Present and Future Computing Requirements for the Centre for Simulation of Wave-Plasma Interactions (CSWPI)

Paul T. Bonoli

Plasma Science and Fusion Center - MIT

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

The SciDAC Center for Simulation of Wave – Plasma Interactions (CSWPI)

L.A. Berry, D.B. Batchelor, D. L. Green,
E.F. Jaeger, E. D`Azevedo

D. Smithe

P.T. Bonoli, J.C. Wright,
J. P. Lee, A. Ram

C.K. Phillips, E. Valeo
N. Bertelli, H. Qin

R.W. Harvey, Y. Petrov
A.P. Smirnov N.M. Ershov

M. Brambilla
R. Bilato

M. Choi

Politecnico di Torino
R. Maggiora

D. D'Ippolito, J. Myra - Lodestar Research

1. Project Description

Principal Investigator	**Institution**
Bob Harvey	CompX
Myunghee Choi	GA
Dan D’Ippolito	Lodestar Research
Paul Bonoli (Project PI)	MIT
Lee Berry	ORNL
Cynthia Phillips	PPPL
David Smithe	Tech-X

1. Project Description (continued)

- **Scientific objectives through 2017 are organized into four major thrusts:**
  - *Coupled core-to-edge simulations:*
    - **Increased understanding of parasitic losses of applied RF power in the plasma boundary between the RF antenna and the core plasma.**
  - *Development of models for core interactions of RF waves with energetic electrons and ions:*
    - **More accurate representation of the particle dynamics in the combined equilibrium and wave fields.**
  - *High-resolution simulations of RF effects on fast-particle driven instabilities driven by fusion alpha particles or NBI ions:*
    - **Will these interactions increase (decrease) the instability drive that can lead to reduced fusion power.**
  - *Development of improved algorithms to achieve the needed physics, resolution, and/or statistics to address these issues:*
    - **Algorithms must take advantage of massively parallel computing platforms up to the multi peta-scale level and beyond.**

1. Project Description (continued)

- Presently we are focusing on development and validation of core and edge models separately.
  - Core models include combined wave solvers and Fokker Planck codes to study:
    - Ion cyclotron resonance heating (ICRH) at $\omega \approx \omega_{ci}$ and $\omega >> \omega_{ci}$ in the presence of energetic ions.
    - Importance of finite ion orbit width effects in ICRH and assess phase coherence effects in ICRF diffusion.
    - Importance of full-wave effects and edge losses in lower hybrid current drive (LHCD).
  - Edge models include:
    - 3D solid geometry of antenna launching structure and nonlinear parasitic loss mechanisms such as RF sheath rectification.

1. Project Description (continued)

- By 2017 we expect to have coupled our core wave / Fokker Planck codes to RF edge models:
  - Will have validated coupled core to edge model against ICRF heating experiments in NSTX, DIII-D and Alcator C-Mod.
  - First predictive simulations of ICRF coupling for ITER.
- By 2017 expect to have validated reduced models:
  - Continuum Fokker Planck code with finite orbit width and phase coherence effects included.
  - Reduced ICRF solver validated for all ICRF regimes.
  - Validated description of LHRF full-wave effects using a beam tracing code.
  - Reduced models expected to be either ready for or implemented in time dependent computational framework (s).

3-D visualization of the ICRF wave fields in ITER shows “hot spots” near the antenna surface where the wave amplitude is high

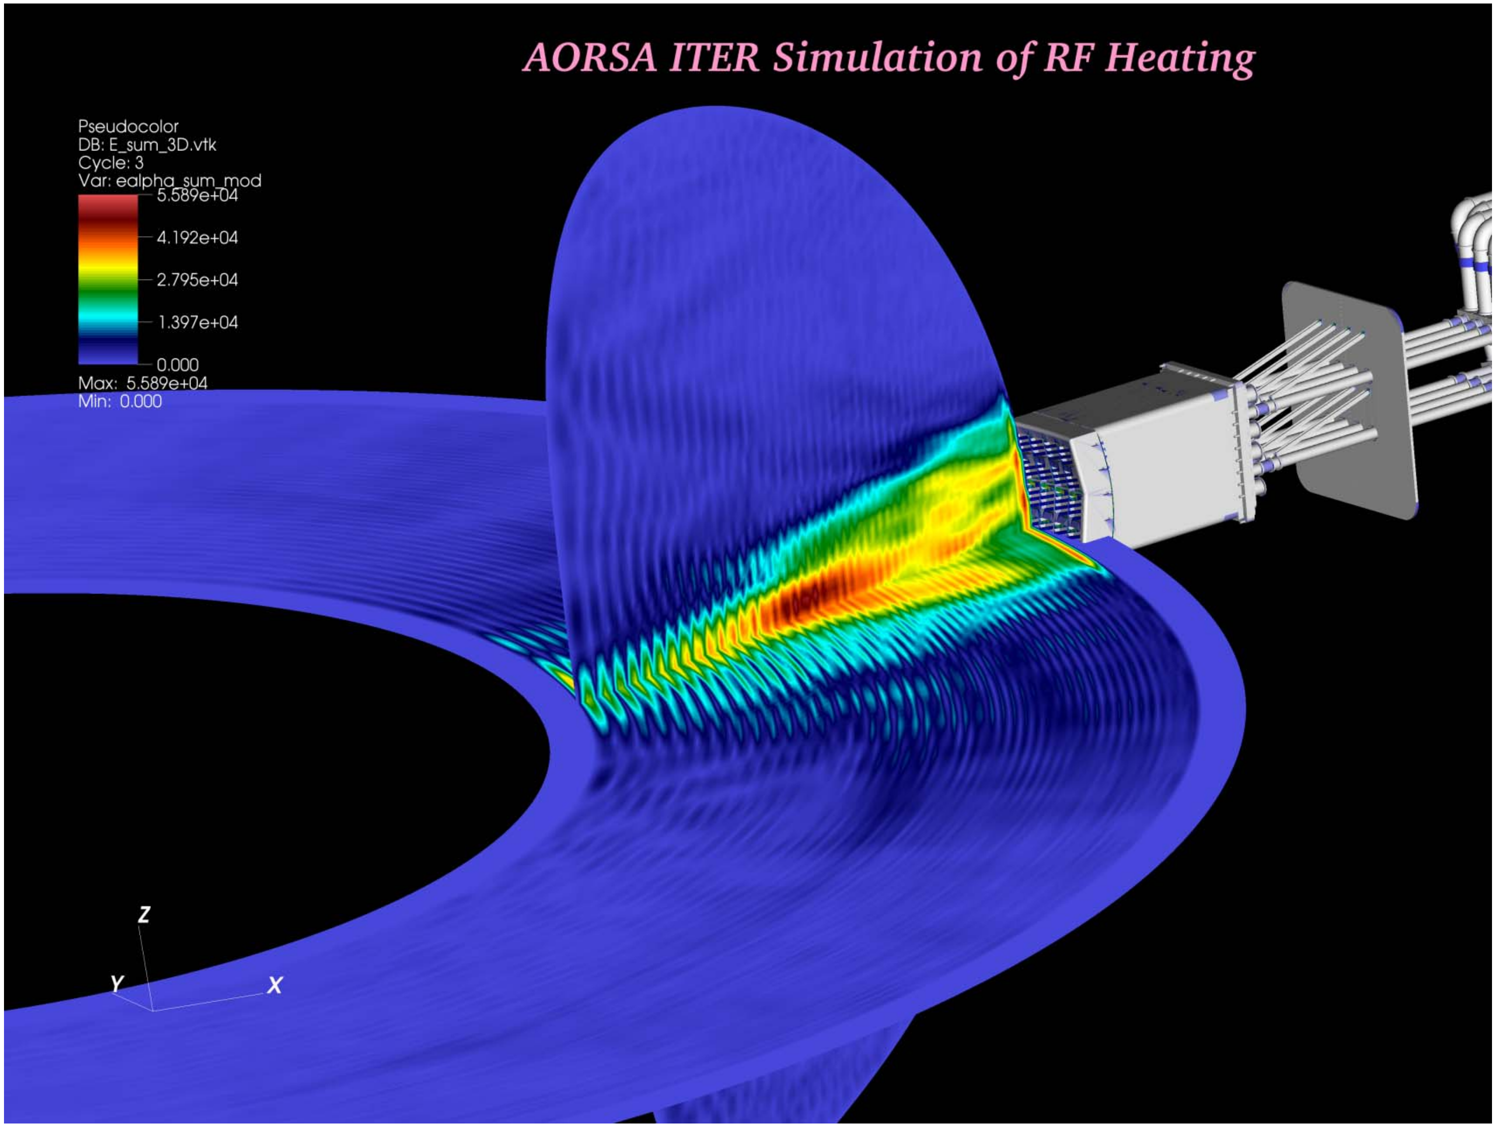

- **AORSA simulation using 100 toroidal modes of the ICRF antenna.**
- **Calculation done on 2048 processor cores in 2 hours on Jaguar facility.**

2. Computational Strategies – Approach

For time harmonic (rapidly oscillating) wave fields E with frequency ω, Maxwell's equations reduce to the Helmholtz wave equation:

$$-\nabla\times\nabla\times\mathbf{E}+\frac{\omega^2}{c^2}\left(\mathbf{E}+\frac{i}{\omega\varepsilon_0}\mathbf{J}_p\right)=-i\omega\mu_0\mathbf{J}_{ant}$$

The plasma current (J_p) is a non-local, integral operator (and non-linear) on the rf electric field and conductivity kernel:

$$\mathbf{J}_p(\mathbf{r},t)=\sum_s\int d\mathbf{r}'\int_{-\infty}^{t}dt'\sigma\left(f_{0,s}(E),\mathbf{r},\mathbf{r}',t,t'\right)\cdot\mathbf{E}(\mathbf{r}',t')$$

$$+\mathbf{J}^{rf}_{sheath}(E_{rf})+\mathbf{J}(E^{rf}_{pump})$$

The long time scale response of the plasma distribution function is obtained from the bounce averaged Fokker-Planck equation:

$$\frac{\partial}{\partial t}(f_0)=\nabla_{\mathbf{u}_0}\cdot\Gamma_{\mathbf{u}_0}+\langle\langle S\rangle\rangle+\langle\langle R\rangle\rangle \quad \text{where} \quad \nabla_{\mathbf{u}}\cdot\Gamma_{\mathbf{u}}=C(f_0)+Q(\mathbf{E},f_0)$$

Wave Solvers:
AORSA
TORIC
TORLH

Evaluate σ(f_0)
p2f →
SIGMAD

RF Coupling
VORPAL

Plasma Response:
CQL3D
ORBIT RF
sMC, DC

Need to solve this nonlinear, integral set of equations for core RF wave fields, velocity distribution function, and RF antenna fields self-consistently. This requires an iterative process to attain self-consistency.

2. Computational Strategies - algorithms

- Wave solvers represent electric field in purely spectral (AORSA) or semi-spectral (TORIC) basis functions:
  - AORSA matrix is completely dense and complex with size ~ $(3\times N_X\times N_Z)^2 \times 16$, where typically (N_X, N_Z) ~ (257, 513) for a size ~ 2.5 TB. [N_X and N_Z are the number of spectral modes, assuming axial (ϕ) symmetry.]

  - TORIC & TORLH matrices are block tri-diagonal with dense, complex blocks of size ~ $3 \times (3 \times 2 \times N_m)^2 \times N_\psi \times 16$, where for N_m ~ 1023 and N_ψ ~ 980 the size is ~ 1.8 TB.

  - Solution is achieved through an LU factorization of the matrix with ScaLAPACK, with inversion time scaling as $(N_Z)^3$ and $N_\psi \times (N_m)^3$.

  - Because of (ϕ) symmetry in tokamk system, toroidal modes of the antenna spectrum are independent and can be solved for separately – in either sequential or concurrent fashion.

2. Computational Strategies - algorithms

- Fokker Planck solvers are continuum and Monte Carlo:
  - **CQL3D uses an implicit solve in velocity-space with 100-200 pitch angle points and 300-1000 velocity space points. Radial solve is done using ADI, with 25-75 flux surfaces typically. Solution is parallel only across flux surfaces.**
  - **Monte Carlo codes (ORBIT RF, sMC) and direct orbit integrators (DC) use RF wave fields to diffuse or "kick" particles in velocity space with excellent parallel scaling.**
- RF wave – edge plasma interaction is modeled using the finite difference time domain (FDTD) VORPAL code:
  - **Cold plasma, accurate antenna geometry (CAD), and nonlinear RF sheath boundary condition.**
  - **Work scales as N^4 for 3 (spatial) + 1 (time) dimension.**

2. Computational Strategies - algorithms

- Coupling between core models (wave solvers and Fokker Planck codes) is expensive:
  - Evaluation of conductivity operator must be done for numerical distributions (SIGMAD) and for statistical particle lists that have been converted to continuum distributions (p2f) .
  - Evaluation of quasilinear diffusion coefficient is 5-D (k_X, k_Y, R, Z, ϕ) in AORSA basis set and is 4-D (N_m, N_m, ψ, ϕ) in the TORIC basis set.
  - Wave solvers and Fokker Planck codes are advanced in time using an explicit method, although vector extrapolation schemes (JFNK) have been useful with TORLH-CQL3D.

2. Computational Strategies - challenges

- Coupling between edge code (VORPAL) and (AORSA) will be a challenge:
  - Considering method of overlapping sub-domains (Alternating Schwarz) to couple solutions.
  - Will need to couple time domain solutions in VORPAL with spectral solutions in AORSA using windowed Fourier transform.
  - May need vector extrapolation to achieve convergence in AORSA-VORPAL iteration.
- Parallel scaling of codes:
  - Wave solvers (AORAS & TORIC) have excellent strong scaling and almost perfect weak scaling (across multiple toroidal modes).
  - Monte Carlo codes and orbit integrators have excellent parallel scaling.
  - Continuum Fokker Planck code (CQL3D) is only parallel across flux surfaces when no radial diffusion operator is used.

2. Computational Strategies - challenges

- Strong scaling of AORSA

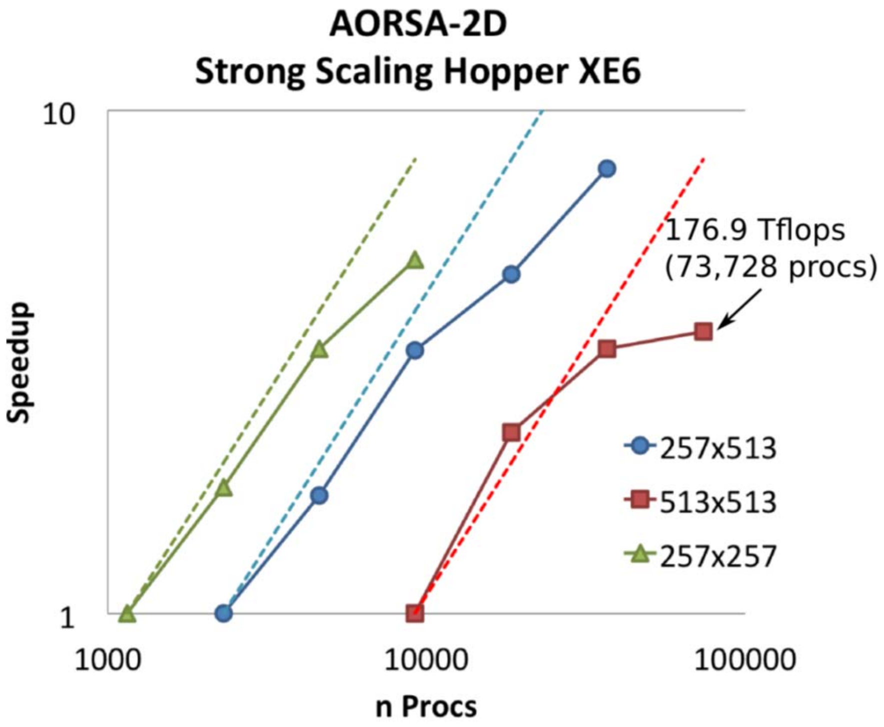

Expect that complex HPL library will help with roll over seen at > 60,000 proc.

- Strong scaling of VORPAL on Hopper (CRAY XE6)

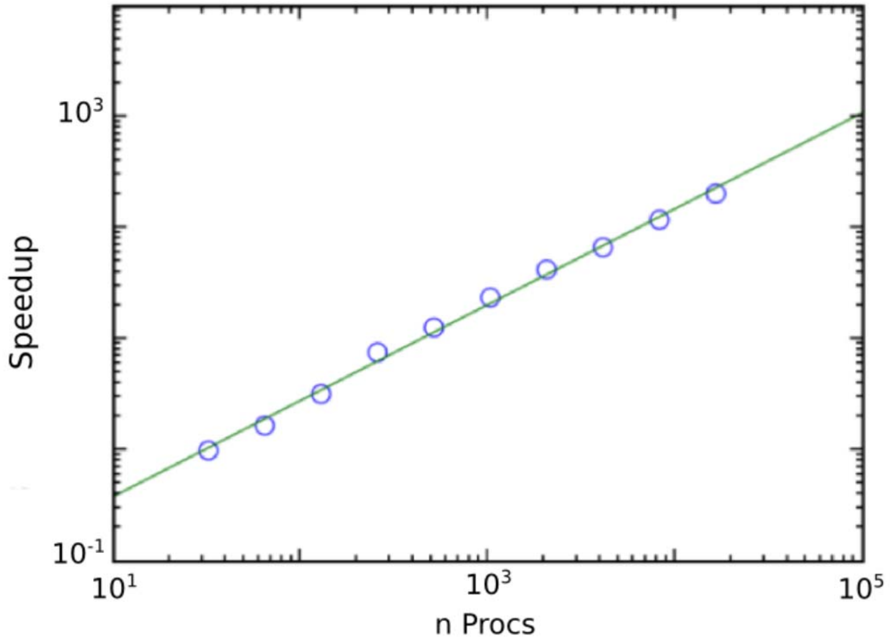

2. Computational Strategies - 2017

- Changes expected in computational approach by 2017:
  - CQL3D will be fully implicit ($v_\perp$, $v_{//}$, r) with a parallel matrix inversion:
    - Will require sparse matrix solver.
  - Utilization of GPU architectures for speed-up of matrix inversion, reconstruction of quasilinear diffusion coefficient and power absorption profiles (already partially demonstrated with AORSA).

2. Computational Strategies - 2017

- Also plan to pursue more speculative approaches:
  - Iterative solution of matrix in full-wave spectral solvers:
    - Needs the proper pre-conditioner.
  - Development of core to edge full-wave solver using pure finite elements:
    - Advantageous for describing complicated 3-D solid geometry of ICRF and LHRF launchers and tokamak vacuum vessel.
    - Requires that the conductivity operator be re-derived in the appropriate finite element basis set.

3. Current HPC Usage

- Machines currently used:
    - Hopper (NERSC)
    - CRAY XK6 (ORNL)

- Hours used in 2012:
    - Hopper (NERSC) – 8,000,000 hours
    - CRAY XK6 (ORNL) – 3,200,000 hours

- Typical parallel concurrency and run time, number of runs per year for three largest codes (taken from 2013 ERCAP Request):
    - AORSA – typically 1 toroidal mode per run, 1203 runs per year
    - TORLH – typically 1 toroidal mode per run, 1000 runs per year
    - VORPAL – 5 runs per year – full 3-D antenna simulations

3. Current HPC Usage

- Data read/written per run
  - Codes read ASCII format input data files.
  - Wave solvers and Fokker Planck codes produce NETCDF format output files with field solutions, RF diffusion coefficients, and distribution functions.
- Memory used per (node | core | globally) – Hopper (32 BG / node and 24 cores per node):
  - AORSA (single toroidal mode on 2040 cores) – (29 GB / node), (1.2 GB / core), 2.5 TB global
  - TORLH (single toroidal mode on 2040 cores) – (21GB / node), (0.88 GB / core), 1.8 TB global
- Necessary software, services or infrastructure
  - Run time - FFTW, MPI, NetCDF, ScaLAPACK
  - Post analysis: IDL, MATLAB, pgplot, PYTHON, VisIT

3. Current HPC Usage

- Data resources used (HPSS, NERSC Global File System, etc.) and amount of data stored:
  - /SCRATCH and /SCRATCH2 for running simulations
  - NERSC Global File System for management of source code and input data
  - HPSS for storing results, although most output data sets are small enough to transfer back to local computers (~ 5TB).

4. HPC Requirements for 2017

(Key point is to directly link NERSC requirements to science goals)

- Compute hours needed (in units of Hopper hours):
  - ~ 60,000,0000 hours
- Changes to parallel concurrency, run time, number of runs per year:
  - Expect full-wave solvers to be simulating ~100-2000 toroidal modes concurrently with approximately 5-7 runs per year.
- Changes to data read/written
  - Data read expected to have same format and size per toroidal mode.
  - Total data written will be larger – but still small (0.1 TB / simulation).

4. HPC Requirements for 2017

(Key point is to directly link NERSC requirements to science goals)

- Changes to memory needed per (core | node | globally):
  - **Memory per node and memory per core unchanged.**
  - **Global memory increases to 700-18000 TB because of parallel currency in simulating toroidal modes.**
- Changes to necessary software, services or infrastructure:
  - **Do not anticipate changes in software for data visualization.**
  - **Plans to do ITER calculation will rely on CPU / GPU architecture (see next section).**
  - **May need next generation of High Performance Linpack (HPL) to maintain strong scaling at > 50,000 processors as platform changes.**

5. Strategies for New Architectures

- Our strategy for running on new many-core architectures:
  - Use mixed CPU/GPU to improve performance of key parts of AORSA and VORPAL., so that we simulate the multi-scale RF power coupling problem self-consistently, in 3D.
  - Use GPU acceleration for LU matrix factorization in AORSA.
  - Use GPU acceleration for RF diffusion coefficient and power absorption profile reconstructions in AORSA.
  - Use GPU acceleration to improve performance of the FDTD algorithm in VORPAL.

5. Strategies for New Architectures

- To date we have prepared for many core by …
  - **Have carried out development work for AORSA on the TITAN supercomputer at the OLCF (16 CPU + 1 GPU per node) using the TITAN-dev partition.**
    - **Matrix factorization algorithm in AORSA solver modified to use an out-of-core (OOC) LU factorization for large dense complex matrices that takes advantage of GPU acceleration.**
    - **The library is designed to be compatible with the ScaLAPACK LU factorization routine PxGETRF.**
    - **External memory (or out-of-core) left-looking algorithm allows significant problems that are larger than available GPU device memory to be factored.**

5. Strategies for New Architectures

- To date we have prepared for many core by …
  - Have carried out development work for VORPAL on the Dirac GPU machine at NERSC (work carried out by Tech-X as this is a commercial code):
    - FDTD algorithm adapted to GPU's, retaining strong scaling.
    - Challenges were creation of dynamically generated kernels for initial/boundary conditions supplied at runtime (using a code generator solution), and efficient hiding of the GPU-to-CPU and CPU-to-CPU data transfer latency which required reordering the execution steps.
- Our plans include the following:
  - Submitted a proposal to acquire resources on TITAN through the ASCR Leadership Computing Challenge (ALCC) Program (D. L. Green, PI and D. N. Smithe).
  - Proposal is titled "Unraveling the Coupling of Radio Frequency Power to Fusion Plasmas".

5. Strategies for New Architectures

- To be successful on many-core systems we will need help with:
  - Maintaining strong scaling with our direct matrix solvers beyond 50,000 cores, by having software such as the HPL available on new platforms.
  - Making transition to GPU architectures as seamless as possible – for example, the compatibility of the GPU based OOCLU library with ScaLAPACK.

6. Summary

- What new science results might be afforded by improvements in NERSC computing hardware, software and services?
  - The capability to simulate multi-scale RF power coupling problem self-consistently, in 3D and at the scale of an ITER sized device.
  - Capability to perform time dependent core to edge simulations in a whole device modeling framework (such as the IPS framework developed in the SWIM Proto-type FSP).
- Recommendations on NERSC architecture, system configuration and the associated service requirements needed for your science:
  - Our experience on the TITAN-dev partition has been favorable thus far with a mixed CPU-GPU architecture.
  - For us, libraries on GPU system (such as the LU matrix factorization) that are compatible with ScaLAPACK are quite useful.

6. Summary

- NERSC generally refreshes systems to provide on average a 2X performance increase every year. What significant scientific progress could you achieve over the next 5 years with access to 32X your current NERSC allocation?
  - The capability to simulate multi-scale RF power coupling problem self-consistently, in 3D and at the scale of an ITER sized device.
  - Capability to perform time dependent core to edge simulations in a whole device modeling framework (such as the IPS framework developed in the SWIM Proto-type FSP).
- What "expanded HPC resources" are important for your project?
  - Shorter queue waits would really help a lot …