Plasma Accelerator Simulation Using Laser and Particle Beam Drivers

C.G.R. Geddes, LBNL
F.S. Tsung, UCLA

27 November 2012

SciDAC-2&3 Compass

BERKELEY LAB
LAWRENCE BERKELEY NATIONAL LABORATORY
UCLA
U.S. DEPARTMENT OF ENERGY
Office of Science

Approach: plasma wave accelerator structure, laser & particle beam evolution excited by laser or particle beam

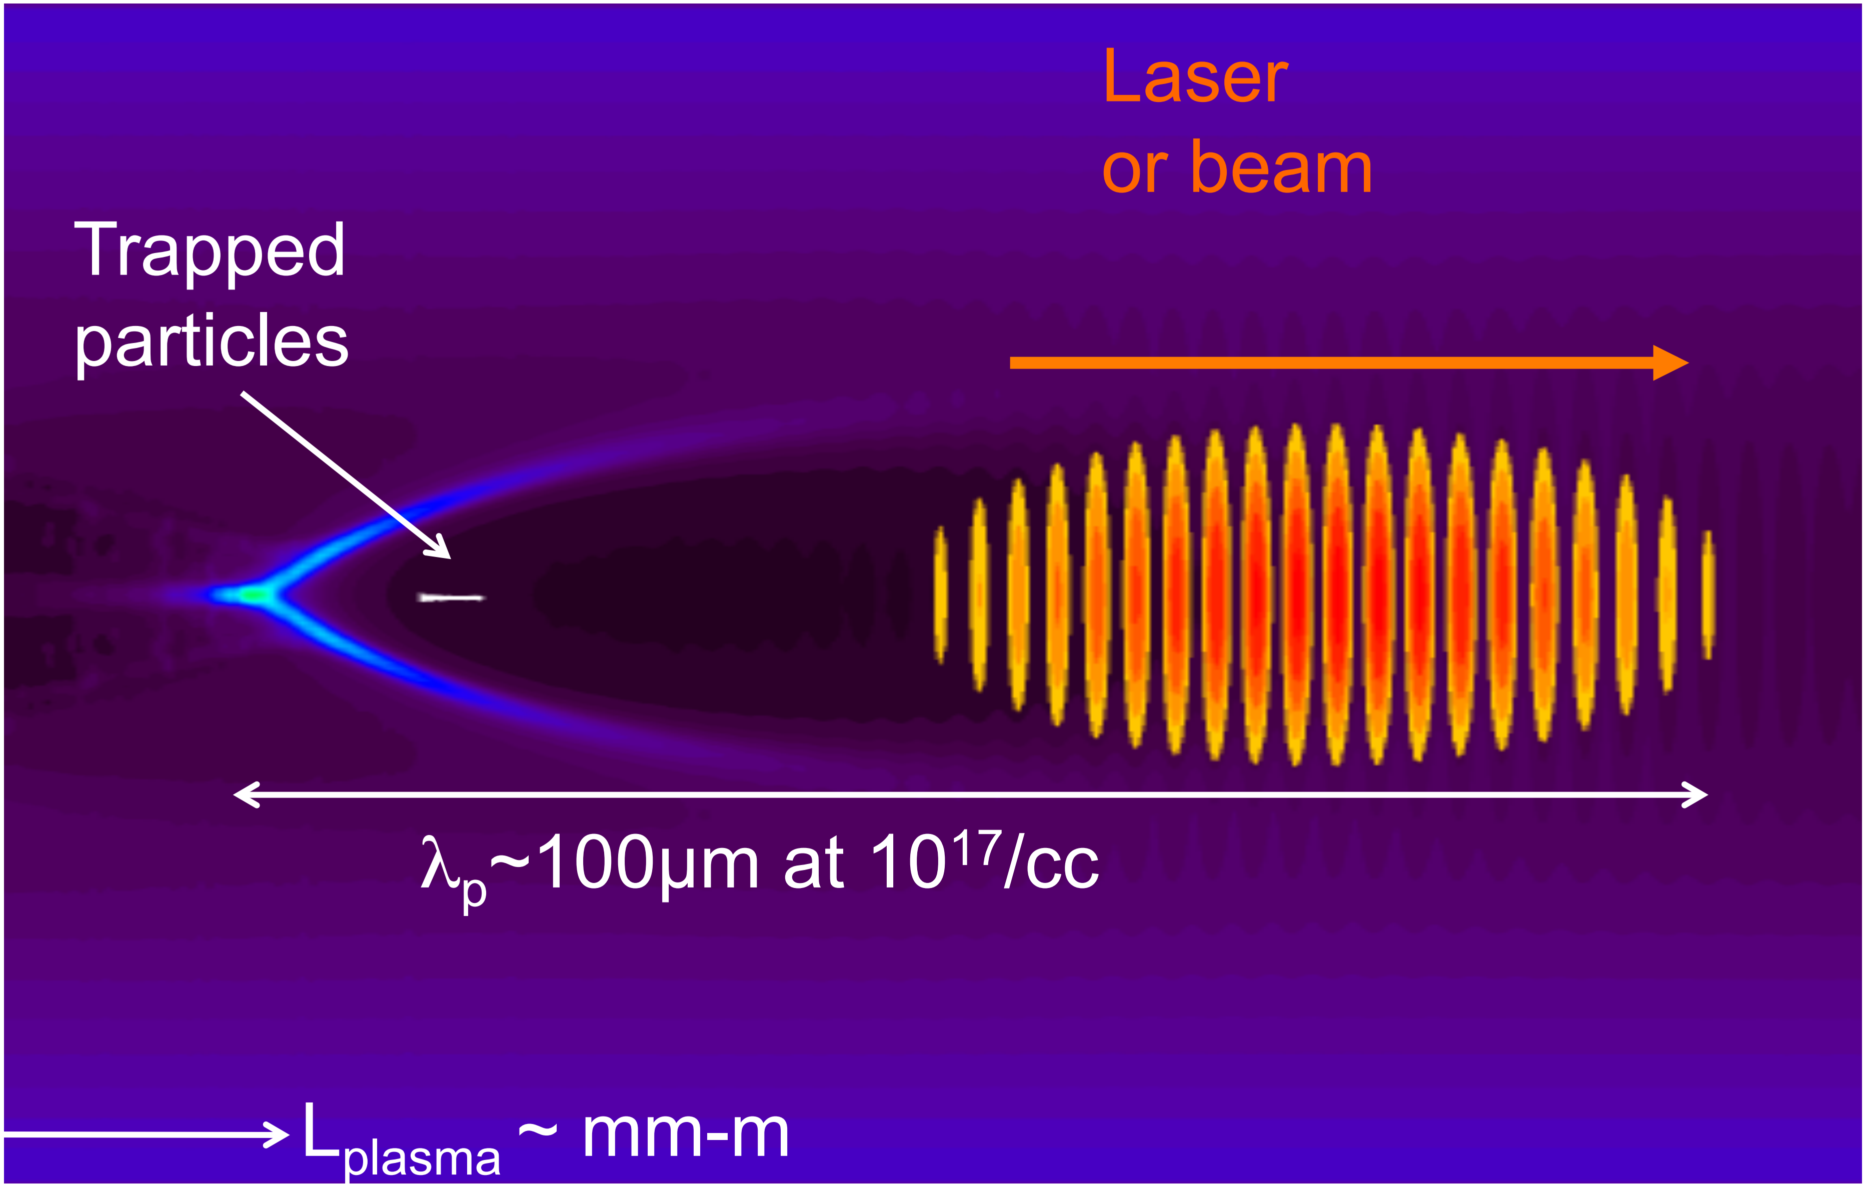

Approach: plasma wave accelerator structure, laser & particle beam evolution excited by laser or particle beam

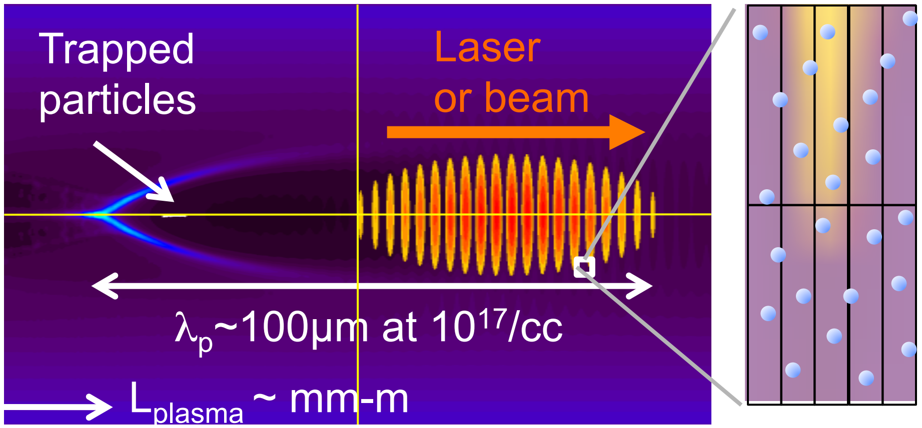

Tajima & Dawson PRL 1979; Esarey et al. RMP 2009

Explicit PIC schematic

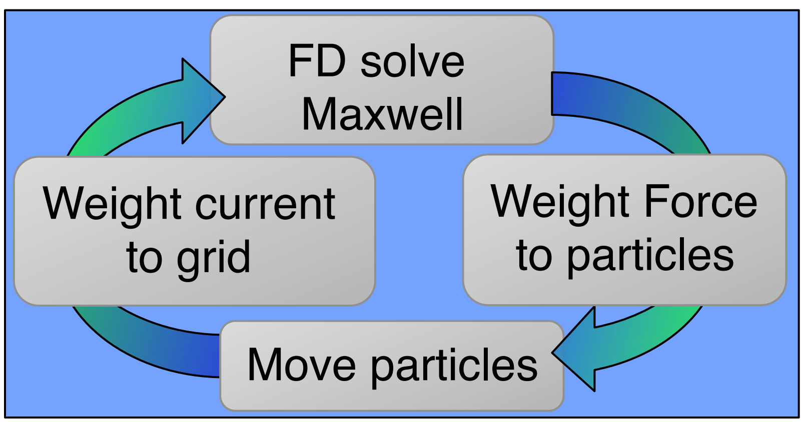

Simulation reviews :ScDAC Review 2009, Huang J. Phys C.S. 2009

- Electromagnetics + particle dist.
- Explicit PIC most common
  - Mhour for cm-scale, base resolution
  - Domain decomp. ~ 50cell^3/core
  - Weak scaling limited by I/O: otherwise OK to 100,000+cores
- Long plasmas – special codes
  - Boost: explicit in moving frame
    - Explicit scaling but more I/O
  - Envelope averages over laser period, Quasistatic adds slow evolution
    - Scaling to 1000's of cores
- For 2017: similar + Vlasov, MHD plasma formation, rad/e+

Plasma Accelerator Simulation Using Laser Drivers
M558 case study

Related: Continuing Studies of Plasma Based Accelerators (MP113)

BERKELEY LAB
LAWRENCE BERKELEY NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY
Office of Science

- LOASIS: C.G.R. Geddes, C. Benedetti, S. Bulanov, J.-L. Vay, E. Esarey, S. Rykovanov, C.B. Schroeder, L. Yu, W.P. Leemans

- Visualization: W. Bethel, O. Rübel, G. Weber, Prabhat

TECH-X CORPORATION

- E. Cormier-Michel, B. M. Cowan, E. Hallman, N. Naseri, K. Paul, J.R. Cary

Objective: model 10 GeV stages for the LBNL BELLA laser[@] – record rep-rated PW laser for accelerator science

Modeling from first principles challenging because of scale separation

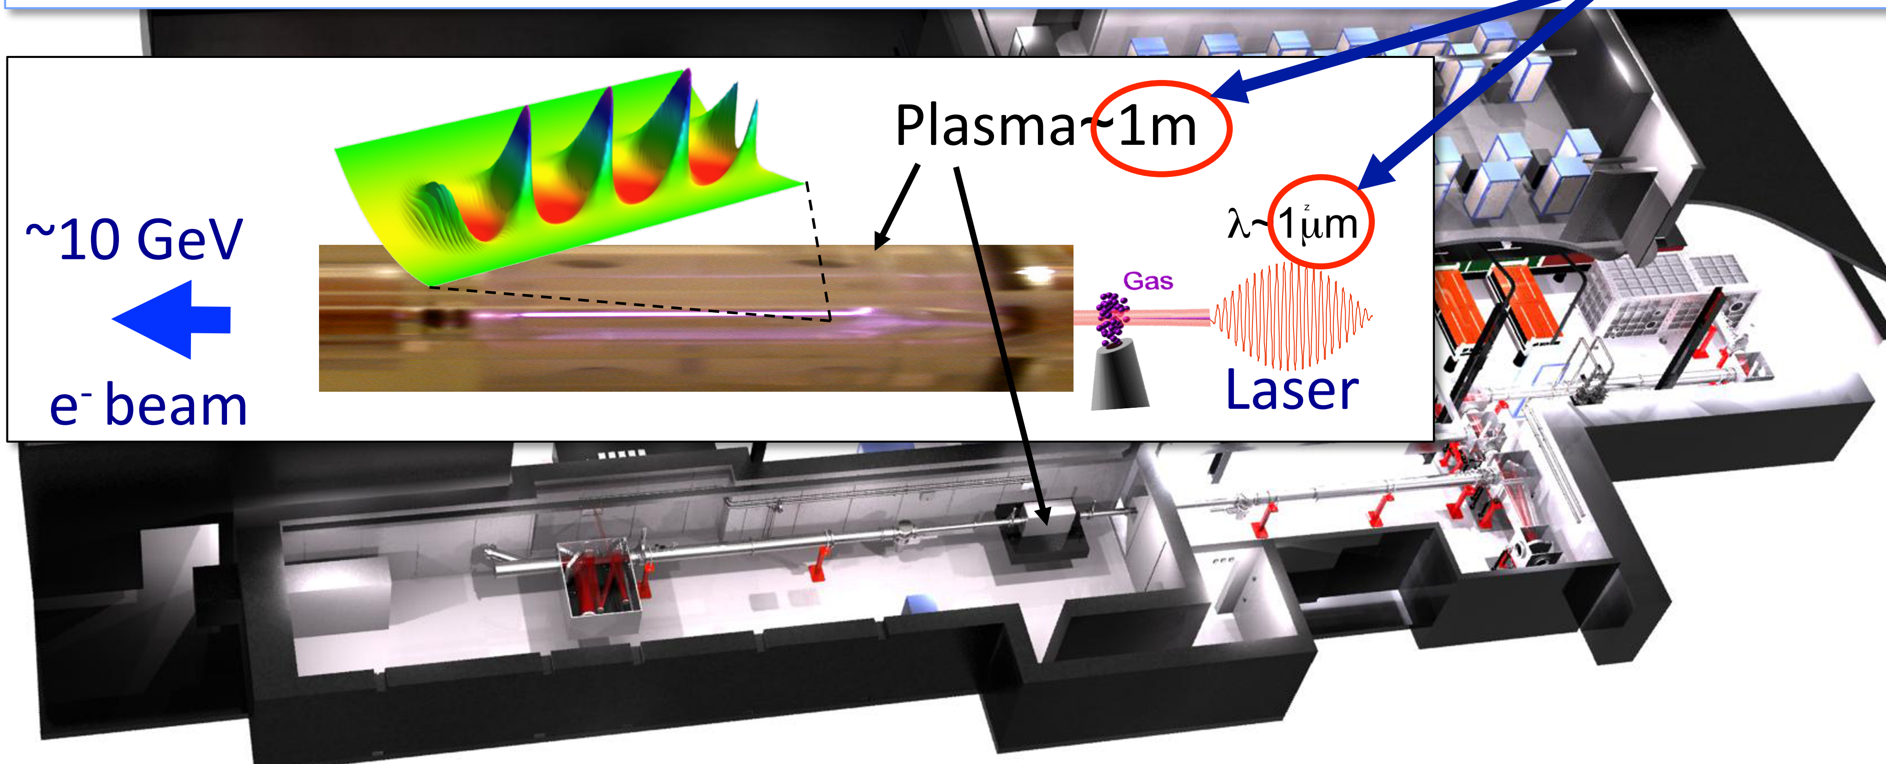

2012: 2/3D simulations @ Mhours on Hopper with advanced models
2017: Modeling low emittance beams*: ~ 3-10x resolution -> 30x cells

[@] Leemans et al, Proc. AAC 2010, *Plateau PRL 2012

Objective: Technology for future laser-plasma collider concepts & FEL/gamma source drivers

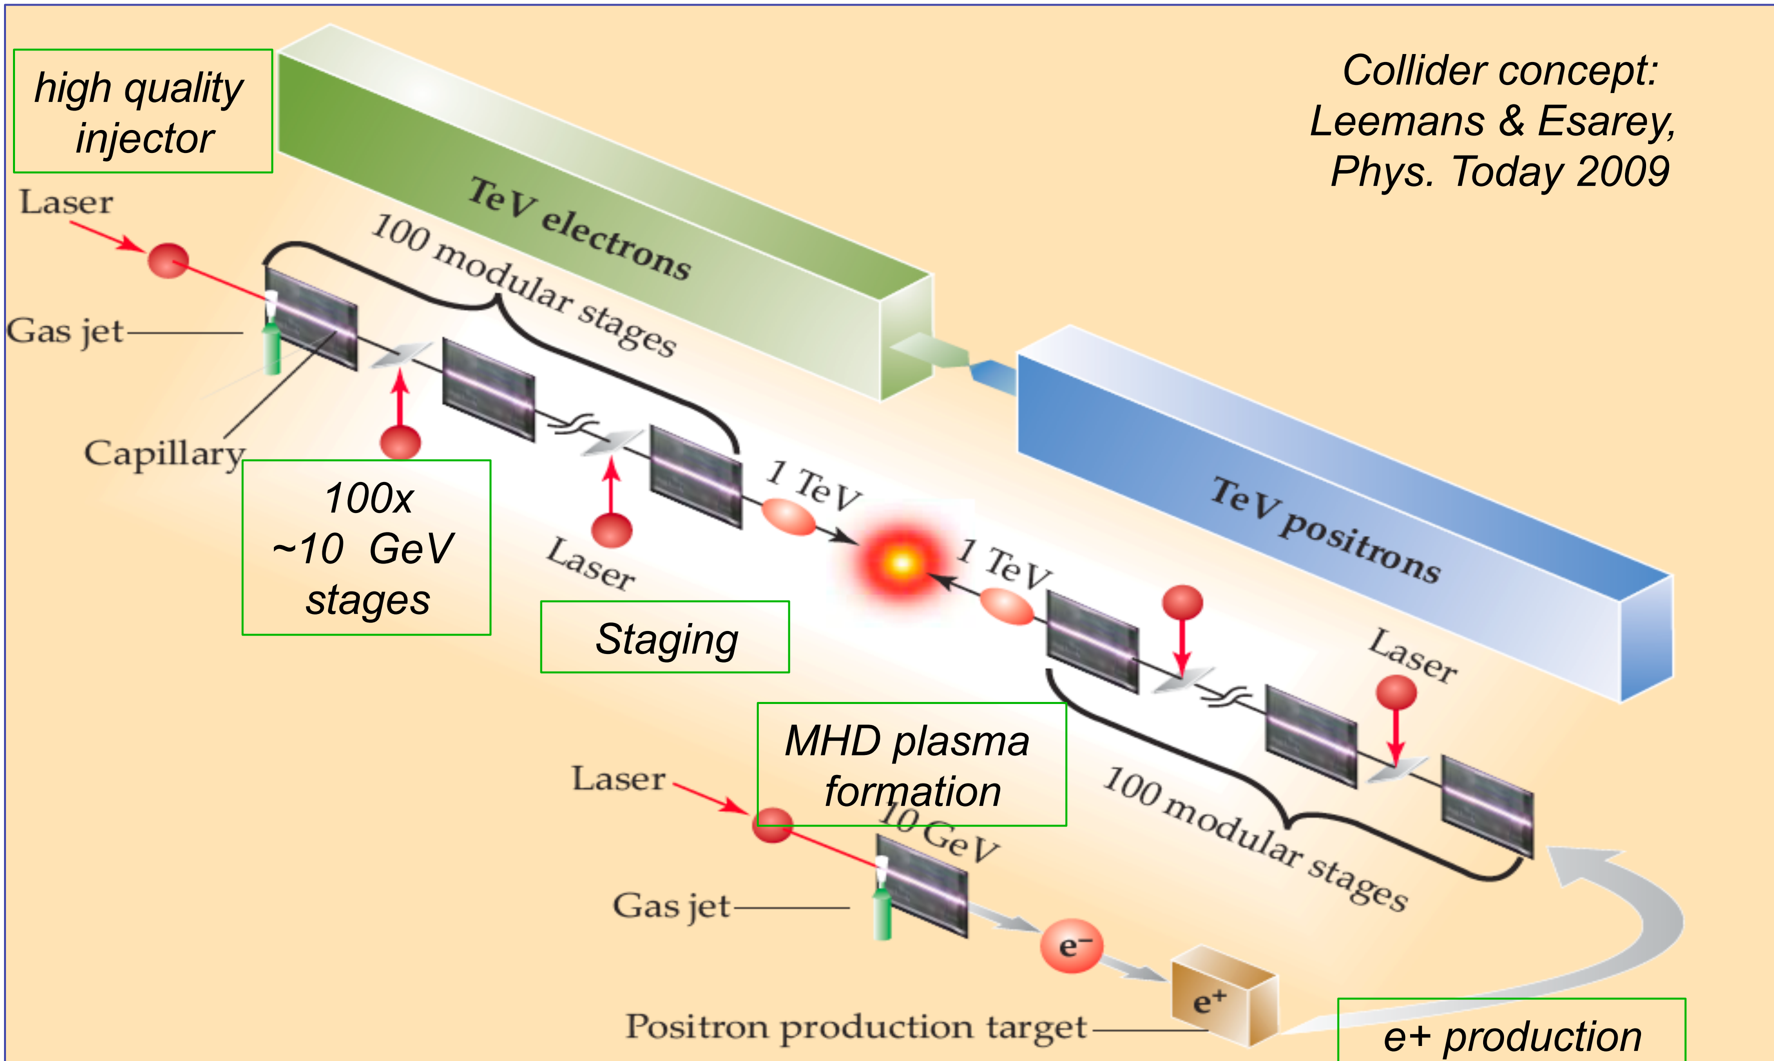

2017: Multiple stages → 30-100X cost

Add physics – e+, radiation, scattering

Codes and Algorithms

Tasks	Use	Other notes
WARP*	1/2/3D Explicit electromagnetic PIC Boosted frame.	ES, FFT, multigrid AMR, RF cavities, surfaces with emission, ionization
VORPAL*	1/2/3D Explicit electromagnetic PIC & fluid, Envelope, boosted frame, beam frame Poisson, Ionization	RF cavities, surfaces, collisions
INF&RNO*	2D cylindrical envelope PIC& fluid (3D-like focusing), boosted frame	
VDSR	Radiation generation and particle tracking. Thomson, Compton, and radiation reaction.	

*Common elements: relativistic, absorbing boundaries, high order spline particles, parallel I/O

- related: OSIRIS (explicit PIC) and QuickPIC (envelope Quasistatic) from UCLA – see Tsung

LPA codes scale well to 100k cores limited by I/O

Warp strong scaling good to >100kcore
3D EM, periodic1024x1024x1024 grid, 8ppc

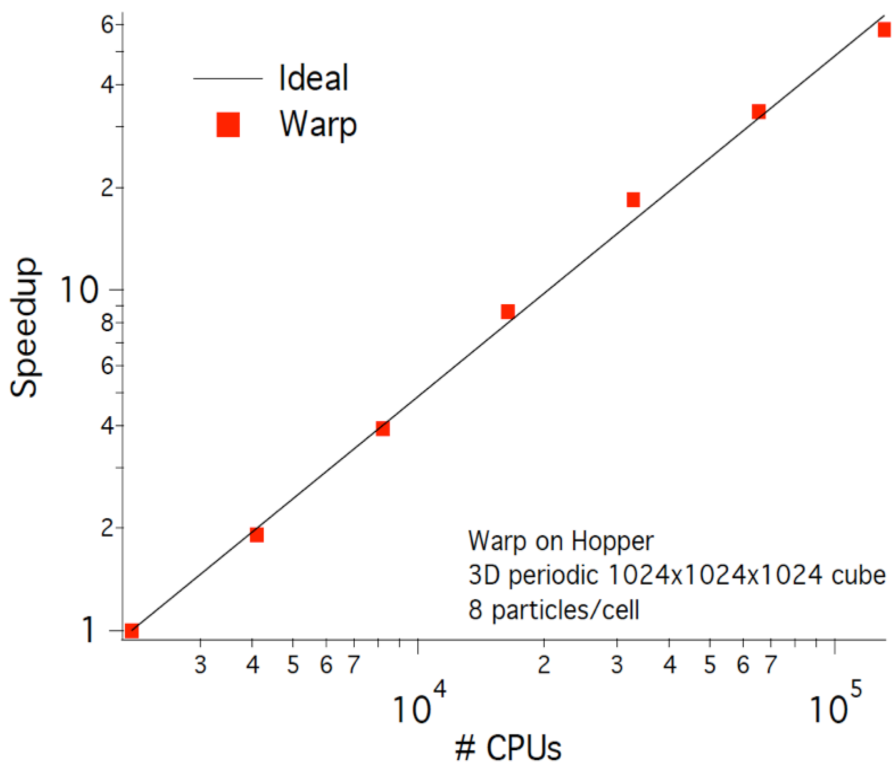

Vorpal HDF5 parallel I/O flat
Franklin data - Hopper is similar

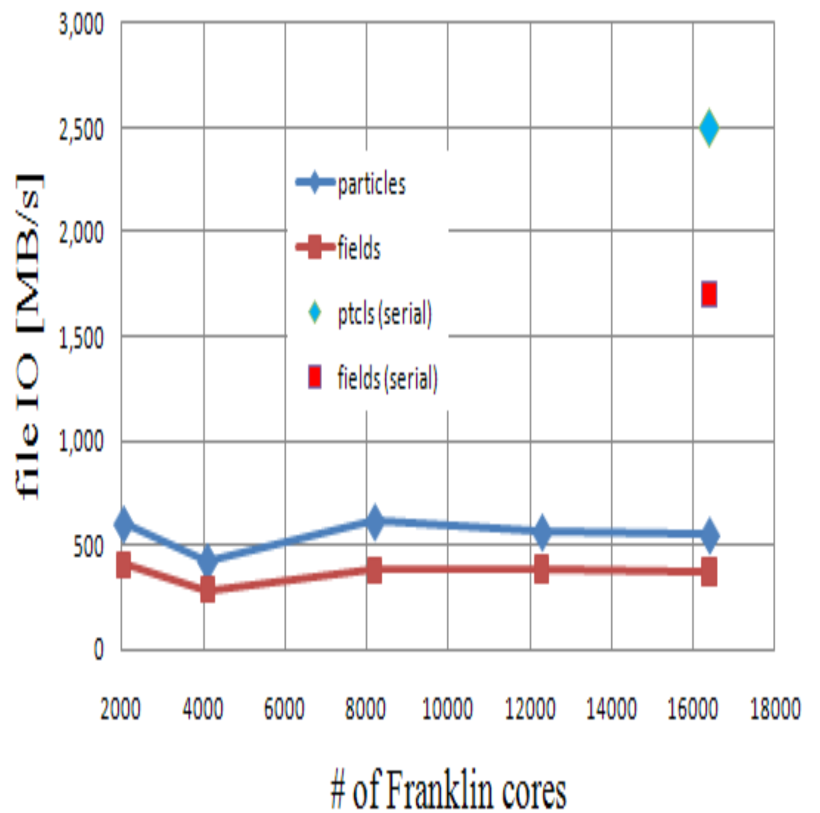

•VORPAL Scales similarly to WARP

•INF&RNO scales to ~4kcores – adequate for 2D r-z
limit is envelope tridiagnoal solve

Current HPC enables leadership LPA science

10 GeV beams
Vay et al, PoP 2011

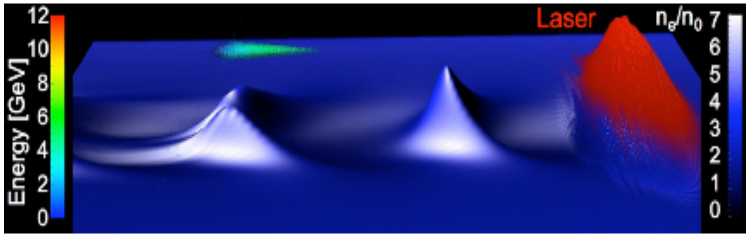

Controlled injection
Cormier – SciDAC vis award 2011

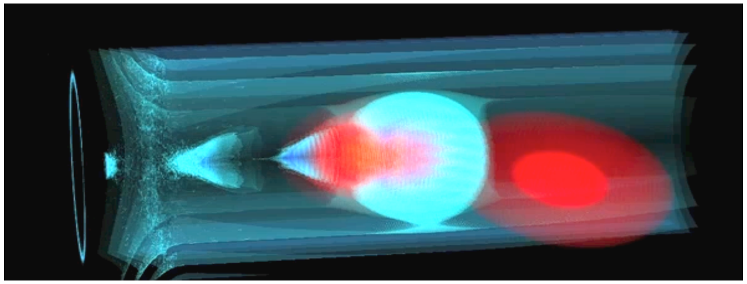

Low emittance beams
Plateau et al PRL 2012

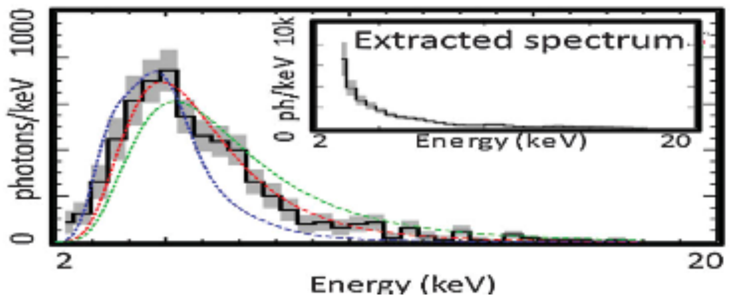

- Hours: 15 Million/year (Hopper)
- Cores:
  - Explicit 2D ~ 1 kCore, 3D ~ 16 kCore
  - Envelope ~ 4 kCore
  - 10's of simultaneous runs
  - 10-100 hr/run
  - Scaling limits: I/O & queue
- 0.1 GB/core, save 4TB/run
- Archival data: 150 TB

HPC needs 2017
Resolve low emittance beams, 10's of stages, transport

- Hours: 500 Million/year. Driven by:
  - Resolution – for low emittance and energy spread
  - Length – for staged approach to collider
  - Added physics – MHD plasma formation, radiation, e+....
- Cores: weak scaling dominant
  - 50-500k, problem and network dependent
  - Staging: long queue@ 10-50 kCore for 100m plasma
  - 10's of simultaneous runs at few kCore – 10's of kCore
- Memory approx. 0.1 GB/core + few w/GB's (analysis)

HPC needs 2017
I/O is a key constraint

- Scalable I/O needed – keep below ~10% time
  - Parallel HDF5 I/O scaling near flat currently
  - Individual processor files / custom solutions cumbersome
- ≥ 10 GB/s needed – determines scaling
  - Shared data: 100 TB/run, total of 600 TB. Fast NGF needed.
  - Archival data: 5000 TB
- Use inline analysis, subsetting to keep growth ~10-30x
  - Require a few full checkpoints for restart
- Common need of high performance tools – NERSC role?

HPC needs 2017
GPU/manycore requires Network bandwidth, preparation

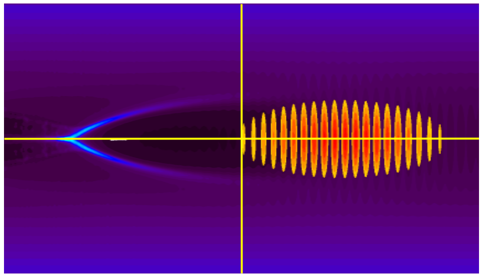

- PIC is not compute intensive:
  - µs/ptcl on CPU, and GPU can be 10x faster
  - Network bandwidth key : at 50^3 cells, exchange ~ 20%
- Codes run on GPU test beds – NERSC, NIU
  - Methods suitable
  - AVX also attractive
- Bottleneck: specialized development for each system
  - Test bed for architecture a year or more ahead
  - Compiler optimization + code work on structured access

HPC needs 2017
Software - Performance of HDF5 and Python are key

- Libraries built with codes:
  - MPI/IO, HDF5 parallel, Trillinos, Mercurial
  - Shared libraries
- Need long queues and memory for serial data analysis
  - IDL, VisIT, file combination
- Desire run scheduling/monitoring access – no logout
  - Limited privileges?

HPC needs 2017
Services – NERSC strength

- General consulting/support is a NERSC strength
- Parallel visualization is required for large runs
  - Ongoing collaboration
  - Need to capture analytics/math from serial tools (IDL, etc.)
- More premium queue options would be of interest
  - Allow only some fraction as premium?
- Support for optimization – especially for GPU
- I/O support needed to avoid effort duplication

Summary

- Advanced LPA concepts need ≥30x resources by 2017
  - Resolution – for low emittance and energy spread ~30x
  - Length – for staged approach to collider ~30-100x
  - Added physics – MHD plasma formation, radiation, e+....
- Codes scale to >100kCore
- Algorithms suited to GPUs but support/testbed needed
- I/O support needed to avoid effort duplication
- Parallel visualization needed for large runs
  - General consulting/support is a NERSC strength to maintain