Needs and challenges for modeling FACET-II and Beyond

W.B.Mori
University of California
Los Angeles (UCLA)

Departments of Physics and
Astronomy and of Electrical Engineering

Institute of Digital Research and Education

mori@physics.ucla.edu

With help from ***W. An, X.Xu,*** *C. Joshi, R. Fonseca, A. Tableman, W. Lu, M. Hogan, PICKSC, and SLAC.*


INSIGHT

nature

THE INTERNATIONAL WEEKLY JOURNAL OF SCIENCE

FULL SPEED AHEAD

Plasma wakefield machines — the particle accelerators of the future? PAGES 40 & 92

CENTRAL EUROPE
LIFE AFTER THE WALL
Science 25 years after the collapse of communism
PAGE 22

ENVIRONMENTAL SCIENCE
CASH, CONFLICT, CONSERVATION
To protect the planet, stop the infighting and fudging
PAGES 27, 28 & 32

INTERACTIVE NOTEBOOKS
SHARE AND SHARE ALIKE
IPython allows on-the-run analysis
PAGE 151

NATURE.COM/NATURE
6 November 2014

Simulations will be critical for FACET-II and PWFA linear collider research

- Need simulation tools that can support the design of experiments at FACET II.
- Need simulation tools that can aid in interpreting experiments at FACET II.
- Need simulation tools that can simulate new physics concepts, e.g., 3D down ramp injection and matching sections.
- Need simulation tools that can simulate physics of a PWFA-LC including the final focus.
- Need simulation tools that aid in helping to design a self-consistent set of parameters for a PWFA-LC.

Simulations are critical for FACET-II and PWFA linear collider research

- Simulations tools need to be continually improved and validated.
- Simulation tools need to run on entire ecosystem of resources.
- Simulation and analysis tools need to be easy to use.
- Relationship between code developers/maintainers and users is critical (best practices are not always easy to document).

Local clusters can be very useful: Dawson2

Dawson2 @ UCLA
- 96 nodes
- Ranked 148 in top 500
- 68 TFlops on Linpack

Node configuration
  - 2× Intel G7 X5650 CPU
  - 3× NVIDIA M2070 GPU

Computing Cores
  - Each GPU has 448 cores
  - total GPU cores: 129,024
  - total CPU cores: 1152

Funded by NSF

Existing leadership class facilities are useful

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

Exascale computing is on the horizon (not needed for FACET II)

High Performance Computing Power Evolution

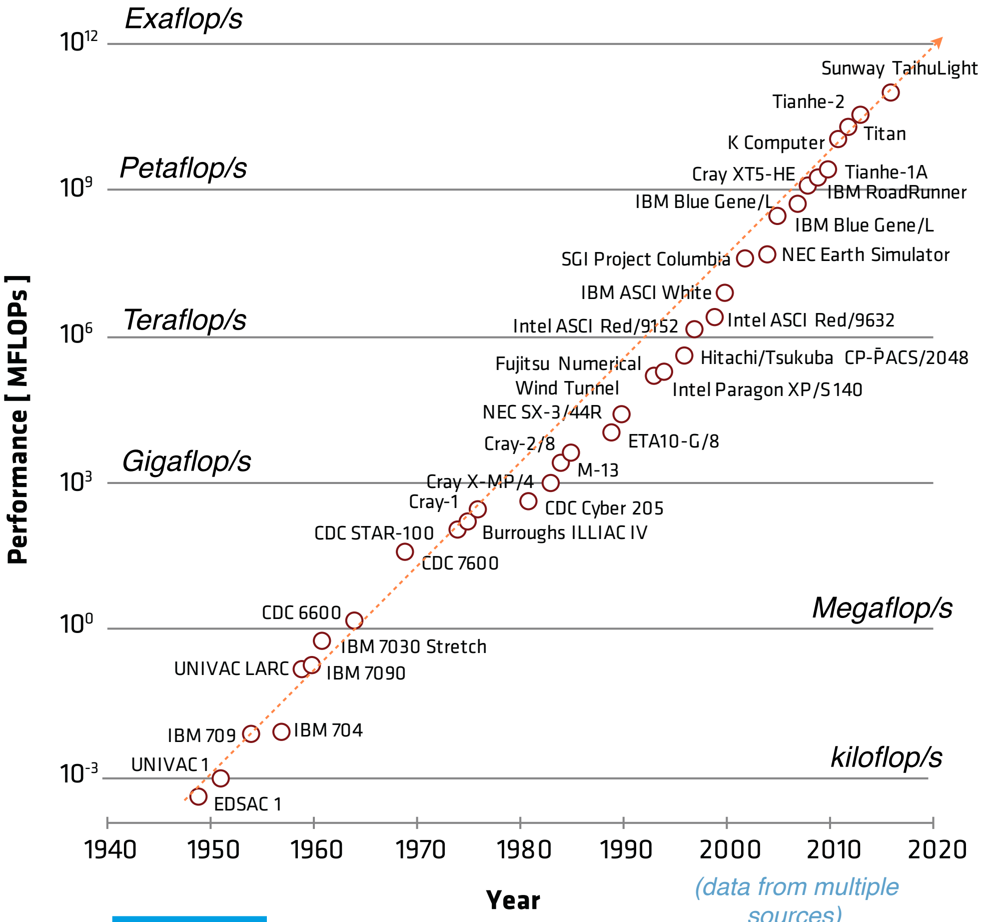

(data from multiple sources)

Sunway Taihulight
NRCPC, China
#1 - TOP500
Jun/16

Sunway Taihulight

- 40 960 compute nodes

Node Configuration

- 1× SW26010 manycore processor
  - 4×(64+1) cores @ 1.45 GHz
- 4× 8 GB DDR3

Total system

- 10 649 600 cores
- 1.31 PB RAM

Performance

- R_{peak} 125.4 Pflop/s
- R_{max} 93.0 Pflop/s

OSIRIS and QuickPIC have used to model FFTB and FACET for past 20 years: Design experiments, interpret experiments, study physics inaccessible to experiments

WarpX next talk

osiris 3.0 (OSIRIS 4.0 is now the development branch)

Osiris 3.0

TÉCNICO LISBOA

UCLA

osiris framework

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

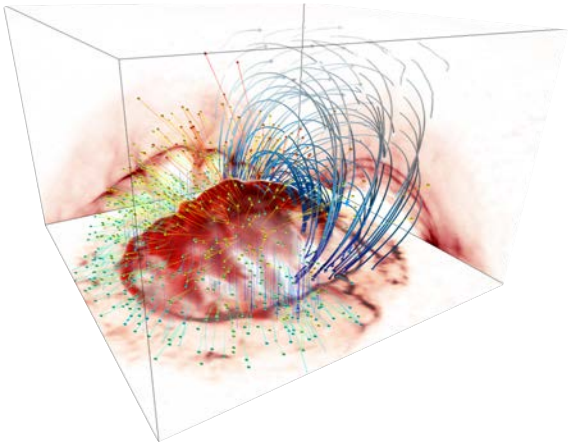

code features

- Scalability to ~ 1.6 M cores
- SIMD hardware optimized
- Parallel I/O
- Dynamic Load Balancing
- PGC
- QED module
- Particle splitting/merging
- Quasi-3D
- Boosted frame/-NCI
- GPGPU support
- Xeon Phi support

Ricardo Fonseca:
ricardo.fonseca@tecnico.ulisboa.pt
Frank Tsung:
tsung@physics.ucla.edu

http://epp.tecnico.ulisboa.pt/
http://picks.idre.ucla.edu/

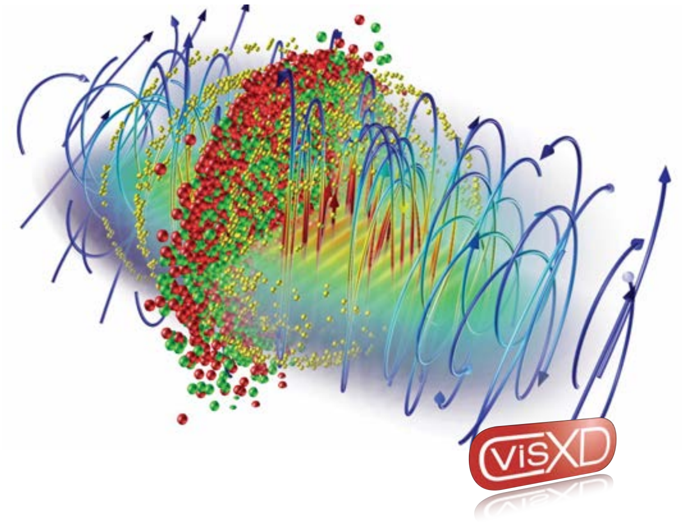

QuickPIC: A 3D quasi-static PIC code

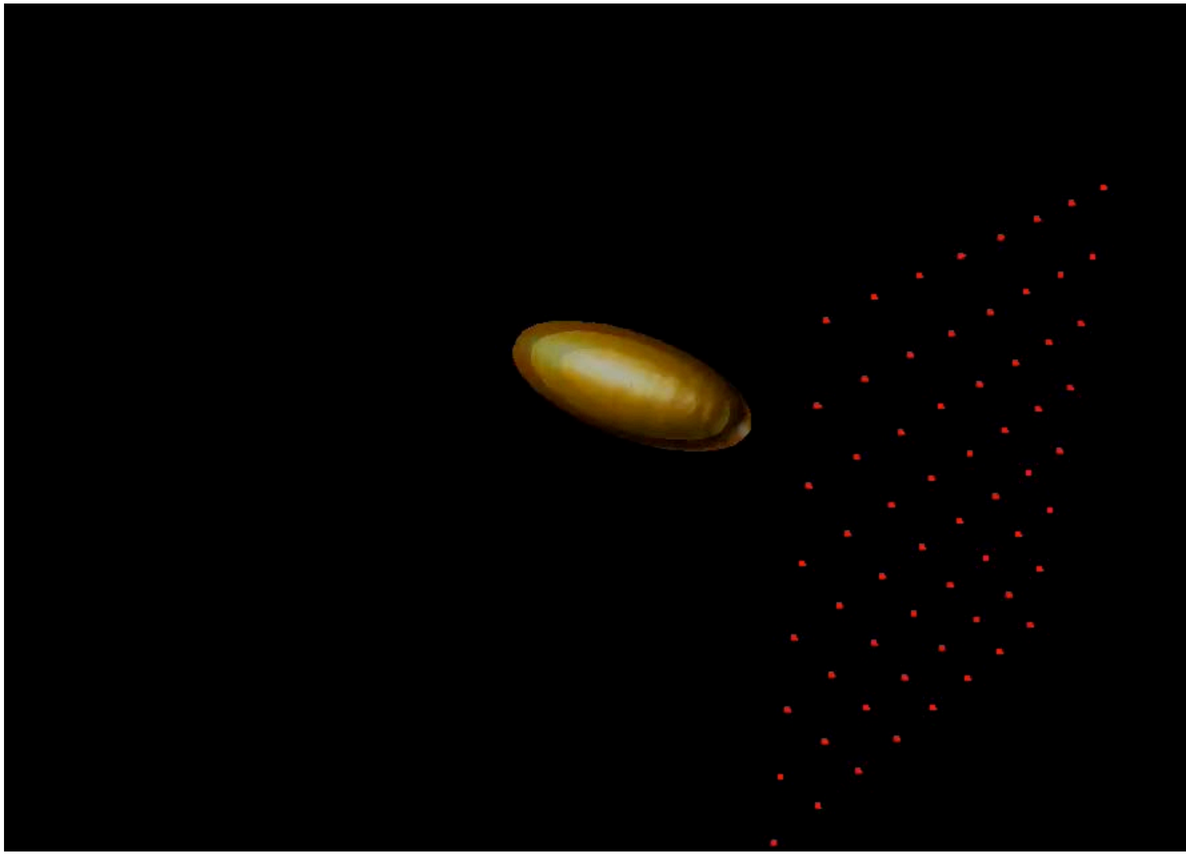

Fully parallelized and scaled to 100,000+ cores

Requires predictor corrector, has some similarities with a Darwin code.

Will be open source soon.

C-K. Huang et al., 2006
W. An et al., 2014

Recently HIPACE (not fully 3D)

Embeds a parallelized 2D PIC code inside a 3D PIC code based on UPIC Framework.

Comparison of Osiris 3D and QuickPIC

PWFA using field ionized plasma

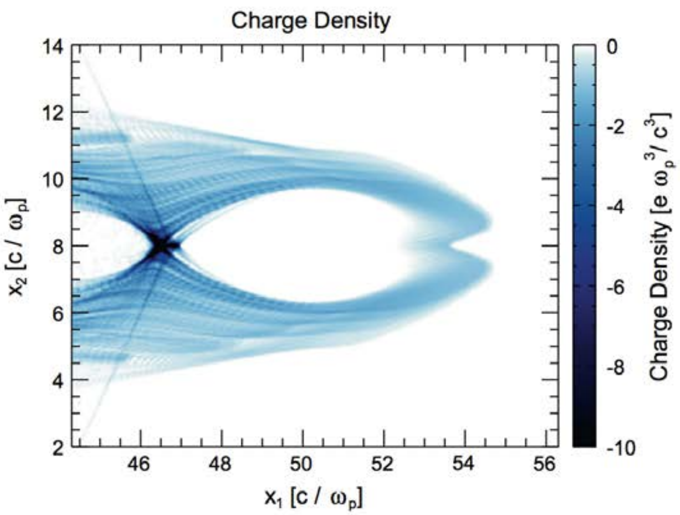

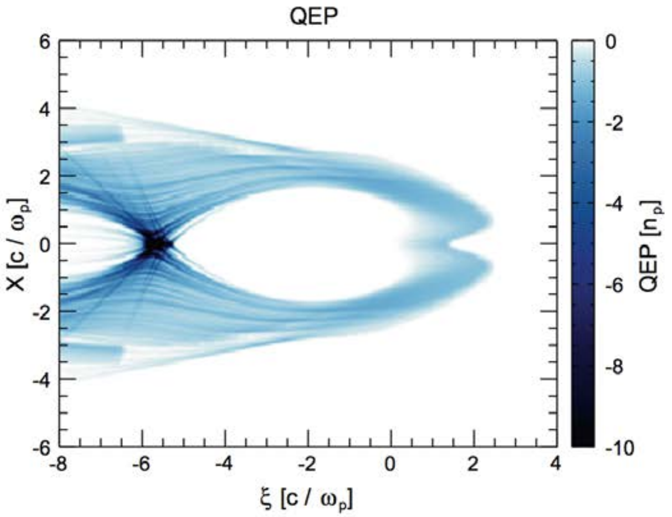

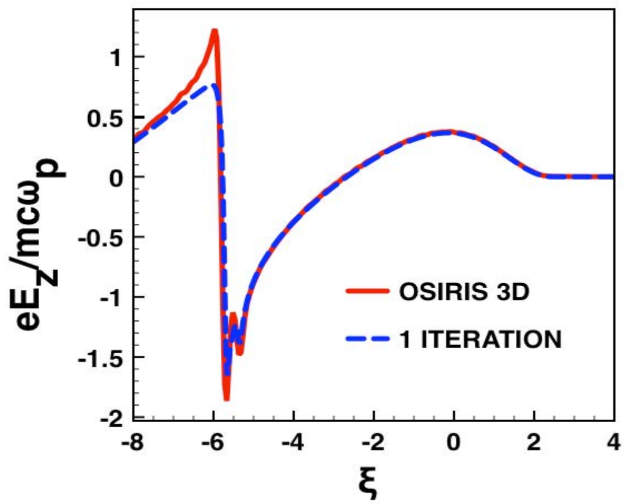

PWFA-LC using preformed plasma

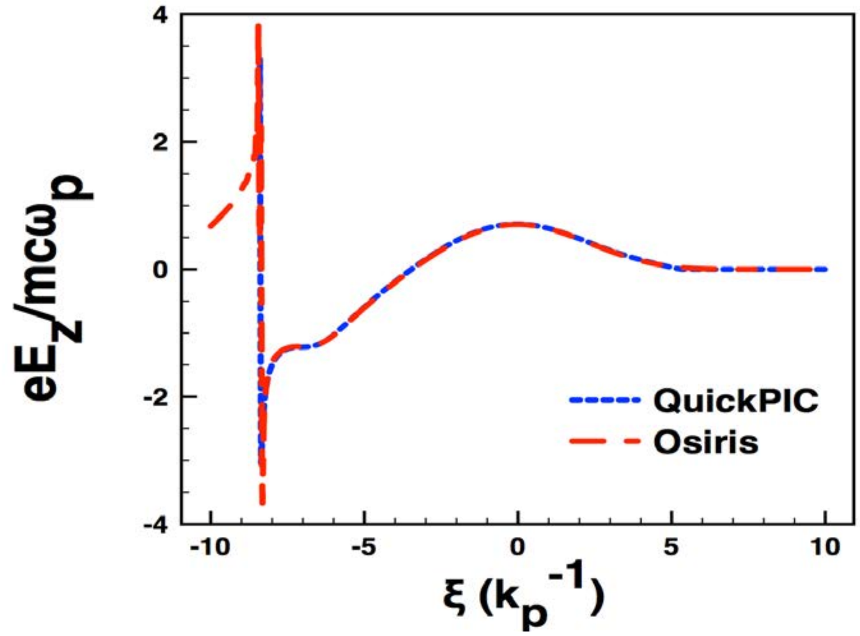

OSIRIS and QuickPIC access is international for HED and AA Science

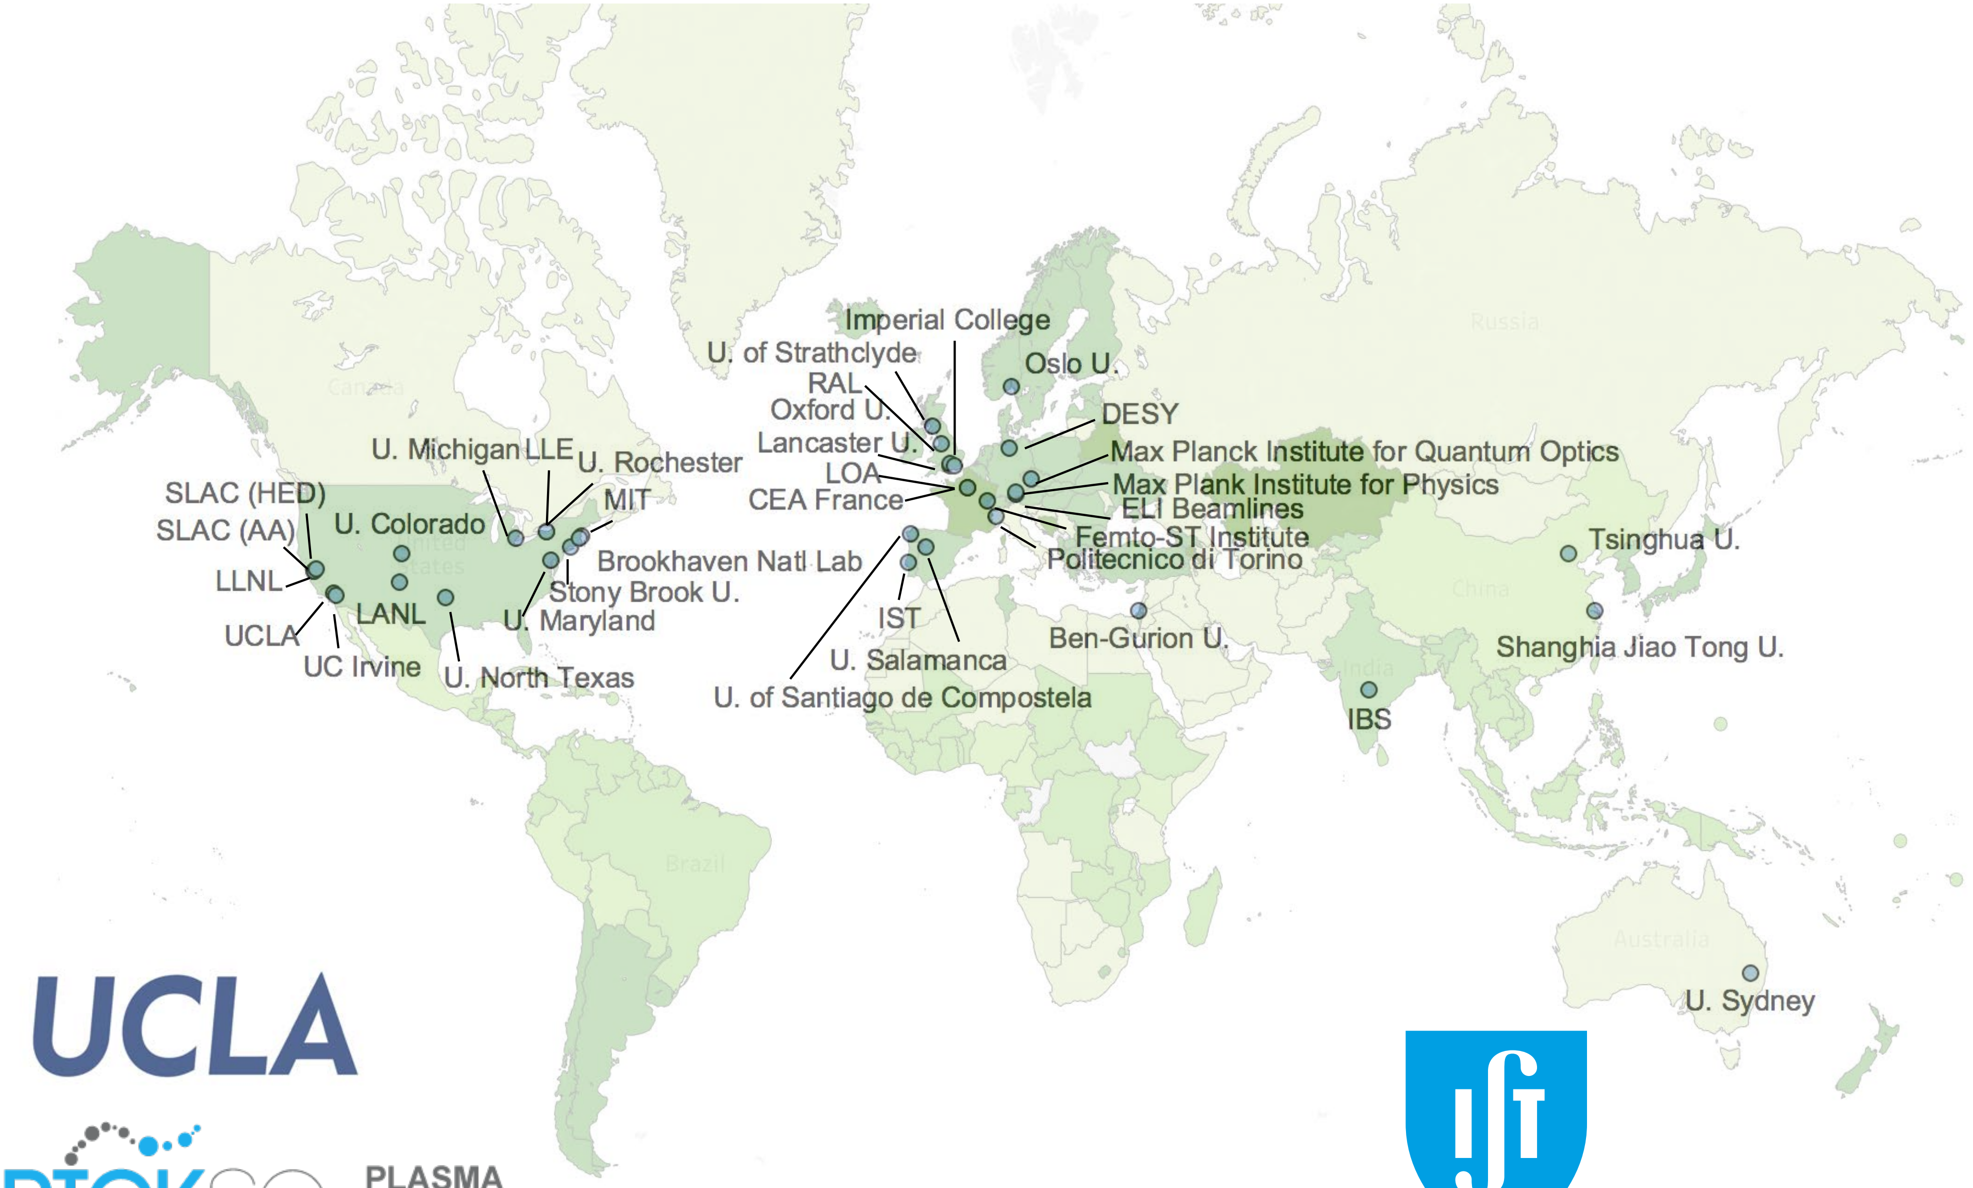

UCLA

PICKSC

PLASMA SIMULATION UCLA

1:Propose a major experiment that is consistent with DOE's one or more strategic goals

Proposal for an experiment at the FACET Science meeting at UCLA

- Deplete the drive beam of its energy
- 50% Energy extraction Efficiency
- 10 GeV energy gain for the trailing beam (TB)
- Minimize the energy spread of TB
- Demonstration of emittance preservation of TB
- (this is the first step towards eventually getting a collider quality beam)

- All at the same time

UCLA

Two Key Concepts for high quality beams from plasma accelerators

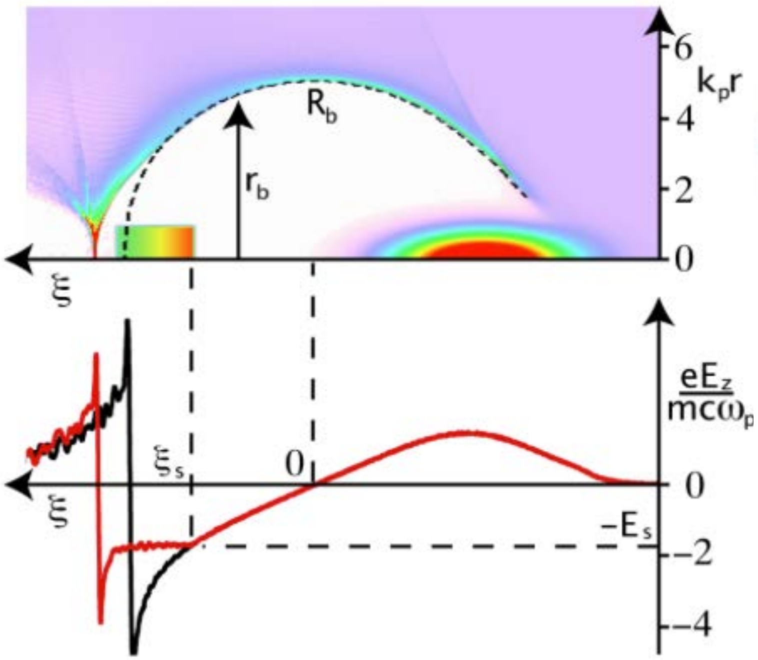

Ref: M. Tzoufras et al PRL

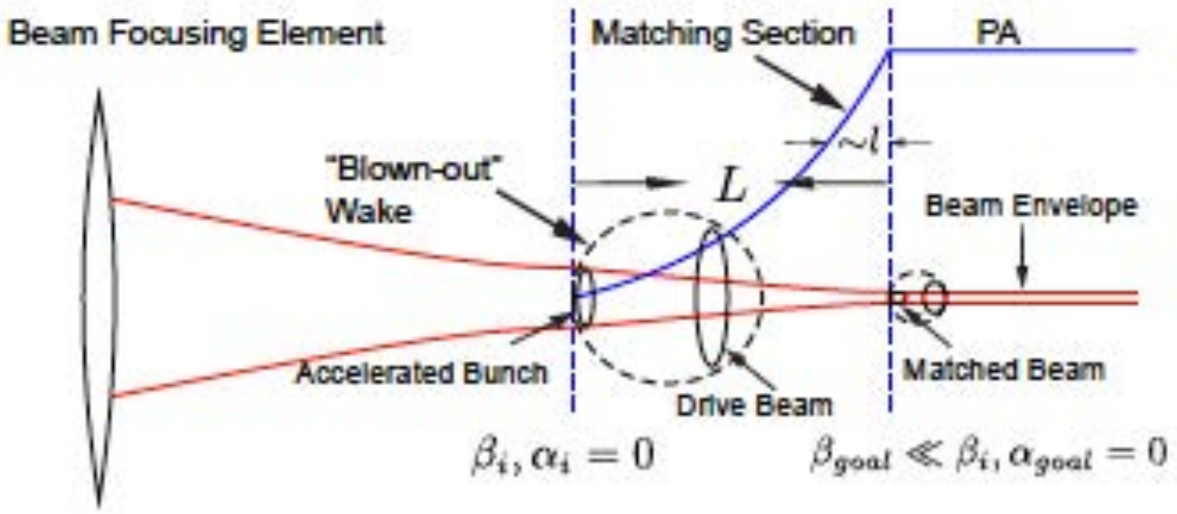

Ref: X. Xu et al PRL 2015

Beam Loading

Energy Spread and Efficiency

Matching Section

Emittance Preservation

Simulations conducted 10 years ago show this is energetically possible for a 25 GeV stage

Preionized

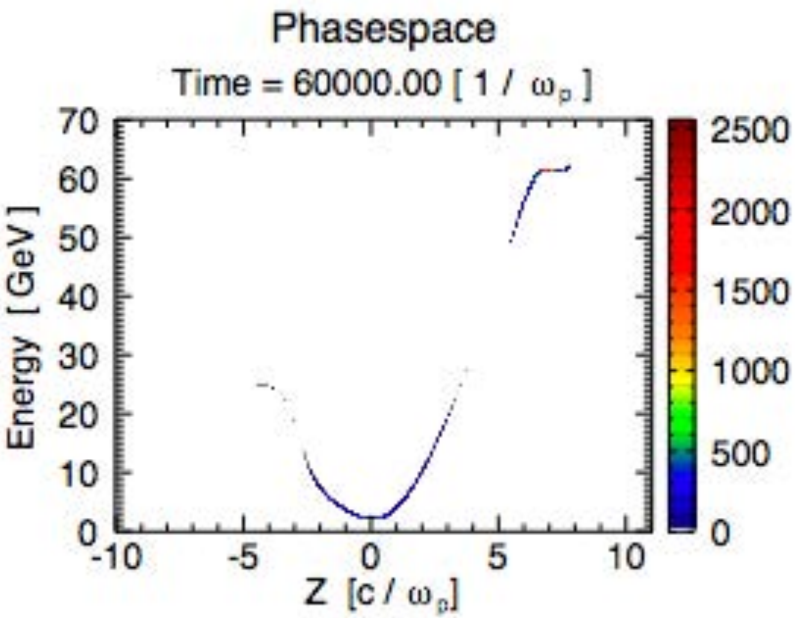

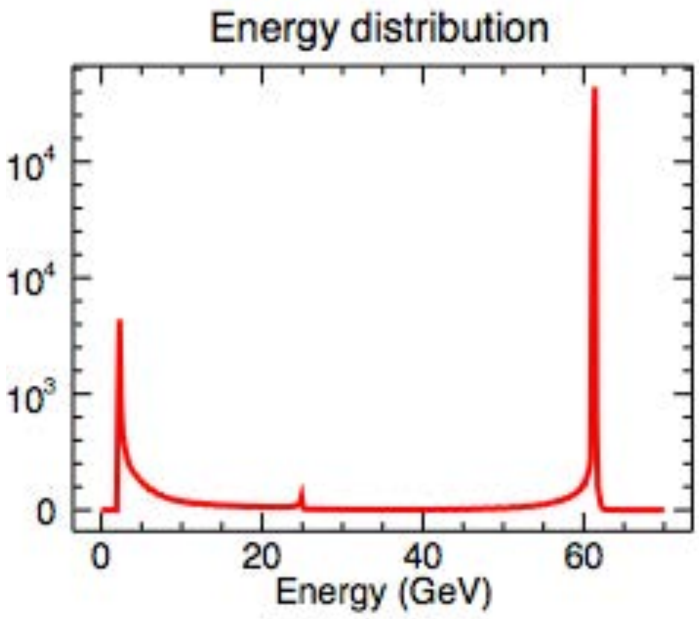

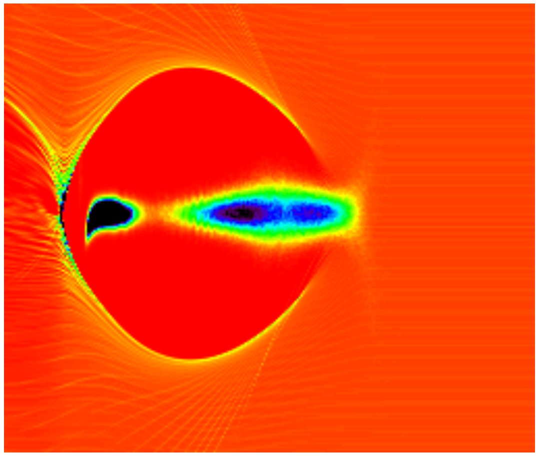

n_p= 1×10^{17}cm^{-3}

N_{driver} = 2.9×10^{10}, σ_r= 3 μ, σ_z = 30 μ, Energy = 25 GeV

$N_{trailing}$ = 1.0×10^{10} , σ_r= 3 μ , σ_z = 10 μ, Energy = 25 GeV

Spacing= 110 μ

R_{trans} = -E_{acc}/E_{dec} > 1 (Energy gain exceeds 25 GeV per stage)

1% Energy spread

Efficiency from drive to trailing bunch ~48%!

Experiment 1:Realizable because of Differences between FACET I and II beams

Parameter	FACET I	FACET II
Drive Beam	20 GeV	10 GeV
Norm. Emittance	50x200 um	< 3x7 um
Pump Depletion	No	Yes
Trailing Beam		
Bunch Charge	>100 pC	> 100 pC
Energy Spread	~5%	~1
Energy gain	max 8 GeV	10 GeV
Efficiency	30%	50%
Emittance Preservation	No	Yes?

We are going to optimize beam loading and demonstrate beam matching.

UCLA

FACET & FACET II Simulations

FACET Two-Bunch

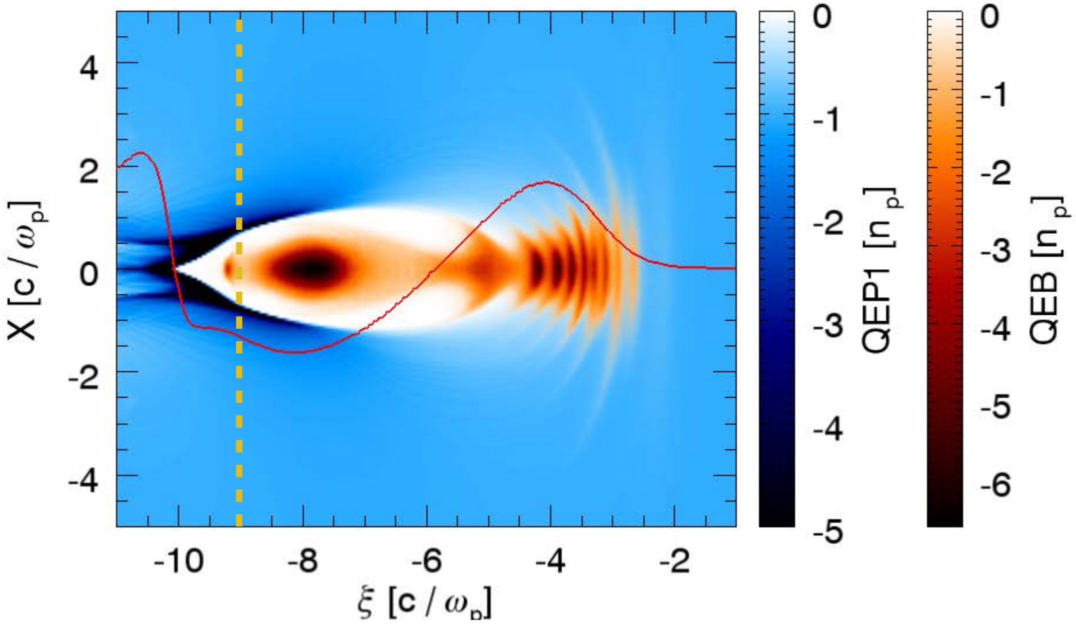
Plasma and Beam Densities

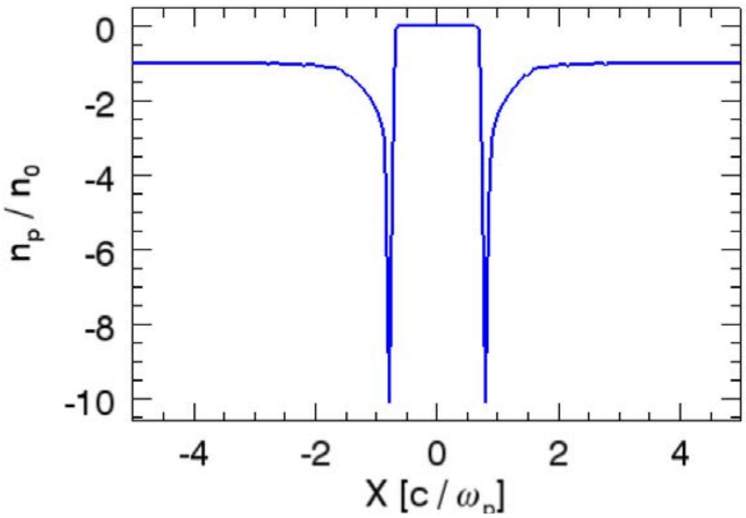
Plasma Density Lineout

FACET II Two-Bunch(Low ε_N)

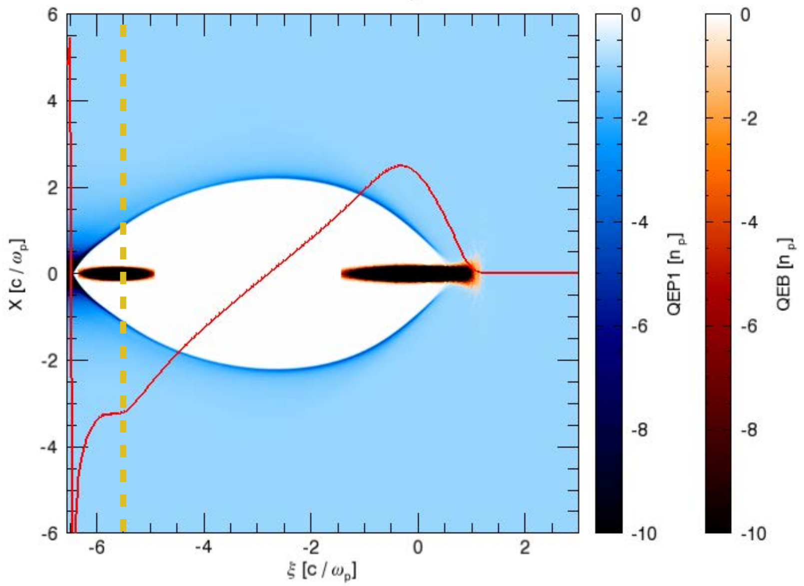
QEP1

Time = 12200.00 [1 / ωp]

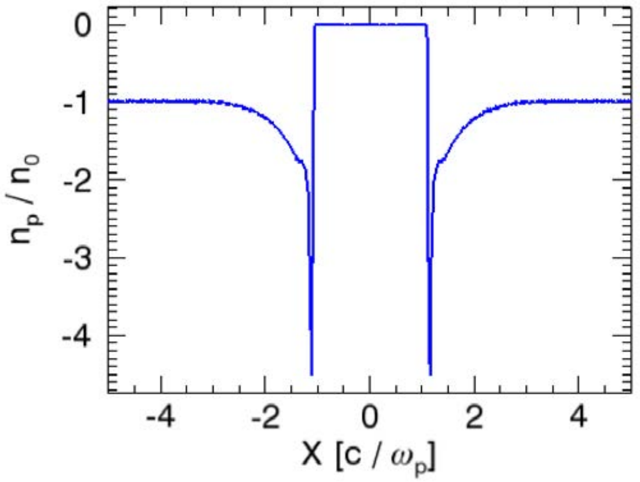
Plasma Density Lineout

Two-bunch PWFA @ FACET II

Simulations helped to understand the proper spacing and current ratios

Drive Beam: E = 10 GeV, I_{peak}=15 kA
β = 89.61 cm, α = 0.0653,
σ_r = 21.17 µm, σ_z = 12.77 µm ,
N =1.0 x 10^{10} (1.6 nC),
ε_N = 10 µm
Trailing Beam: E = 10 GeV, I_{peak}=9 kA
β = 89.61 cm, α = 0.0653,
σ_r = 21.17 µm, σ_z = 6.38 µm ,
N =0.3 x 10^{10} (0.48 nC),
ε_N = 10 µm
Distance between two bunches: 150 µm
Plasma Density: 4.0 x 10^{16} cm^{-3} (with ramps)

OSIRIS Simulations are useful for determining if there is ionization self-injection

Plasma Density Profile

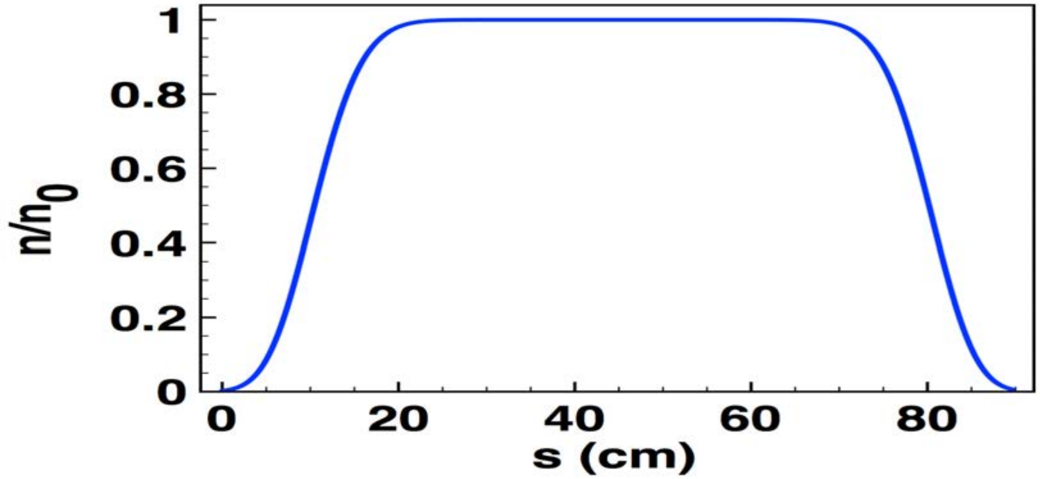

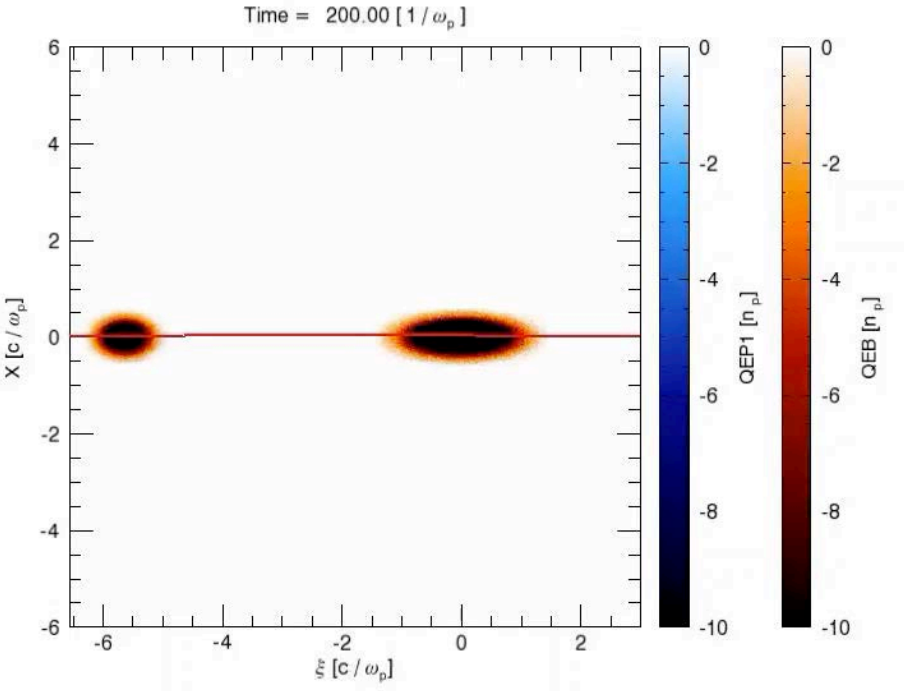

To understand the computational needs to introduce the concept of pump depletion: Transformer ratio

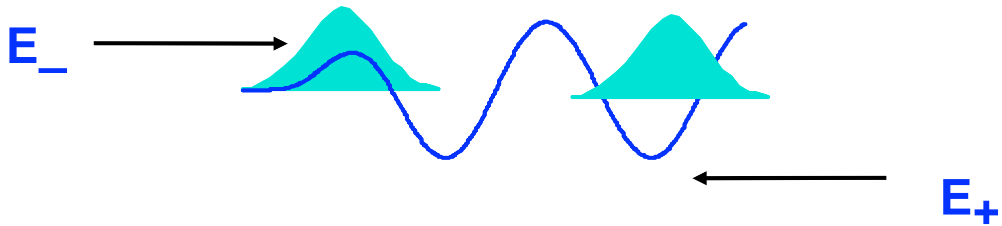

$$E = \partial_\xi \psi$$

$$E_- = \partial_\xi \psi_- \qquad eE_- L_{pd} = \gamma_b mc^2$$

$$\Delta W = E_+ L_{pd} = \frac{E_+}{E_-}\gamma_b mc^2 \qquad R \equiv \frac{E_+}{E_-}$$

Transformer ratio

You want particles in drive bunch to slow down together (Larger Lpd):

$$E_- = \partial_\xi \psi_- = Constant$$

You want particles in witness bunch to slow down together (beam loading):

$$E_+ = \partial_\xi \psi_+ = Constant$$

Simulation needs for two bunch experiment

$L_\perp$

$\approx 20\frac{c}{\omega_p}$

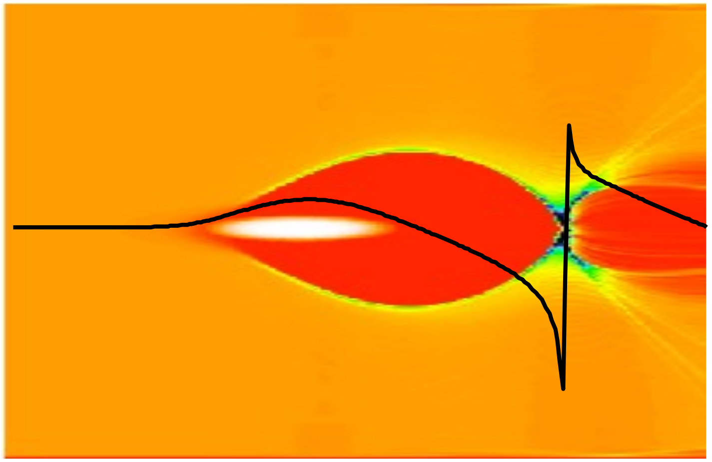

L_z

$\approx 25\frac{c}{\omega_p}$

$\Delta_t \approx min(\Delta_z, \Delta_\perp)$

OSIRIS

$\Delta_t \approx \gamma_b^{1/2}\Delta_z$

QuickPIC

$\gamma_b = 20,000$

FACETII

$\gamma_b = 50,000$

LC Stage

$\frac{L_z L_x L_y}{\Delta_z \Delta_x \Delta_y} N_{pcell} \frac{\gamma_b}{\Delta_t}$ = Particle Pushes

$\Delta \approx (.025 - .1)\frac{c}{\omega_p} \quad N_{pcell} = 2 - 20$

Cost=particle pushes x cost/push

UCLA

Osiris and QuickPIC: Rough estimates

Total Number of Particle Pushes

	Osiris 3D (8ppc)	QuickPIC (8ppc)
FACET II	7×10^{15}	1×10^{13}
PWFA-LC	1×10^{21}	5.6×10^{16}

Total CPU-Hours: assuming no load imbalance

	Osiris 3D (8ppc)	QuickPIC (8ppc)
FACET II	5.9×10^{5}	2.8×10^{3}
PWFA-LC	8.7×10^{10}	1.5×10^{7}

Exascale is not needed for FACET II

2:Experiment aligned with early application

- General consensus at present is the next generation of coherent light source.
- Need to produce electron bunches with brightness orders of magnitude larger than the brightest beams available today.
- What are the beam and plasma requirements?

Extreme Bunches a la Brendan

- Good news and Bad news

- First the bad news
- The charge in the witness too low to beam load the wake, need 2:1 ratio of beam currents

- Now the Good News: we can use just the drive beam to do the following
- Can operate at high density 2x10^{20} cm^{-3}
- No dephasing and TeV/m gradients
- Ionization injection and downramp injection possible
- May be possible to generate collider quality ultrabright beams.
- May be the easiest route to a first application-generation of coherent x-ray radiation
- We have developed or are developing plasma sources for such beams

Simulating downramp injection

When using normalized parameters a single simulation corresponds to a family of simulations with different densities

So the good news is that what has been simulated for lower densities can be simply “scaled” to higher densities: Brightness scales with n_0

3D simulations of LWFA and PWFA (e and p) can be expensive, but “r-z” can be useful for parameter scans

- 2D cylindrical r-z simulations can get the geometric scaling correct: Used extensively for PWFA
- EM waves are radially polarized in r-z simulations, so cylindrical r-z simulations not used for LWFA studies.
- Expand in azimuthal mode number and truncate expansion! [1]: LASER is an m=1 mode. This is PIC in r-z and gridless in Φ.
- A charge conserving current deposit was developed and incorporated into OSIRIS [2].

[1] A.F. Lifshitz et al., JCP 228, pp.1803 (2009).

[2] A. Davidson et al., JCP 281, pp. 1063 (2014).

[3] R. Lehe et al., submitted (20

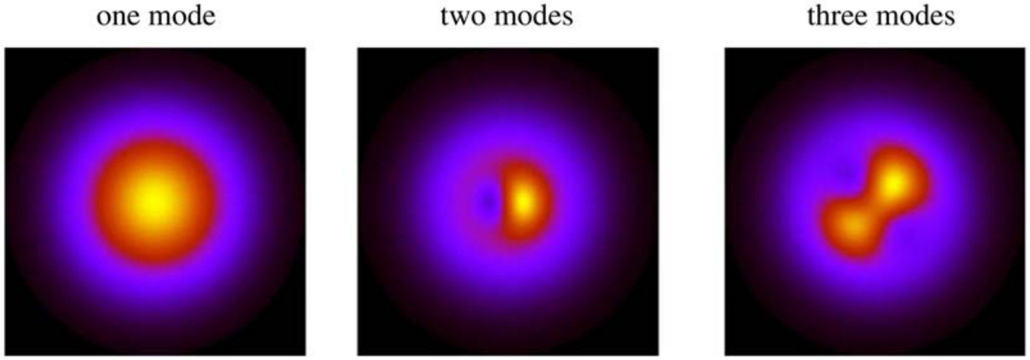

quasi-3D agrees with full 3D for symmetric cases with CPU savings of ~100 or more

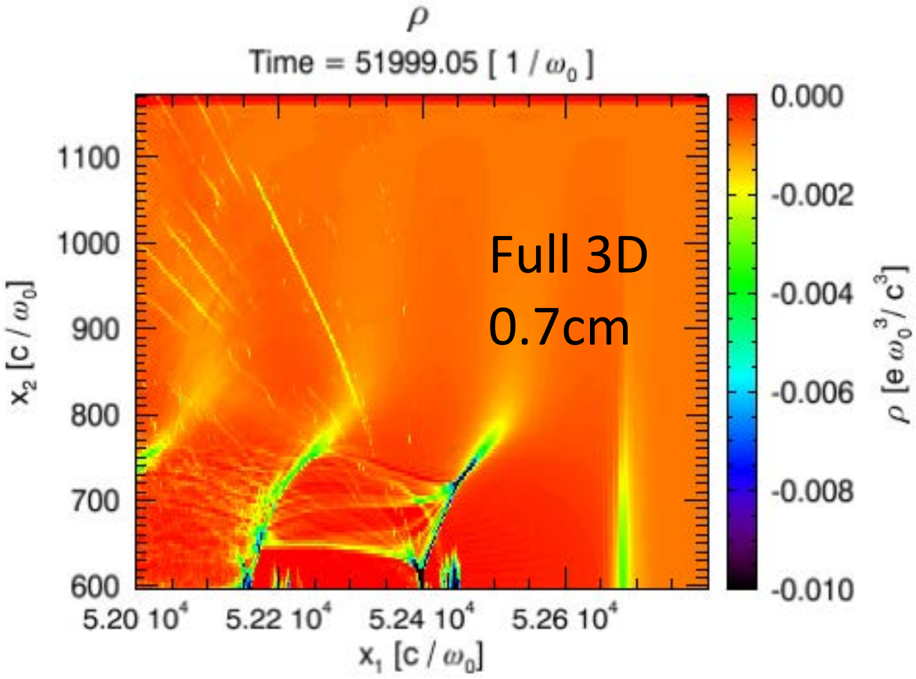

340pC 1.57 GeV

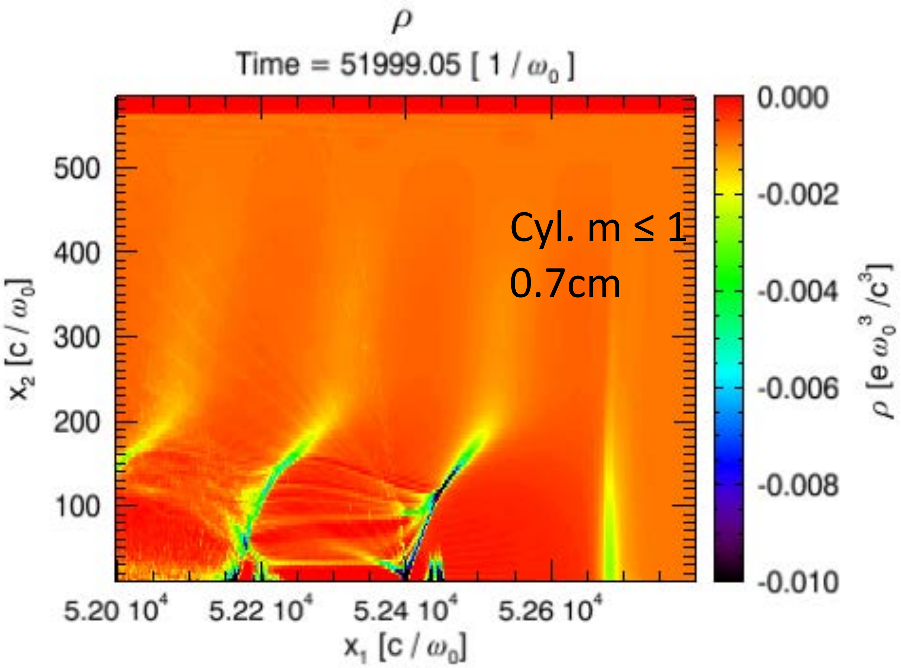

328pC 1.55 GeV

UCLA

PWFA-LC Stage: Matched beams lead to ion motion

Drive Beam : σ_r = 3.45 μm , σ_z = 30.0 μm , N_1 = 3.0 x 10^{10} , ε = **100 mm·mrad**
Trailing Beam: σ_r = 0.1 μm (0.006 k_p^{-1}), σ_z = 10.0 μm, N_2 = 1.0 x 10^{10}, ε = **0.1 mm·mrad**
Distance between two beams : 115 μm; Plasma Density : 1.0 x 10^{17} cm^{-3}

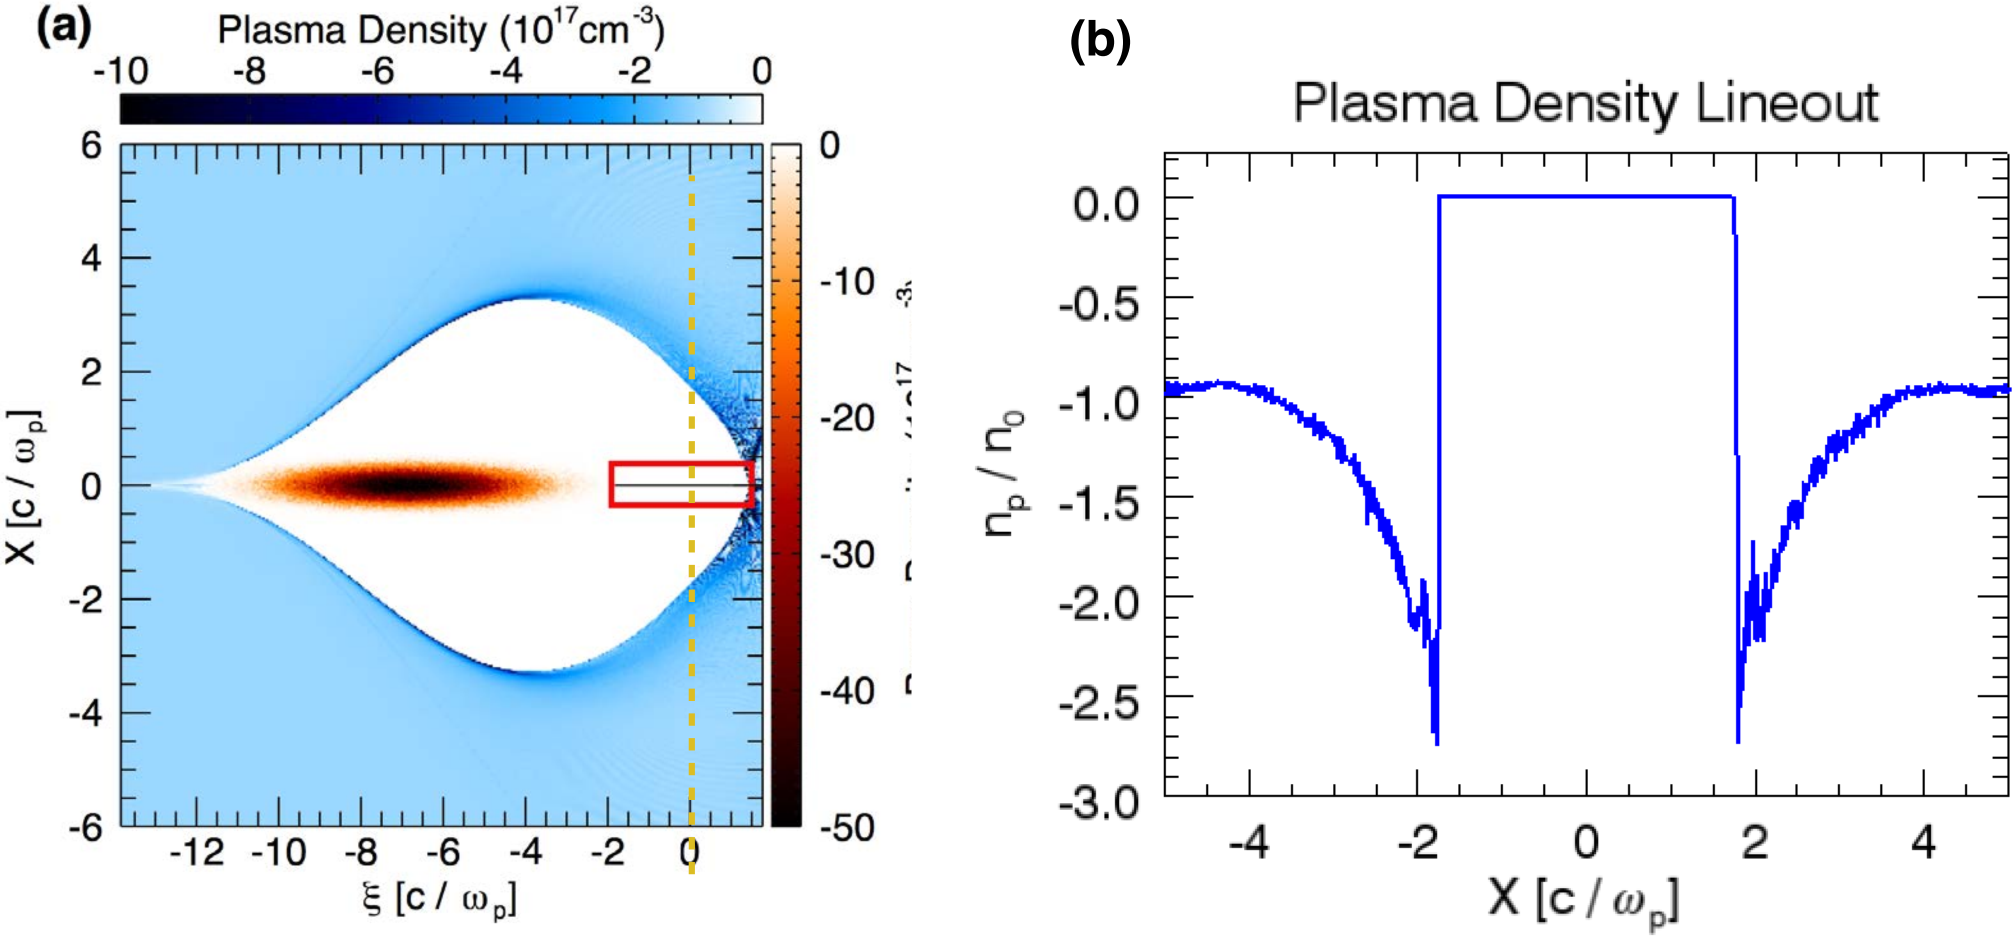

$$16384x16384x1024 = N_xN_yN_z$$

Simulation needs for PWFA LC stage

$L_\perp$

$\approx 20\frac{c}{\omega_p}$

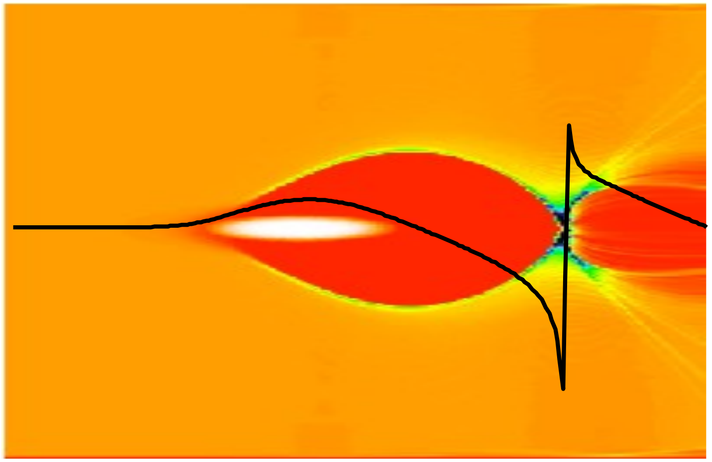

$$\Delta_t \approx min(\Delta_z, \Delta_\perp)$$

OSIRIS

$$\Delta_t \approx \gamma_b^{1/2}\Delta_z$$

QuickPIC

L_z

$\approx 25\frac{c}{\omega_p}$

$$\gamma_b = 50,000$$

LC Stage

$$\frac{L_z L_x L_y}{\Delta_z \Delta_x \Delta_y} N_{pcell} \frac{\gamma_b}{\Delta_t}$$ = Particle Pushes

$$\Delta_\perp \approx 5x10^{-4}\frac{c}{\omega_p} \qquad N_{pcell} = 2 - 20$$

Ion Motion Driven by Asymmetric Trailing Beam

Trailing Beam: σ_z = 10.0 μm , N = 1.0 x 10^{10} ,
σ_x = 0.463 μm , ε_{Nx} = **2.0 mm·mrad** , σ_y = 0.0733 μm , ε_{Ny} = **0.05 mm·mrad**
Y = 48923.7 (25 GeV), Plasma Density : 1.0 x 10^{17} cm^{-3}

$\sigma_x / \Delta_\perp$ = 75.9
$\sigma_y / \Delta_\perp$ = 12.0

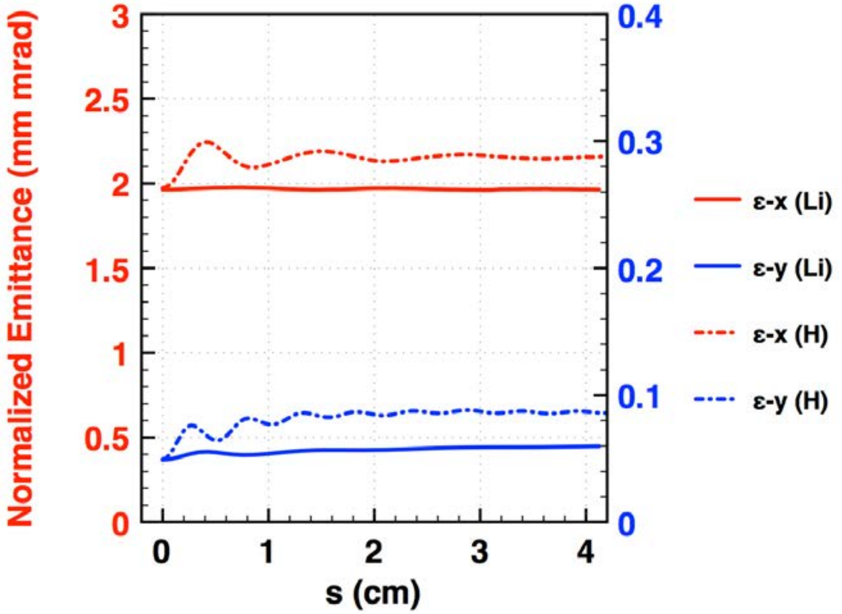

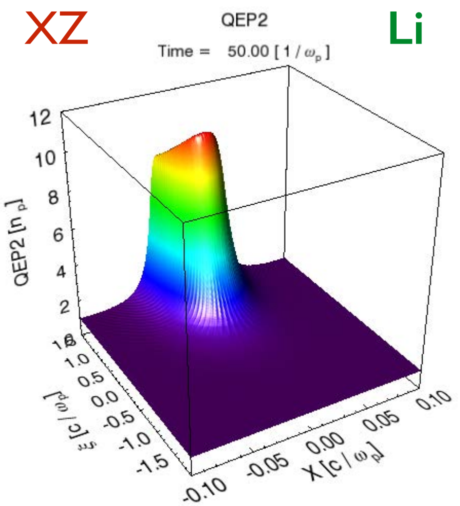

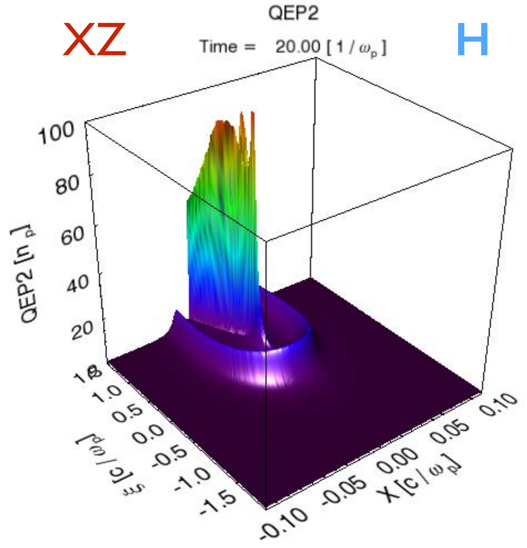

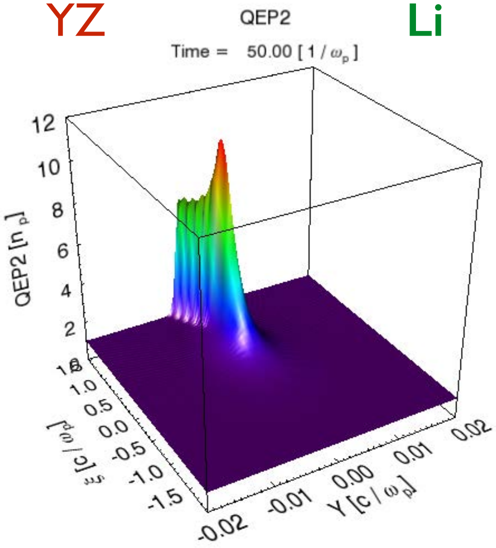

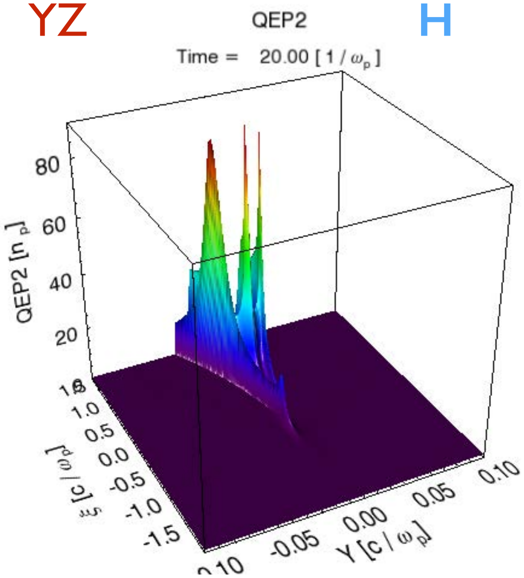

In Li, the emittance in x does not change, and in y direction it only increase by 20%.

In H, the emittance in x increase by 10%, and in y direction it increases by 70%.

UCLA

Osiris and QuickPIC: Rough estimates

Total Number of Particle Pushes

	Osiris 3D (8ppc)	QuickPIC (8ppc)
FACET II	7×10^{15}	1×10^{13}
PWFA-LC	1×10^{21}	5.6×10^{16}

Total CPU-Hours: assuming no load imbalance

	Osiris 3D (8ppc)	QuickPIC (8ppc)
FACET II	5.9×10^{5}	2.8×10^{3}
PWFA-LC	8.7×10^{10}	1.5×10^{7}

Exascale is not needed for FACET II

If ion motion does not lead to emittance growth then lower resolution simulations are possible

LC examples can be simulated with ~10-50 times less resources

UCLA

Accelerating positrons in nonlinear regime: Much harder than electrons but not impossible

Corde et al, 2015

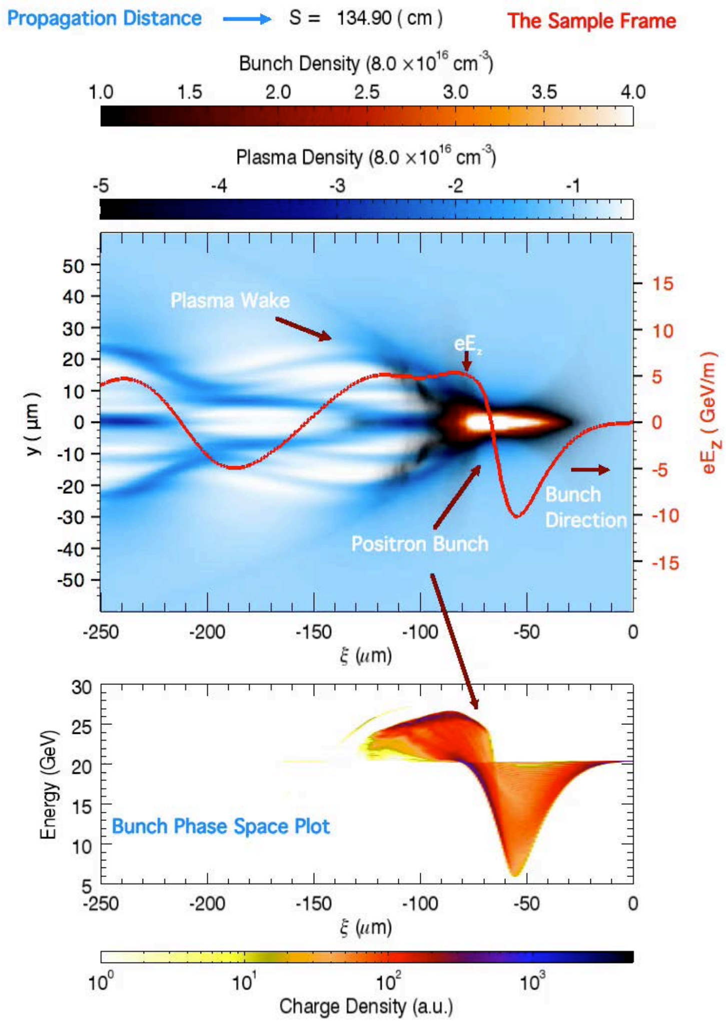

Drive Beam: σ_r = 70.0 μm , σ_z = 30.0 μm , N_2 =1.4 x 10^{10}, ε_N =(50,200) mm·mrad

Plasma Density: 8.0 x 10^{16} cm^{-3} (1.5 meters long)

Positrons in tail are accelerated in the wake of the head of the same beam.

Need a column of electrons on axis: Other recent work, Vieira et al. and Jain et al.

SLAC
NATIONAL ACCELERATOR LABORATORY
New SLAC National Accelerator Laboratory logo

Incomplete list of directions for future algorithmic work to reduce turn around time

- Adaptive mesh refinement (quasi-static/full PIC)
- Adaptive particle loading:Vary Npcell and/or particle merging and splitting (quasi-static/full PIC)
- Dynamic load balancing (quasi-static/full PIC)
- NCI mitigation (e-beams in lab frame and boosted frame) (full PIC)
- Boosted frame (more challenging than for lasers) (full PIC)
- Adaptive 2d and 3d time steps (quasi-static)
- Intel Phi and GPUs (quasi-static/full PIC)

Take advantage of new computer hardware

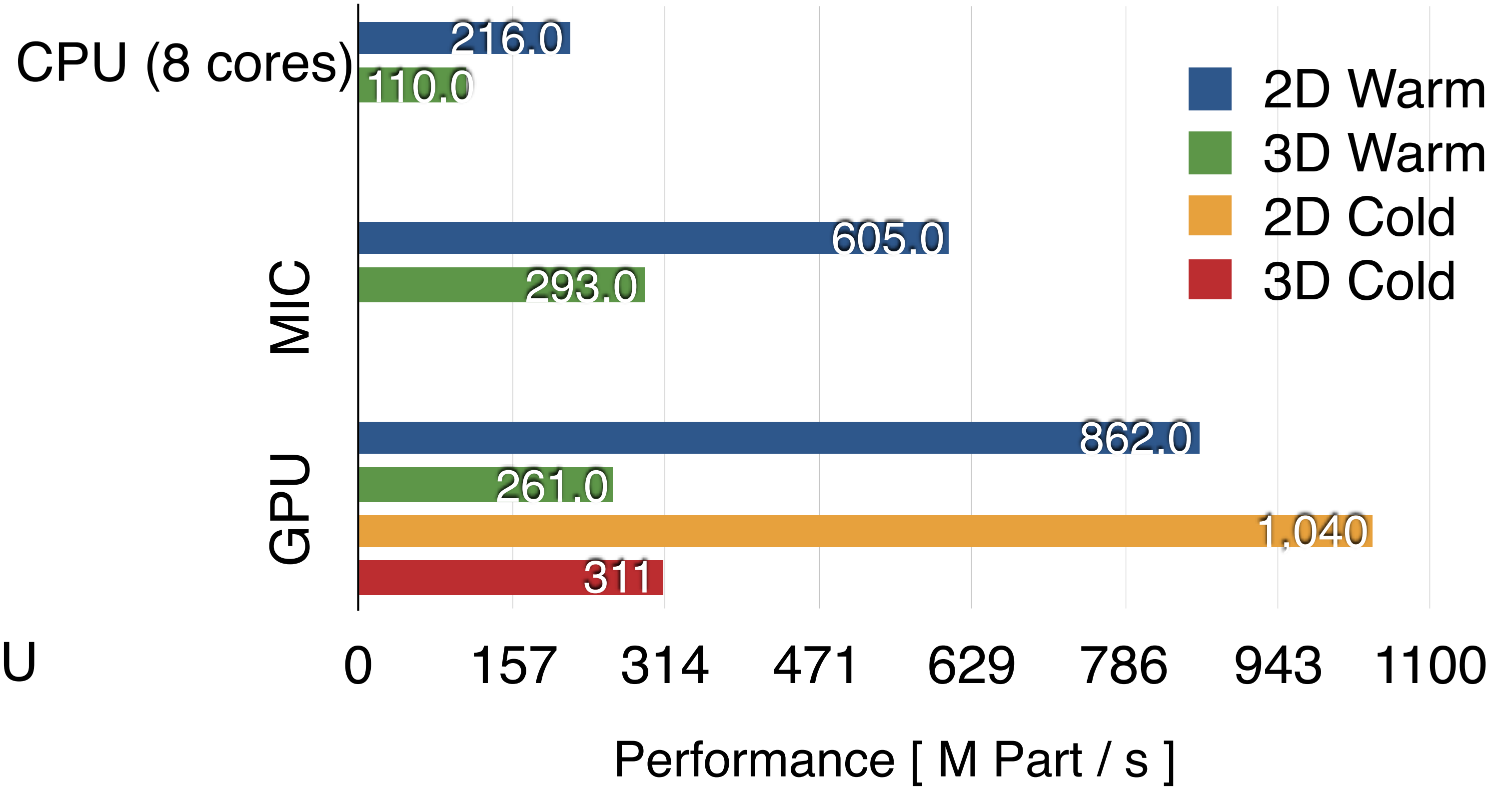

PIConGPU
PSC

e.g., OSIRIS is
already GPU and Intel
Phi enabled

Challenges (Opportunities): From a talk at LBNL Workshop

- PWFA and LWFA research are now focused on collider concepts that have multiple stages (10-100) that are each ~1meter in length.
- The challenges fall into a variety of areas:
  - Driver (particle beams and/or lasers)
    - Need development and design such they have a low cost for high average power and are efficient.
    - May need to develop methods to shape them (axially, transversely, chirp them etc.)
    - There analogies but also key differences.
- Interstage transport of the particle beams (emittance preservation) and injection of new drive beams.
- Final focus and interaction point: Oide limit, disruption, beamstrahlung, QED (OSIRIS?)

In my opinion the biggest challenge remains developing self-consistent beam loading scenarios for electrons and positrons (they don't have to be the same, e.g., use electron beam to accelerate electrons in blowout regime and lasers to accelerate positrons in a hollow channel) in a single stage.

- There are many options with decisions that are inter-related.
- Any scenario needs to be tested self-consistently over meter distances (including the evolution of the driver).
- It is my sense that the two scenarios being discussed most seriously are: 1. Nonlinear wakes in the blowout regime and 2. Linear wakes in a fully or nearly hollow channel.