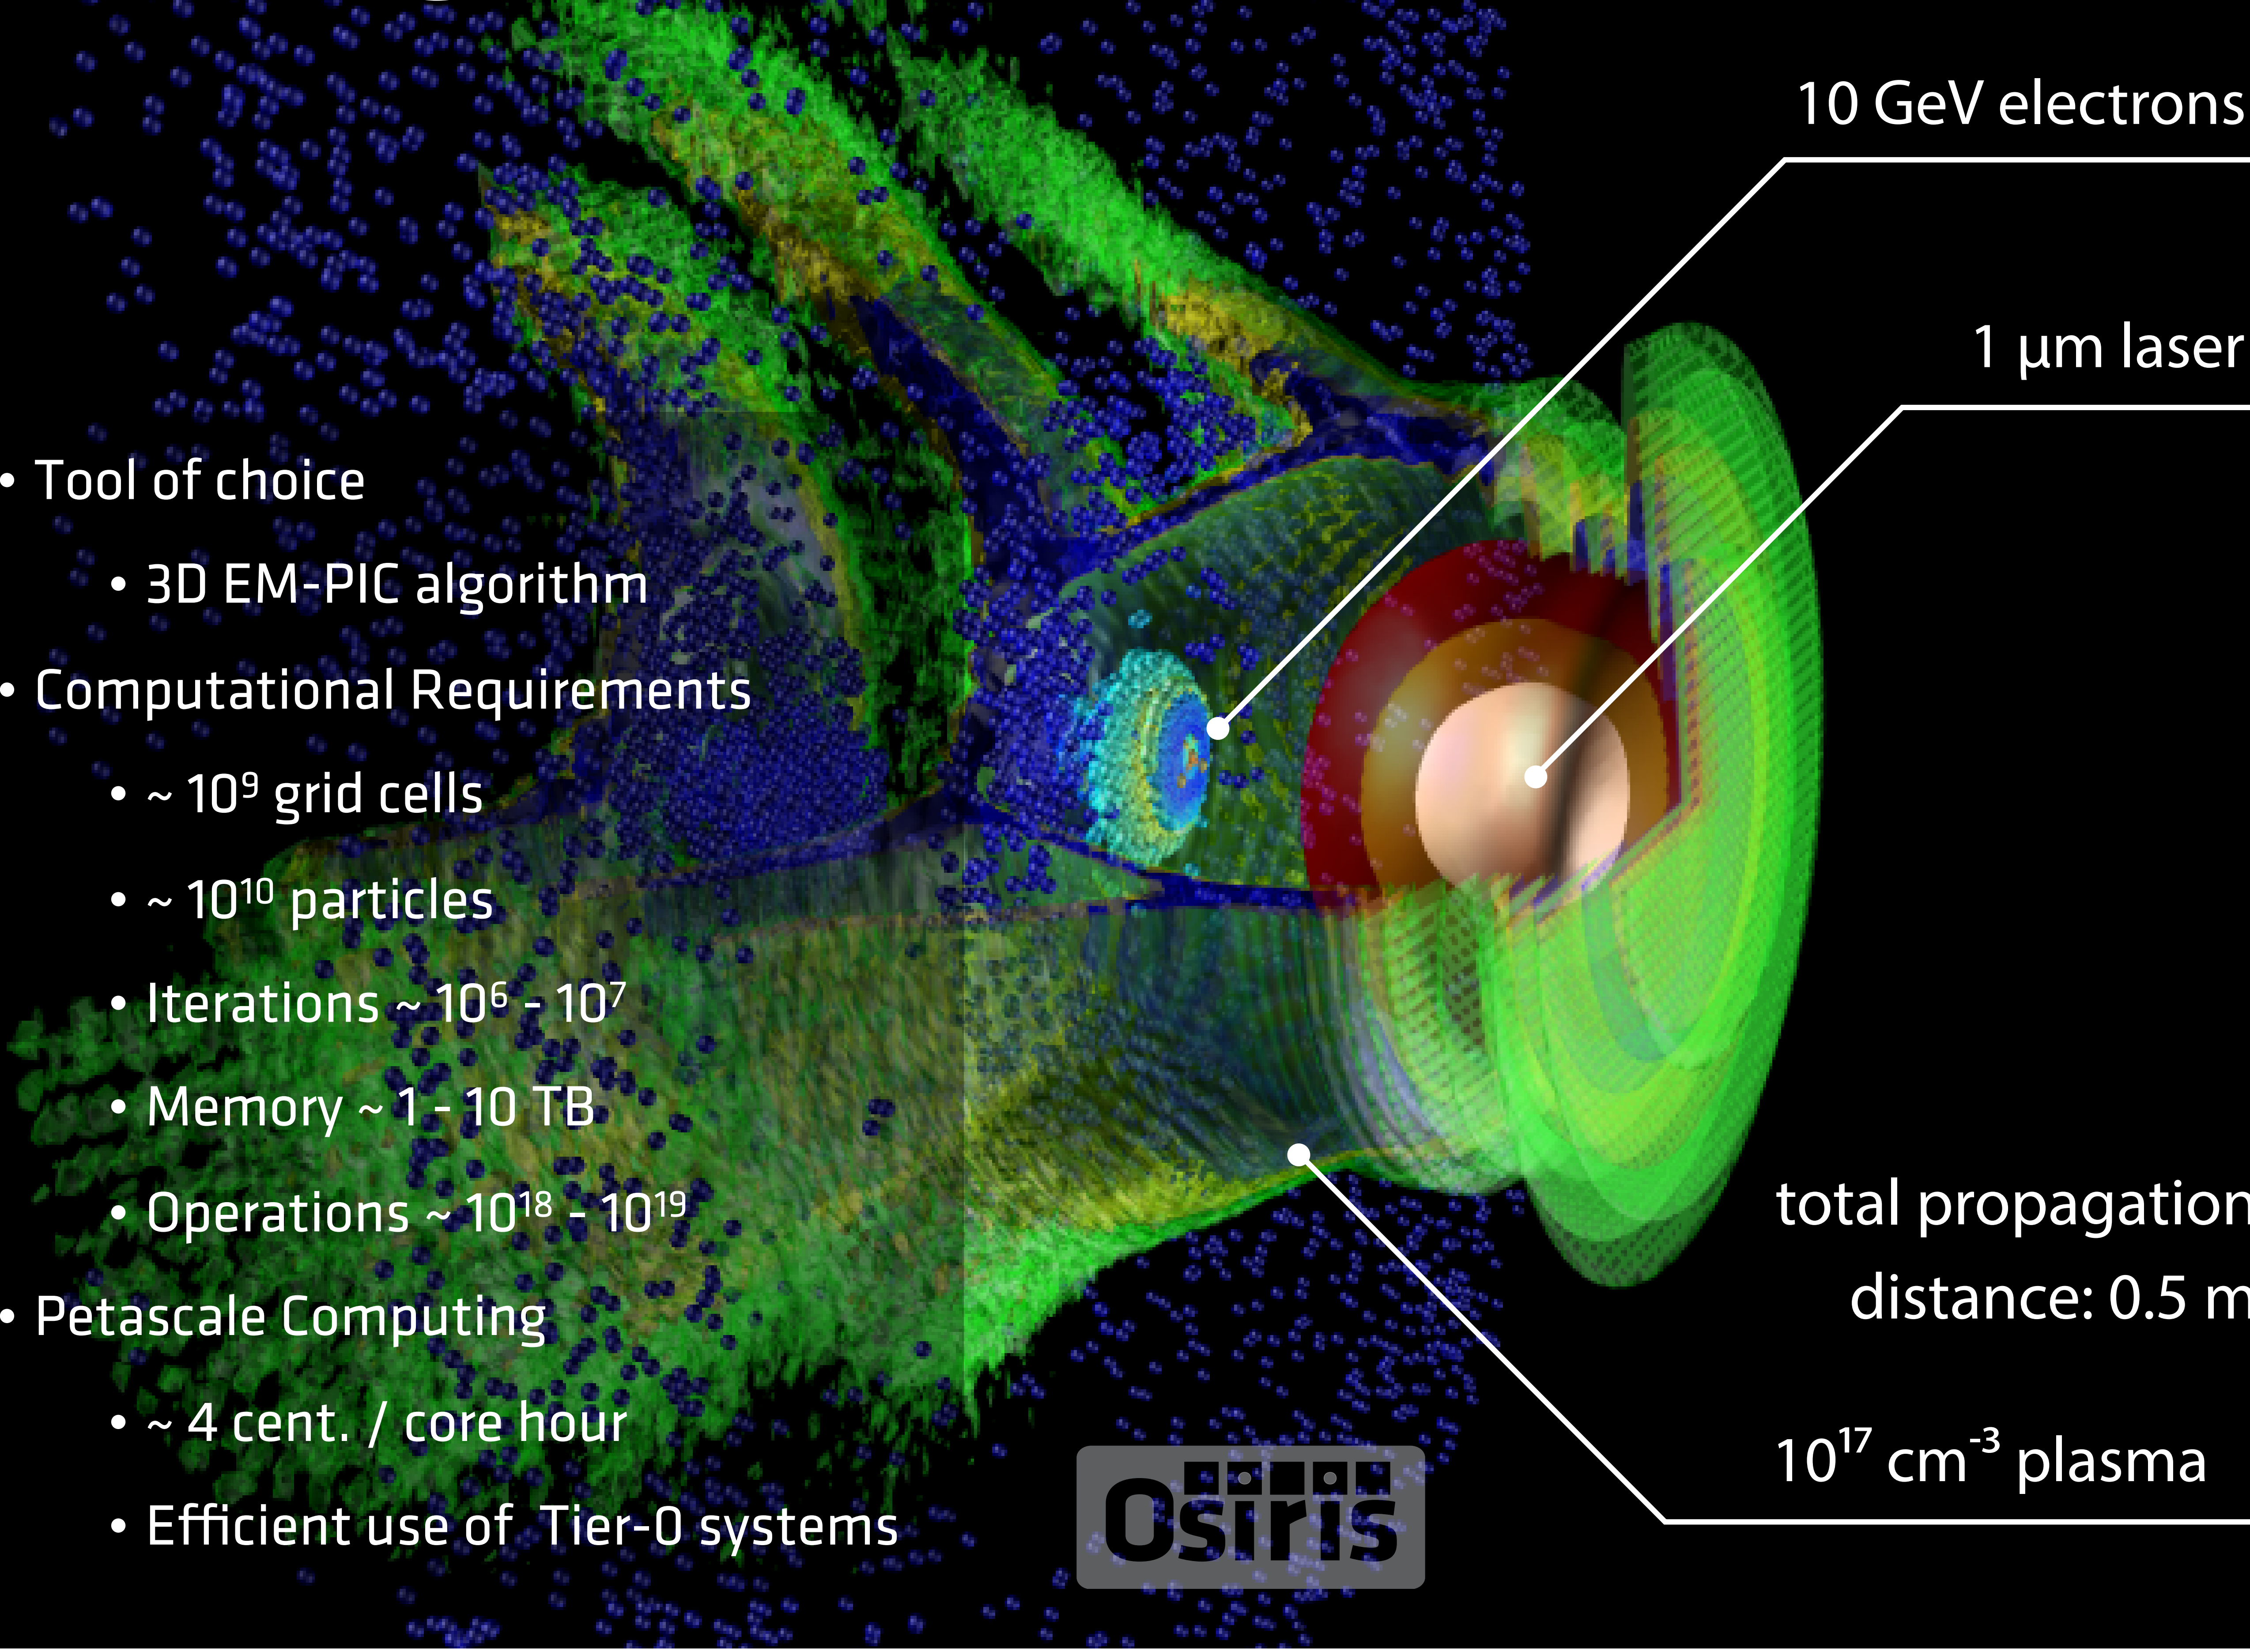

Outline

- Modeling Plasma Based Accelerators
  - The OSIRIS framework
  - Algorithm
  - Computational Requirements
- Extending the algorithm
  - Lorentz boosted frames
  - Ponderomotive guiding center
  - High density hybrid
  - LWFA in cylindrical geometry
- Scaling to Petascale Systems
  - Multi-level parallelism
  - Multi-dimensional dynamic load balance
- Plasma Acceleration on Tier-0 systems
  - Scalings @ Sequoia & BlueWaters
  - 20 m long beam-plasma interaction @ SuperMUC
  - High-res 1 GeV LWFA @ Jaguar
  - Full scale target LWFA modeling @ SuperMUC
- Overview

Ponderomotive Guiding Center (PGC)*

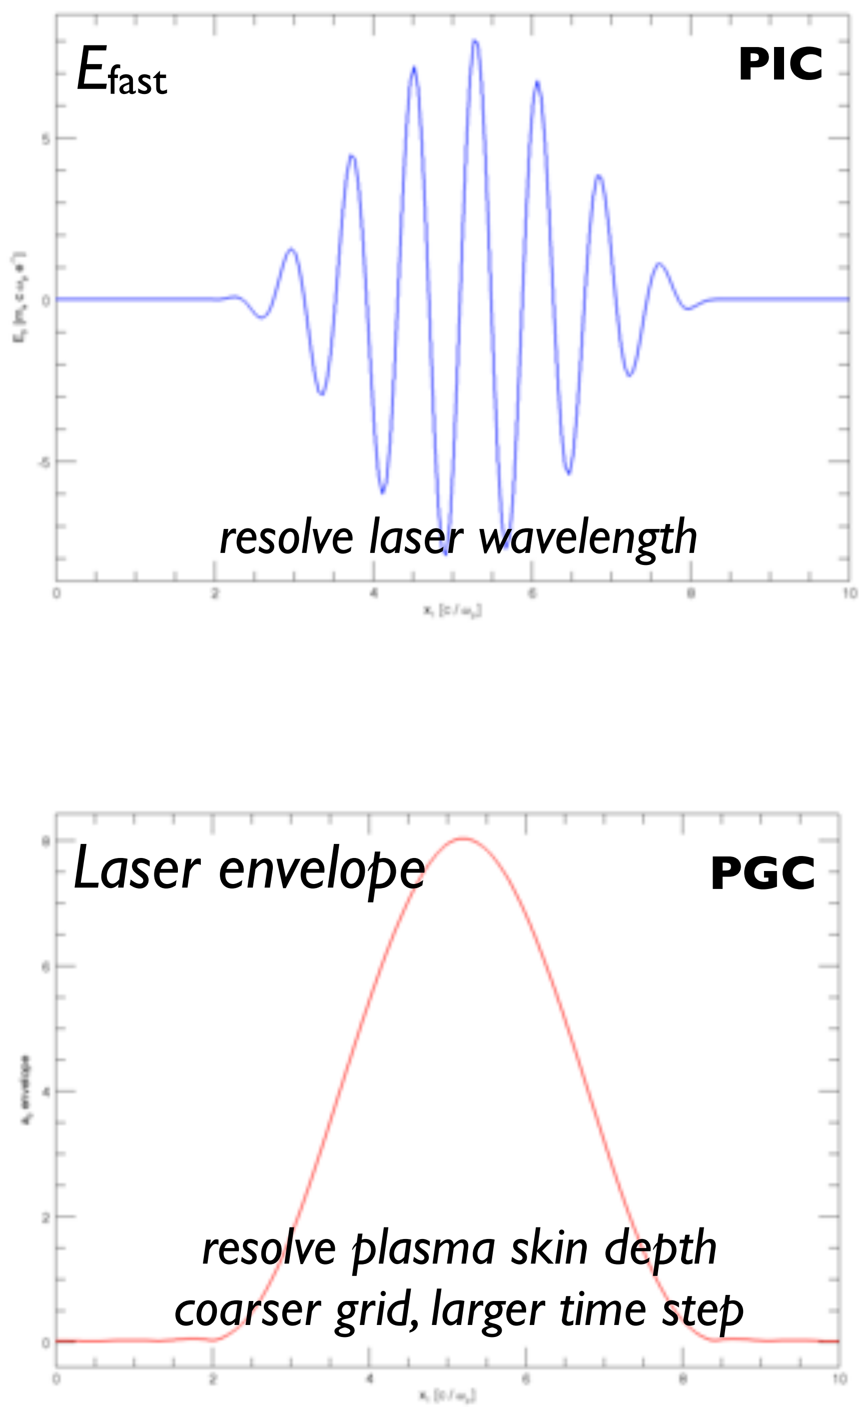

Long distance (10-100m) LWFA

- **Significant challenge for PIC codes**
  - Large disparity of spatial scales
    - propagation distance / laser wavelength
  - Algorithm must resolve the smallest scale in the simulation
  - High resolution, large iteration count
- **Ponderomotive guiding center approximation**
  - Model laser envelope propagation
  - Push particles using self consistent plasma fields and ponderomotive force

* D. Gordon, W. Mori, T. Antonsen, IEEE-TPS, **28** 1135-1143 (2000)

PGC - PIC simulations of LWFA

Algorithm

Laser pulse envelope equation:

$$\frac{\partial a}{\partial \tau} = \frac{1}{2i\omega_0}\left(1 + \frac{1}{i\omega_0}\frac{\partial}{\partial \xi}\right)\left(\chi + \nabla_\perp^2\right) a$$

τ=t; laser frequency; ξ =x-ct; laser envelope

Ponderomotive guiding center pusher:

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}\tau} = q\left(\mathbf{E} + \mathbf{v} \times \mathbf{B} - \frac{1}{4}\frac{q}{\langle m \rangle}\nabla |a|^2\right)$$

slow varying momentum; slow varying electric (**E**) and magnetic (**B**) fields; relativistic mass in laser fields

Features

- Speedup ~ $(\omega_0/\omega_p)^2$
- Large boosted frame like computational savings
- Ultra-fast LWFA simulations
- Ionization energy depletion
- 2D slab/2D cylindrical

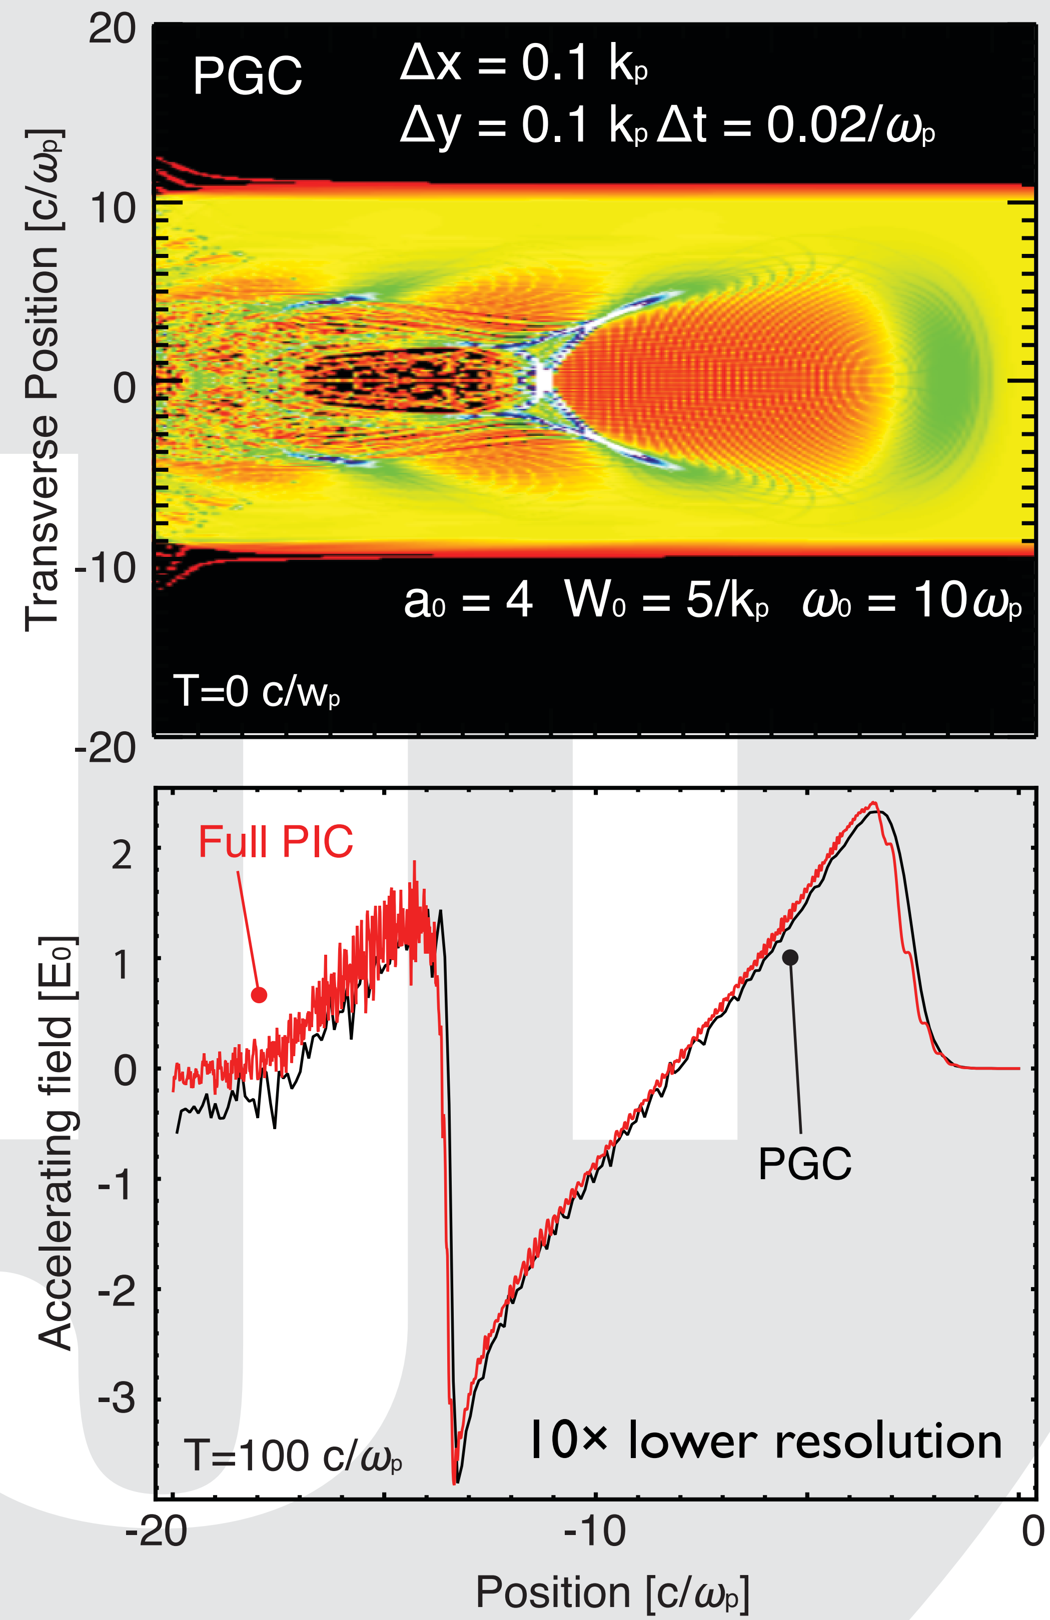

D. Gordon, W. Mori, T. Antonsen, IEEE-TPS, **28** 1135-1143 (2000)

Lorentz Boosted Frame*

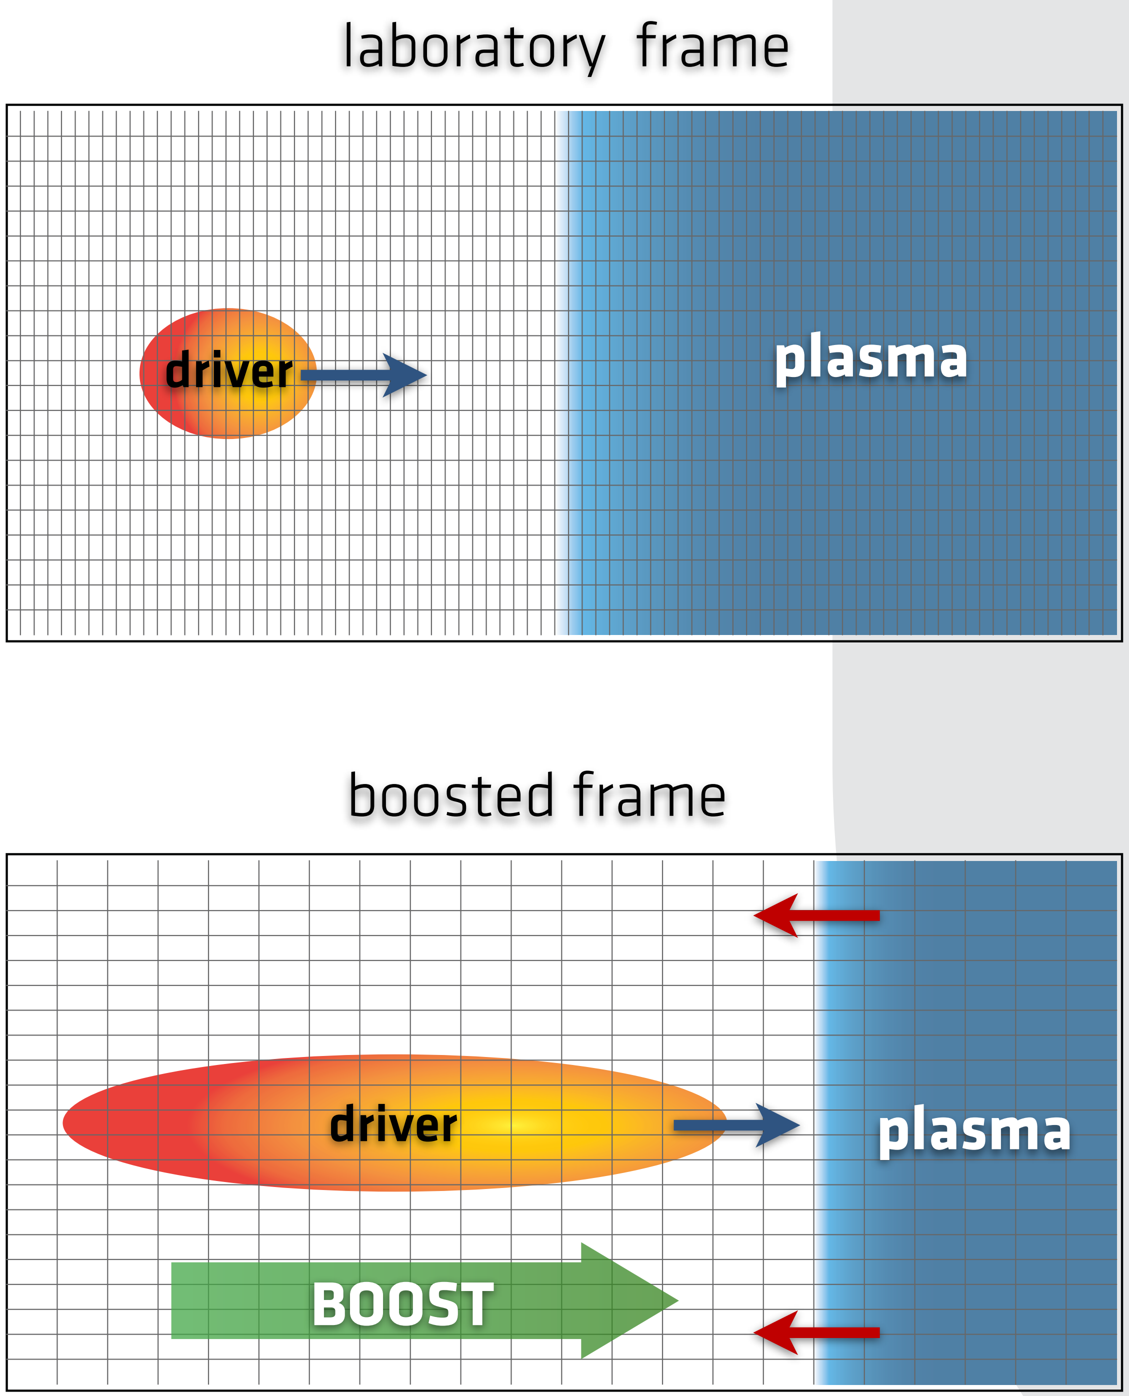

Long distance beam/plasma Interaction

- **Different reference frame**
  - Work in a reference frame moving at γ to reduce scale disparity
  - Driver (laser/beam) is elongated
    - Also lowers requirements on grid resolution
  - Plasma target contracts
    - Shorter propagation distance
- **High performance gains**
  - Lower resolution / plasma contraction
    - Fewer particles/Less iterations
    - Speedup ~ $\gamma^2(1+\beta)^2$
- **Difficult simulations**
  - Numerical Čerenkov
  - Backscattered radiation / self-injection not accurately modeled

*J.-L. Vay, PRL **98,** 130405 (2007).

LWFA in boosted frames

+40 GeV external injection

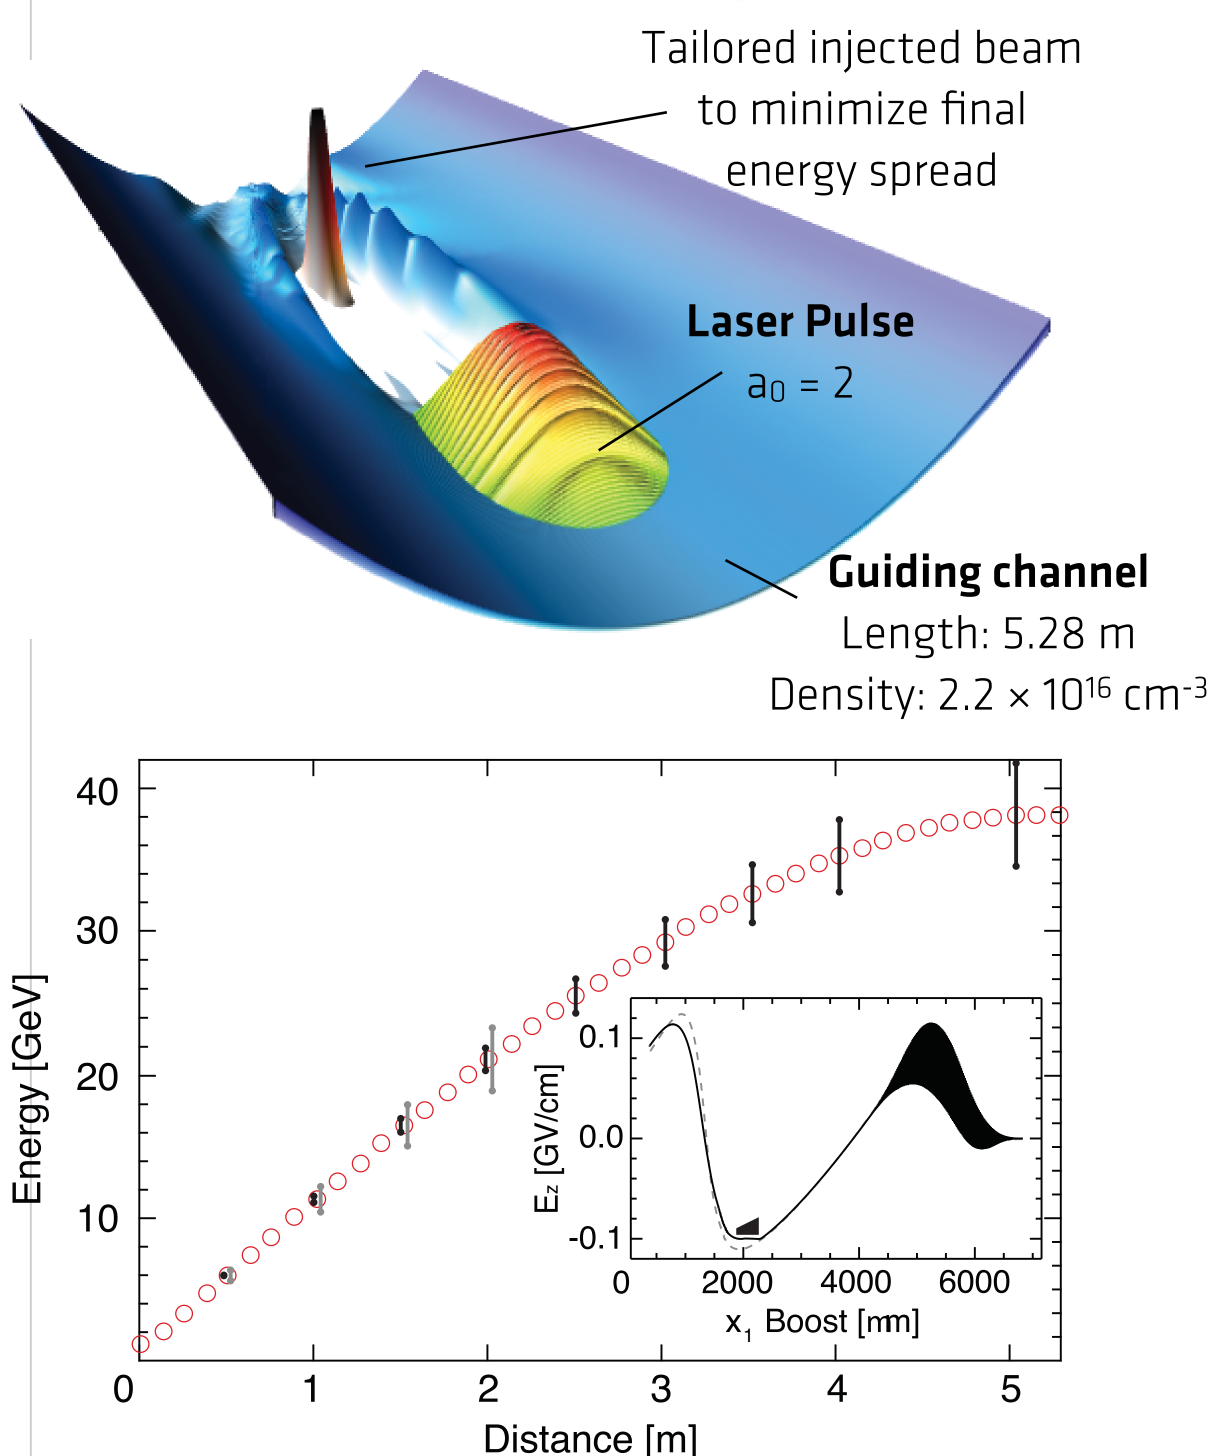

+10 GeV self-injection

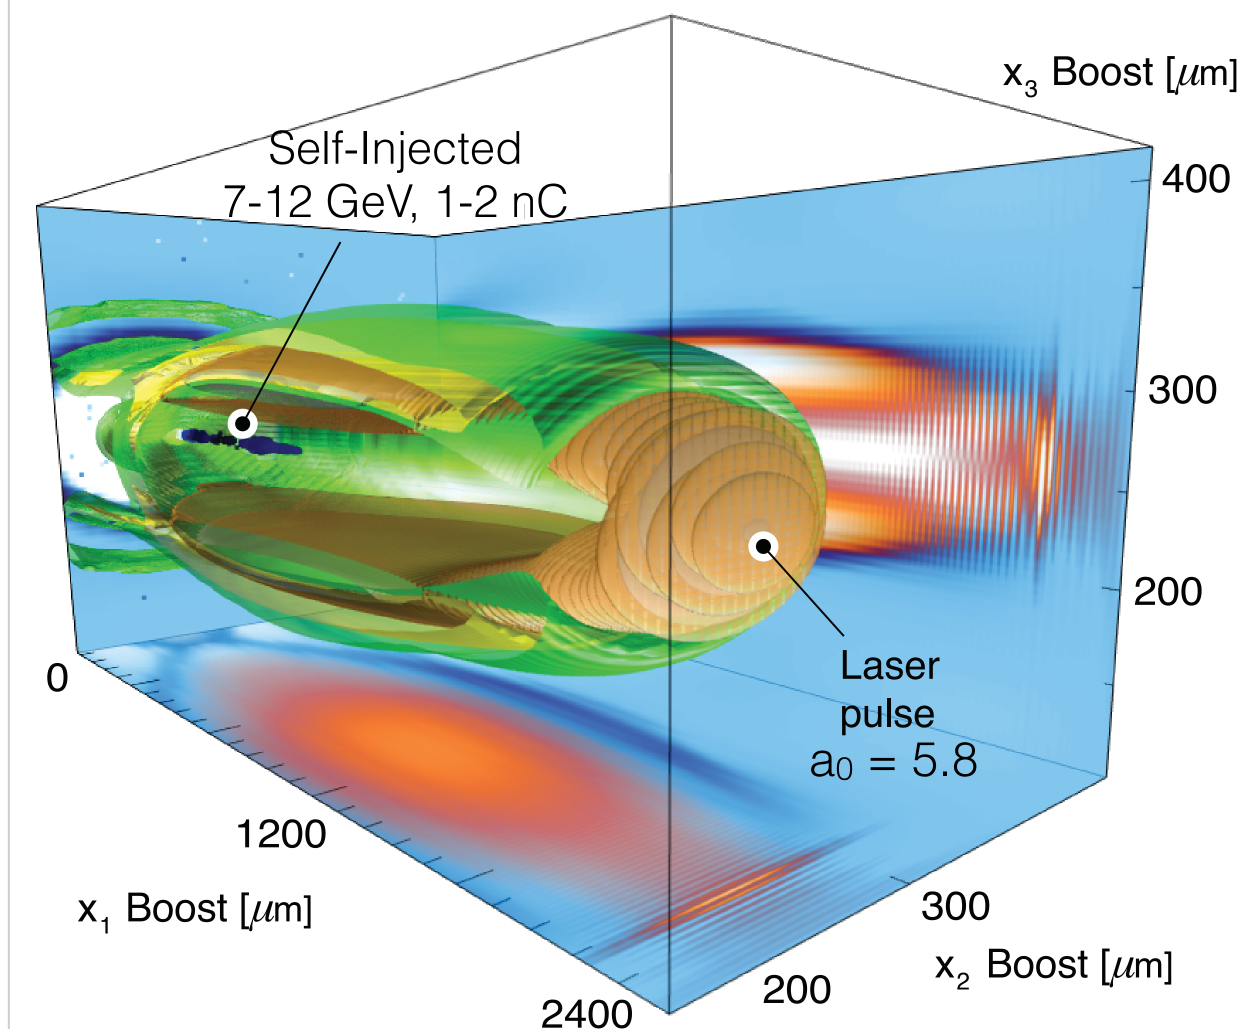

Boost γ = 10

~ 300 × faster than lab. simulation

S. Martins et al., Nat. Phys. **6**(4), 311, (2010)

Speedups up to ~10^5

Boost γ ~ 10^2 - 10^3

- **Control Numerical Noise**
  - Advanced filtering
  - Alternative field solver
    - Finite difference: 4th order, CK, Lehe
    - Spectral
  - “Magic” time step
- **Improved laser injection**
  - Moving antenna
  - Reduce transverse box size
- **Good agreement with theory**
  - Accelerating gradient
  - Energy gain
  - Depletion length

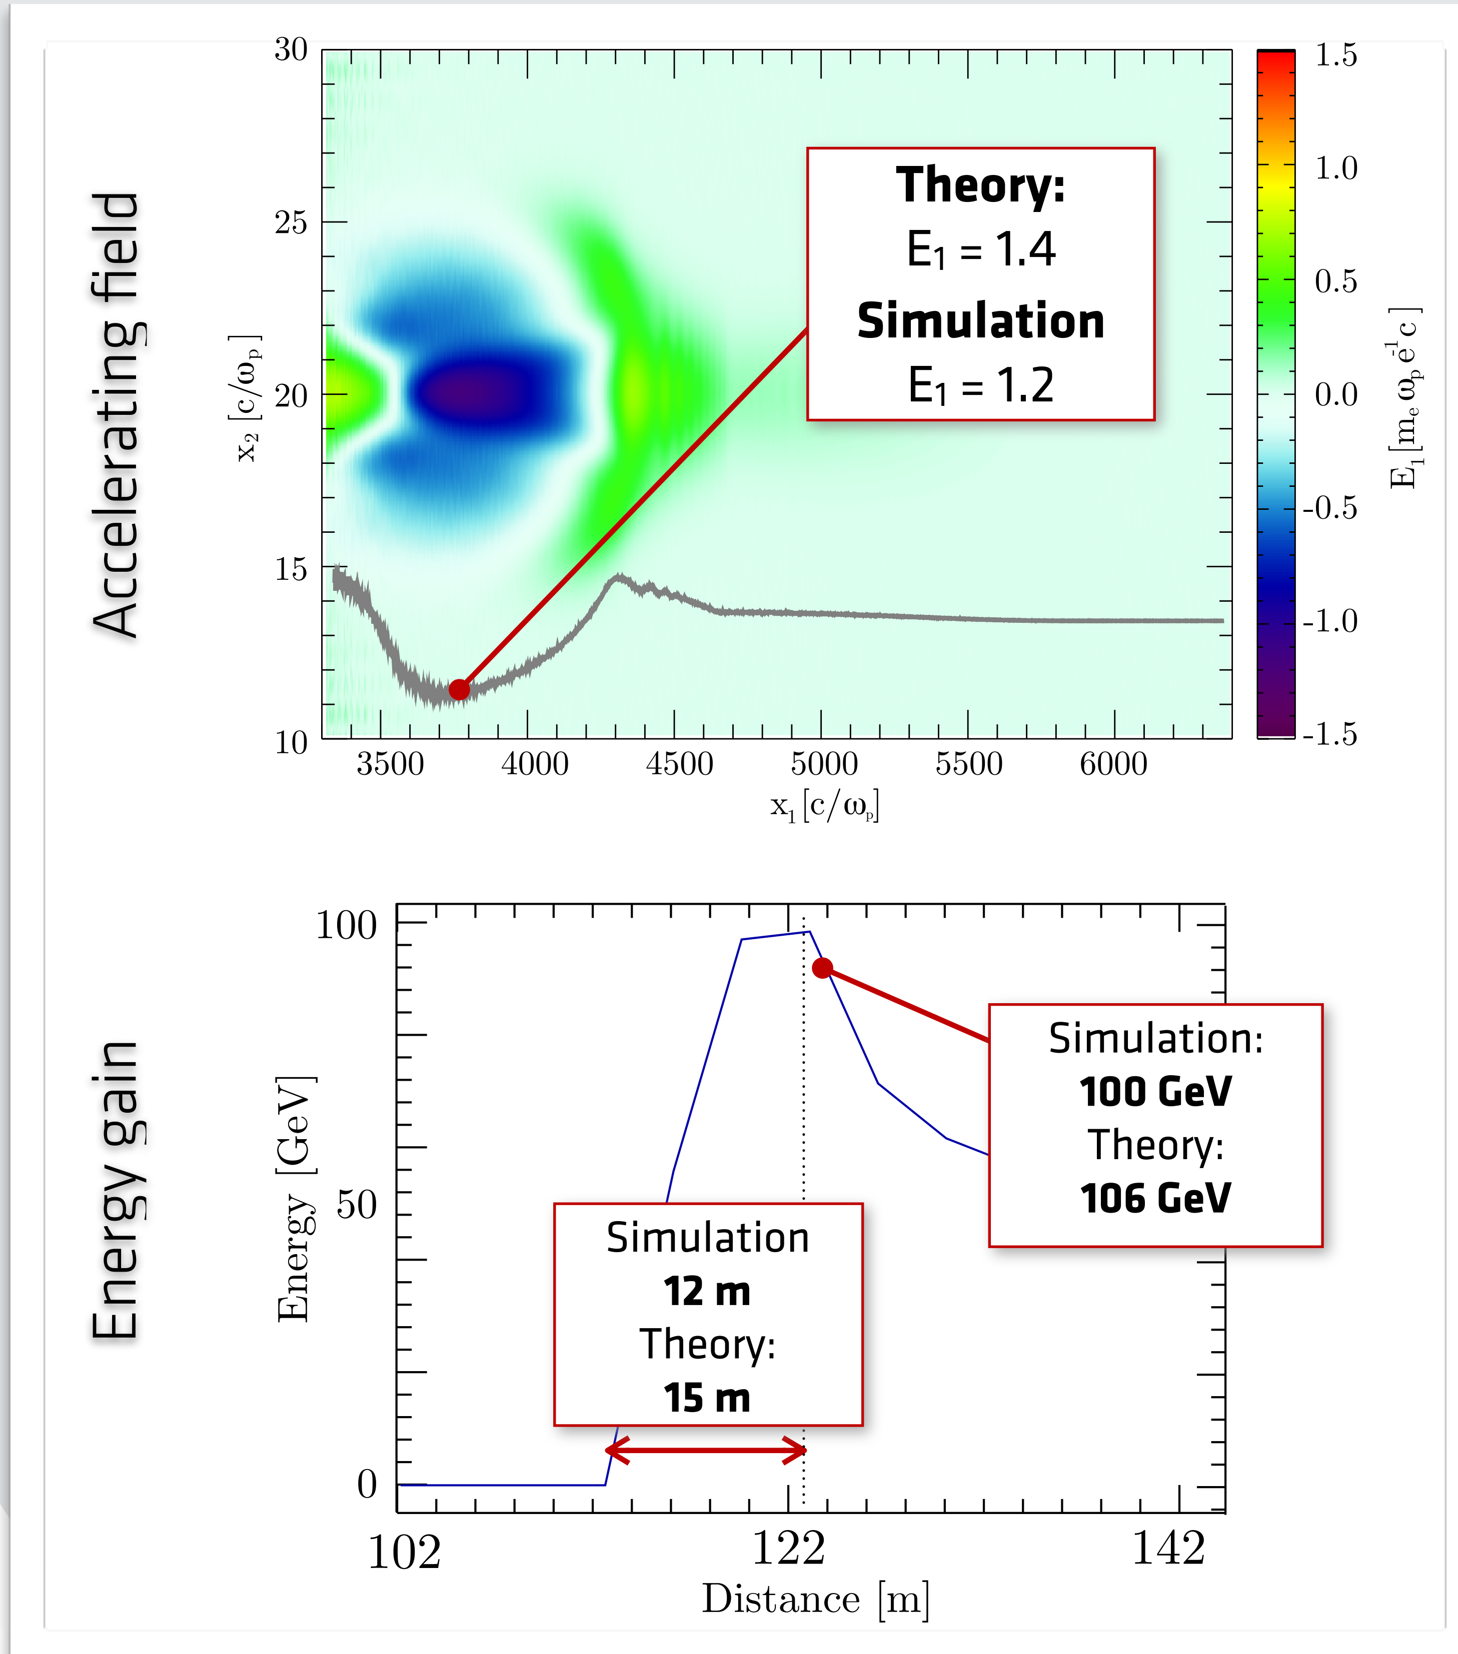

S.Martins et al, Comp. Phys. Comm. **181**(5), 869 (2010)

New hybrid-PIC algorithm for HEDP modeling*

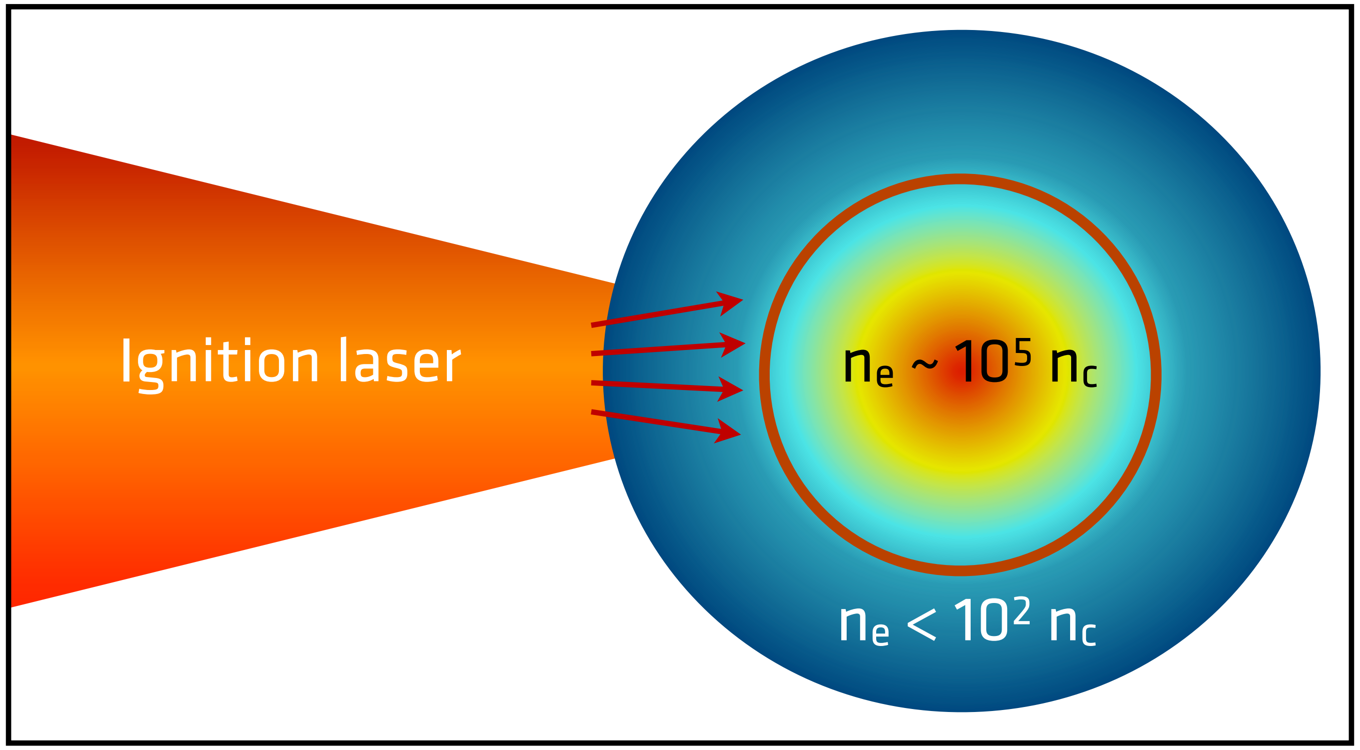

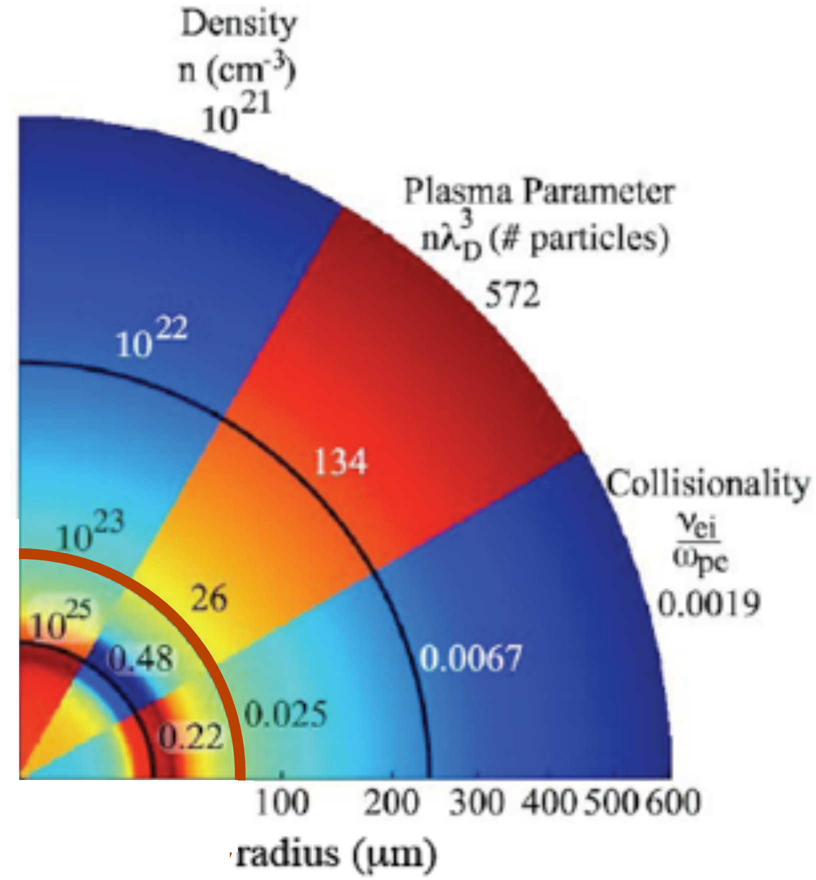

Full-PIC code
• Full Maxwell's equations
• Kinetic species
• $n_0 < 10^{23}$ cm^{-3}
• $\omega_p \Delta t < O(1)$
• $\Delta x\ \omega_p / c < O(1)$
• $c\Delta t / \Delta x < 1$

If resistivity (Ohm's law) matches collisional model transition is natural and self-consistent

Hybrid-PIC code
• Maxwell's equations + Ohm's law (inertialess)
• Kinetic species
• $n_0 > 10^{23}$ cm^{-3}
• $\nu_{ei} \Delta t < O(1)$
• $c\ \Delta t / \Delta x < 1$

* B. Cohen, A. Kemp, and L. Divol, JCP **229**, 4591 (2010), F. Fiuza et al., PPCF **53**, 074004 (2011)

Ion acceleration from laser-solid interactions

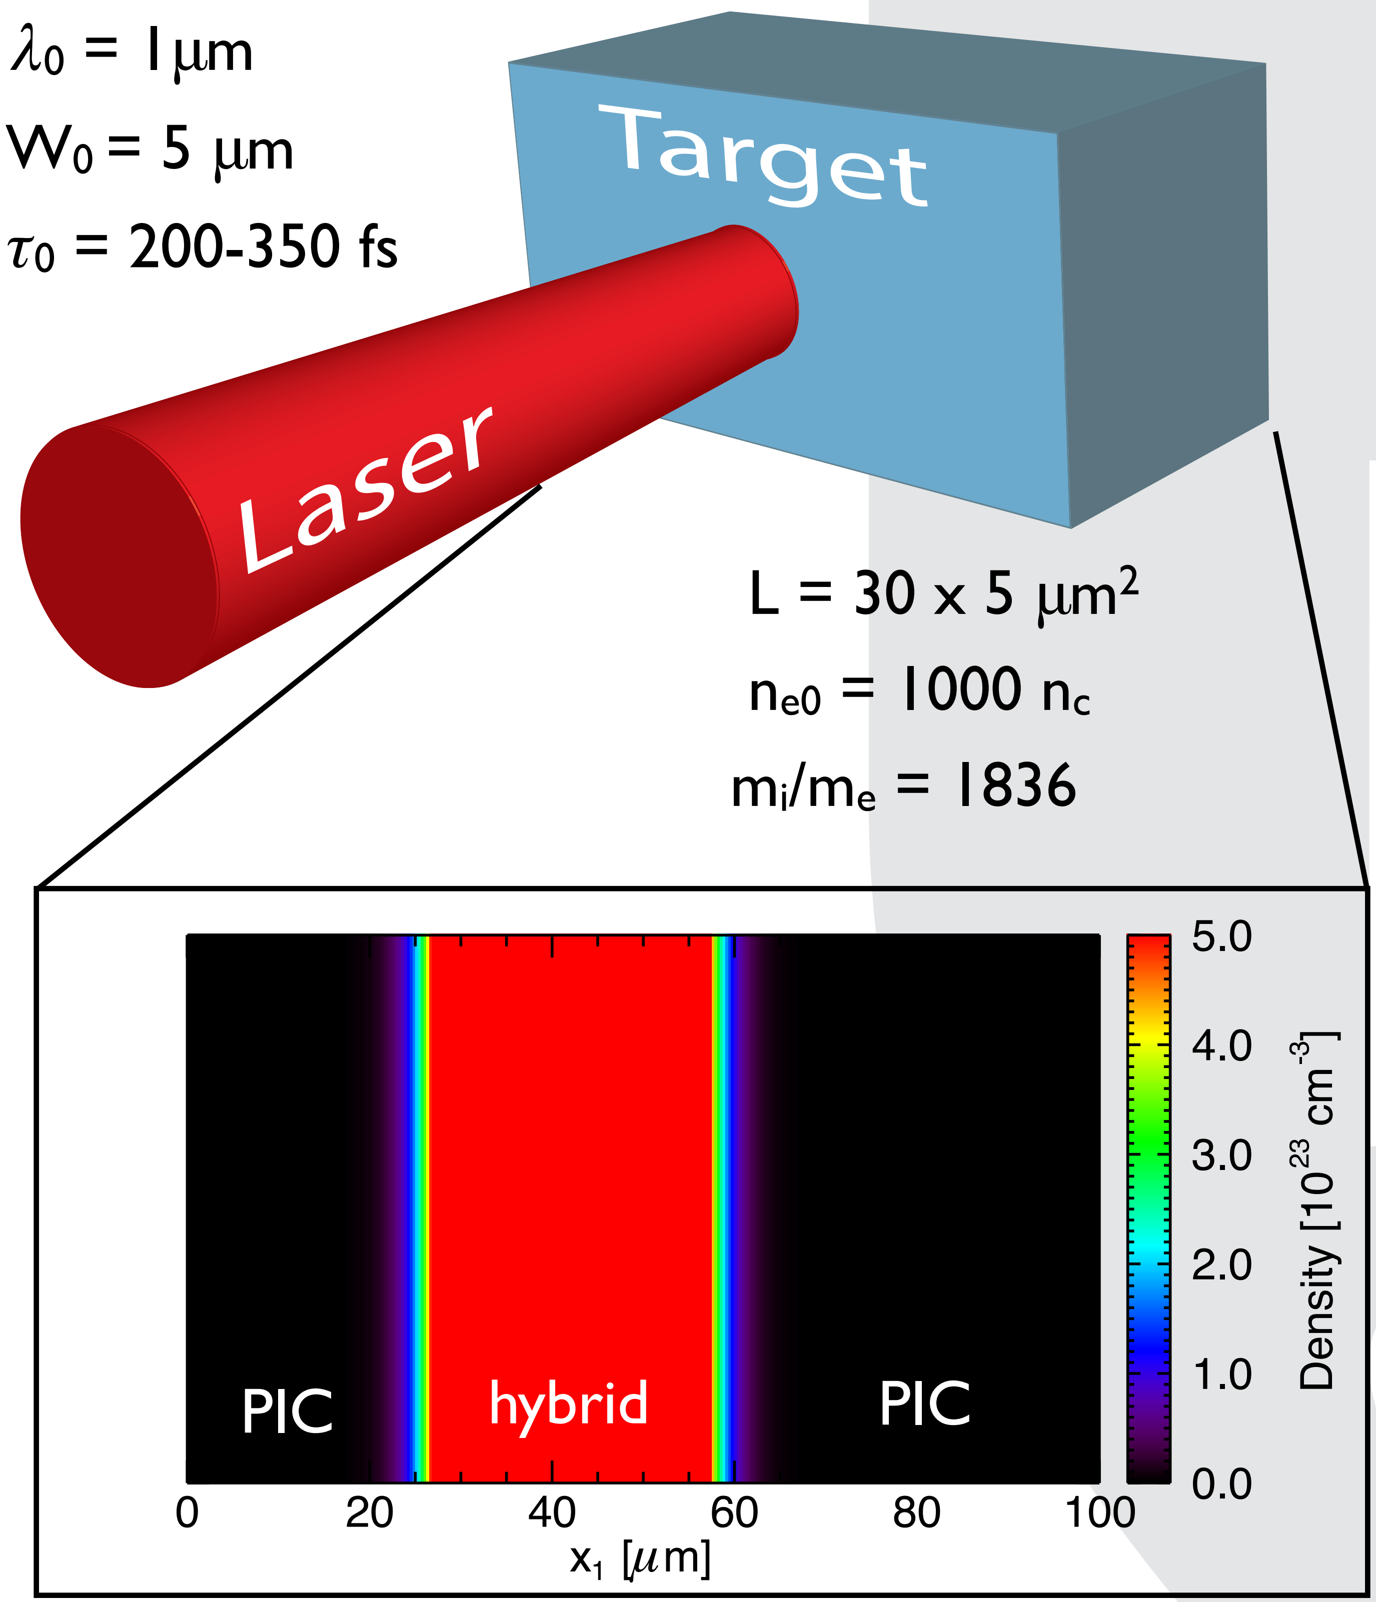

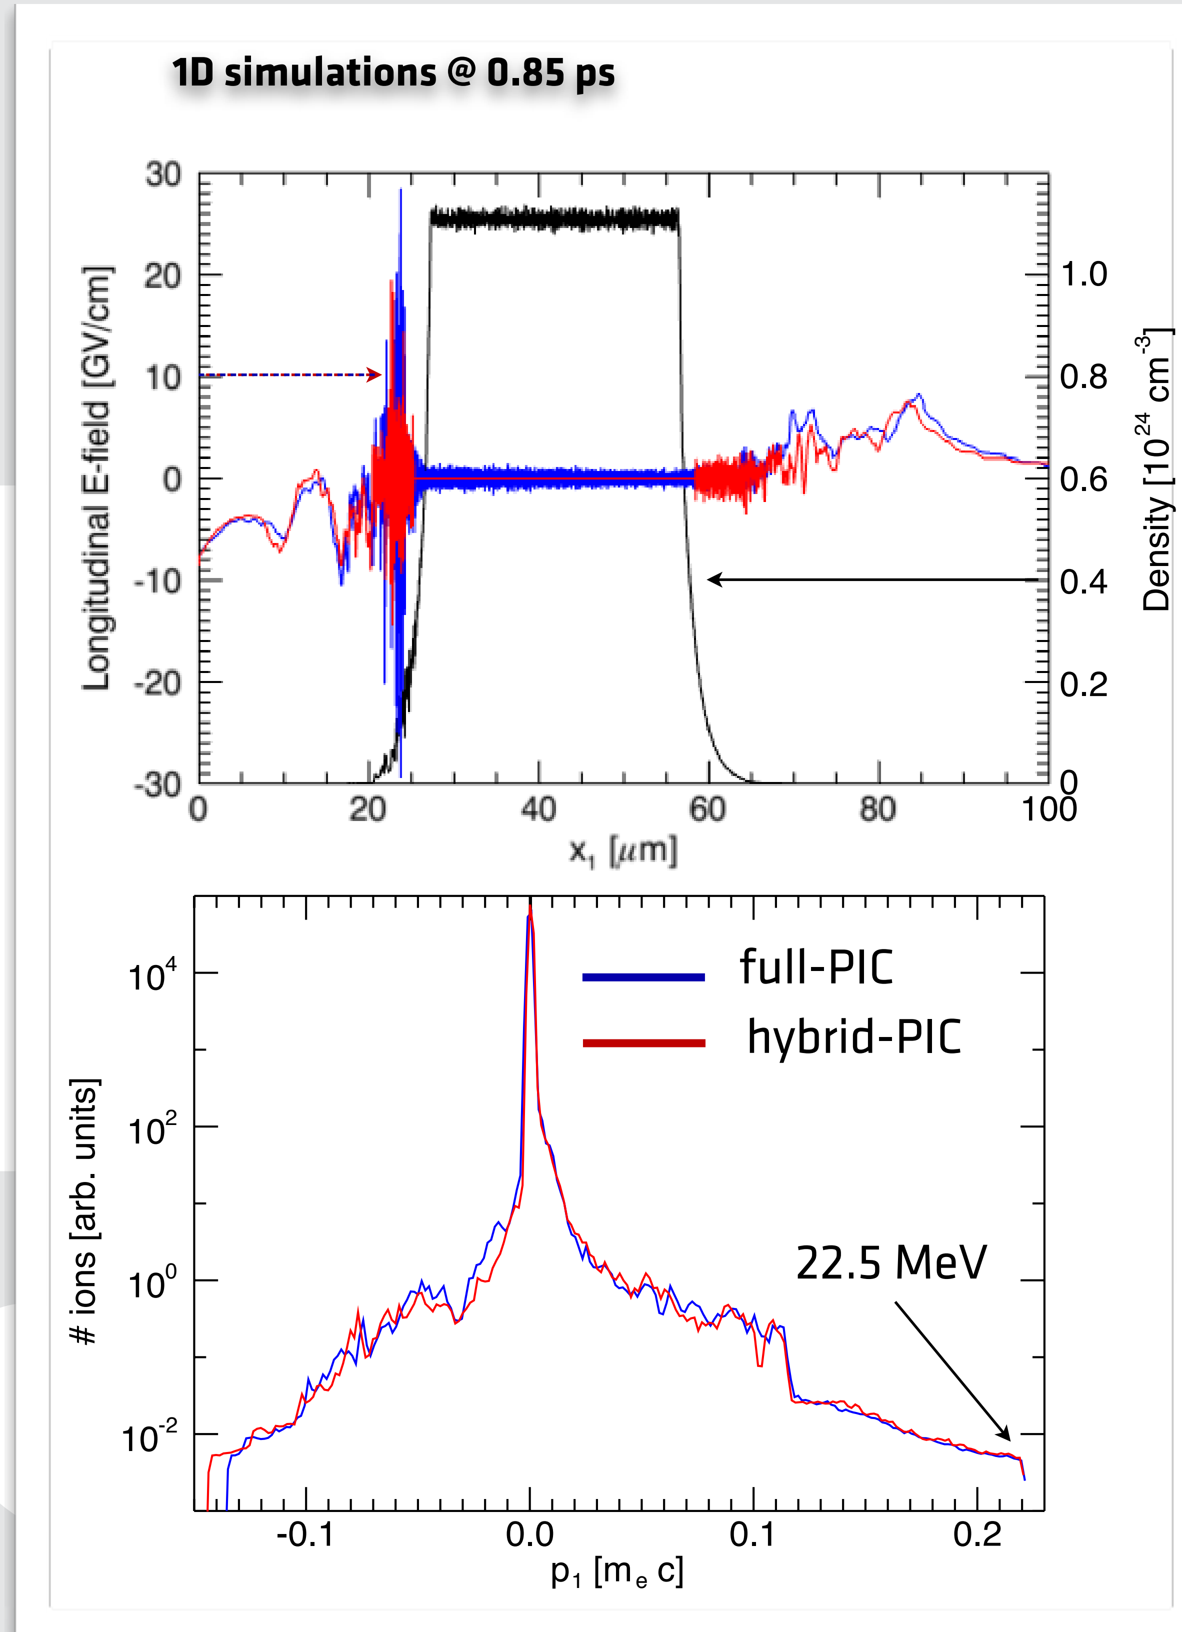

Monte Carlo Coulomb collisions modeled in all simulations

High speedups and accuracy

- New algorithm allows for smooth hybrid - PIC transition and accurate results
- Very large speedups can be achieved while maintaining the accuracy
  - 1D - 90× speedup
  - 2D - 300× speedup
- FI runs can achieve ~10^5 speedups

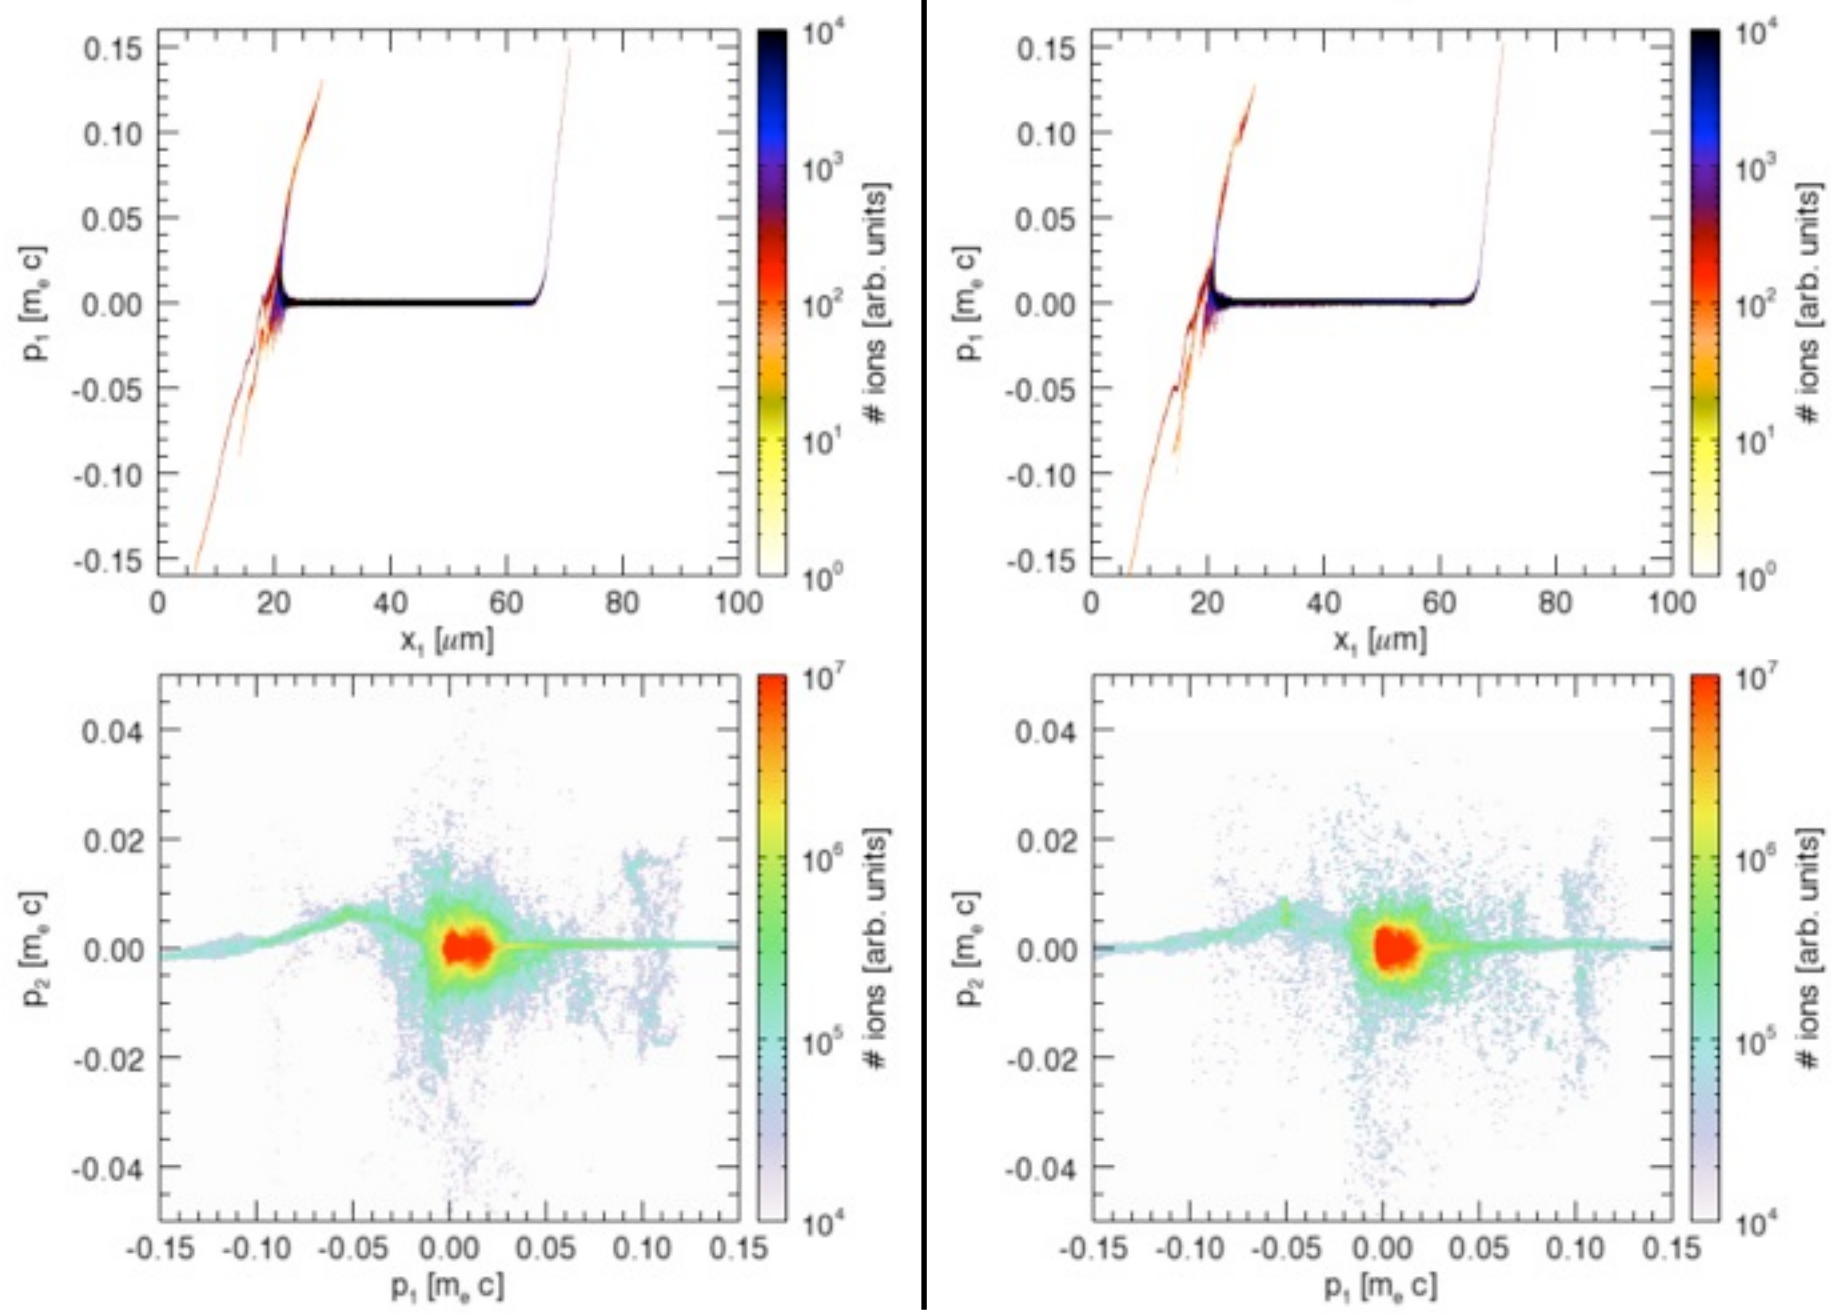

2D simulations @ 0.50 ps

F. Fiuza et al., PPCF **53**, 074004 (2011)

Quasi-axissymetric PIC for LWFA*

- Use cylindrical coordinates for the fields
- Do a Fourier decomposition in the poloidal direction
- The physics of LWFA is quasi-axisymmetric
  - Only a few modes needed, higher modes can be neglected
  - Only a few "2D" modes are calculated
    - faster simulations
- Mode *m*=0
  - No dependence on θ
  - Correspond to a 2D cylindrical case
- The laser field is described on mode *m*=1
  - Only 2 modes are enough to describe the laser field and axisymmetric wakefield

Calder-Circ

$$\mathbf{E}(r,\theta,x) = E_r(r,\theta,x)\mathbf{e_r} + E_\theta(r,\theta,x)\mathbf{e_\theta} + E_x(r,\theta,x)\mathbf{e_x}$$

$$F(r,\theta,x) = \sum_{m=-\infty}^{+\infty} \hat{F}^m(r,x)e^{-im\theta}$$

$$\hat{F}^m(r,x) = \frac{1}{2\pi}\int_0^{2\pi} F(r,\theta,x)e^{im\theta}d\theta$$

$$\begin{aligned}\mathbf{E}(r,\theta,x) &= a_0(r,x)\mathbf{e_y} = a_0(r,x)\cos(\theta)\mathbf{e_r} - a_0(r,x)\sin(\theta)\mathbf{e_\theta}\\ \mathbf{B}(r,\theta,x) &= a_0(r,x)\mathbf{e_z} = a_0(r,x)\sin(\theta)\mathbf{e_r} + a_0(r,x)\cos(\theta)\mathbf{e_\theta}\end{aligned}$$

*A. Lifschitz et al, JoCP **228**, 1803 (2009)

Similar results to full 3D

Calder-Circ

- $\lambda_0 = 0.8\ \mu m$
- $\tau_0 = 30$ fs
- $w_0 = 9\ \mu m$
- $a_0 = 5$
- $n_e = 0.007\ n_c$

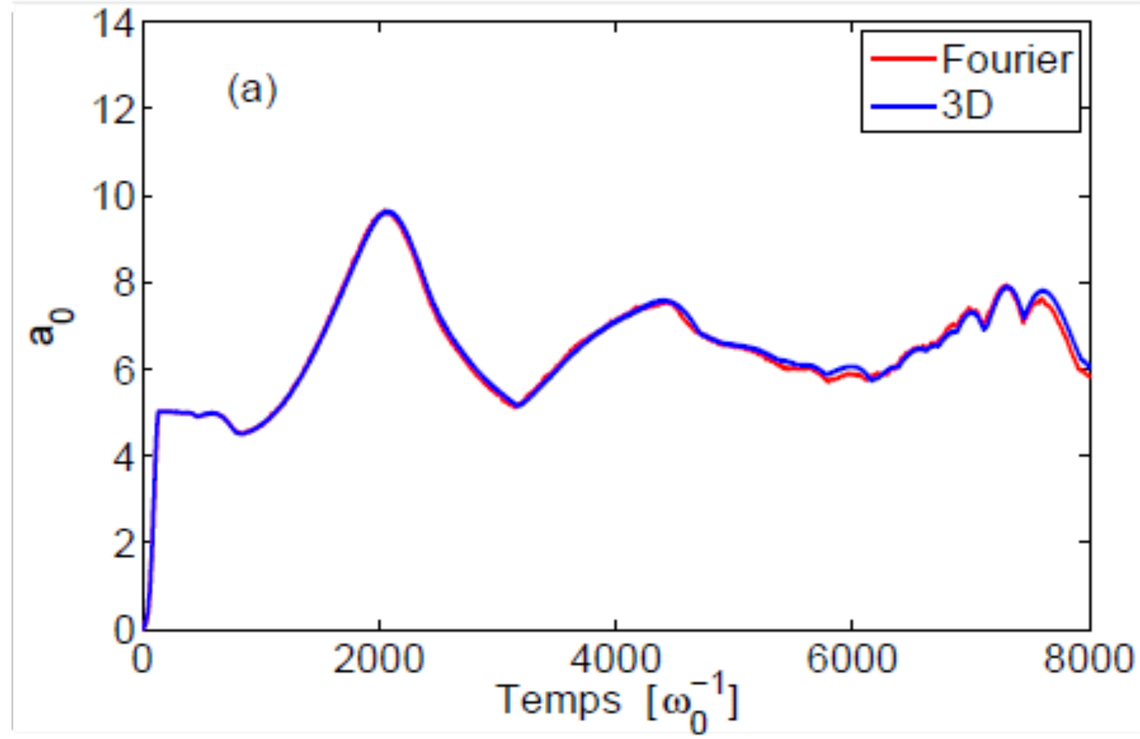

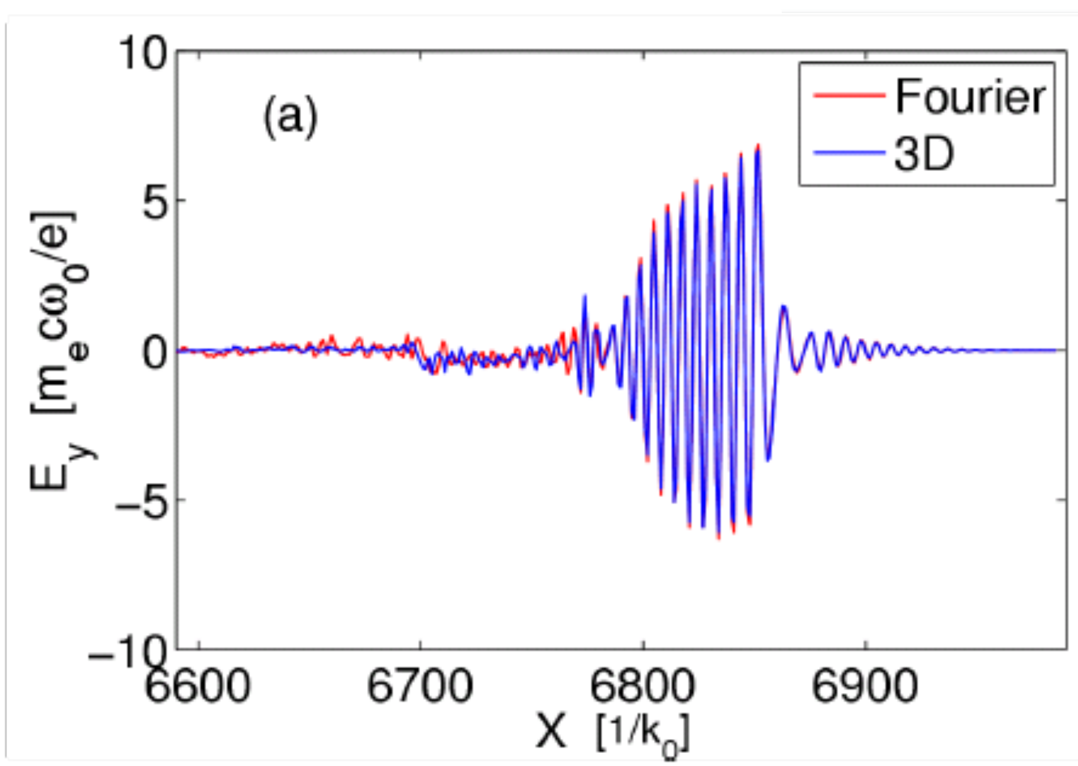

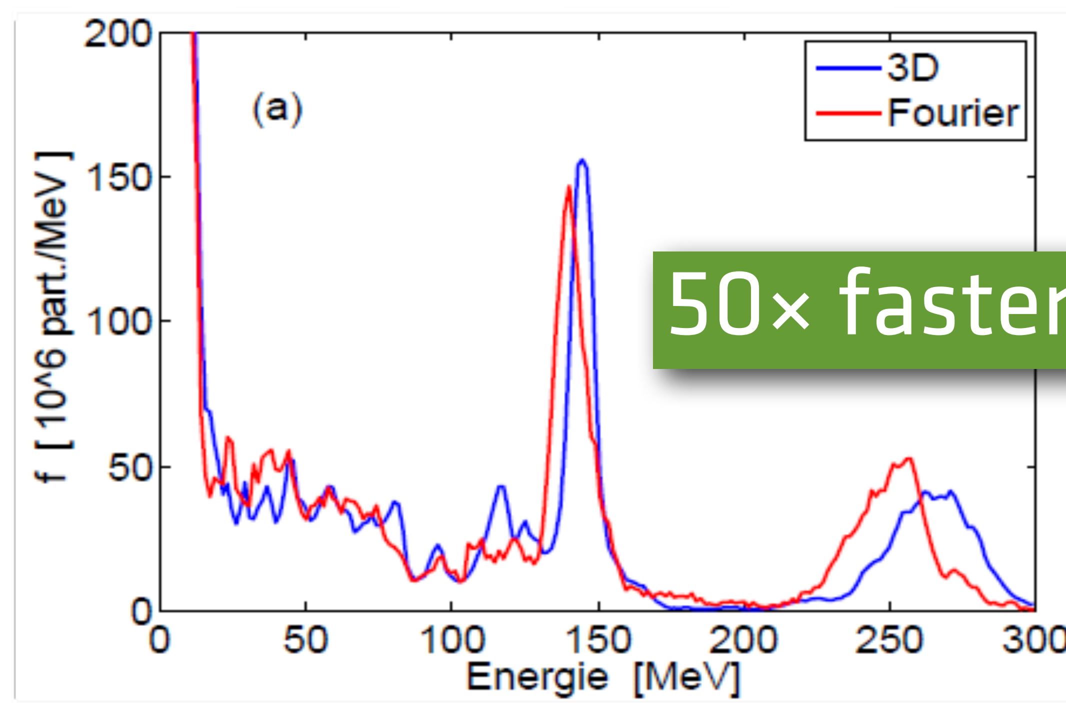

Outline

- Modeling Plasma Based Accelerators
  - The OSIRIS framework
  - Algorithm
  - Computational Requirements
- Extending the algorithm
  - Lorentz boosted frames
  - Ponderomotive guiding center
  - High density hybrid
  - LWFA in cylindrical geometry
- Scaling to Petascale Systems
  - Multi-level parallelism
  - Multi-dimensional dynamic load balance
- Plasma Acceleration on Tier-0 systems
  - Scalings @ Sequoia & BlueWaters
  - 20 m long beam-plasma interaction @ SuperMUC
  - High-res 1 GeV LWFA @ Jaguar
  - Full scale target LWFA modeling @ SuperMUC
- Overview

Achieving Petascale Performance

Multiscale Parallelism

- Modern HPC systems present a hierarchy of parallelism
  - At the highest level they are a network of computing nodes
  - Each node is a set of CPUs / cores (GPUs/MICs) sharing memory inside the node
  - Most processing cores have a vector SIMD unit (AMD, Intel, Power7)
- **Efficient HPC system use requires taking advantage of all these levels of parallelism**

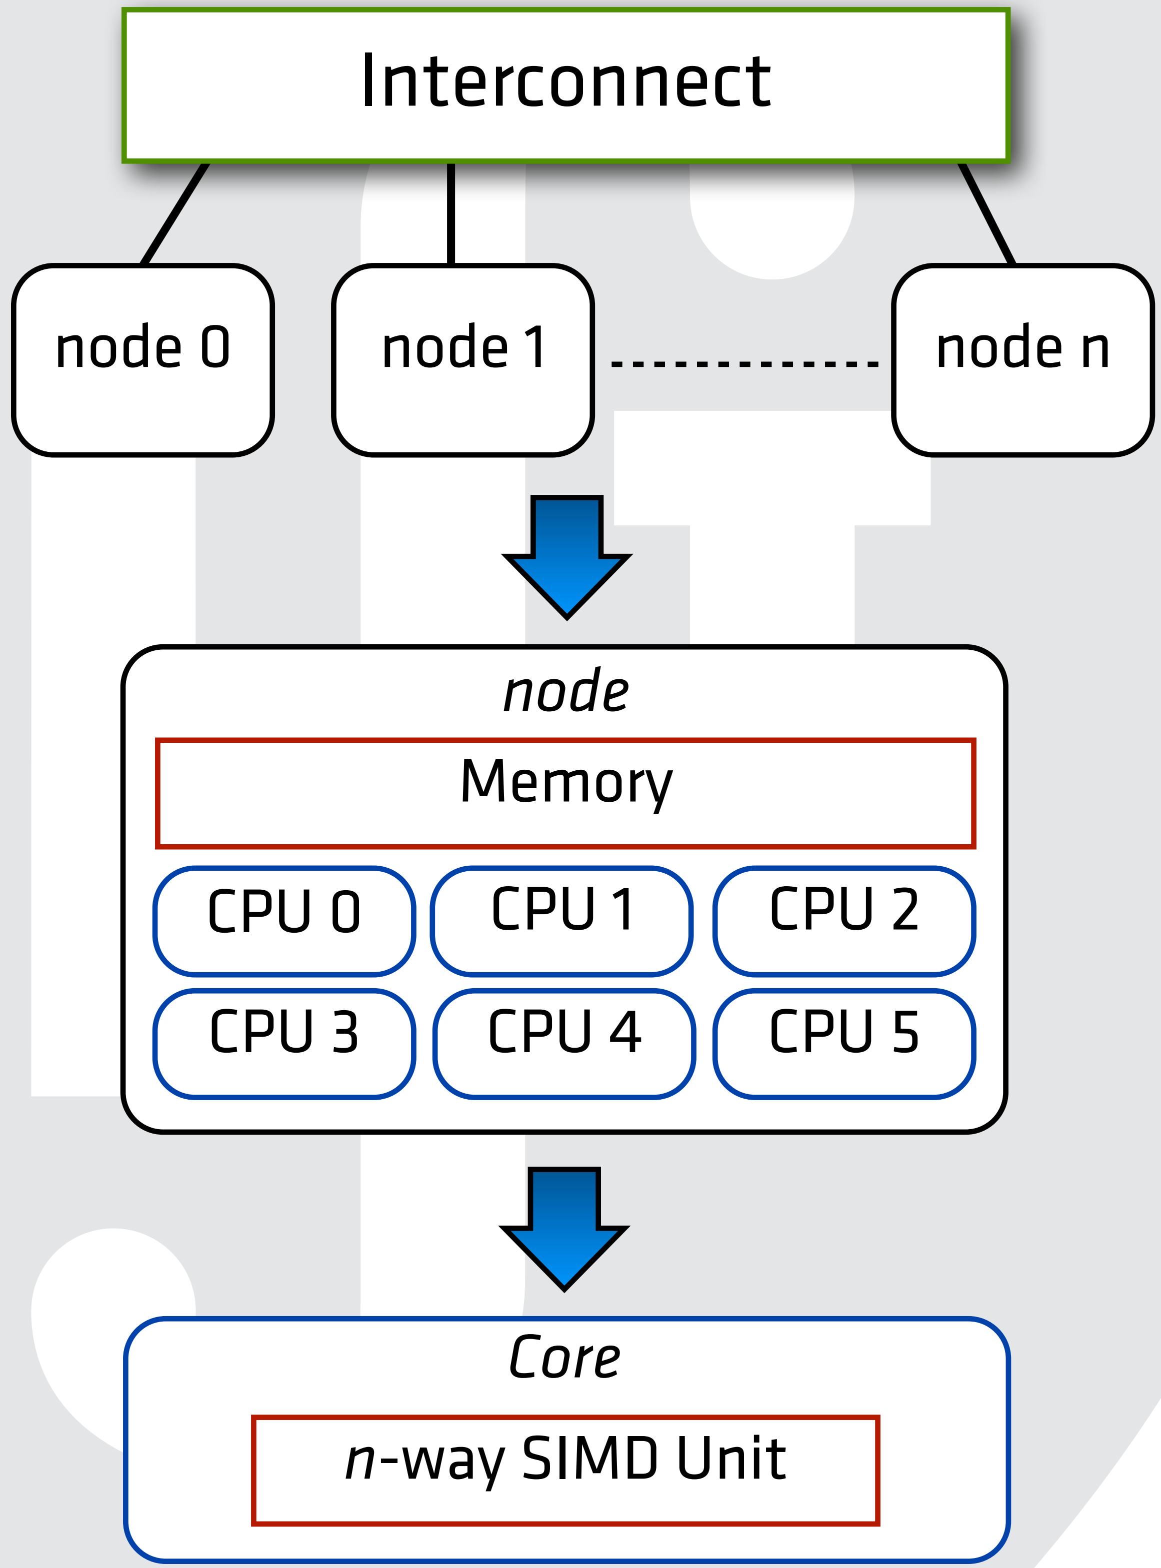

HPC system level parallelism

Distributed memory

- Each process cannot directly access memory on another node:
  - Information is exchanged between nodes using network messages (MPI)
- Standard parallelization uses a spatial decomposition:
  - Each node handles a specific region of simulation space
- Works very well also on multi-core nodes
  - Message passing inside a node is very efficient
- Very efficient for uniform plasmas
- Susceptible to imbalance if particles accumulate in a subset of nodes

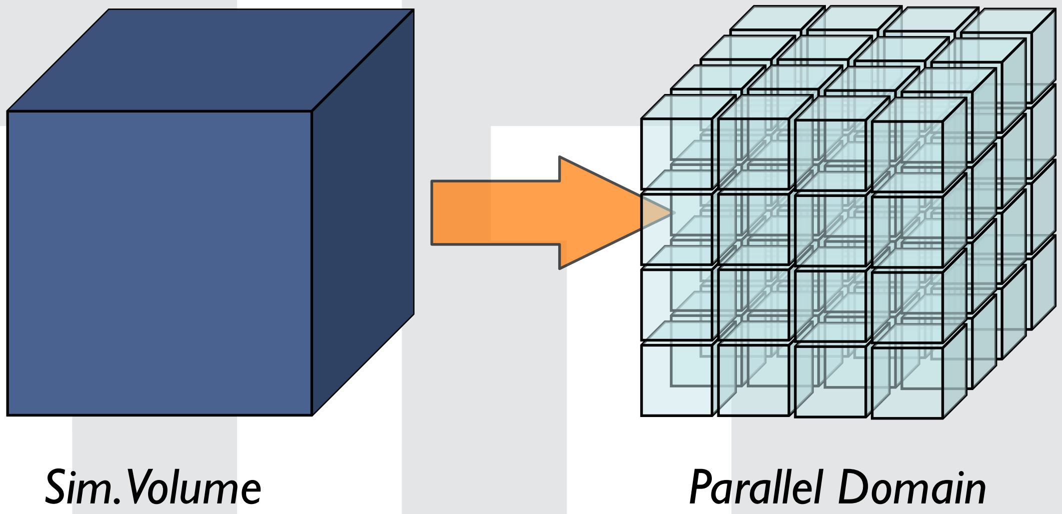

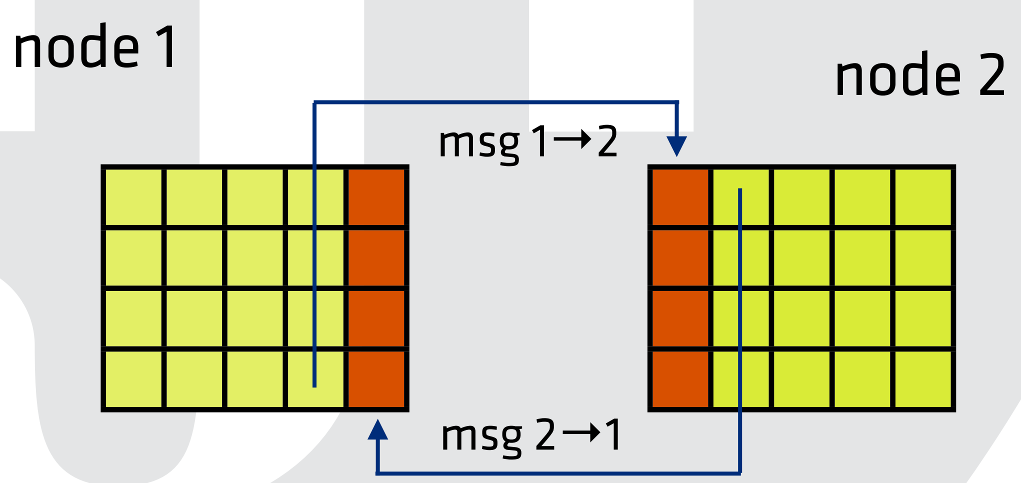

Load Imbalance

Imbalance Problem

- For non-uniform plasmas many particles may accumulate in a single process
- Processes don't share memory access and only communicate through messages

- This process with the most particles will dominate overall simulation time

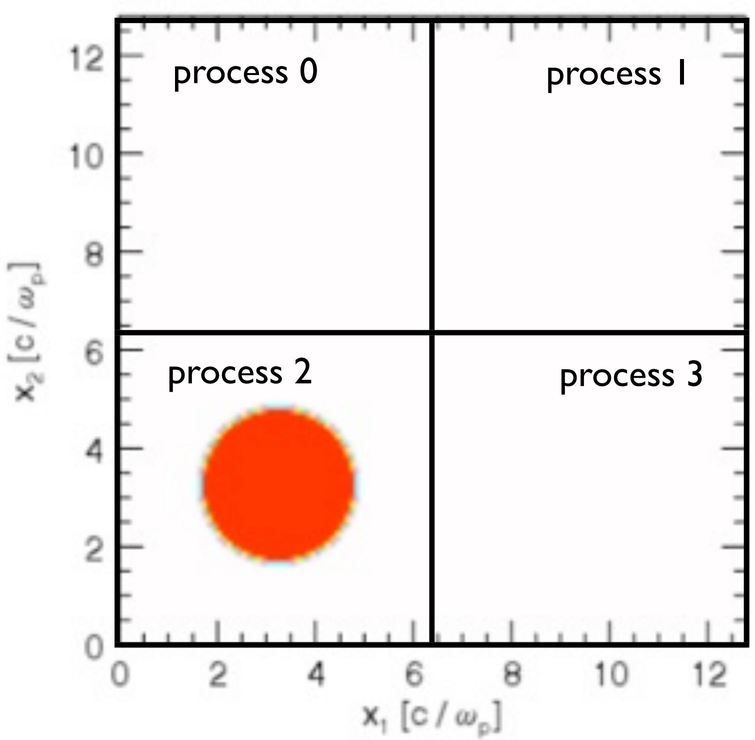

All work handled by process 0 4 cores

Dynamic Load Balance

- The code can change node boundaries dynamically to attempt to maintain a even load across nodes:
  - Determine best possible partition from current particle distribution
  - Reshape the simulation

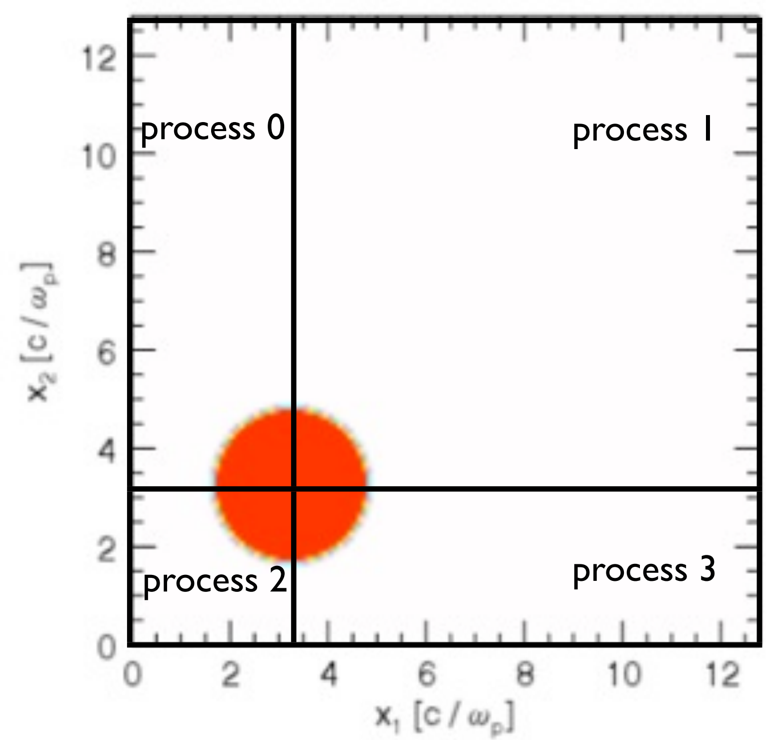

Even particle load across processes 4 cores

Multi-Dim. Dynamic Load Balance

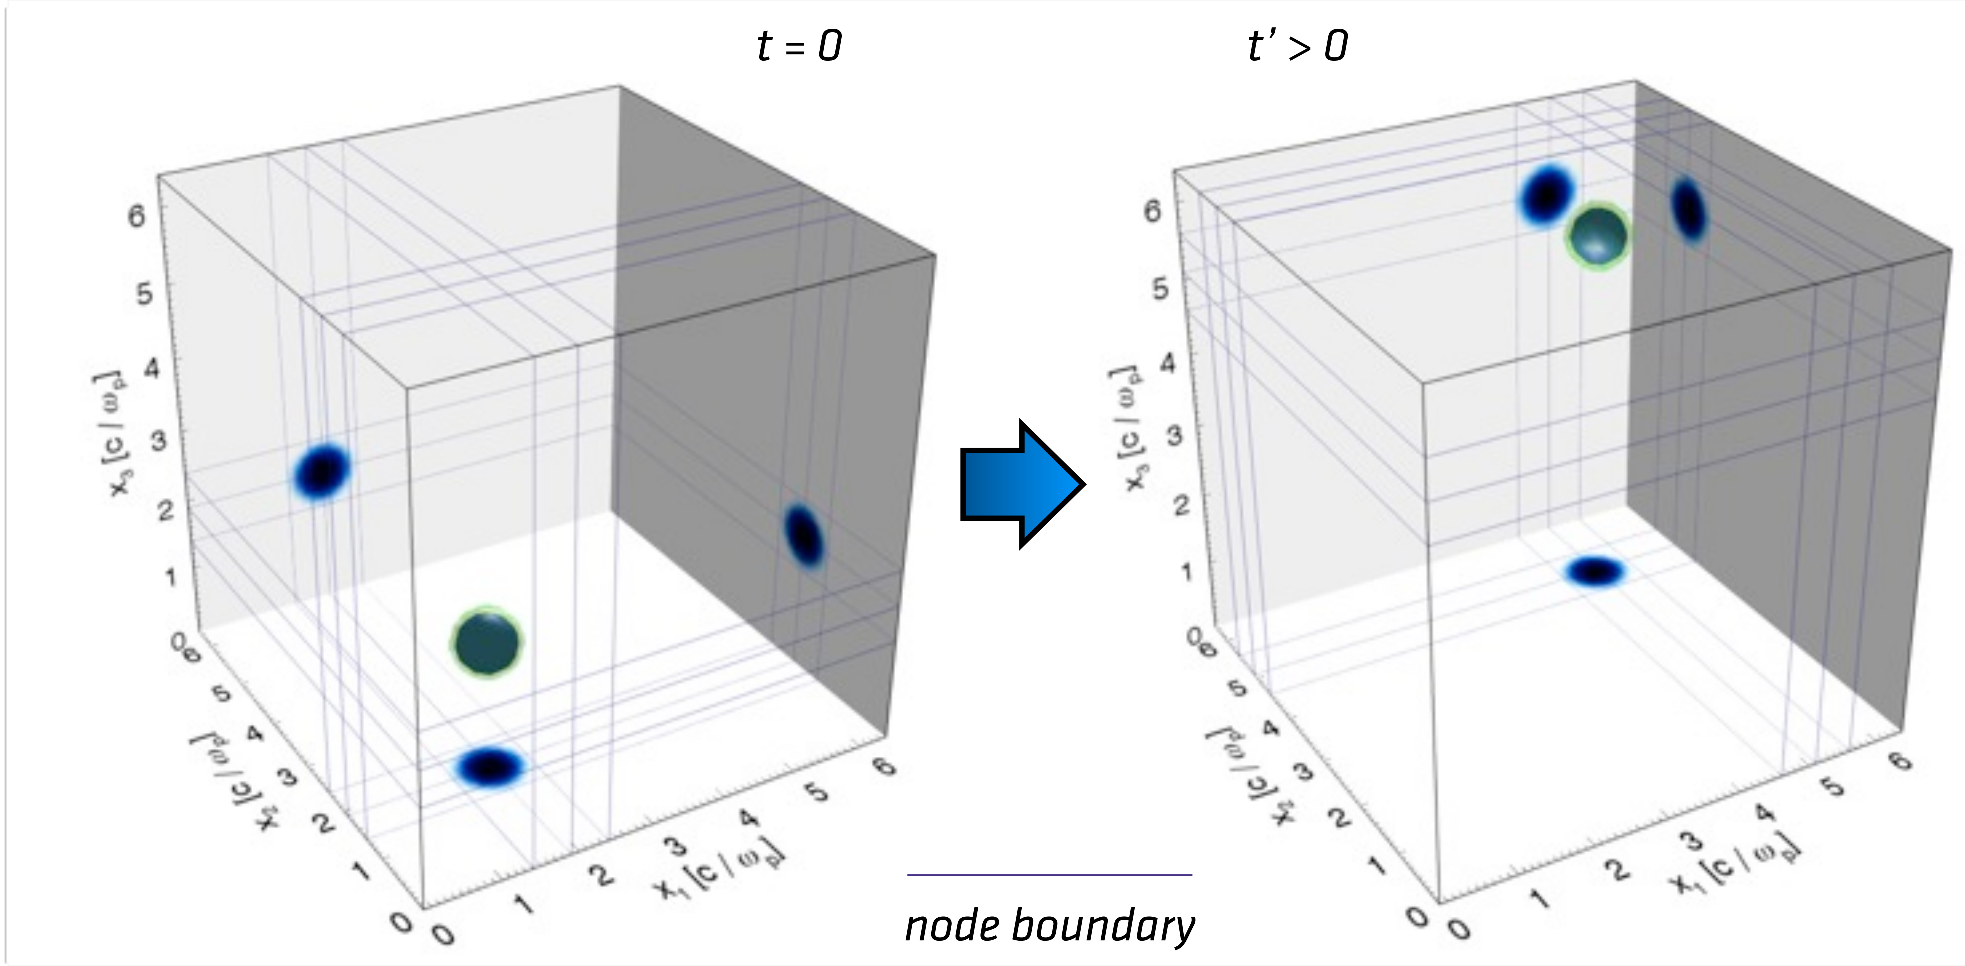

Best Partition

- Difficult task:
  - Single solution exists only for 1D parallel partitions
  - Improving load balance in one direction might result in worse results in another direction

Multidimensional Partition

- Assume separable load function (i.e. load is a product of $fx(x) \times fy(y) \times fz(z)$)
  - Each partition direction becomes independent
- Accounting for transverse directions
  - Use max / sum value across perpendicular partition

CPU / node level parallelism

Imbalance Problem

- For non-uniform plasmas many particles may accumulate in a single process
- Processes don't share memory access and only communicate through messages
- This process with the most particles will dominate overall simulation time

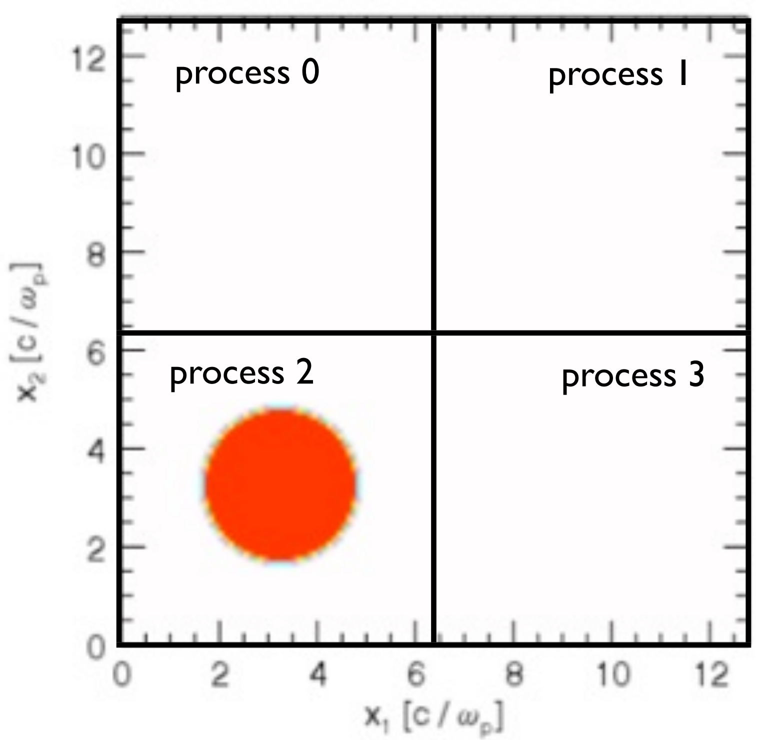

4 cores

Using Shared Memory

- A group of cores will share memory access / simulation region:
  - Distribute particles evenly across these cores
- The workload for these cores will always be perfectly balanced
- Implement using OpenMP
  - Simple API for shared memory parallelism
  - Use 1 thread / core

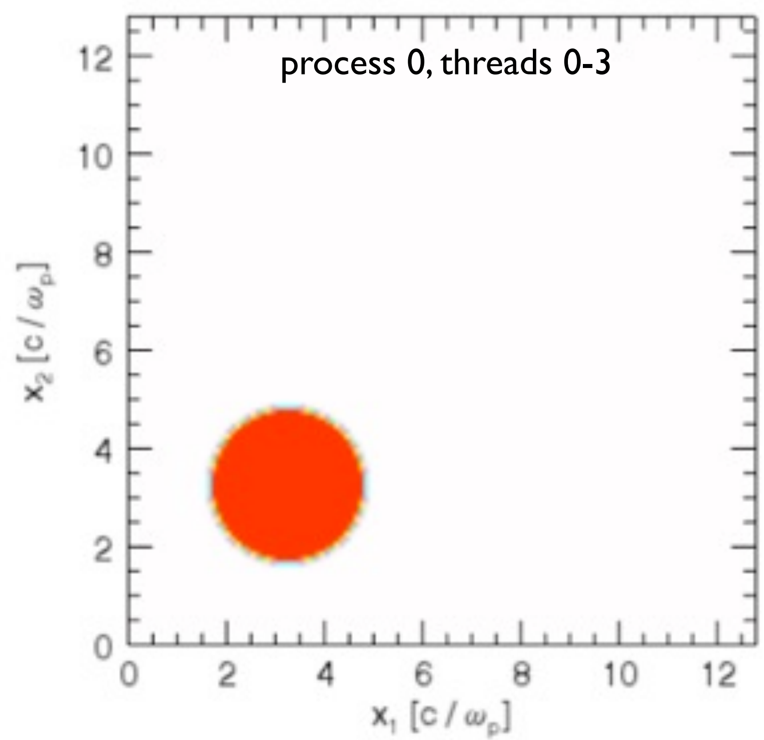

4 cores

SMP particle pusher

Algorithm

- An MPI process handles a given simulation region
- This process spawns n_T threads to process the particles:
  - Use n_T copies of the current grid
  - Divide particles evenly across threads
  - Each thread deposits current in only 1 of the grid copies
- Accumulate all current grids in a single current grid
  - Divide this work also over *nt* threads
- Algorithm overhead
  - Zeroing the additional grid copies
  - Reduction of multiple current grids

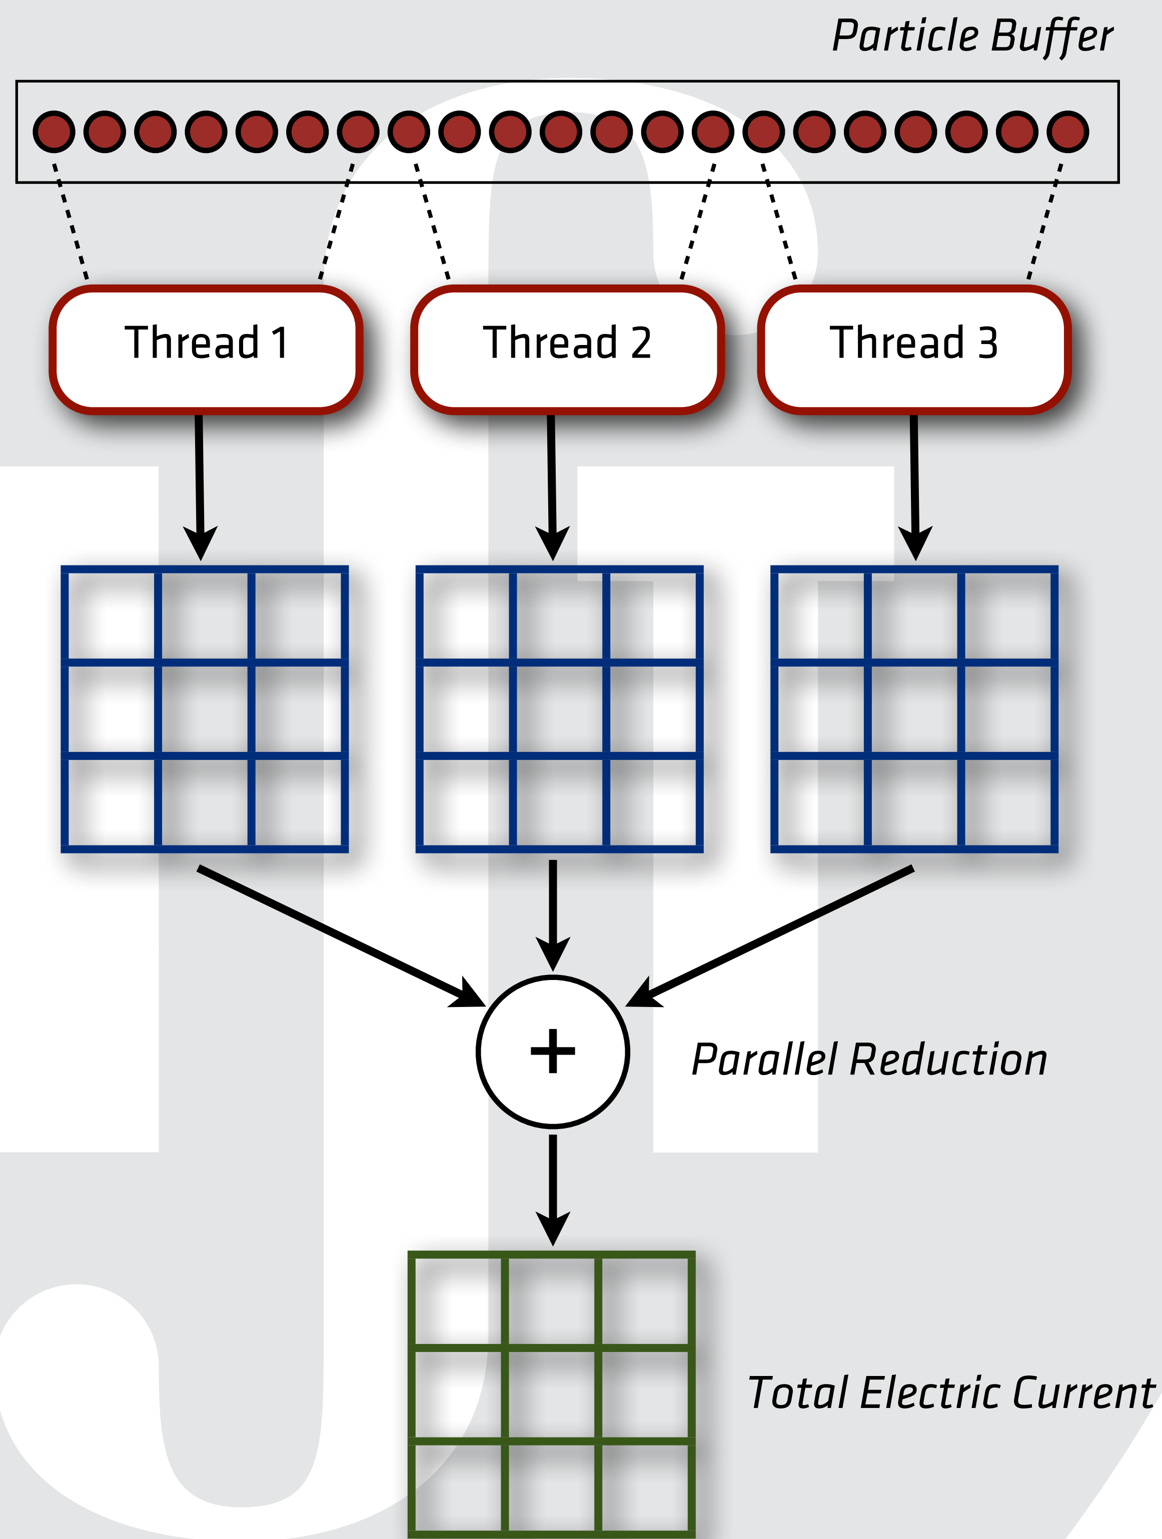

SMP Algorithm Performance

Simulation Time

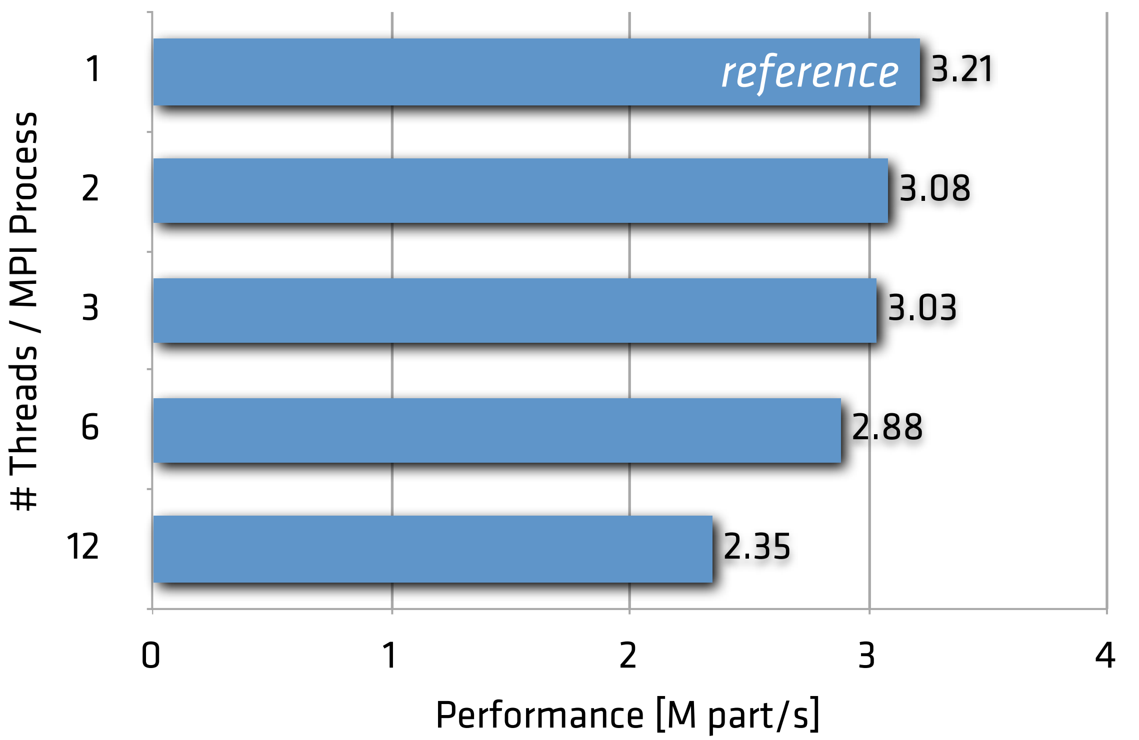

- Additional sections were parallelized:
  - Boundary processing of particles / fields
  - Particle sorting (rearrange stage)
- Lower node communication volume
- Only 10% slowdown for 6 threads / MPI process from reference MPI parallelization
- Large Drop from 6 to 12 threads / MPI process

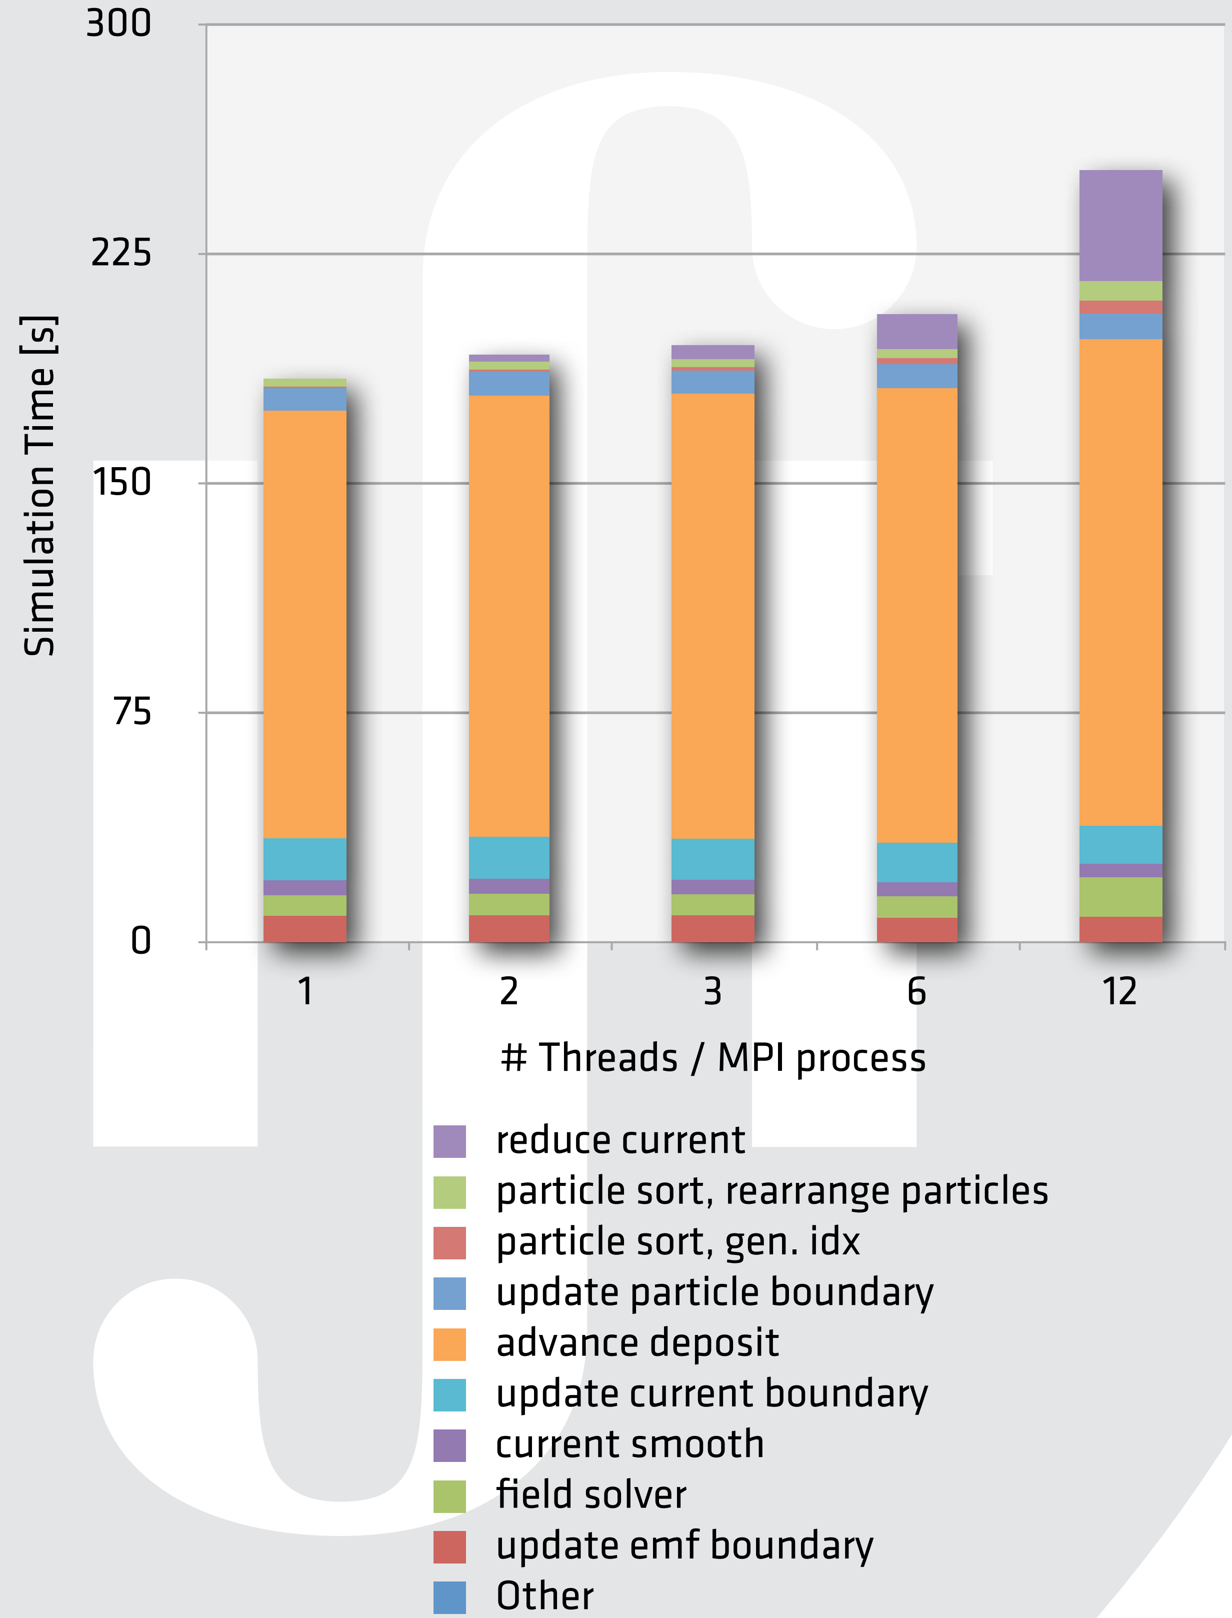

Core Level Parallelism

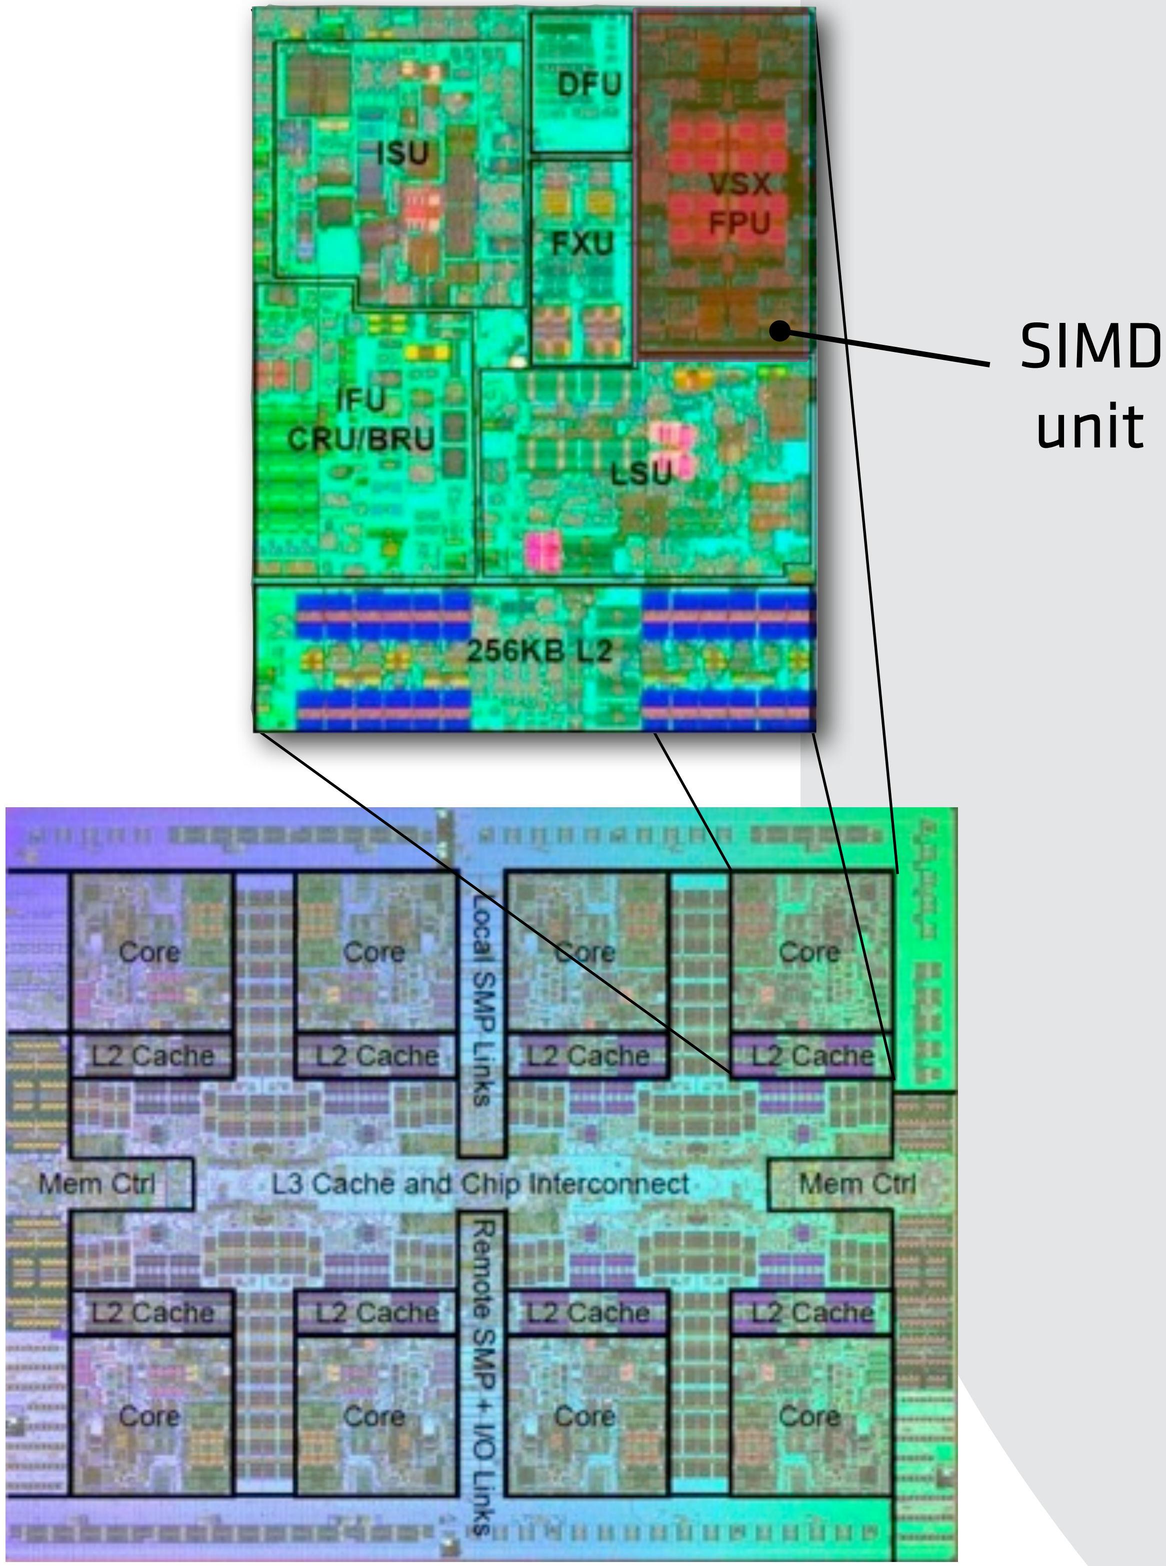

IBM Power 7 CPU

Single Instruction Multiple Data (SIMD)

- Modern CPUs (Intel / AMD / PowerPC) include a SIMD vector unit:
  - Modern cpus include a single instruction multiple data (SIMD) vector unit
  - Vector registers ($n \times 32$ bit int/float)
  - Instructions act on vector registers
  - Same operation on *n* different values simultaneously
- Programming
  - Require ASM or C intrinsics
  - Some compilers will generate SIMD code automatically
- System Architecture
  - Same overall architecture
  - Cache based memory system
  - More data being processed per cycle:
  - Optimize memory access

SIMD Algorithm

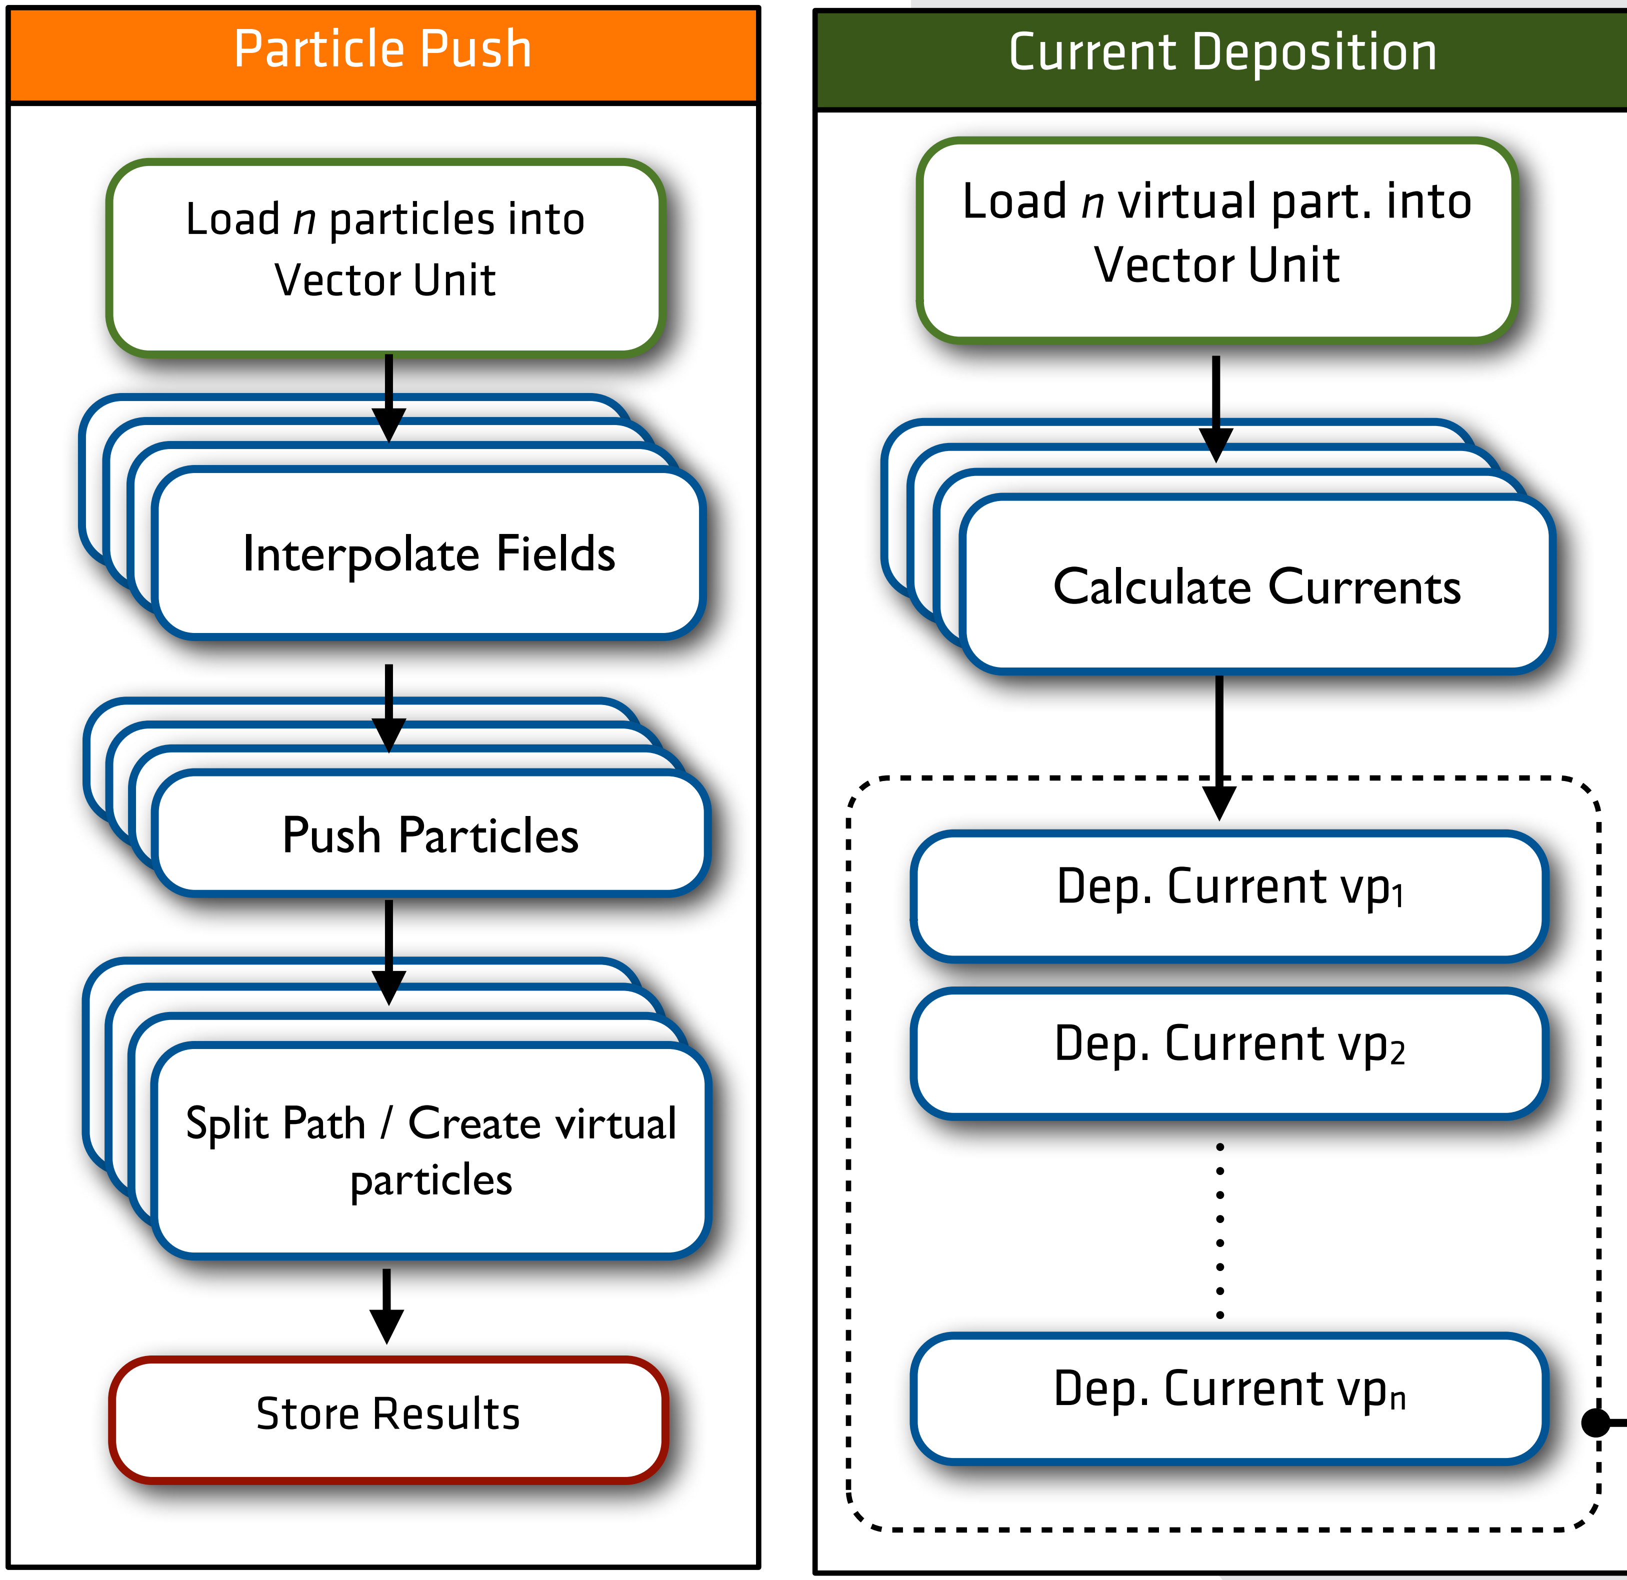

- PIC codes are good candidates for optimization:
  - Operations on each particle independent from each other…
  - except for current deposition
  - For most cases work well in single precision
- Process *n* (vector width) particles at a time
- Memory access much more expensive than calculation
  - Avoid temp buffers
  - Optimize for cache

- Particles may deposit to same cell
- Process each *n* particles sequentially

Performance on modern CPUs

- Intel E5-2680 @ 2.7 GHz
  - 8 cores / 16 threads
  - 256 bit wide SIMD (AVX)
  - 32 kB L1i + 32 kB L1d + 256 kB L2 per core
  - 20 MB L3 per cpu
  - (up to) 2 cpu per node
- Benchmark
  - 3D Warm plasma
  - 1 Mcells / 8 Mparticles
- Floating point efficiency: ~ 50%

LRZ SuperMUC
IBM iDataPlex
#7 - TOP500 Nov/12
147456 cores
R_{max} 2.9 PFlop/s

Performance on Intel E5-2680@2.7 GHz

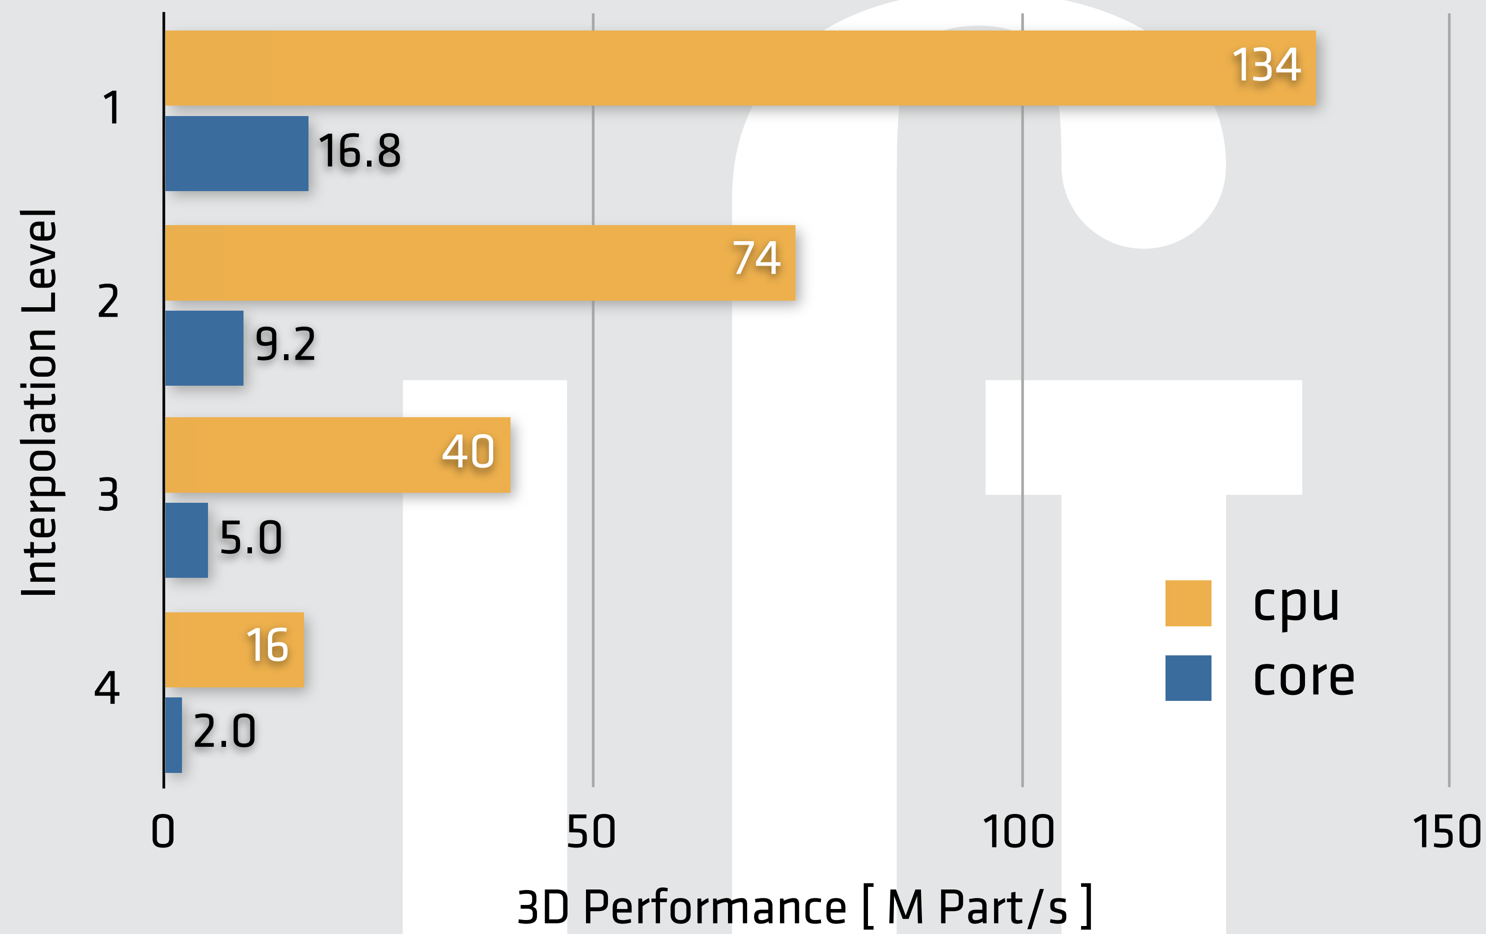

Core

Level	Push [ns]	Performance [MPart/s]
1	59.5	16.79
2	108.7	9.20
3	198.3	5.04
4	492.2	2.03

CPU

Level	Push [ns]	Performance [MPart/s]
1	7.4	134.34
2	13.6	73.62
3	24.8	40.34
4	61.5	16.25

Outline

- Modeling Plasma Based Accelerators
  - The OSIRIS framework
  - Algorithm
  - Computational Requirements
- Extending the algorithm
  - Lorentz boosted frames
  - Ponderomotive guiding center
  - High density hybrid
  - LWFA in cylindrical geometry
- Scaling to Petascale Systems
  - Multi-level parallelism
  - Multi-dimensional dynamic load balance
- Plasma Acceleration on Tier-0 systems
  - Scalings @ Sequoia & BlueWaters
  - 20 m long beam-plasma interaction @ SuperMUC
  - High-res 1 GeV LWFA @ Jaguar
  - Full scale target LWFA modeling @ SuperMUC
- Overview

Scaling to 1.6 million cores

Scaling Tests

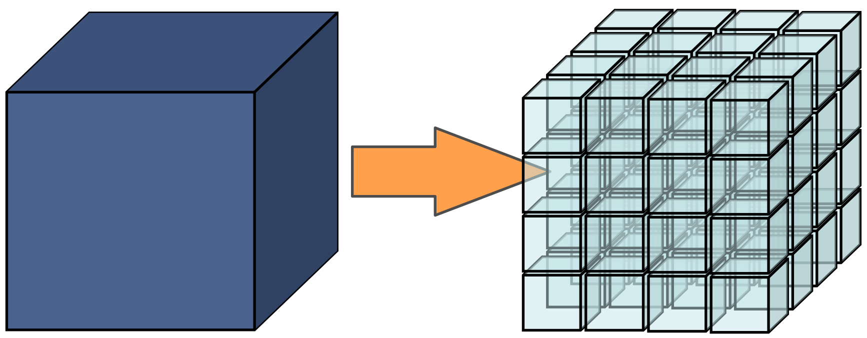

Sim. Volume → *Parallel Domain*

- Scaling tests on LLNL Sequoia
  - 4096 → 1572864 cores (full system)
- Warm plasma tests
  - Quadratic interpolation
  - $u_{th} = 0.1\, c$
- Weak scaling
  - Grow problem size
  - cells = $256^3 \times (N_{cores} / 4096)$
  - 2^3 particles/cell
- Strong scaling
  - Fixed problem size
  - cells = 2048^3
  - 16 particles / cell

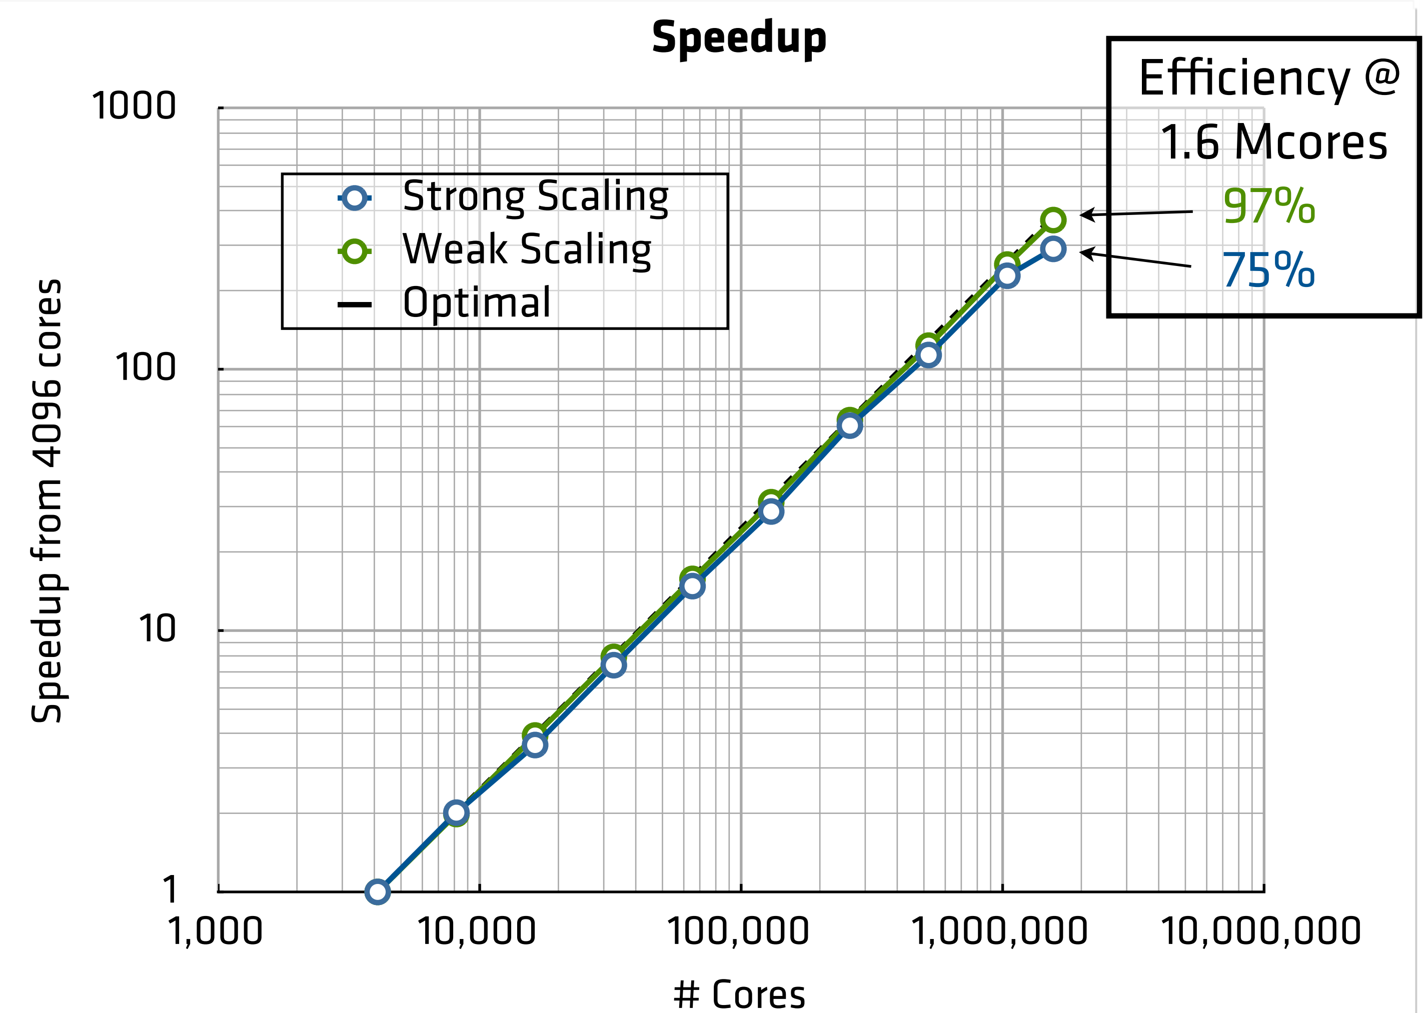

LLNL Sequoia
IBM BlueGene/Q
#2 - TOP500 Nov/12
1572864 cores
R_{max} 16.3 PFlop/s

Sustained Petascale performance

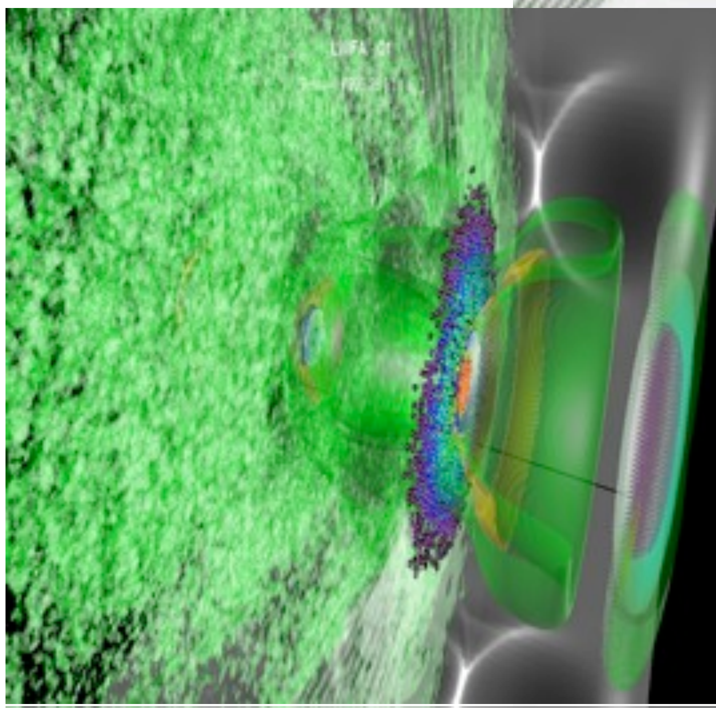

- Performance tests on Blue Waters
  - 772 480 cores (XE partition)
- Warm plasma tests
  - Quadratic interpolation
  - u_{th} = 0.1 c
- Problem size
  - cells = 38624 × 1024 × 640 (~2.5×10^{10})
  - 400 particles/cell (~ 10^{13})
- Computations
  - **2.2 PFlop/s performance**
  - 31% of R_{peak}

Blue Waters - Cray XE/XK hybrid

24140 XE Compute Nodes
2× 16 core AMD 6276 @ 2.3 GHz
R_{peak} 7.1 PFlop/s

3072 XK Compute Nodes
1× 16 core AMD 6276 @ 2.3 GHz
1 × Nvidia Tesla K20 GPU
R_{peak} 4.51 PFlop/s

$R_{peak\ aggr}$ 11.61 Pflop/s

Self-modulated plasma wakefield acceleration in 10-20 m long plasmas with ionizing laser

TÉCNICO LISBOA

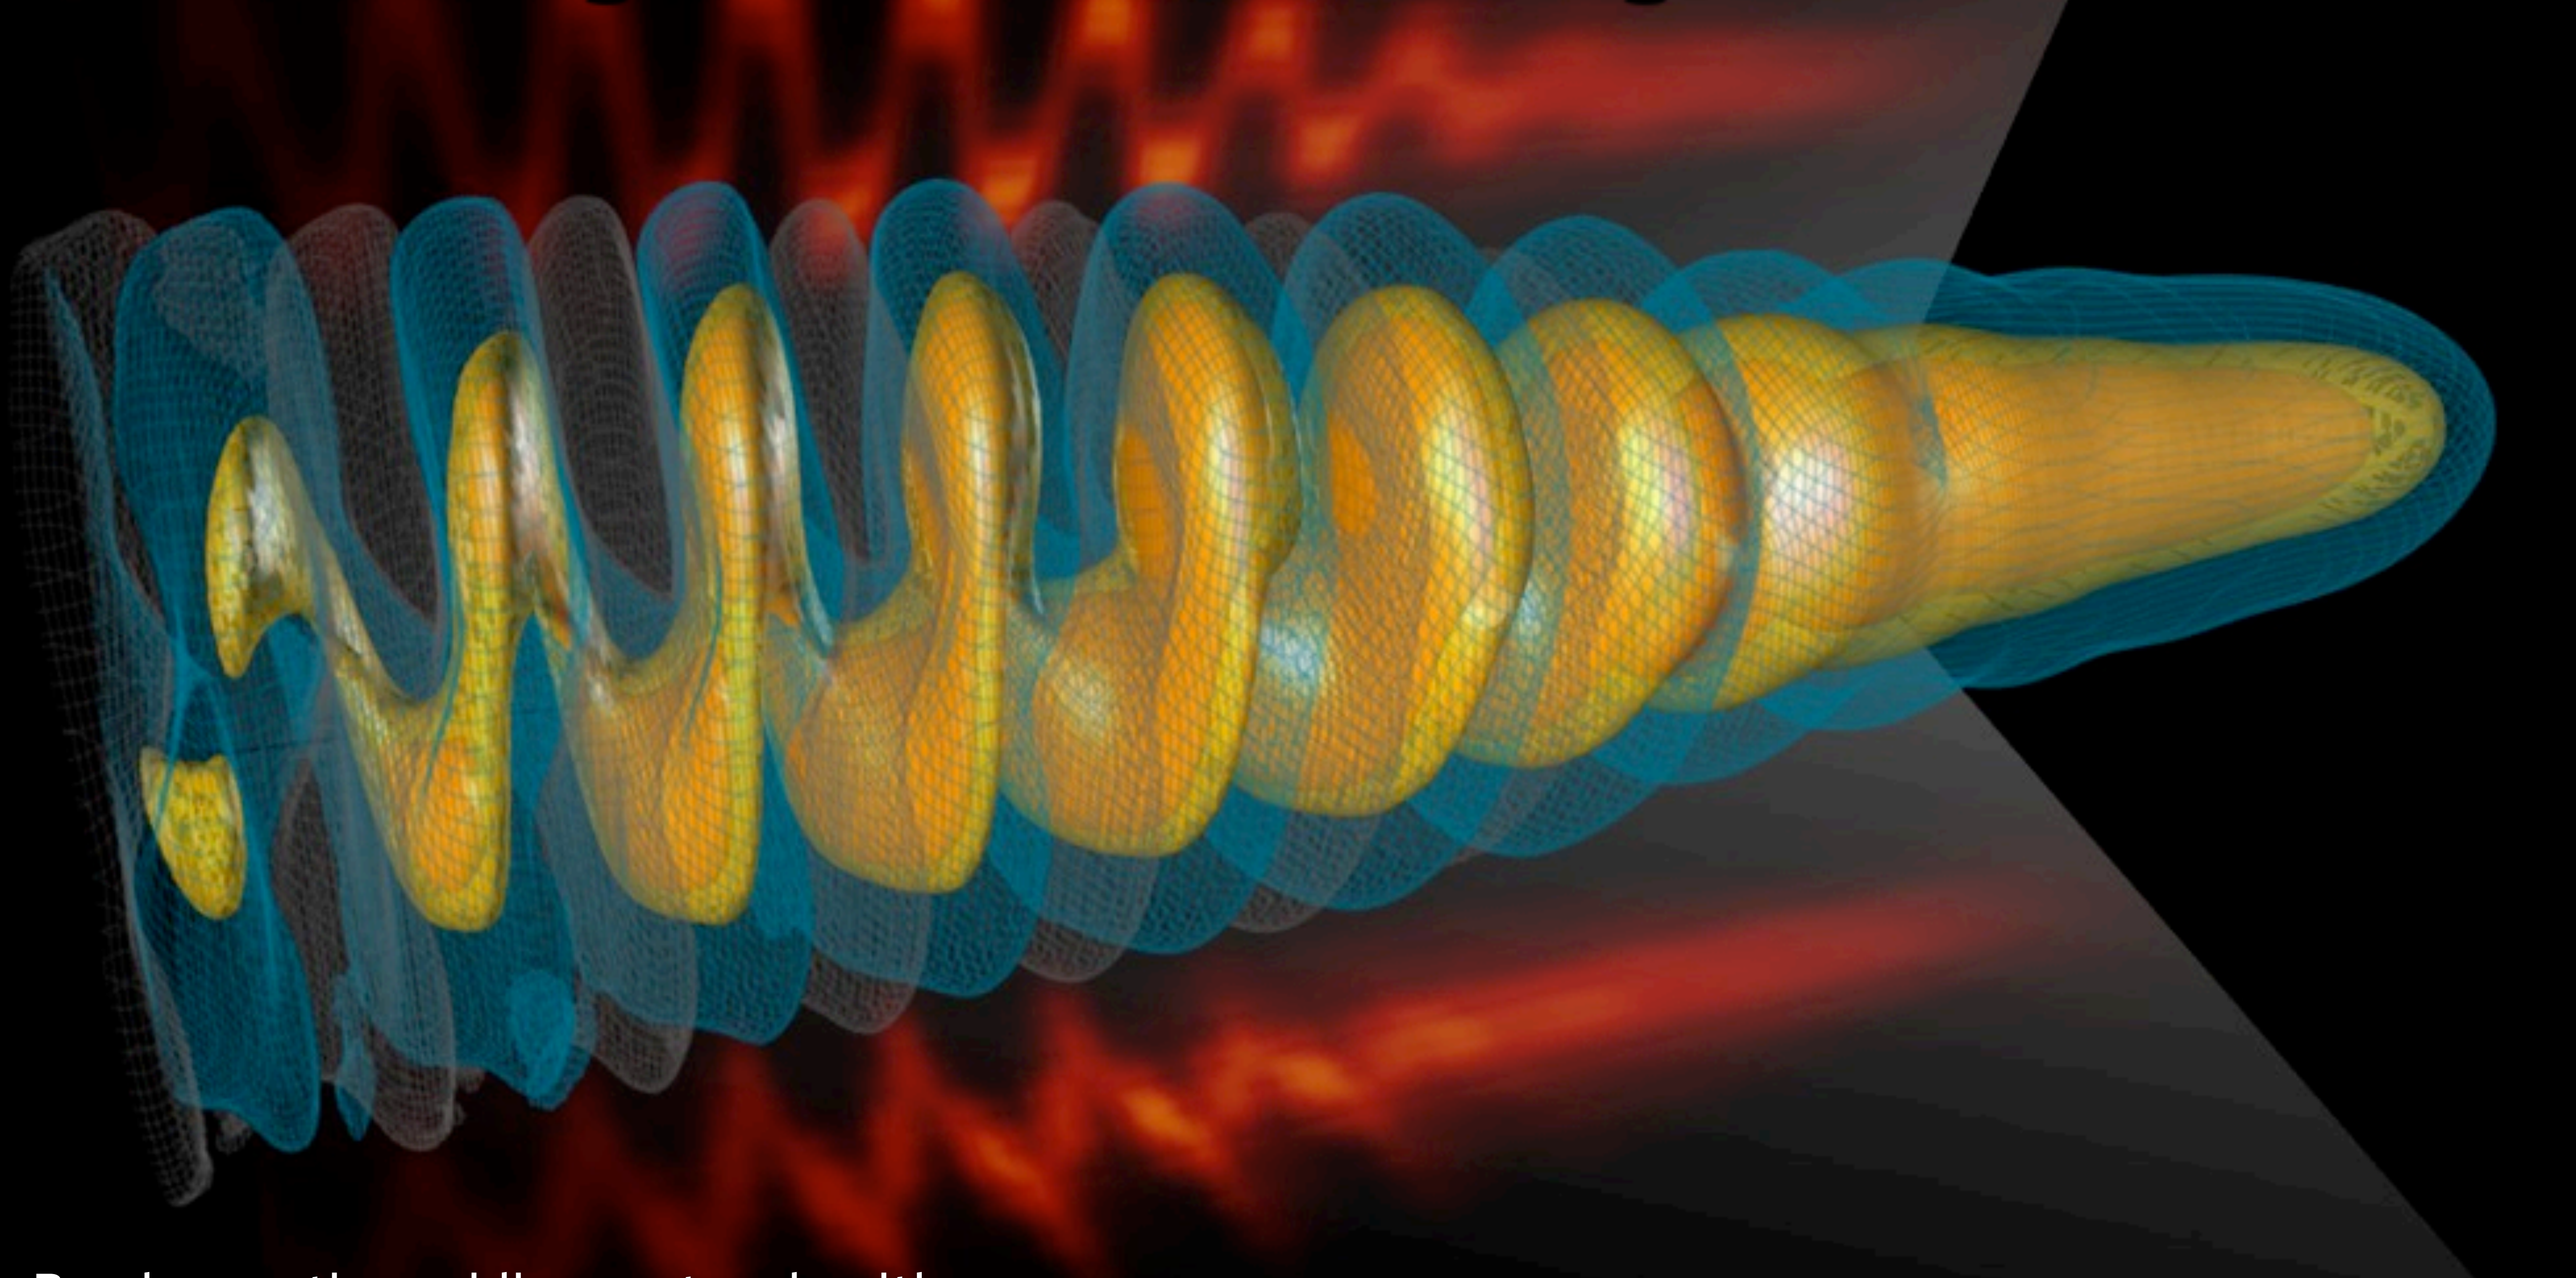

- Ponderomotive guiding center algorithm
- full-scale modeling of 10-20 meter long plasmas
- 20 cm long simulation window
- Ionization seeding for self-modulation and ionization laser depletion
- ~10^6 time-steps per simulation

Osiris

LWFA on a Tier-0 system

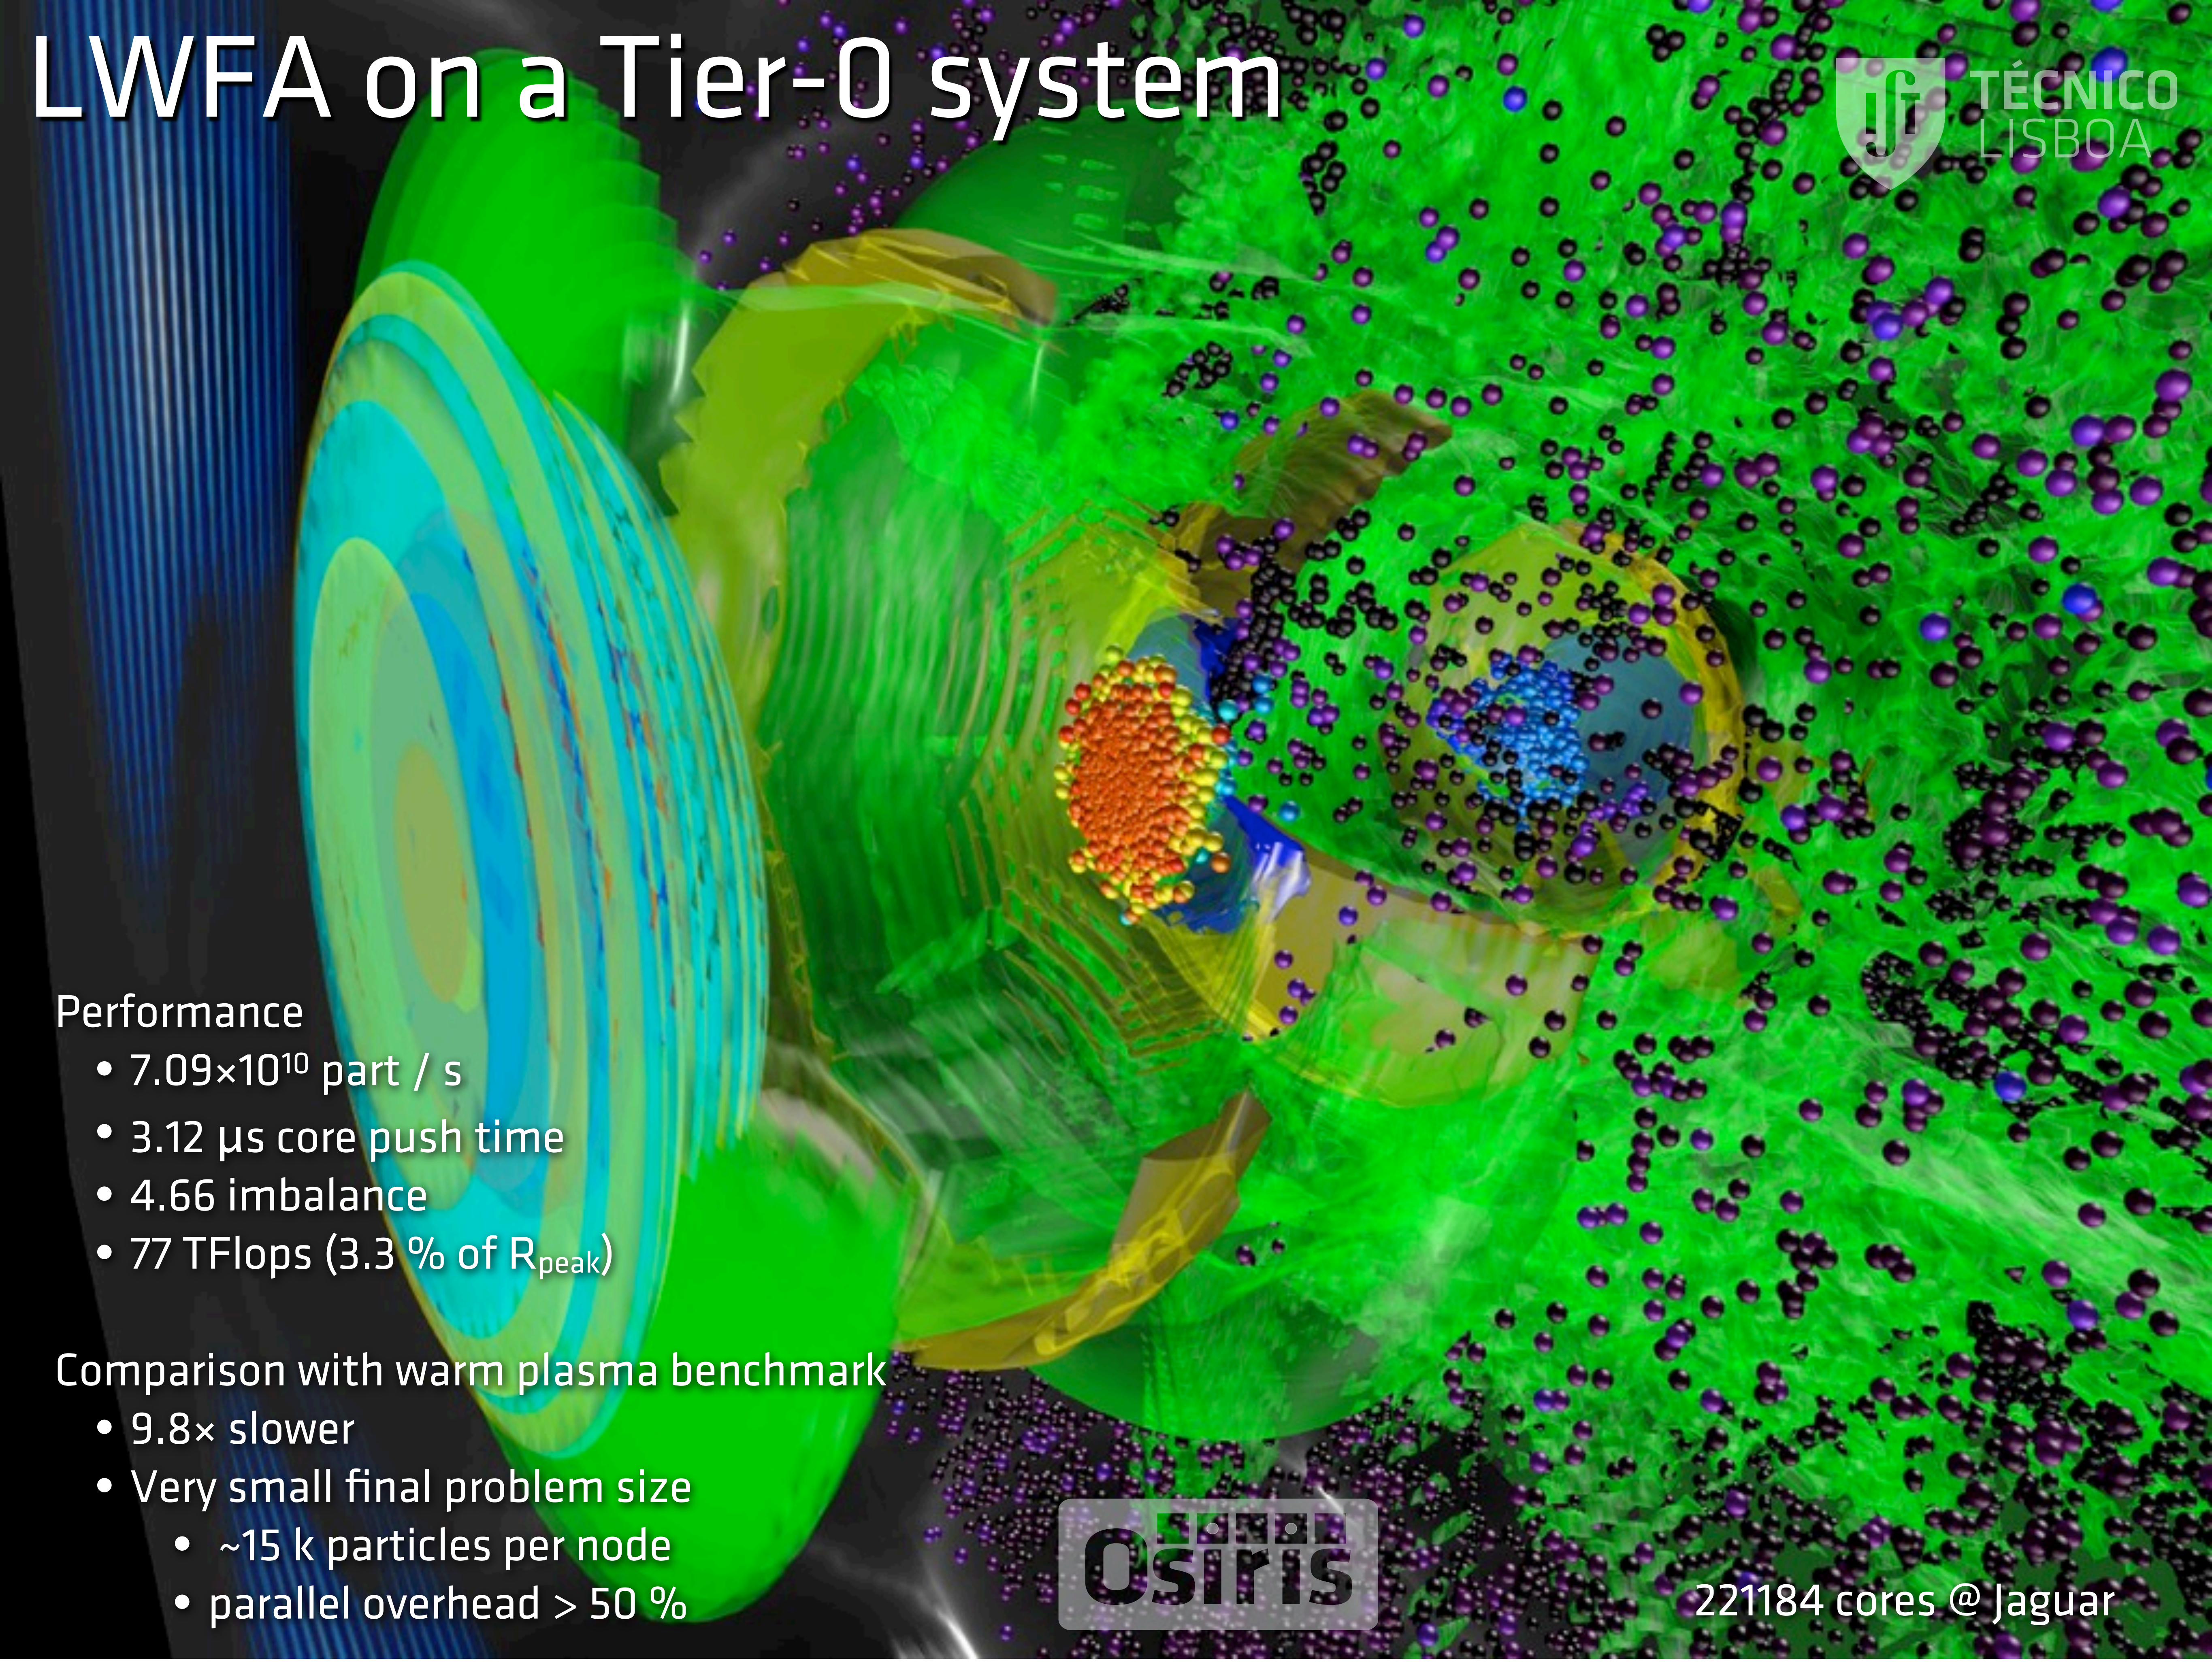

Performance

- 7.09×10^{10} part / s
- 3.12 μs core push time
- 4.66 imbalance
- 77 TFlops (3.3 % of R_{peak})

Comparison with warm plasma benchmark

- 9.8× slower
- Very small final problem size
  - ~15 k particles per node
  - parallel overhead > 50 %

221184 cores @ Jaguar

Laser wakefield acceleration
Full-scale 3D modeling of entire gas jet

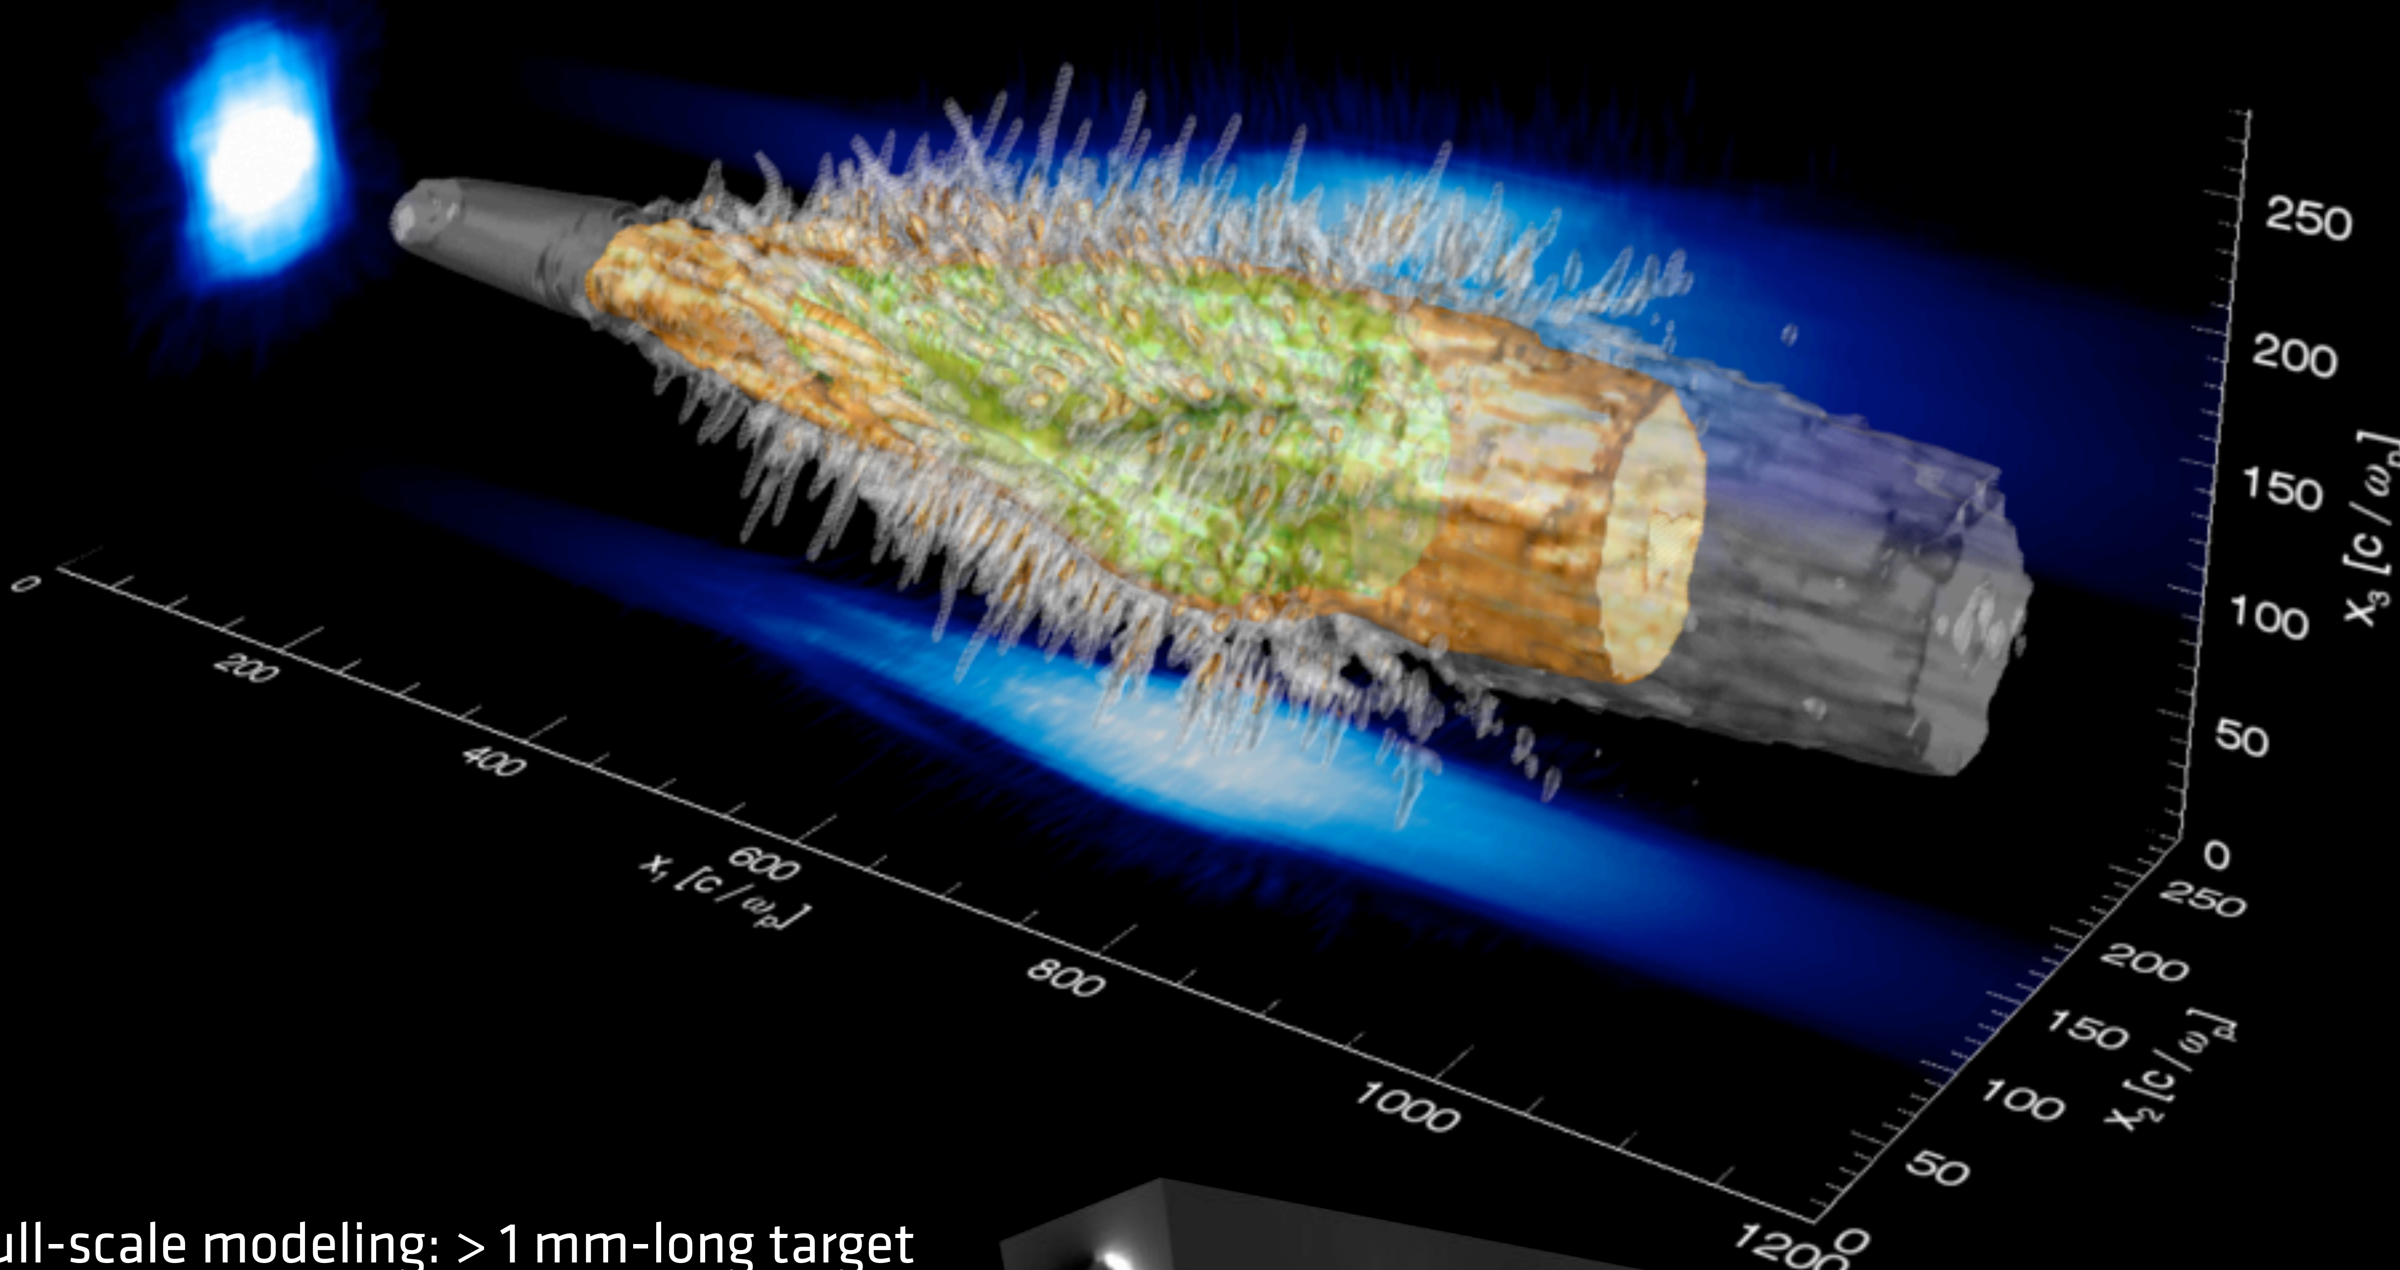

- full-scale modeling: > 1 mm-long target
- Multi-level ionization
- ~ 3×10^{10} cells
- ~ 6×10^{10} particles
- ~ 3×10^{4} time-steps
- ~ 0.4 million core h (~ 16 k€)

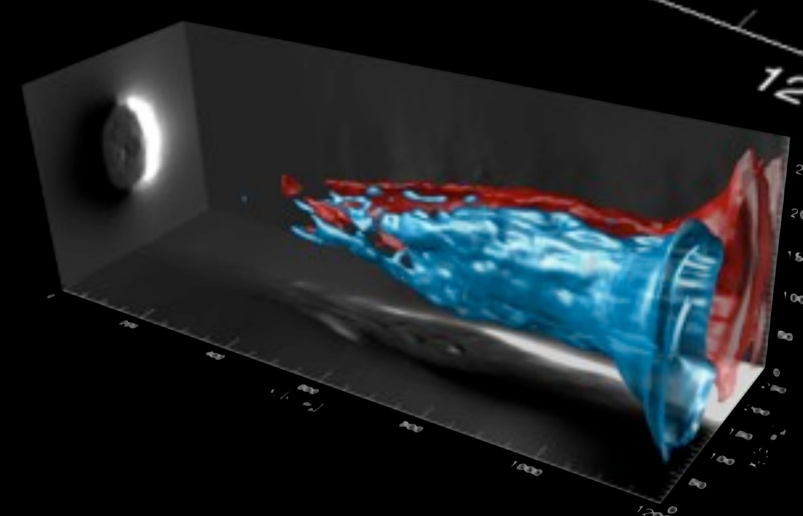

Outline

- Modeling Plasma Based Accelerators
  - The OSIRIS framework
  - Algorithm
  - Computational Requirements
- Extending the algorithm
  - Lorentz boosted frames
  - Ponderomotive guiding center
  - High density hybrid
  - LWFA in cylindrical geometry
- Scaling to Petascale Systems
  - Multi-level parallelism
  - Multi-dimensional dynamic load balance
- Plasma Acceleration on Tier-0 systems
  - Scalings @ Sequoia & BlueWaters
  - 20 m long beam-plasma interaction @ SuperMUC
  - High-res 1 GeV LWFA @ Jaguar
  - Full scale target LWFA modeling @ SuperMUC
- Overview

Overview

OSIRIS in numbers

- Scalable to ~ 1.6 Mcores
- > 2 PFlop/s sustained perf.
- > 10^{10} cells
- > 10^{13} particles
- < 8 ns / part / cpu

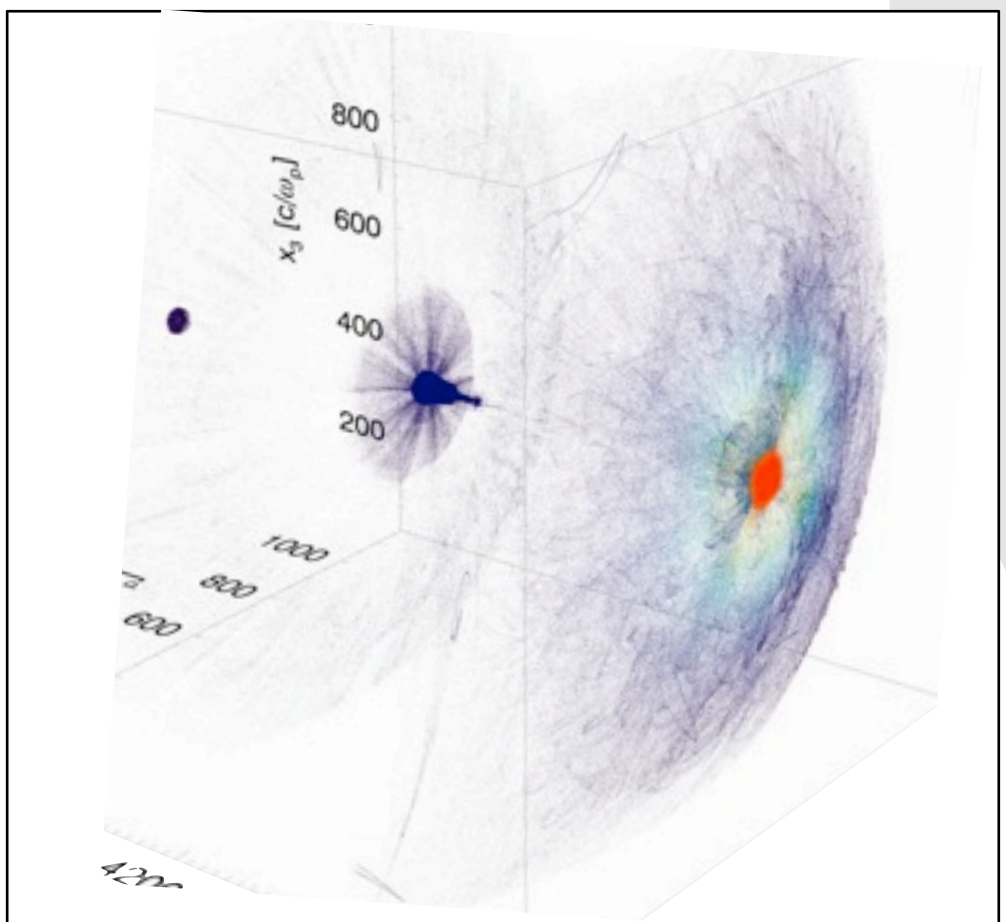

- **Petascale HPC systems online**
  - Outstanding tools for scientific research
  - Leverage on an hierarchy of parallelism
- **Improve our existing models**
  - Explore multiscale scenarios
  - Over/under critical laser plasma interaction
- **Improve our codes**
  - Exploit all levels of parallelism available
  - Make every cycle count
- **One to one kinetic modeling of plasma based accelerators within reach**
  - Multiscale simulations
  - High-Fidelity full scale models