UCLA

Case Study: Continuing Studies of Plasma Based Accelerators (mp113)

- PI: W. B. Mori (UCLA)
- Presenter: F. S. Tsung (UCLA)

Users: W. An, A. Davidson, V. K. Decyk, (UCLA), J. Vieira, L. Silva (IST), W. Lu (UCLA/ Tsinghua)

INSTITUTO SUPERIOR TÉCNICO

UCLA

HEP Requirements:
Continuing Studies of Plasma Based Accelerators (mp113)
(PI: W. B. Mori, Presenter: F. S. Tsung)

- An alternate scheme to accelerate particles using plasmas is the Plasma WakeField Accelerator (PWFA) concept where a particle beam is used to drive a plasma wave which is used to accelerate particles.

- Both laser and particle beams excite a highly nonlinear plasma wave (in which all of the plasma electrons are blown out by the driver), and particle methods are ideally suited to study these systems. The injection process, where some of the background plasma electrons become resonant with the plasma wave and become accelerated is another process that can only be studied by PIC codes (more on this later).

- In the past few years, there have been many high impact experiments, and particle-in-cell codes have given insights to these experiments andled to many high profile publications (see right)

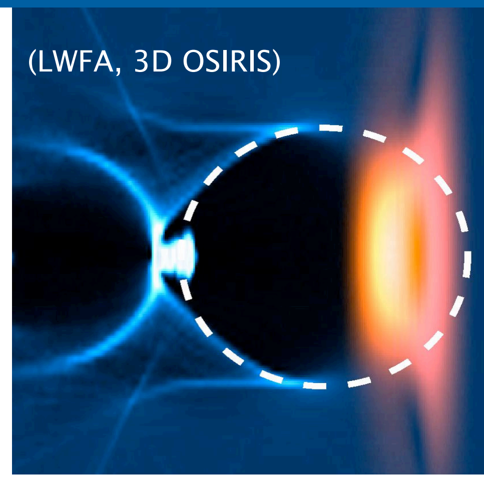

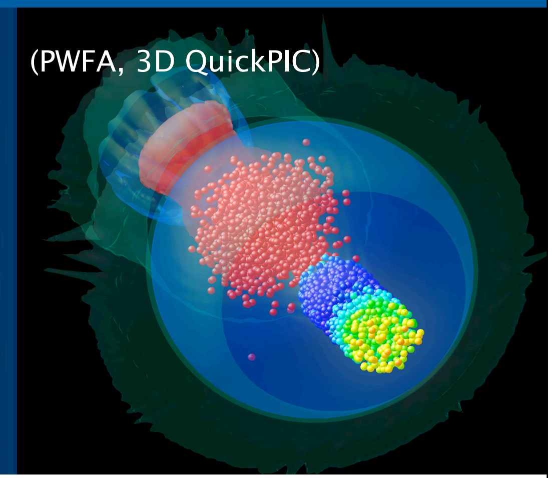

Facilities for ACcelerator science and Experimental Test

Beams

Energy	24 GeV
Charge	3 nC
Sigma z	17 μm
Sigma r	< 10 μm
Peak Current	22 kAmps
Species	e^- & e^+

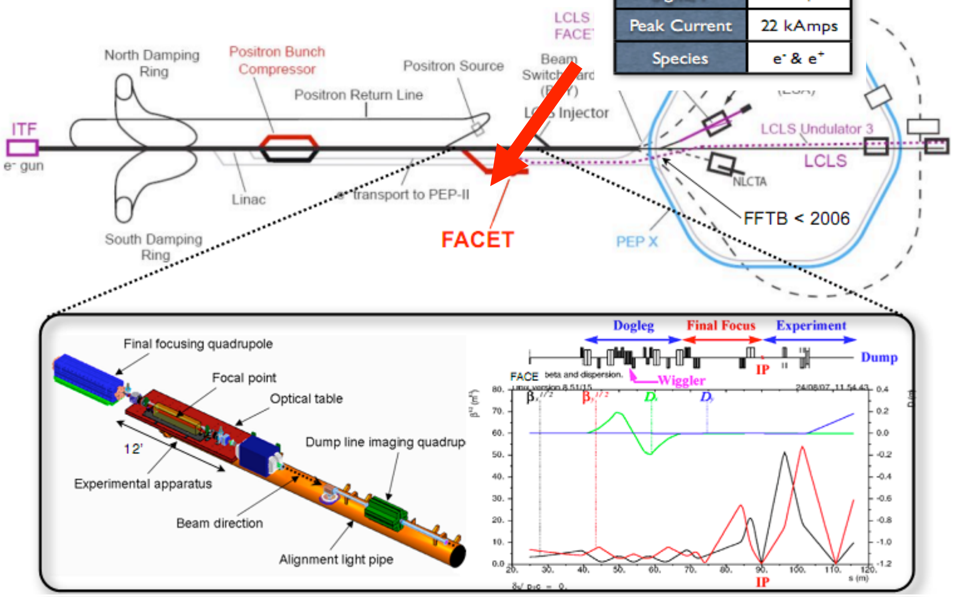

The PWFA-LC illustrates the key questions that must be answered:

- High beam loading with both e- and e+.
- Small energy spread (required to achieve luminosity and luminosity spectrum)
- Small emittance and small emittance dilution (required to achieve luminosity).
- Staging of multiple PWFA's
- Source of "dark current" in current FACET experiments

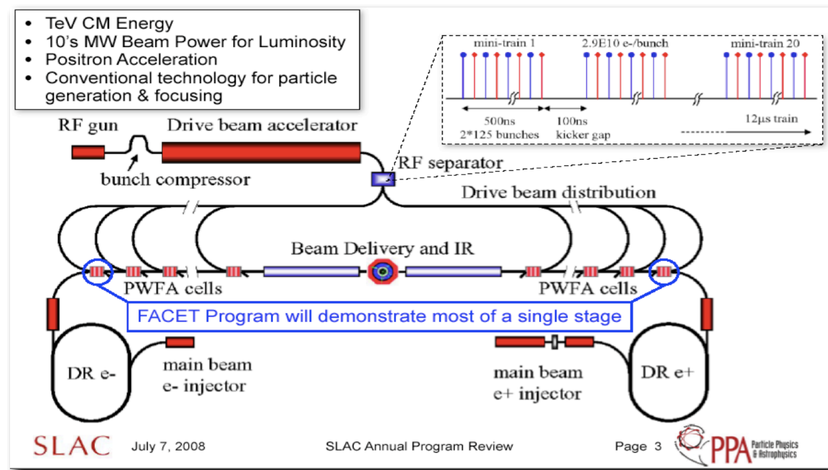

FACET is a new facility to provide high-energy, high peak current e⁻ & e⁺ beams for PWFA experiments at SLAC, the goal is to achieve high efficiency, with low energy spread and low emittance. (In 2006 this facility demonstrates energy doubling in 1 meter using a long beam, and it is in the process of creating a two bunch experiment to produce high quality e- beams)It is also a testbed for the multi-stage collider concept (shown on right)

Proton driven wakefield accelerator (PPA) collaboration

PDPWFA experiment at CERN

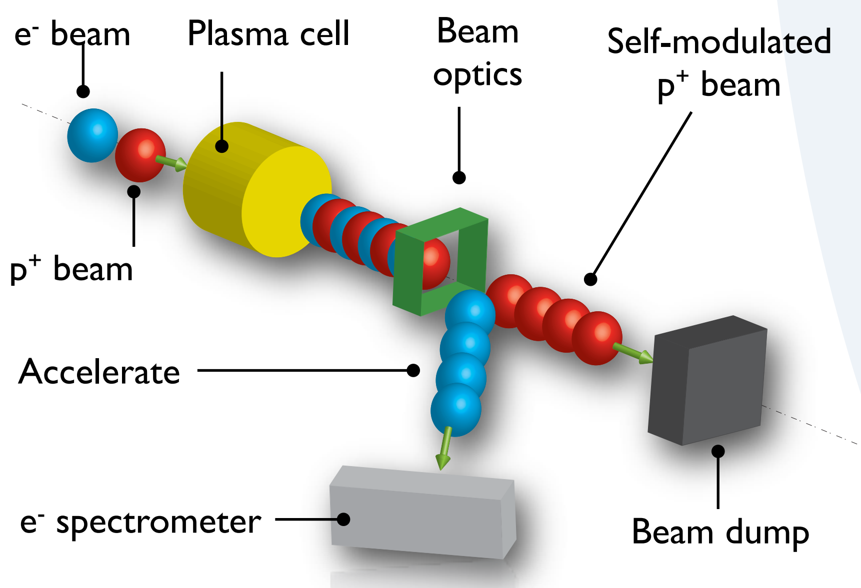

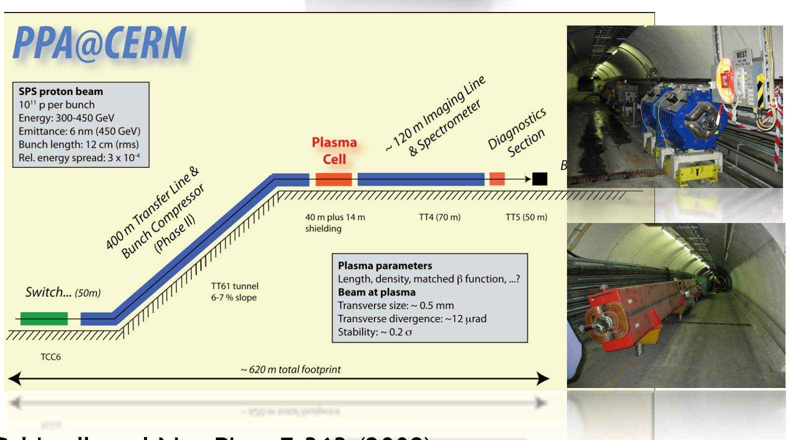

*A. Caldwell et al Nat. Phys. **5** 363 (2009).

Proton Driven Accelerator (PPA):

- Proton beams @ CERN is the highest energy beam on earth (7TeV), and an attractive driver for PWFA's.
- Existing proton beams have much longer pulse lengths (10cm) than the plasma skin depth (~100μm). These beams can undergo self-modulation instability (and hosing, which cannot be studied in r-z) (see below). In PWFA experiments @ SLAC, both the driving beam and the witness beam sits within the first plasma period.
- The UCLA and the IST group studying the possibility of using the proton beam @ Fermilab (120GeV) and CERN (respectively) to accelerate particles.

Full 2D PIC in 10-20 meter long plasma

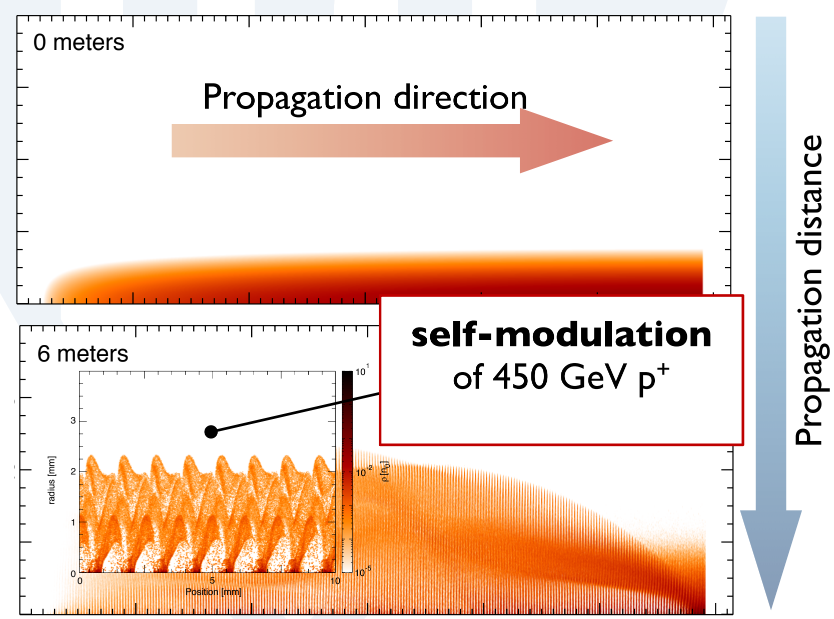

osiris 2.0

osiris framework

- Massivelly Parallel, Fully Relativistic Particle-in-Cell (PIC) Code
- Visualization and Data Analysis Infrastructure
- Developed by the osiris.consortium
- UCLA + IST
- >80% efficient for ~300,000 cores.
- Achieved >30% peak speed on full Jaguar in 2011

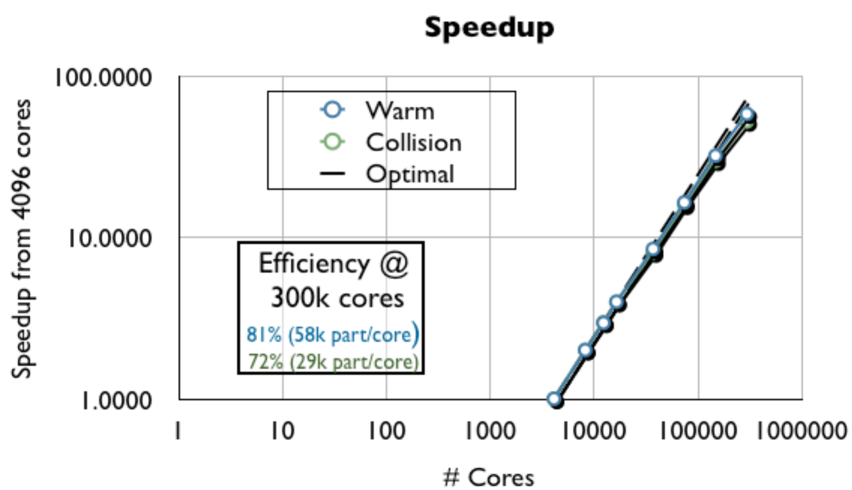

Ricardo Fonseca: ricardo.fonseca@ist.utl.pt
Frank Tsung: tsung@physics.ucla.edu

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

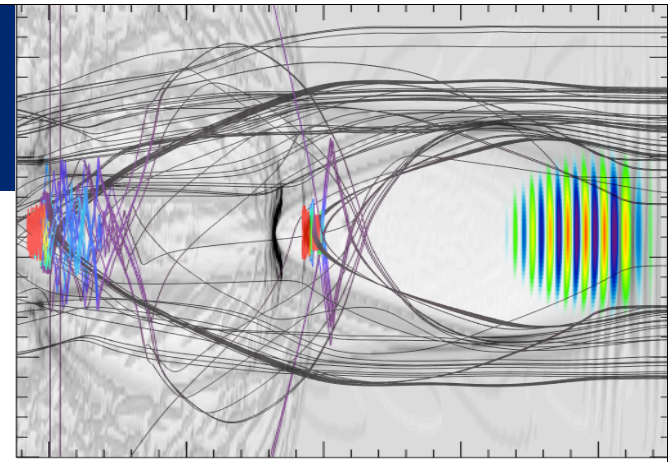

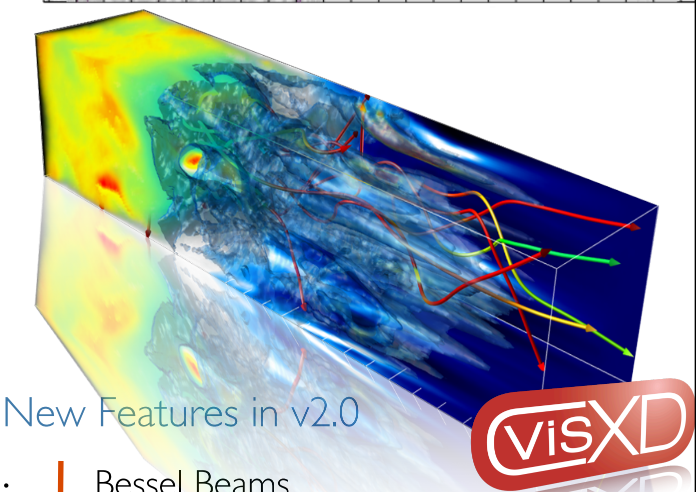

New Features in v2.0

- Bessel Beams
- **Energy Conserving Algorithm**
- **Multi-dimensional Dynamic Load Balancing**
- **OpenMP/MPI hybrid parallelism**
- PML absorbing BC
- **Higher order splines**
- Parallel I/O (HDF5)
- Boosted frame in 1/2/3D

UCLA

QuickPIC

- QuickPIC[1] is a 3D parallel Quasi-Static PIC code, which is developed with the framework UPIC[2].
- Pondermotive guiding center model for laser driver (laser solved on the Rayleigh length scale).
- Realisms such as field ionization and radiation damping are included.
- The pipeline algorithm and 2D domain decomposition ensure the code scale more than 30,000

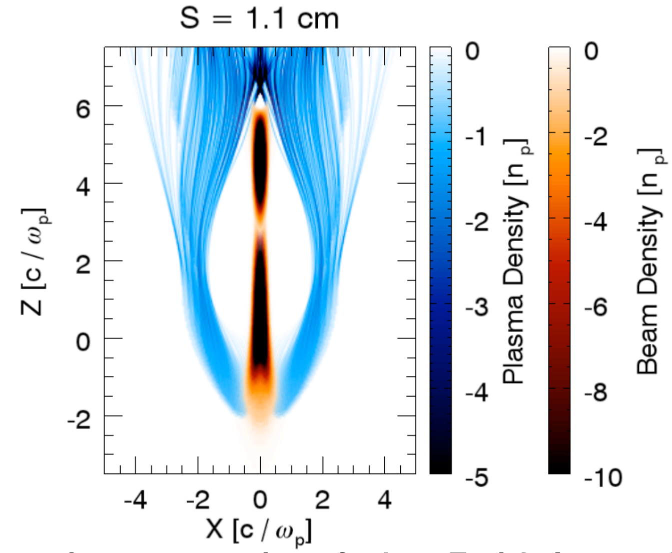

Simulation result of the Field Ionized Plasma Density in a Two-Bunch PWFA using QuickPIC

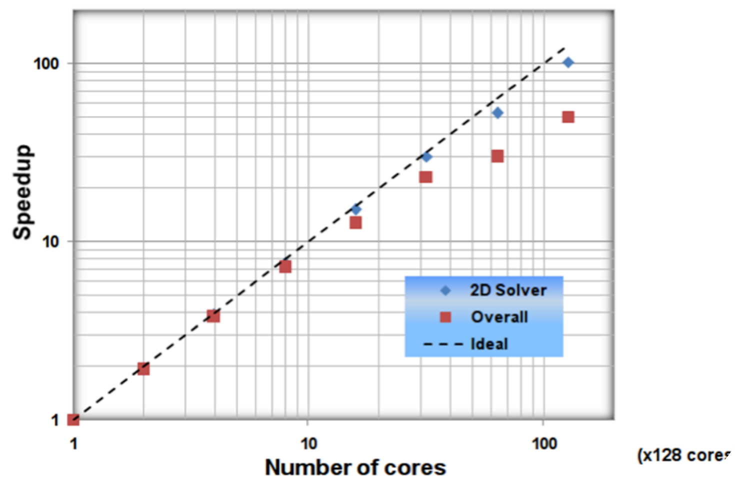

[1] C. Huang et al., J. Phys.: Conf. Ser. 46, 190 (2006).
[2] V. K. Decyk, Computer Phys. Comm. 177, 95 (2007).

Current HPC Requirements (3D single-stage FACET PWFA Simulation)

- 3D high resolution QuickPIC Simulations of a single FACET stage:
  - 137billion grids (2,048x8,192x8,192) -- 1TB/ grid quantity
  - 268 million particles per 3D slice (for plasmas)
  - 30 million beam particles ~1GB/particle data
  - 1 million CPU hrs total (. 5m-1m long plasma)
  - Running on 16-32K cores on Hopper/Jaguar

	2012 Usage
Total Computation Hrs	26M total, 12M for PWFA
Typical # of cores	>16,000
Maximum Number of Cores That Can Be Used for Production Runs	~300,000
Data read/written per run	2TB data 5-10TB restart
I/O bandwidth	5-10GB/sec (very rough estimate, it is not a bottleneck now)
Percent of runtime for I/O	5-10% (mostly due to restart)
Shared filesystem space	
Archival data	10-15TB/year
Memory per core	0.3GB
Aggregate Memory	5-10TB

2017 HPC Requirements (Ionization Seeded Proton PWFA in 3D)

1 meter, 3D ionization seeded Proton PWFA.

plasma density = 6.9 10^{15}

- [500,000 x 133 x 133] cells (71GB/grid quantity)
- 280 billion particles (14.5 TB particle data, 15TB total)
- 6.25 million timesteps.
- 1.78*18 particle pushes --> 49.4 million hours (using 100ns/(particle*step) (using the SSE)

We hope to address:

- Ionization trapping/injections & beam loading in proton driven PWFA's
- Instabilities associated with the proton beam (e.g., self modulation, hosing)

	2012 Usage
Total Computation Hrs	50M each (300M-400M total)
Typical # of cores	>200,000
Maximum Number of Cores That Can Be Used for Production Runs	*
Data read/written per run	~10TB data 15-20TB restart
I/O bandwidth	> 10GB/sec
Percent of runtime for I/O	5-10% (mostly due to restart)
Shared filesystem space	
Archival data	> 100 TB/year
Memory per core	0.1GB
Aggregate Memory	> 15 TB

* depending on whether CPU or GPU's are available in 2017

OSIRIS on SSE -- Velocity Integration & Current Deposition

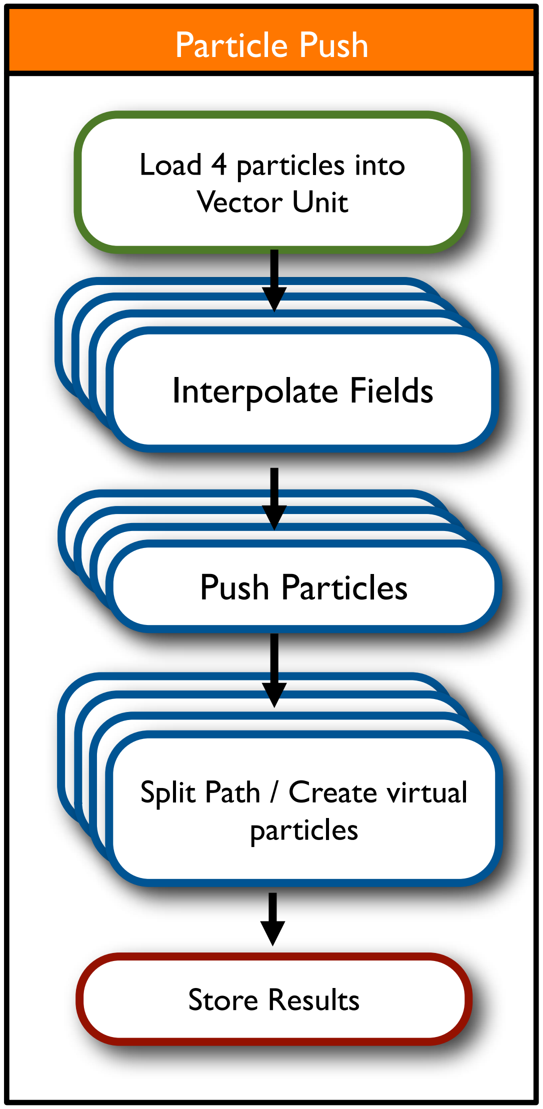

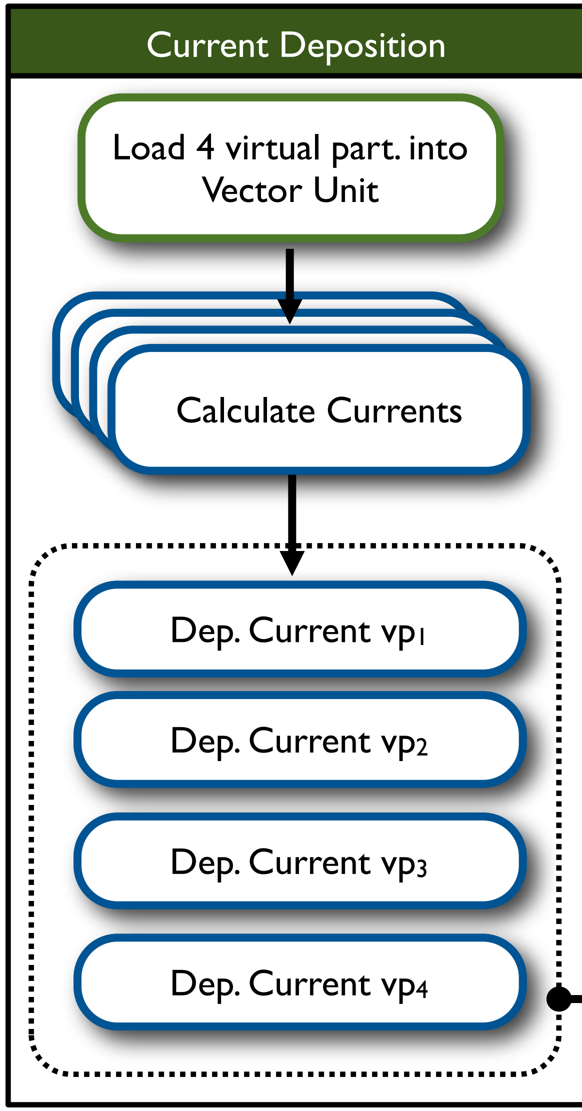

- Particles may deposit to same cell
- Process each 4 particles sequentially (to avoid memory collision)

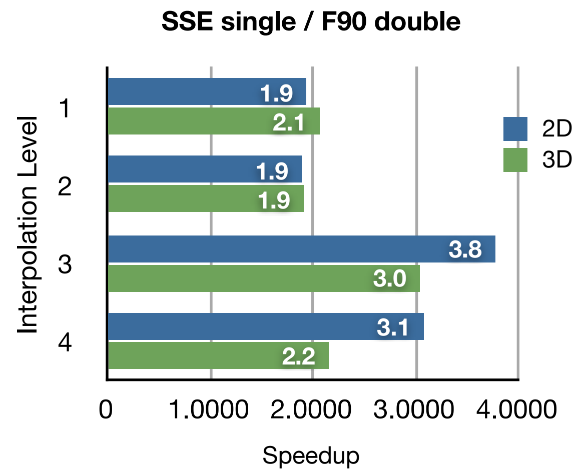

- Excellent speedup on all dimensionalities / Interpolation level
- For lower interpolation levels compiler does vectorization on its own (consistent with comments from the previous talk)
- 3rd order interpolation has optimal ratio of computation over memory fetch and yielded the largest speedup on the SSE

Porting UPIC on GPU-MPI Systems (using CUDA)

	CPU: Intel i7	GPU: Fermi M2090	GPU: Tesla C1060
Push	18.9 ns	557 ps	735 ps
Deposit	8.7 ns	254 ps	232 ps
Reorder	0.4 ns	134 ps	818 ps
Total	28.0ns	944 ps	1785 ps

Dawson2 at UCLA: 96 nodes, ranked 384 in top 500, 70 TFlops on Linpack
- Each node has: 12 Intel G7 X5650 CPUs and 3 NVIDIA M2090 GPUs.
- Each GPU has 512 cores: total GPU cores=147,456 cores, total CPU cores=1152

```
The time reported is per particle/time step.
The total speedup on the Fermi M2090 was 30x,
on the Telsa C1060 was 16x.
```

CPU: Intel i7	1 core (ns)	12 core (ns)
Push	20.3	1.80
Deposit	8.34	0.75
Reorder	0.34	0.04
MPI Move	0.01	0.04
Total	28.94	2.64

GPU: Fermi M2090	1 GPU	3 GPUs
Push	345	135
Deposit	266	97
Reorder	478	187
MPI Move	36	88
Total	1125	506

```
The time reported is per particle/time step.
The total speedup on the 3 Fermi M2090s compared to 12 cores was 5.2x,
Speedup on 3 M2090s compared to 1 M2090 was 2.2x (2 PCI bus per node?)
* OSIRIS on GPU (treating Viktor's routines as blackboxes)
* UPIC on Blue Waters via OpenACC
```

Uncertainties And Future HPC Needs

- I/O needs -- will future simulation codes contain postprocessing capabilities to address the I/O bandwidth limitations or trade computation for I/O bandwidth?

Example: Particle orbit tracking in OSIRIS. In OSIRIS, we study the injection of electrons into the plasma wave by running the same simulation twice, the first time to identify the injected particles, and the 2nd time to record the detailed orbits of all injected particles.

As computation gets cheaper relative to I/O, this way of extracting data where we extract data/physics by doing the identical simulation repeatedly may be more common. (VORPAL and WARP has the same feature)

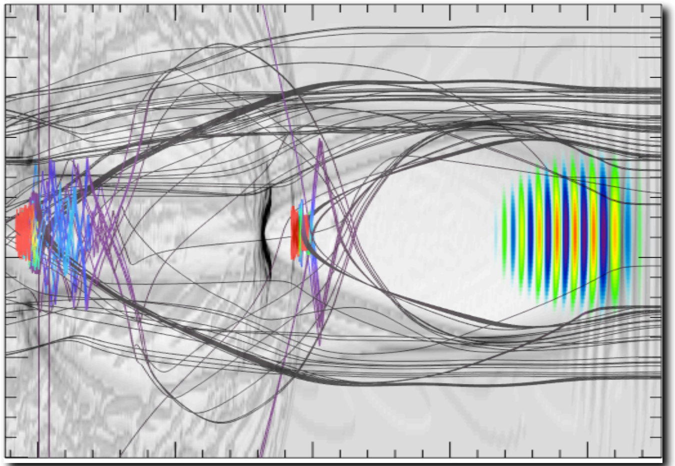

- How does NERSC (and NERSC users) pressure vendors to provide parallel post-processing tools? (e.g., IDL, Matlab) We can write our data using parallel HDF5 but only VisIT can read data in parallel.

- Other.......

4. Summary

- There are interesting PWFA physics that can be studied using exascale supercomputers
  - PWFA's for positrons
  - beam loading
  - small emittance and small emittance dilution
  - proton drivers (PPA's)
- Development Plans
  - GPU's and multi-core CPU's share some common features, specifically both hardware are limited by memory bandwidth. In PIC code, we have begun to eliminate random memory fetch to gather/scatter particle data on grids (using particle re-ordering) and this will work on both multi-core CPU's and GPU's.
- Uncertainties and future needs:
  - As computation become cheaper, simulation codes may have to handle some of the analysis and post-processing that is done after the simulation to reduce the I/O load. Will this become part of the future metrics?
  - Restart on large machines will be difficult. What are some promising new technologies for high bandwidth file servers?
  - How can NERSC (and NERSC users) help in getting commercial softwares (e.g., IDL and Matlab) to be more responsive to the needs of HPC users?

GPU Particle Reordering (analogous to message passing of particle data via MPI)

An important part of the GPU version of UPIC is the particle reordering, which insures that the particle data needed by the small GPU tile (analogous to the domain under domain decomposition) is situated locally and can be easily accessed. This will eliminate the memory collisions shown in the previous slide. We feel this technique is portable to all future SIMD architectures where there are a large number of cores and the bottleneck is memory bandwidth.

The 2D particle reordering, implemented by Dr. V. K. Decyk of UCLA is shown on the right. We hope to use this as a "black box" for all of our PIC codes.

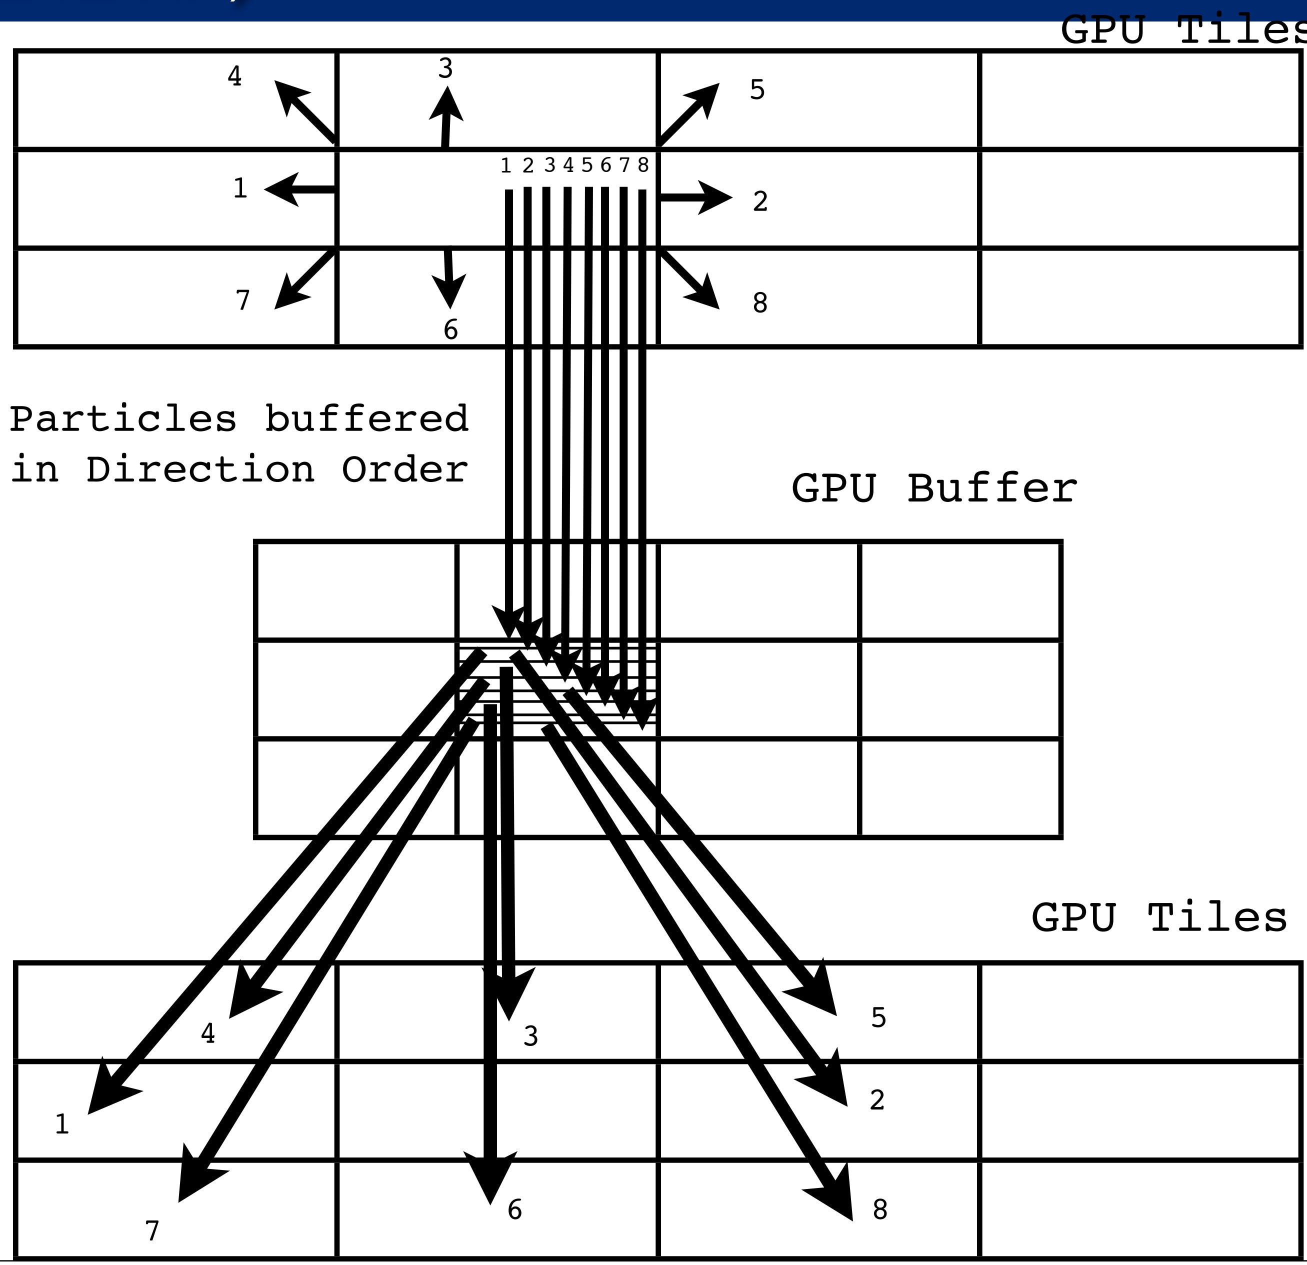

UPIC: UCLA Particle-in-Cell Framework

Features of UPIC:

- Provides trusted components for rapid construction of new parallel PIC codes (You-PICK)
- Support multiple physics models, levels of accuracy, optimizations, computer architectures.
- Supports both MPI and threaded programming models.
- Hides parallel processing by reusing communication patterns: Physicists only need to know the data layout.
- Components used in wide variety of applications: Magnetic Fusion, Space Physics, **Plasma Accelerators (QuickPIC)**, Cosmology, Quantum Plasmas, Ion Propulsion (DRACO).

(V. K. Decyk, Comp. Phys. Comm. **17**, 95 (2007).)

Recently UPIC has been ported to the MPI-GPU systems and we will show some preliminary results and discuss the move to new multi-core architectures.

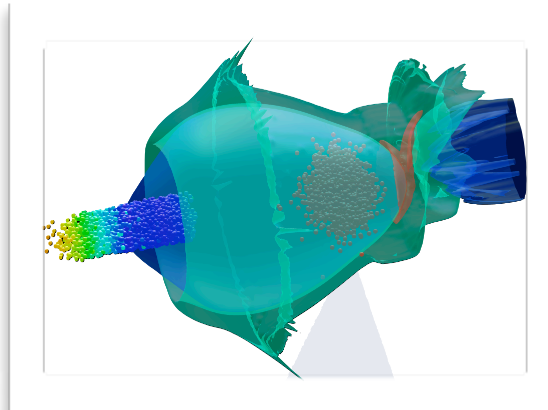

QuickPIC: Plasma Accelerators
(C. K. Huang, *et al*)

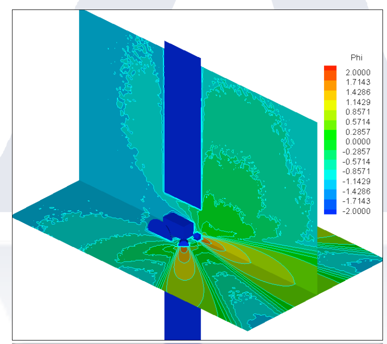

DRACO: Ion Propulsion
(J. Wang, *et al*)