Present and Future Computing Requirements for Kinetic Modelling in ICF

Chuang Ren
University of Rochester

NERSC FES Requirements for 2017
March 19-20, 2013
Rockville, MD

Investigation of LPI and Advanced Ignition Physics in ICF

C. Ren (m792), W. B. Mori (m412), T. S. Tsung (m1157)

- Develop a first principle-based understanding of LPI and other HED physics that can be included in design codes in a computation-efficient way.

- Our present focus is understand relevant physics under idealized conditions (e.g. laser as plane waves or a few speckles)

- By 2017 we expect to model the physics under realistic conditions and compare with experiments

LPI are important to laser-target coupling and target preheating in ICF

Courtesy D. Froula

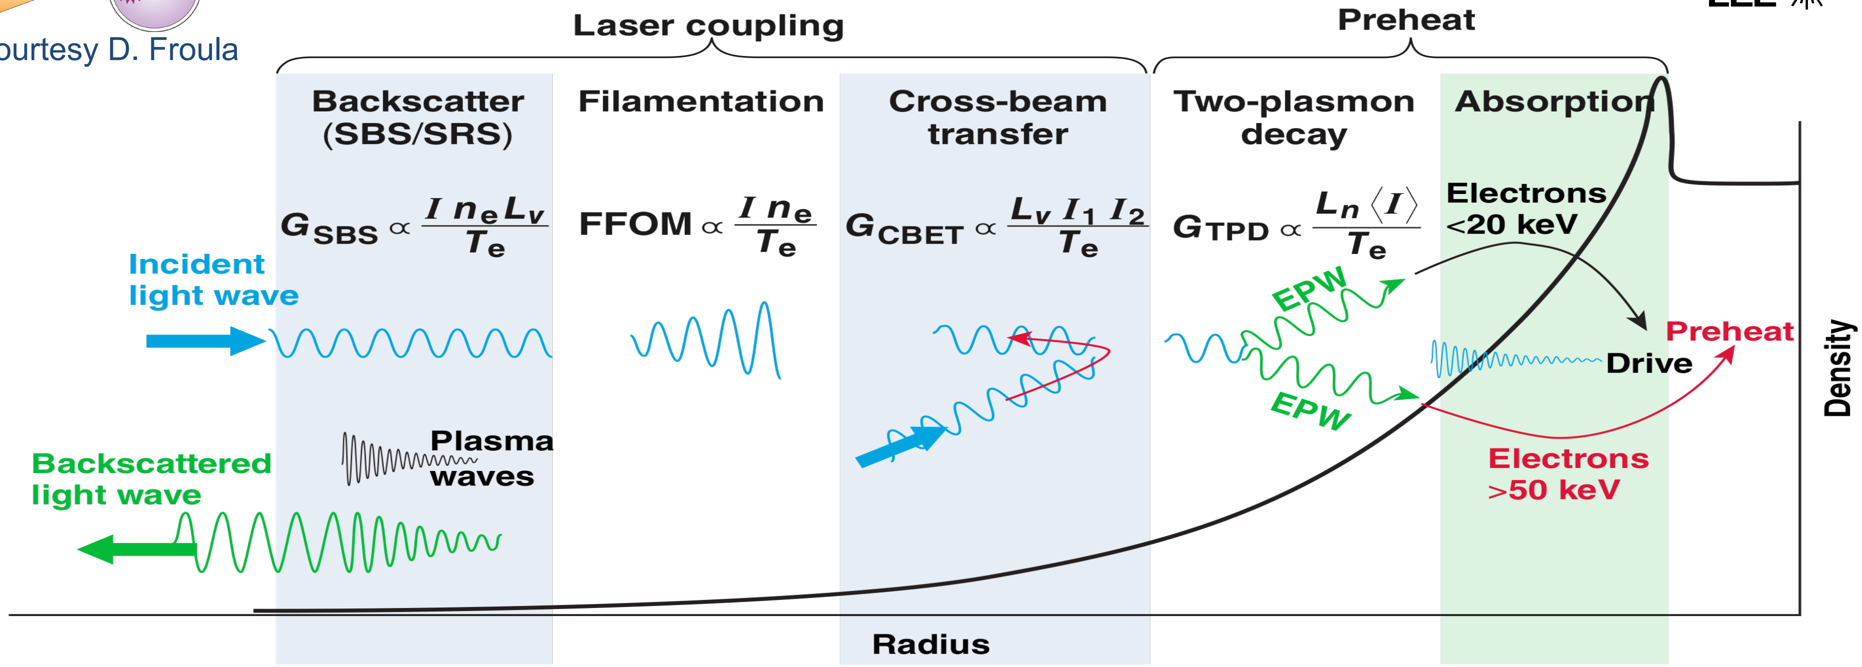

- Laser-plasma instabilities can change laser-target coupling through scattering and ‘collisionless’ heating
  - Backscattering (SRS & SBS) and CBET (SBS)
  - Additional laser absorption away from the n_c-surface (SRS & TPD)
- LPI may generate hot electrons that can preheat the target (TPD)
- Hot electrons can also help ignition in shock ignition
- LPI physics is currently not fully incorporated in target design codes

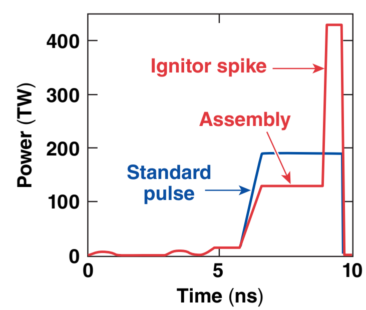

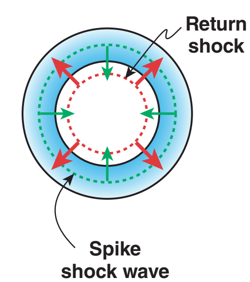

Computation Challenges: Now and 2017

- Our present focus is understand relevant physics under idealized conditions
  - 2D Plane wave TPD simulations (Yan et al. PRL '12) discovered staged-acceleration of hot e- and collisions can reduce the acceleration
  - But hot e- fraction predicted (5% of laser energy) is much larger than experimental values (0.1-1%, Froula et al. PRL'12)

- In 2017 we expect to simulate LPI in realistic beam conditions
  - An actual beam consists of tens of thousands of speckles of different polarizations
  - A 3D simulation of 10-100 speckles requires 18-180 M hours

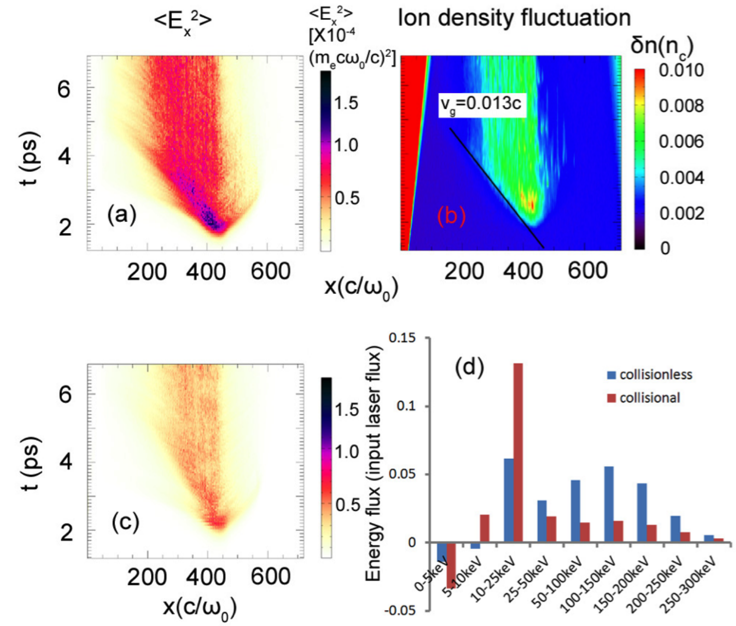

From Yan PRL'12. Need 100,000 CPU-hours to simulate 10 ps (Grid: 3600 x 6000, 100ppc)

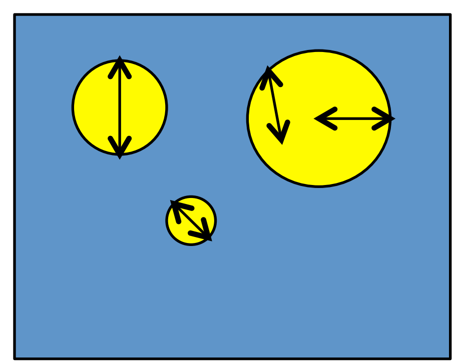

speckles in a beam

Recent Highlights from m412 and m792

- We have studied TPD evolution and hot electron generation in direct-drive ICF under Omega conditions (Yan et al. PRL '09, '12, PoP '10)
- We have studied many aspects of fast ignition scenario, including laser channeling, laser-plasma interactions inside coned-targets, and laser hosing in relativistic plasmas (Li et al. PRL '08, '13, PoP '11, '13)
- We have studied hot-electron generation and their role in assisting Giga-bar shock launching in shock ignition
- We have studied collisonless shock generation from counter-flowing plasmas and particle-energization in the shocks, which can be realized on the Omega facility (Workman et al. PoP '11, Park et al. PoP '12, ApJ '13)

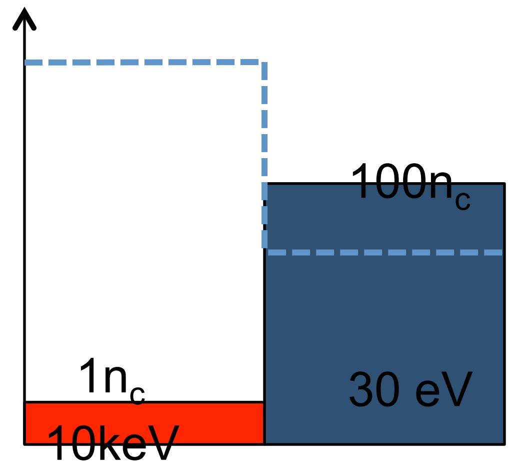

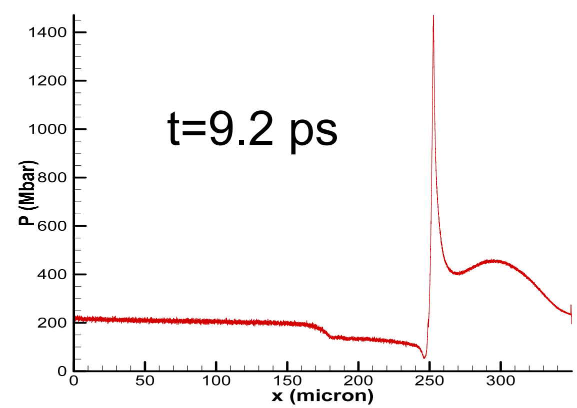

~1M hours on 20k cores for a 1D run.
Hybrid runs can save 100x

2D and 3D simulations of SRS under IFE Relevant Conditions (PI: Tsung/repo: m1157)

The m1157 repo studies both the TPD instability and SRS (stimulated raman instability). In SRS, the laser decays into a back-going EM wave and and a forward going plasma wave. The SRS instability (and LPI in general) is an important issue for the National Ignition Facility (NIF) because:

- The backward going light cannot couple to the target
- The plasma wave can accelerate electrons and cause the target to preheat, and thus degrades compression.

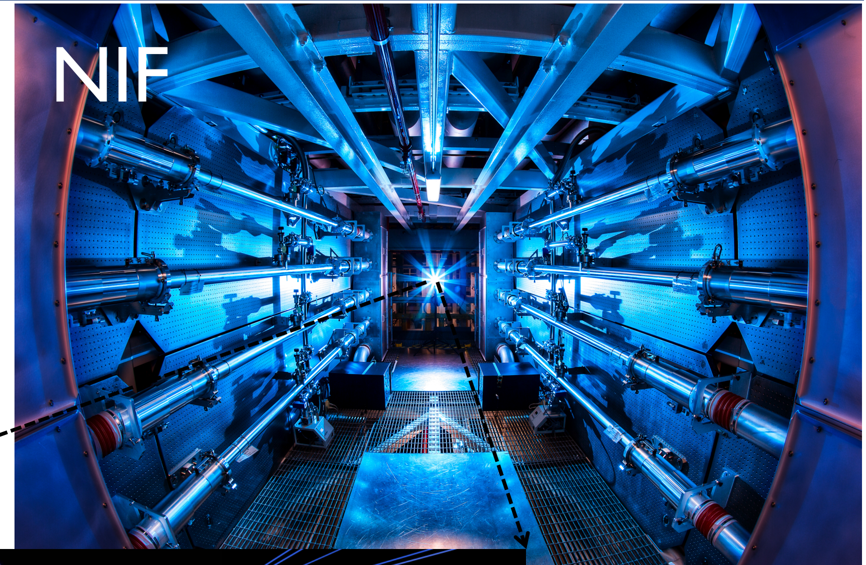

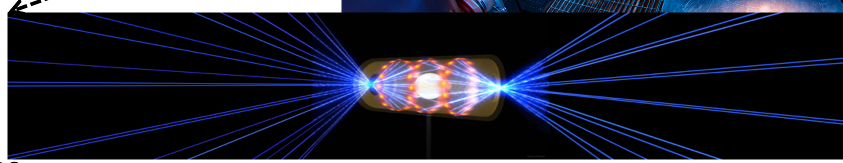

This problem is well suited to large scale supercomputers because of all the spatial and temporal scales involved. Some of these scales are shown below.

Lengthscales

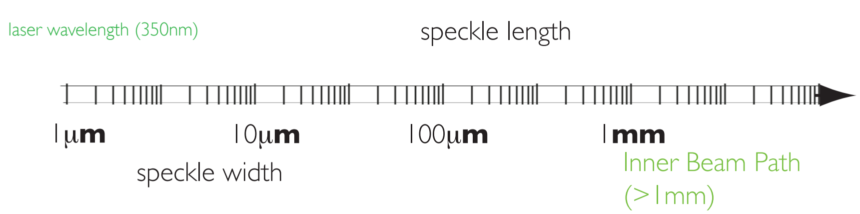

Timescales

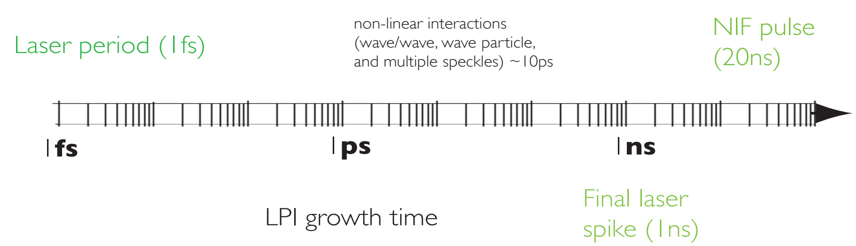

2012 Highlights: 1. NIF-Relevant Simulations of Stimulated Raman Scattering

- We have simulated stimulated Raman scattering (SRS) in 1D and 2D over 0.5-1.5 mm lengths with NIF-relevant density profiles

- 1D simulations can be done quickly (<500 cpu hours each) and allow for methodical parameter scans and comparisons with linear theory
  - Hydro conditions ⟶ NIF scientist uses 1D fluid postprocessing tools such as SLIP/NEWLIP:

    Predict the frequency and reflectivity of the most unstable LPI
  - Hydro conditions ⟶ 1D OSIRIS simulations:

    Similar capabilities + detailed information about energy partition, backscattered light, and energetic electrons (see below)

The figure here shows several snapshots of the electron distribution function from one of our 1D simulations (along the red path, see right). PIC simulations can extract the information about the scattered light as well as the hot electrons, and the underlying physics associated with these hot electrons can be extracted by looking at the plasma waves.

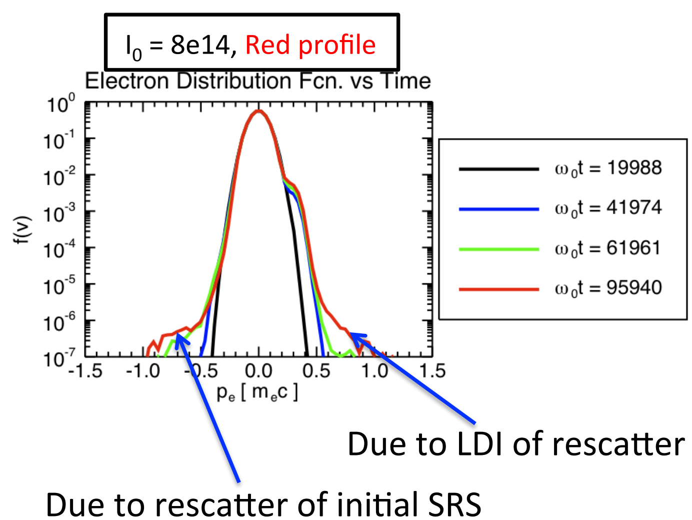

I_{laser} = 2 – 8 x 10^{14} W/cm^2
l_{laser} = 351nm,
T_e = 2.75 keV,
T_i = 1 keV, Z=1,
t_{max} up to 20 ps
Length = 1.5 mm
Density profiles from NIF hydro simulations

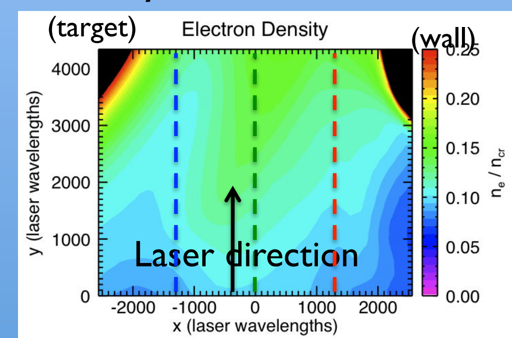

14 million particles
~400 CPU hours per run
~1 hr on modest size supercomputer

2012 highlight 2: 2D effects of SRS under NIF-relevant conditions

- The SRS problem is not strictly 1D -- each "beam" (right) is made up of 4 lasers, called a NIF "quad," and each laser is not a plane wave but contains "speckles," each one a few microns in diameter.

- We have been using OSIRIS to look at SRS in multi-speckle scenarios. In our simulations we observed the excitation of SRS in under-threshold speckles via:
  - "seeding" from backscatter light from neighboring speckles
  - "seeding" from plasma wave seeds from a neighboring speckle.
  - "inflation" where hot electrons from a neighboring speckle flatten the distribution function and reduce plasma wave damping.

- The interaction of multiple speckles is a highly complex process and is ideally suited for PIC simulations

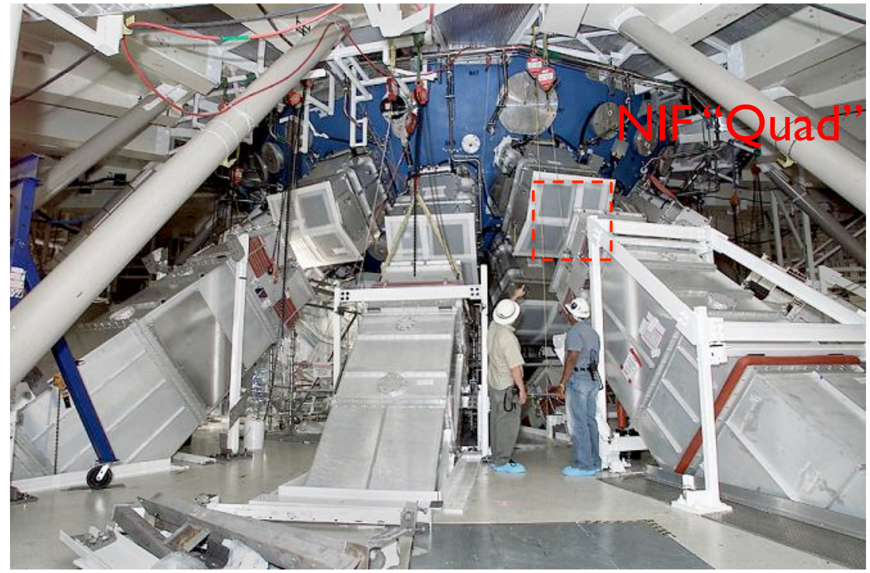

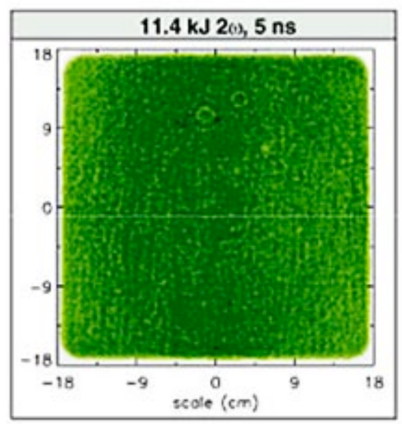

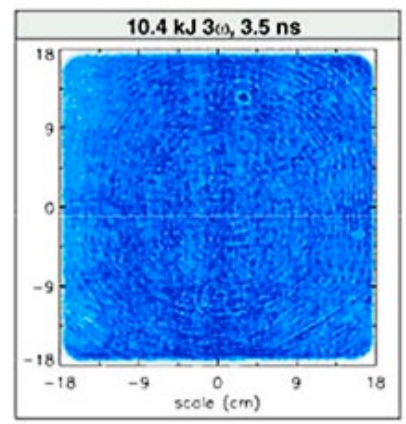

(cont) Two-speckle simulations allow for controlled investigations of each inter-speckle intermediary

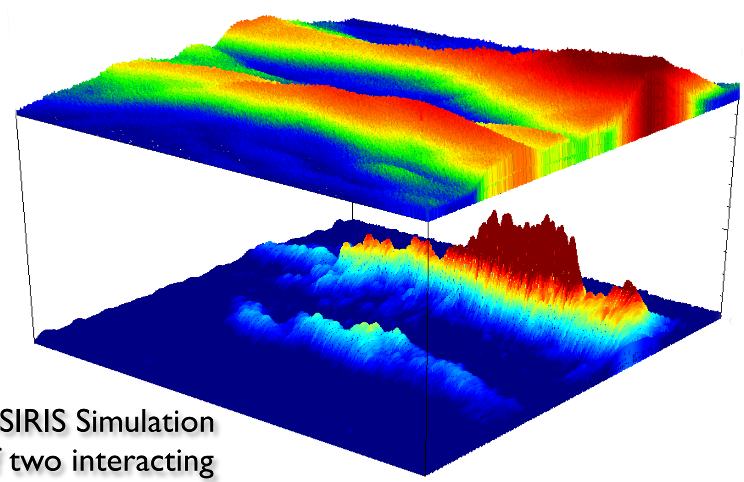

OSIRIS Simulation of two interacting speckles

Two-speckle simulations allow for controlled investigations of different types of inter-speckle interactions:

- ~125mm x 30mm, 10 billion particles
  ~100,000 CPU hours for a >10ps simulation.
- Simulations show the excitation of SRS for speckles below LPI threshold (through kinetic inflation in the case shown on the left). A number of 2 speckle simulations were done to look at the different types of inter-speckle interactions.

Multi-Speckle Simulations

- ~500mm x 80mm, 10 billion particles
  ~500,000 CPU hours for a 4ps simulation
- 2D simulations like those on the right will quantify the effects laser speckles have on LPI, and will provide better predictive capabilities for LPI models used for NIF. We have begun to study the interaction of many (>~10) speckles under NIF relevant plasma conditions.

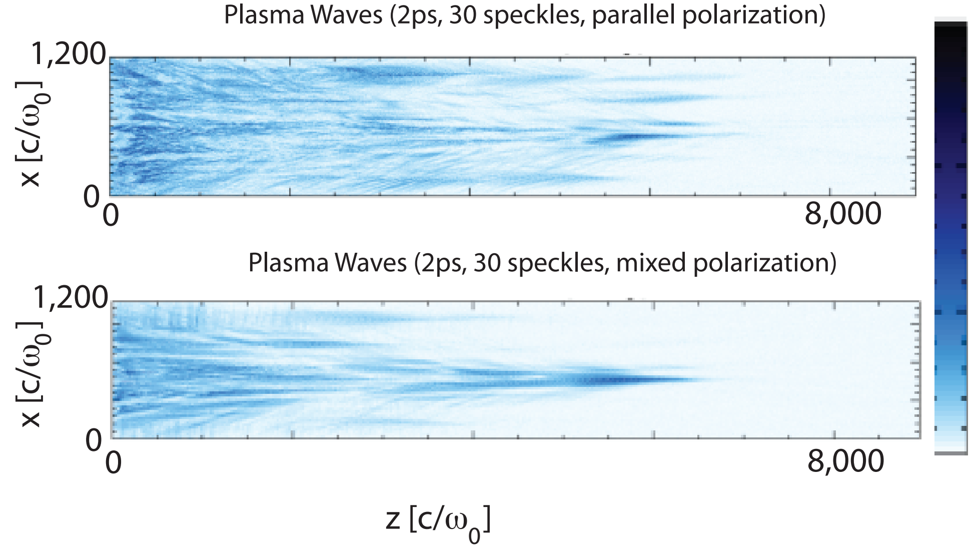

With increased hardware capabilities, we would like to address the following physics relevant to the SRS problem:

Outstanding Issues in LPI/IFE for 2017:

- 2D simulations of speckled laser beams in millimeter-length NIF relevant plasmas, this is import both scientifically & programmatically (100 billion particles, 5 million Hopper hours/run).
- 3D simulations of single- and multiple-speckle SRS and related nonlinear plasma wave phenomena, a typical 3D simulation will take ~2 trillion particles > 40 million hours/run.
- Understand higher dimensional effects in the development and evolution of nonlinear plasma waves for TPD and SRS. (See right)

	Used at NERSC in 2012	Needed at NERSC in 2017
Computational Hours (Hopper core-hour equivalent)	5 Million	250Million
Scratch storage and bandwidth	7.5-10TB	300 TB
	(not sure, whatever the current bandwidth is adequate)GB/sec	50GB/sec(this assumes it would take 30 minutes to write checkpoint files)
Shared global storage and bandwidth (/project)	< 1 TB	10-20TB
	(adequate)* GB/sec	GB/sec
Archival storage and bandwidth (HPSS)	15-20 TB	1000TB
	(adequate)* GB/sec	50-100GB/sec
Number of cores* used for production runs	16,000-32,000	100,000-250,000
Memory per node	.1-.2GB	.5-1GB
Aggregate memory	4TB	50-150TB

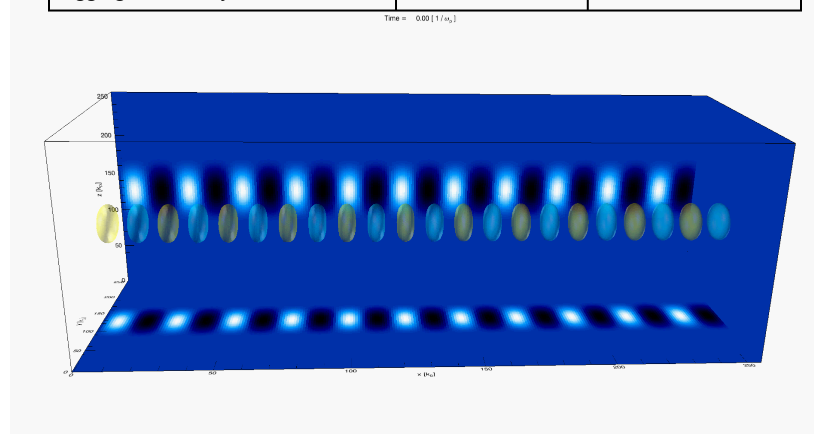

osiris 2.0

UCLA

osiris framework

- Massivelly Parallel, Fully Relativistic Particle-in-Cell (PIC) Code
- Visualization and Data Analysis Infrastructure
- Developed by the osiris.consortium
- UCLA + IST
- #8 most used code @ NERSC in 2012

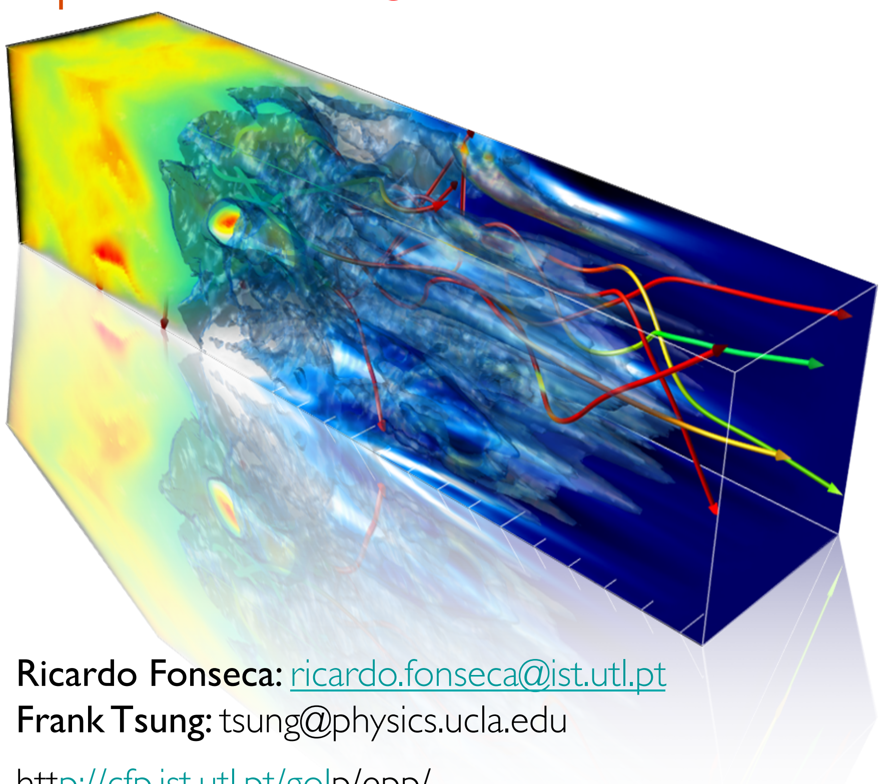

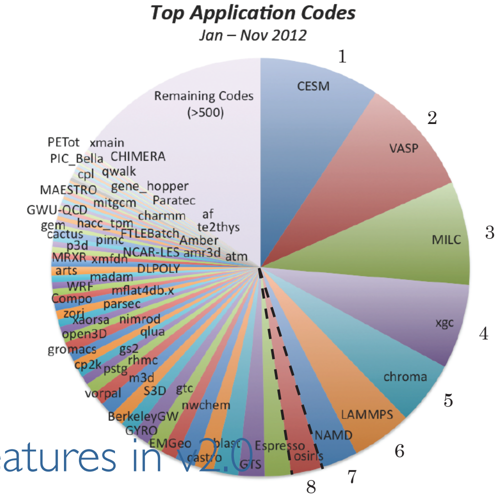

Ricardo Fonseca: ricardo.fonseca@ist.utl.pt
Frank Tsung: tsung@physics.ucla.edu

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

New Features in v2.0

- Bessel Beams
- Binary Collision Module
- **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

VisXD

Profile of OSIRIS

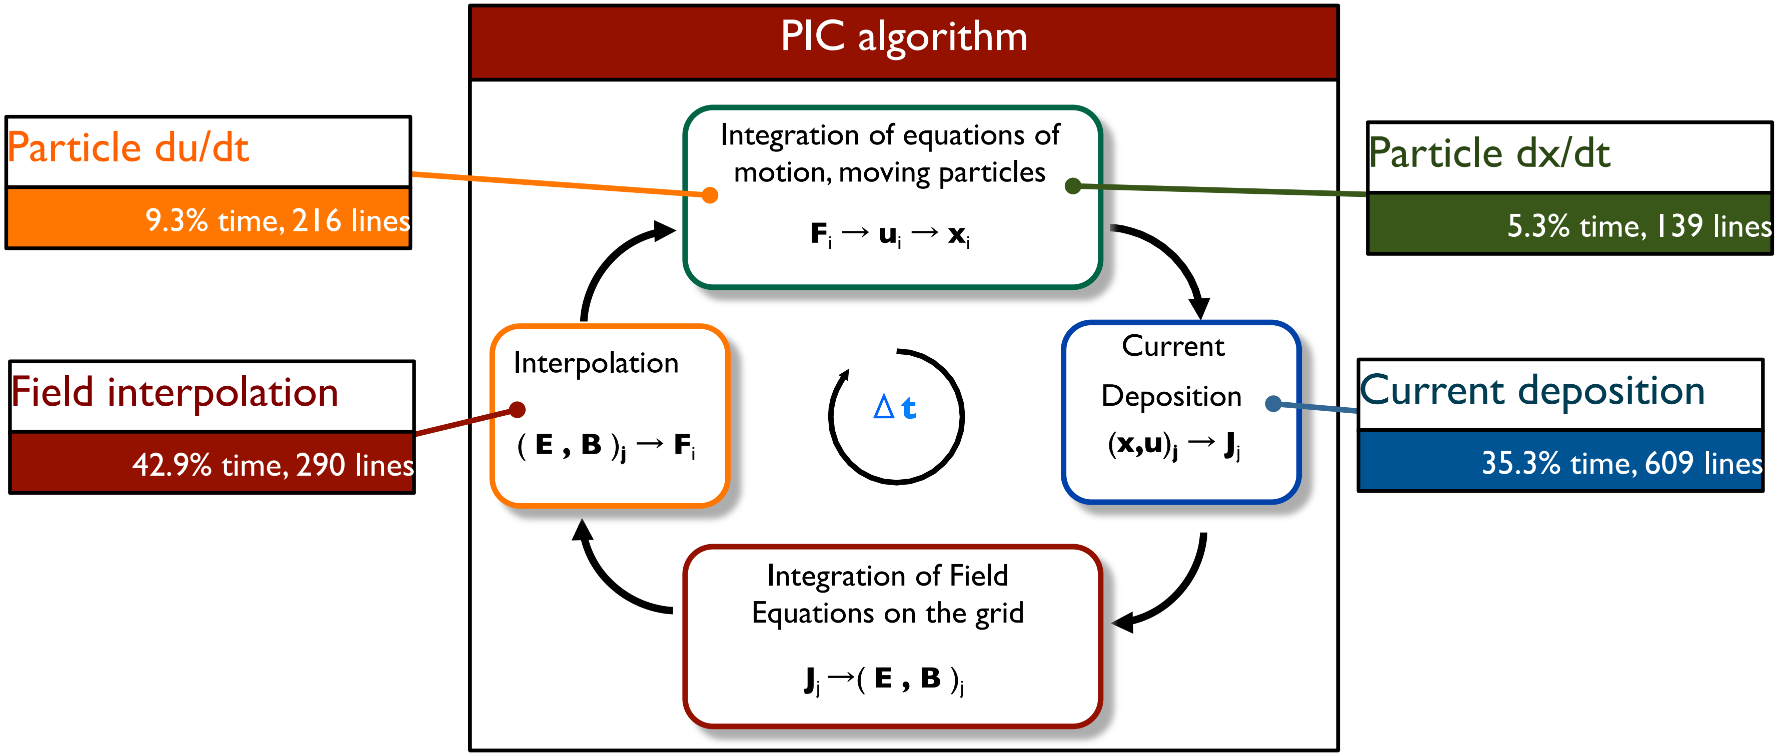

- The code spends over 90 % of execution time in only 4 routines
- These routines correspond to less than 2 % of the code
- Optimization:
  - Focus where cycles are spent
  - Keep existing C/Fortran code structure

Source Code

- Basic PIC algorithm (with various implementations) accounts for only 18% of the code
- Parallel communications correspond to less than 10% of the code

- Most of the code is devoted to:
  - Diagnostics
  - Additional physics
  - User Interface

OSIRIS source distribution

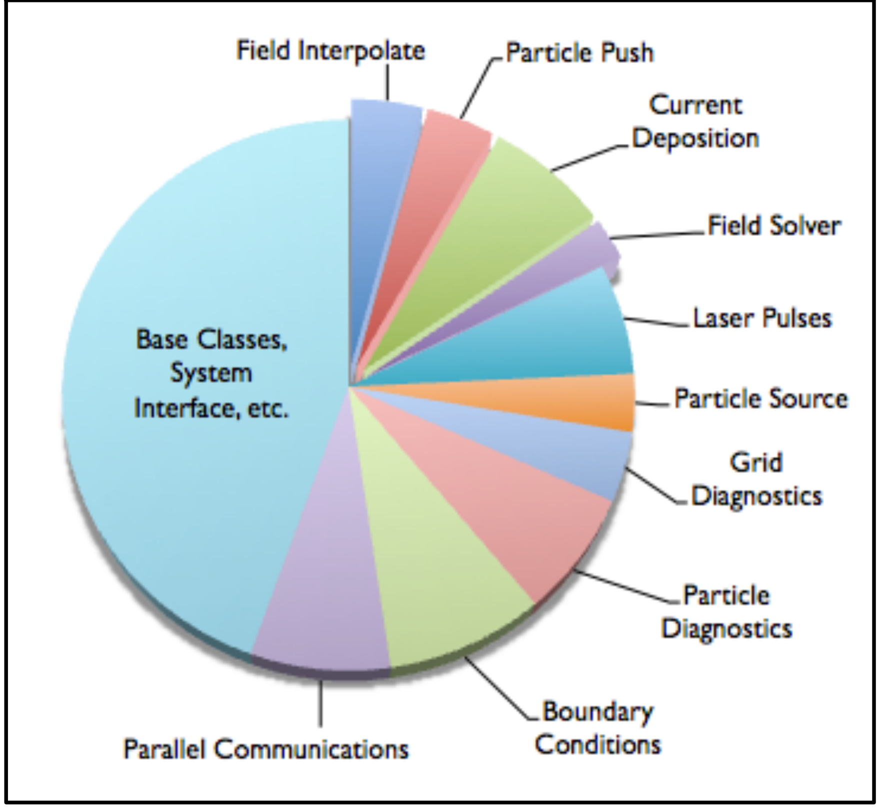

2. Computational Strategies

- We anticipate to use 25x hours in 2017: 20 M -> 500 M
- We meet the computation challenges by
  - Use more cores available: 20k -> 1M (OSIRIS currently scalable to 1.5 M cores on Sequoia and achieved 2.2 pflops on Blue Waters)
  - Adapt to new platforms such as GPU
  - Develop new reduced models such OSIRIS-h

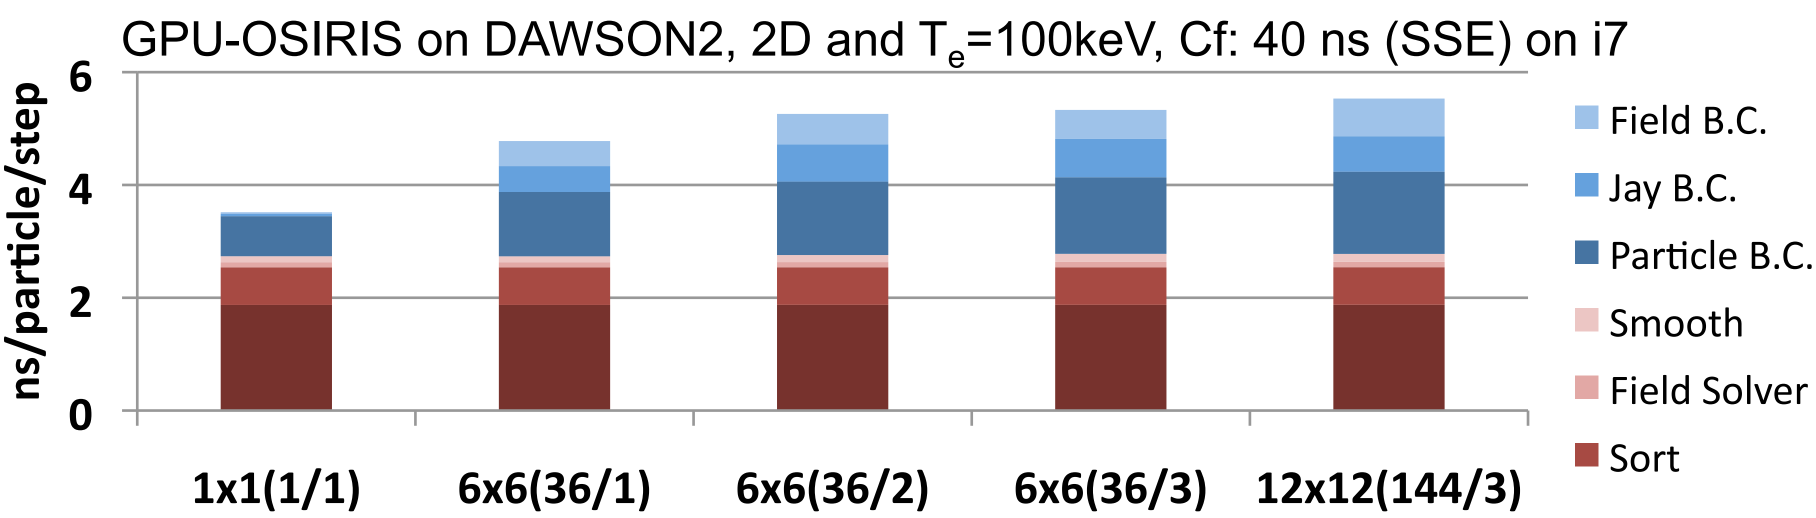

OSIRIS-H

OSIRIS-H framework

- Hybrid Particle-in-cell (PIC) code
- Full-PIC and MHD algorithms
- Massivelly Parallel and Fully Relativistic
- Based on OSIRIS 2.0
- Visualization and Data Analysis Infrastructure
- Developed by the osiris.consortium
  $\Rightarrow$ UCLA + IST

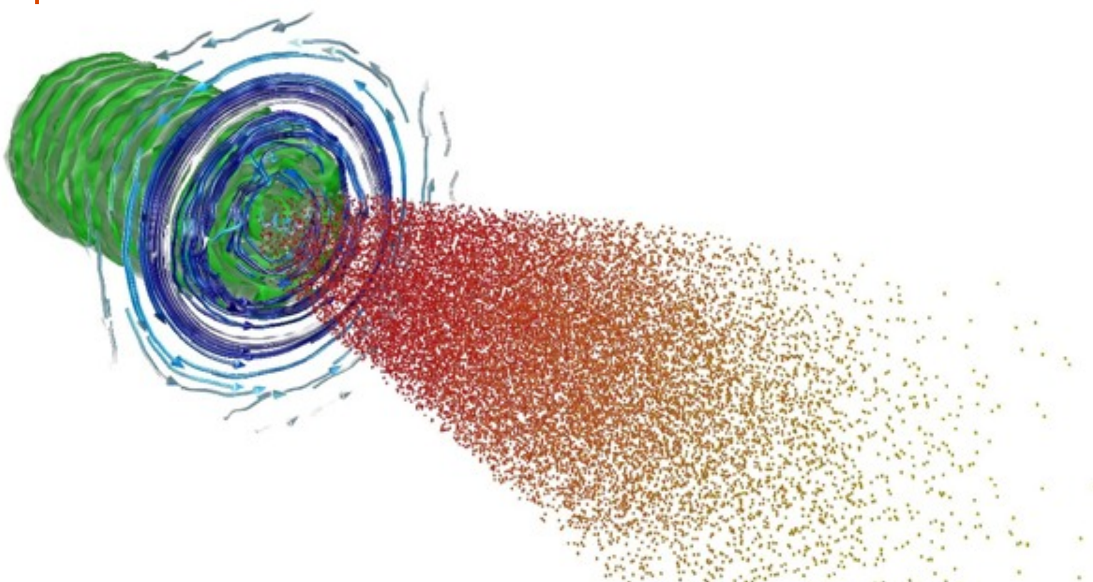

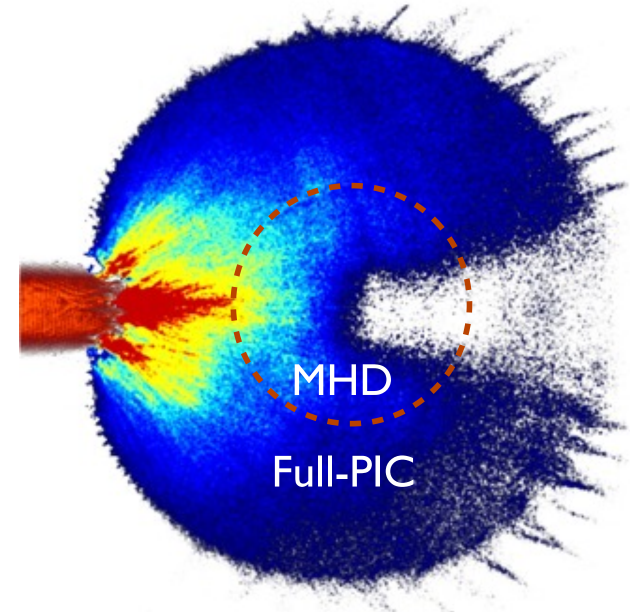

Advanced features

- High-order splines
- Binary Collision Module
- PML absorbing BC
- Tunnel (ADK) and Impact Ionization
- Dynamic Load Balancing
- Parallel I/O

Frederico Fiuza: fiuza1@llnl.gov
Ricardo Fonseca: ricardo.fonseca@ist.utl.pt
Warren Mori: mori@physics.ucla.edu

F. Fiuza *et al.* Plasma Phys. Control. Fusion 53, 074004 (2011)

3. Current HPC Usage

- Machines currently using: Hopper and ACLF
- Hours used in 2012 (list different facilities): 19.5 M
- Typical parallel concurrency and run time, number of runs per year: 2-20k cores for 40-100 hours, ~10 runs/year
- Data read/written per run: 2TB
- Memory used per (node | core | globally): 0.2-1 GB/core, 1-10 TB globally
- Necessary software, services or infrastructure: MPI, HDF5, IDL
- Data resources used (HPSS, NERSC Global File System, etc.) and amount of data stored: 70 TB on HPSS

4. HPC Requirements for 2017

- Compute hours needed (in units of Hopper hours): 200-500 M
- Changes to parallel concurrency, run time, number of runs per year: 20k -> 200k-1M cores, 40-100 hours, 3-5 big runs / year
- Changes to data read/written: 2 TB -> 50-100 TB
- Changes to memory needed per (core | node | globally) GB/core: 0.2-1 (no change) and global memory: 1-10 TB -> 10 – 500 TB
- Changes to necessary software, services or infrastructure: parallel data analysis infrastructure

5. Summary

- We anticipate a ~25x resource demand increase, through more and better cores
  - A large GPU cluster is highly desirable

- Such an increase will allow us to study LPI and other HED physics under realistic conditions and can compare with experiments