TECH-X
SIMULATIONS EMPOWERING
YOUR INNOVATIONS

Simulation of intense beams with exascale-ready Vorpal

Presented by John Cary
FAST/IOTA meeting
June 12, 2019

Performance team
Scott Sides
Jarrod Leddy
Ben Cowan
Sergey Averkin
Ilya Zilberter

IOTA: large nonlinear tune for stability

- Nonlinear tune: smaller resonances from errors = Immunity from nonlinear errors due to Landau damping
- Landau stabilization of otherwise instabilities
- Damping of oscillations due to matching errors

- The goal is to work at high intensity

Stellarators provide (imperfect) analogy for intense, nonlinear, integrable beams (INLB)

- Need magnetic lines to have rotation (like needing tune in an accelerator lattice)
- Rotation is nonlinear (rate varies with distance from axis)
- Self fields are important
  - Stellarator: from confinement currents
  - NLB: from net of charge - current forces
- Equilibrium: a state with the periodicity of the underlying systems
- Instabilities (and stable oscillations): time dependent or static
- Transport
  - Stellarator: collisional + orbit mechanisms, turbulence
  - NLB: collisional + orbit mechanisms, turbulence
- What can we learn?

But the analogy is not perfect

Stellarators	**Nonlinear, integrable accelerators**
Has vacuum rotational transform due to magnetic fields from coils	Has betatron tune due to magnetic fields from coils.
Self-consistent fields important	Self-consistent fields important
Local, PDE for the equilibrium	Integro-differential equation for the equilibrium
Conditions well understood for good confinement (quasihelicity, omnigenity) with self-consistency	No general principles for integrability
Multiple methods for computing equilibrium	No methods for computing equilibrium
Large orbit effect on transport: both theory and computations	No calculations of transport including effects of modified orbits.

Stellarator timeline may temper expectations

Year	Accomplishment
1966	Model-C stellarator so poor, tokamaks adopted, experimental stellarator research dropped for 30 years
1984	Local, PDE for the equilibrium
1982	Discovery of integrable vacuum fields
1984	Discovery of quasihelical symmetry (self-consistent)
1997	Restart of stellarator program with HSX (Wisconsin)
1997	Discovery of omnigenity symmetry (self-consistent)
1997	Design for NCSX initiated
~2003	NCSX construction begins
2008	NCSX cancelled after $90M spent
2017	First plasma in Wendelstein (Greiswald)

← we are here

Worldwide research effort led to performant stellarators

Need to start moving to self-consistent (large tune depression) studies

- Resonance reduction well known
  - 2D resonances studied on Tevatron (E778) in 1988,89
  - Could be studies in 4D
- Intense beams (large space charge)?
  - Problems due to not matching
  - Will resonances open up?

Mismatch oscillations lead to halo

- Core-halo model (Gluckstern): R. L. Gluckstern, Phys. Rev. Lett. **73**, 1247 (1994)
- Cylindrically symmetric
- Add oscillation with associated oscillation of the potential

- Get very large amplitude oscillations for the particles at the edge

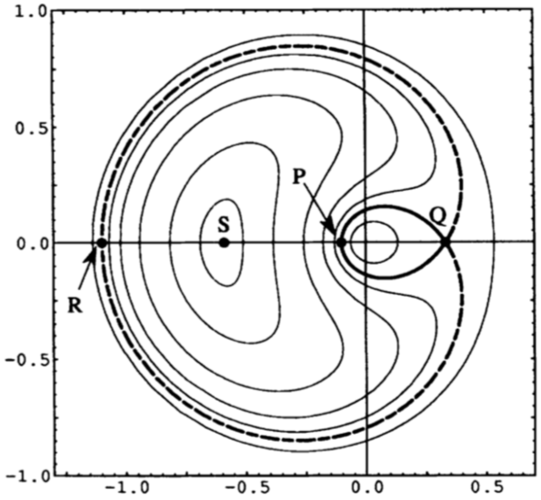

FIG. 2. Polar plot of w vs Ψ for the trajectories corresponding to the parametric resonance using $\Delta = 0.35$, $\epsilon = 0.1$, and the simplified model.

Will integrable nonlinearity prevent this?

- Sonnad, Cary PR-ST/AB 8, 064202 (2005)
- Cylindrically symmetric, with
  - Linear imposed forces $F = -k_0 r,$
  - Nonlinear imposed forces $F = -k_1 r - k_2 r^3.$
  - Self-consistent (space charge) forces

$$\frac{d^2 r}{ds^2} = F + F_{sc} + \frac{L^2}{r^3},$$

Nonlinearity definitely causes damping

- Launched with mismatch of 30%
- After oscillations have died down

- Oscillations decrease significantly, but they never go away

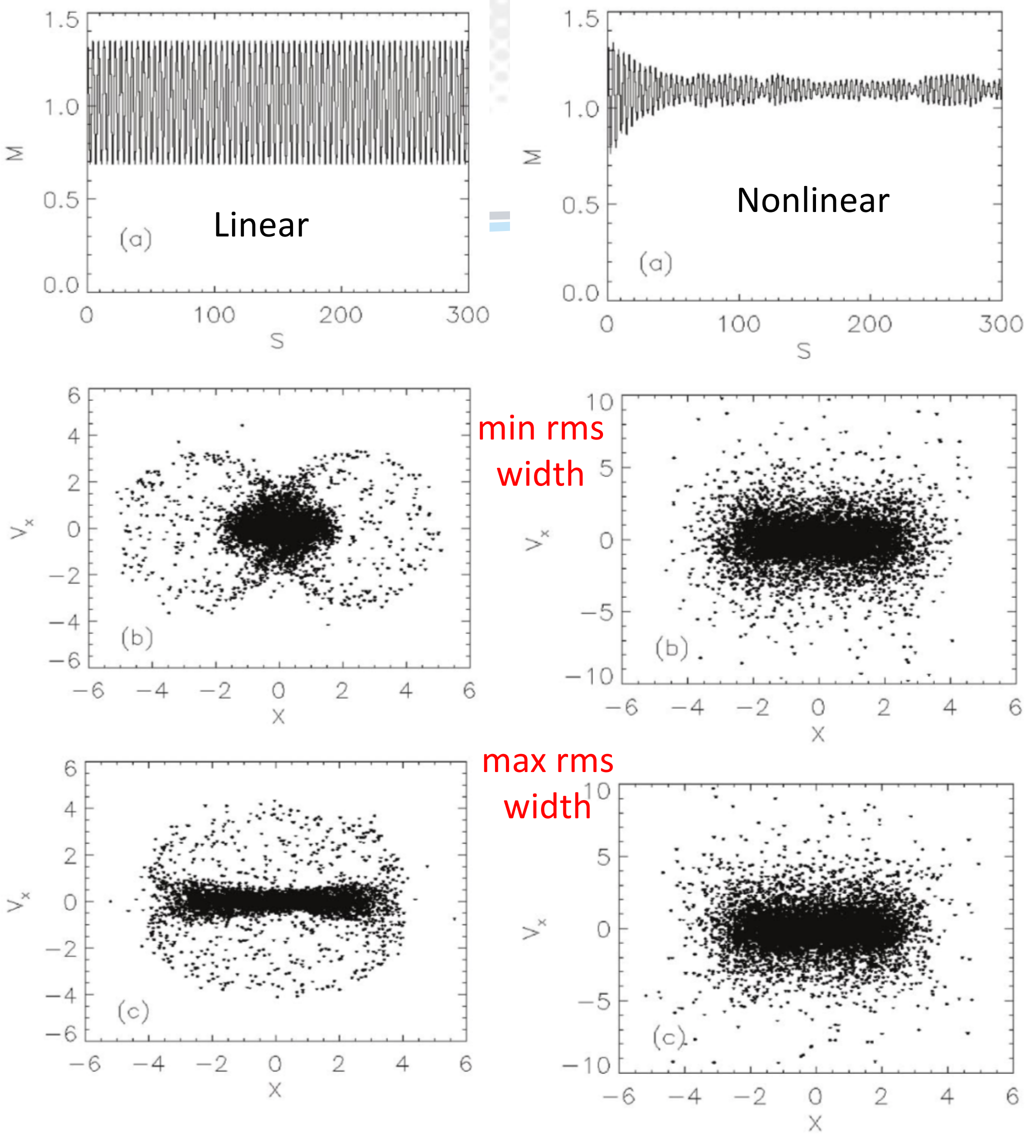

However, halo particles remain: NLB not enough

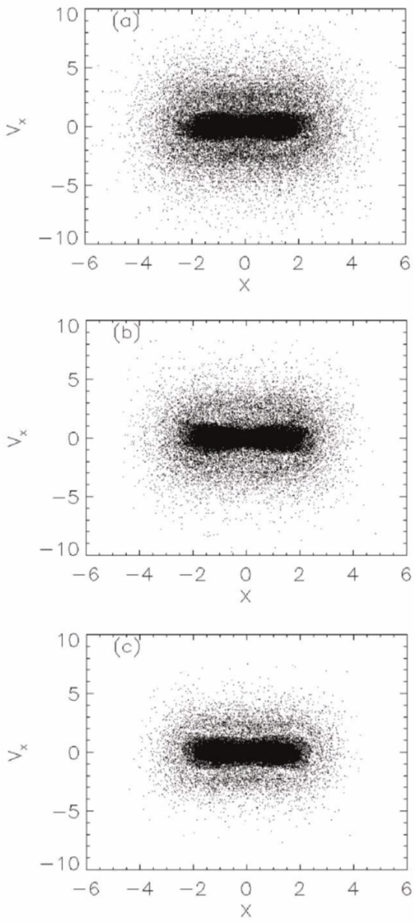

- But still large effect, perhaps too much?
- To avoid this, need to load (paint) beam consistent with the equilibrium, but what is the equilibrium?

Equilibrium calculations by expansion done previously at Colorado

- Finding integrable systems

  W. Wan and J. R. Cary, “Finding Four Dimensional Symplectic Maps with Reduced Chaos,” Phys. Rev. ST/AB 4, 084001 (2001).

  K. Sonnad and J. R. Cary, "Finding a nonlinear lattice with improved integrability using Lie transform perturbation theory," Phys. Rev. E. 69, 056501 (2004)
- Nonlinear systems for halo control

  K. Sonnad and J. R. Cary, "Control of beam halo formation through nonlinear damping and collimation," Phys. Rev. ST/AB **8**, 064202 (2005).
- Equilibria through perturbation theory

  K. G. Sonnad and J. R. Cary, “Near equilibrium distributions for beams with space charge in linear and nonlinear periodic focusing systems,” Phys. Plasmas 22, 043120 (2015); http://dx.doi.org/10.1063/1.4919033.
- See also (and cites within)

  S. M. Lund, S. H. Chilton, and E. P. Lee, “Efficient computation of matched solutions of the kapchinskij-vladimirskij envelope equations for periodic focusing lattices,” Phys. Rev. ST Accel. Beams, vol. 9, p. 064201, Jun 2006.

Previous calculations indicate no-halo equilibria possible

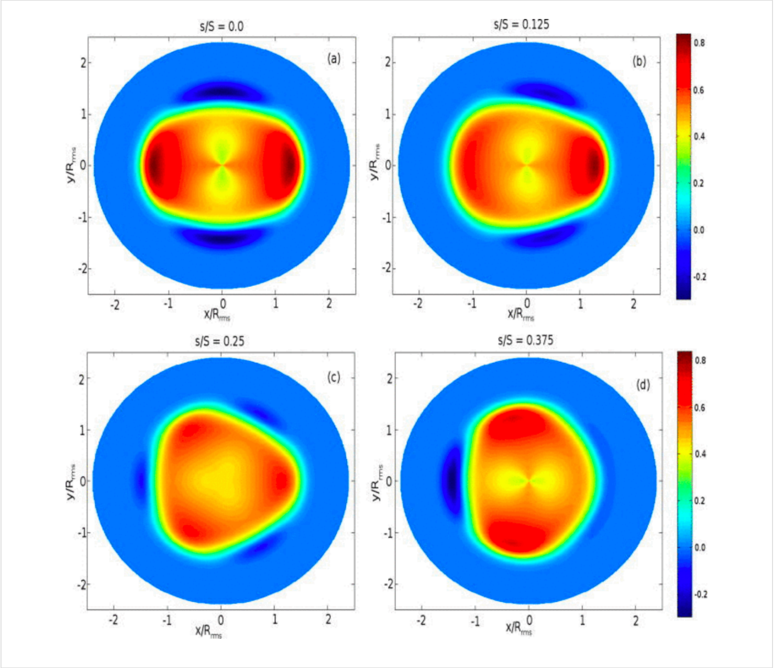

FIG. 4.

Plots of the density function $\pi R_{rms}^2(n_0(r)+n_2(r,\theta,s))/N_{tot}$ of a beam under nonlinear focusing, at different points along the lattice.

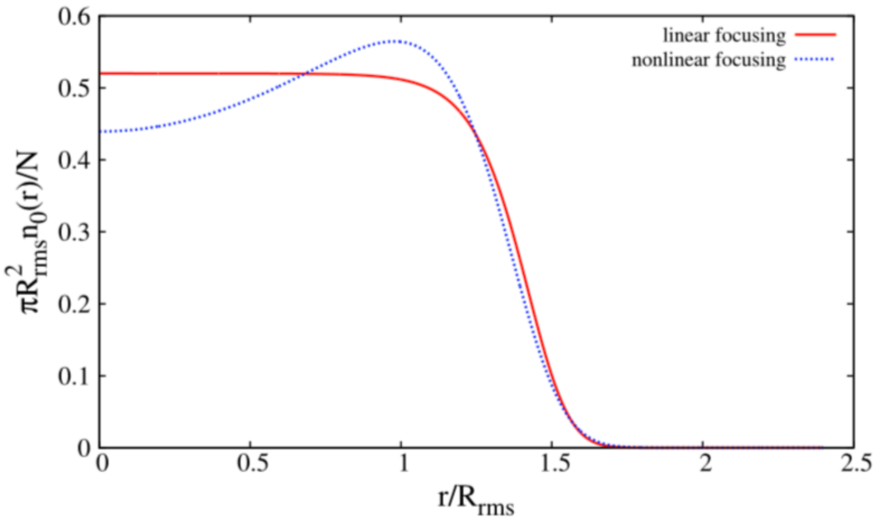

FIG. 6. The unperturbed number density n_0 for the linear and equivalent nonlinear focusing cases.

K. G. Sonnad and J. R. Cary, “Near equilibrium distributions for beams with space charge in linear and nonlinear periodic focusing systems,” Phys. Plasmas 22, 043120 (2015); http://dx.doi.org/10.1063/1.4919033.

Can clean with collimation, but at cost of beam loss

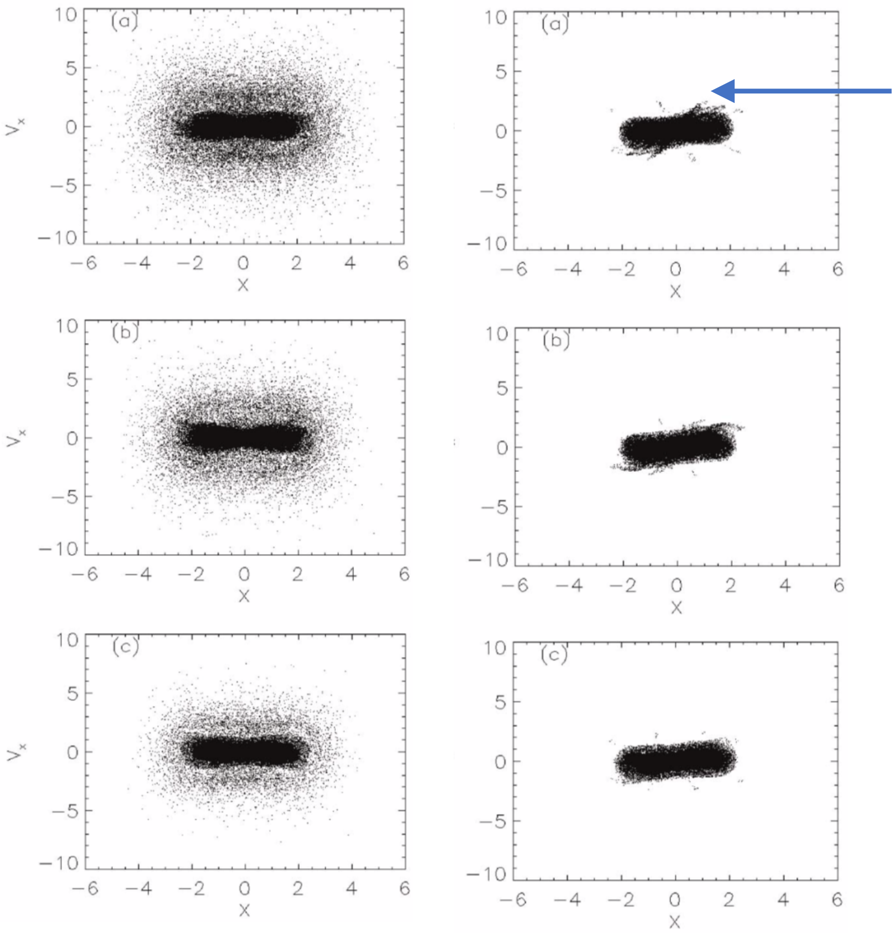

- Cleaning needs to continue, as tendrils keep forming
- Beam is lost with collimation
- Probably cannot afford to lose so much beam

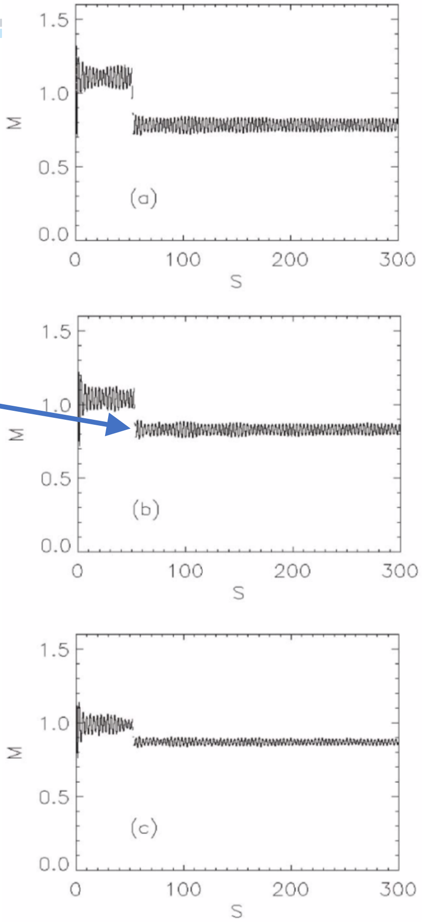

But leads to method of computing beam equilibria

- Launch arbitrary beam into nonlinear lattice
- Let beam relax
  - Due to phase mixing of nonlinearity
  - Due to scraping off of large-orbit particles
- Result: a beam equilibrium with no halo
- Use that for programming the beam painting

Accurate *intense* beam dynamics modeling requires full PIC

Exascale Vorpal coming on line for this purpose

- The computational engine of VSim (https://www.txcorp.com/vsim)
- Multiphysics for electromagnetics, electrostatics, (magnetostatics soon) of structures, kinetic and fluid species
- Cross platform: supercomputers to desktops, including Windows
- User friendly, well documented
- With about 100 FTE-years of investment
- With 100's of licensing agreements in >15 countries since 2012, including multiple labs in US, UK, Germany, Russia
- The most frequently cited computational plasma app (at last check)

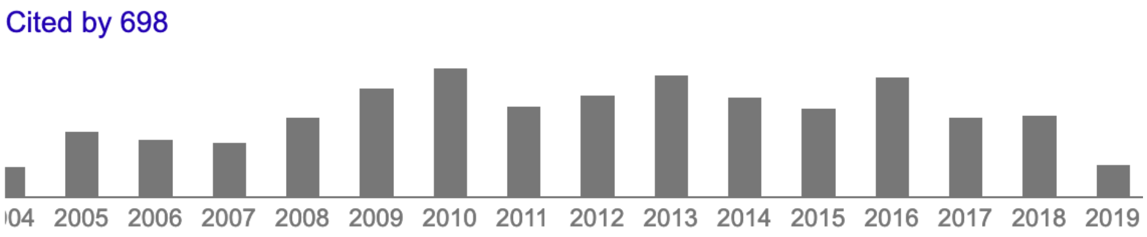

Full package:	VSim
Comp. engine:	Vorpal
Front end:	VSimComposer

Vorpal has a different business model

Code	Method of support	Access
VSim	SBIR, Sales, Grants	Commercial or collaboration
OSIRIS	DOE, SciDAC	MOU
WARPX	DOE, SciDAC, ECP	FOSS

Commercial drives ease of use

Vorpal and Exascale, what gives?

- Vorpal is **not** part of the Exascale Computing Project
  - In HEP, only WarpX is, so if you need beam equilibrium solves, collisions, cut-cell accuracy, MADX parser, sit tight until 2023
- Exascale is inclusive of
  - ✓ Multiple levels of hierarchy: distributed memory, multiple device, threads, and vector instructions (as the case may be)
  - ✓ Running on Cori, other computers as we get access
  - ○ Running on some future computers not yet built
- Vorpal funded by DARPA to be ported to GPUs [but $1.5M over 3 years << $20M ($100M?) over 5 years]
- Tech-X has used the opportunity to get Vorpal ready for device, threaded, vector computing (as well as distributed memory)
- Vorpal success now being built upon by FES
- Vorpal offered to IOTA as part of collaboration

New DOE supercomputers all rely on multi- and heterogeneous device computing

- Summit (2018)
  - 4,608 nodes, each with
  - 2 IBM Power 9 CPUs/node
  - 6 Nvidia Volta GPUs/node
  - Code via CUDA
  - https://www.olcf.ornl.gov/summit/
- Perlmutter (2020)
  - AMD Epyc CPUs
  - 4 NVidia GPUs per node
  - Code via CUDA
  - https://www.nersc.gov/systems/perlmutter
- Frontier (2021)
  - AMD Epyc CPUs
  - 4 Radeon Instinct GPUs per node
  - Code via HIP (designed to be CUDA compatible)
  - https://www.olcf.ornl.gov/frontier
- Aurora (2021)
  - Intel Xeon
  - Intel’s Xe compute architecture (vapor?)
  - Code via SYCL (vapor?)
  - https://aurora.alcf.anl.gov

All require multiple-device coding as is available in VSim

Why use Vorpal?

- Can work collaboratively
  - Available at NERSC for collaborators
- Commercial brings
  - User-friendly interface
  - Variables, parsing
  - CAD capabilities
  - Extensive documentation
  - User support
  - Cost reduction (commercial customers paying for GUI, CAD)
  - Affordable HPC
- Scientific collaboration brings
  - Largest scale
  - Latest algorithms

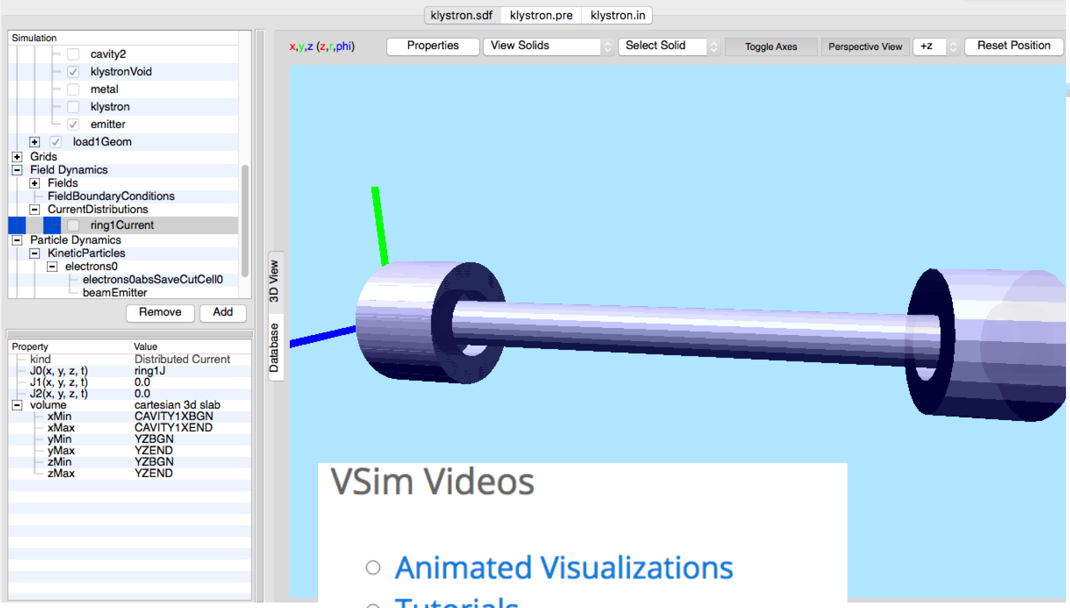

Property	Value
kind	Distributed Current
J0(x, y, z, t)	ring1J
J1(x, y, z, t)	0.0
J2(x, y, z, t)	0.0
volume	cartesian 3d slab
xMin	CAVITY1XBGN
xMax	CAVITY1XEND
yMin	YZBGN
yMax	YZEND
zMin	YZBGN
zMax	YZEND

VSim Videos

- Animated Visualizations
- Tutorials

Documentation

www.txcorp.com/vsim

- Online Documentation
- Installation and Release Notes
- VSim Examples
- VSim User Guide
- VSim Reference Manual
- VSim Customization

Vorpal's basic assumptions align well with exascale - 2

- Use of embedded boundary methods gives accuracy and can be used with Richardson extrapolation
- *Unique in the DoE portfolio.*

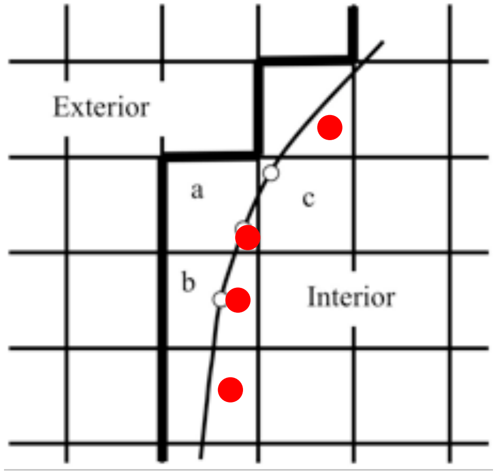

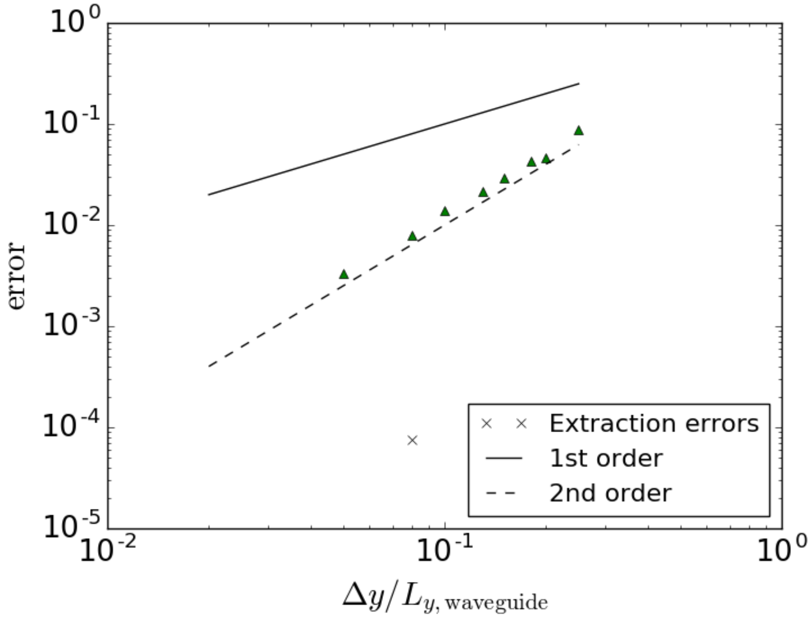

2^{nd} order accuracy

Thermal equilibration now simulated on GPU

- 100x100 cells, 10 PPC, isotropic
  - Argon at 300K
  - Hydrogen at 1000K
  - Binary elastic collisions
  - Ar-Ar, H-H, and Ar-H
- Same code, compile-time option for GPU use (will eventually be run-time)
  - 1-core CPU (i7-6700, 3.4GHz 4 cores): 175s (44s if 4 cores?)
  - GTX 745 GPU (384 cores, 1GHz): 22s
- Next steps: optimization and profiling to get even more speed

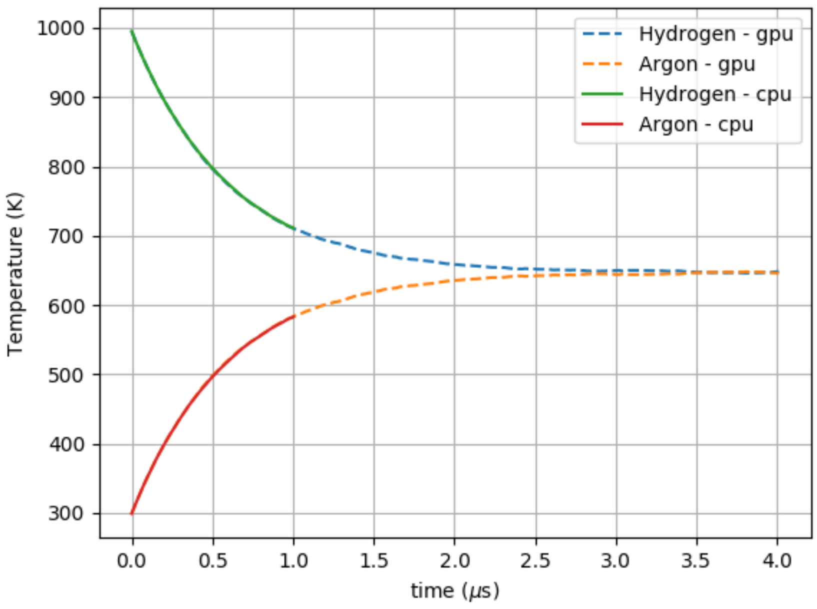

Maxwell equation updates getting near perfect scaling

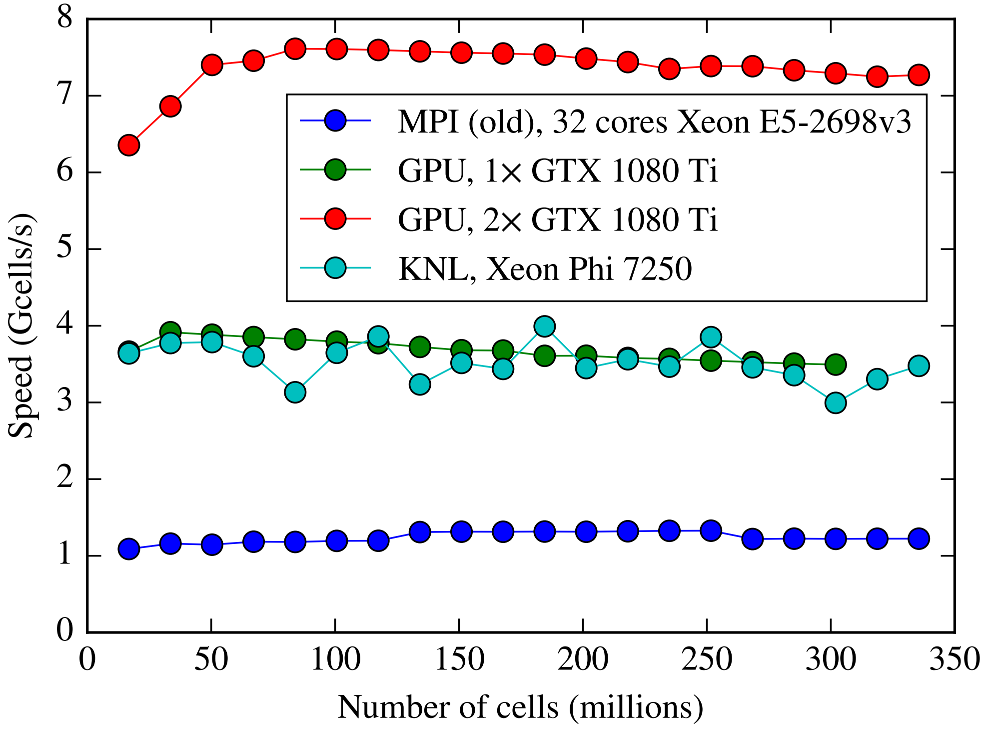

Initial results, not tuned,
not using NVLink

How many particles?

- 100k particles/GPU
- ~1 GP/sec (10^9 particle/sec)
- 10,000 steps/sec

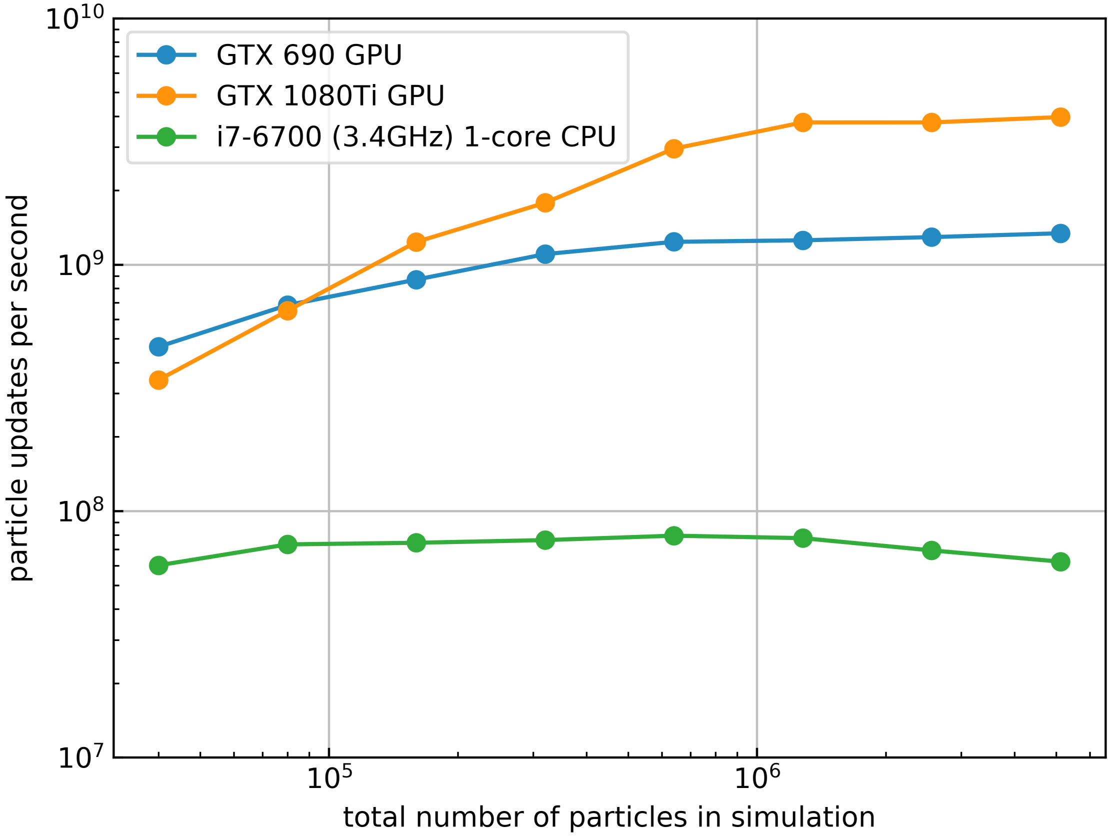

Memory for a moving window simulation modest by current standards

- Longitudinal variations
  - 7 m sections, 10 cm elements, 10 cm beam, so 1 cm cells.
  - 2.9e9 particles (3000 GPUs)
- Looks like cells are 1cm longitudinal,
- Width:
  - Circulating beam size, 1-5mm (protons)
  - Tube radius of 12 mm.
  - Well resolved with 0.12 mm cells, so $\pi 200^2$ or 4e4 cells/plane
- Total volume for moving window: $20\pi 1.2^2 = 90\text{cm}^3$
- Cells could be $1\text{x}0.012^2\text{cm}^3$ (630k cells) but with poor dispersion, 0.012^3cm^3 (52M cells) with good dispersion.
- Neither case seems particularly challenging in terms of memory

Time steps pose stronger condition

- 40m circumference, 0.012 cm cells, 300k steps/turn
- 10k steps/sec, so 30 s/turn
- 100 turns (3e7 steps) is 30m computation.

- Simulating the full ring costs no more except memory.

Proposed research: compute equilibria, enhance algorithms

- Compute equilibria by full PIC plus large-orbit particle removal
  - Bring in elements defined by MAD-X files
  - Start with FODO + nonlin elements from Sonnad/Cary
  - Launch particles as expected experimentally
  - Upon demonstration of method, move to IOTA lattice
  - Work with IOTA to test code at each step
- As we approach very long time (1M turns) simulations, need to prepare for highly stable simulations with space charge
  - Self-consistent, pic-scaling symplectic simulations (wave-particle introduced in Cary, Doxas): requires C2, generalize to EM
  - Structure preserving, perhaps symplectic particle integration for E&B

Vorpal available to collaborators at NERSC

- Contact us
- Exascale capable not yet ready

Stellarators: Now have the Simons Institute on Hidden Symmetries

- https://hiddensymmetries.princeton.edu/
- 10 institutions
- $2M/yr
- Theory ONLY!

- Hidden symmetries also for intense NLB

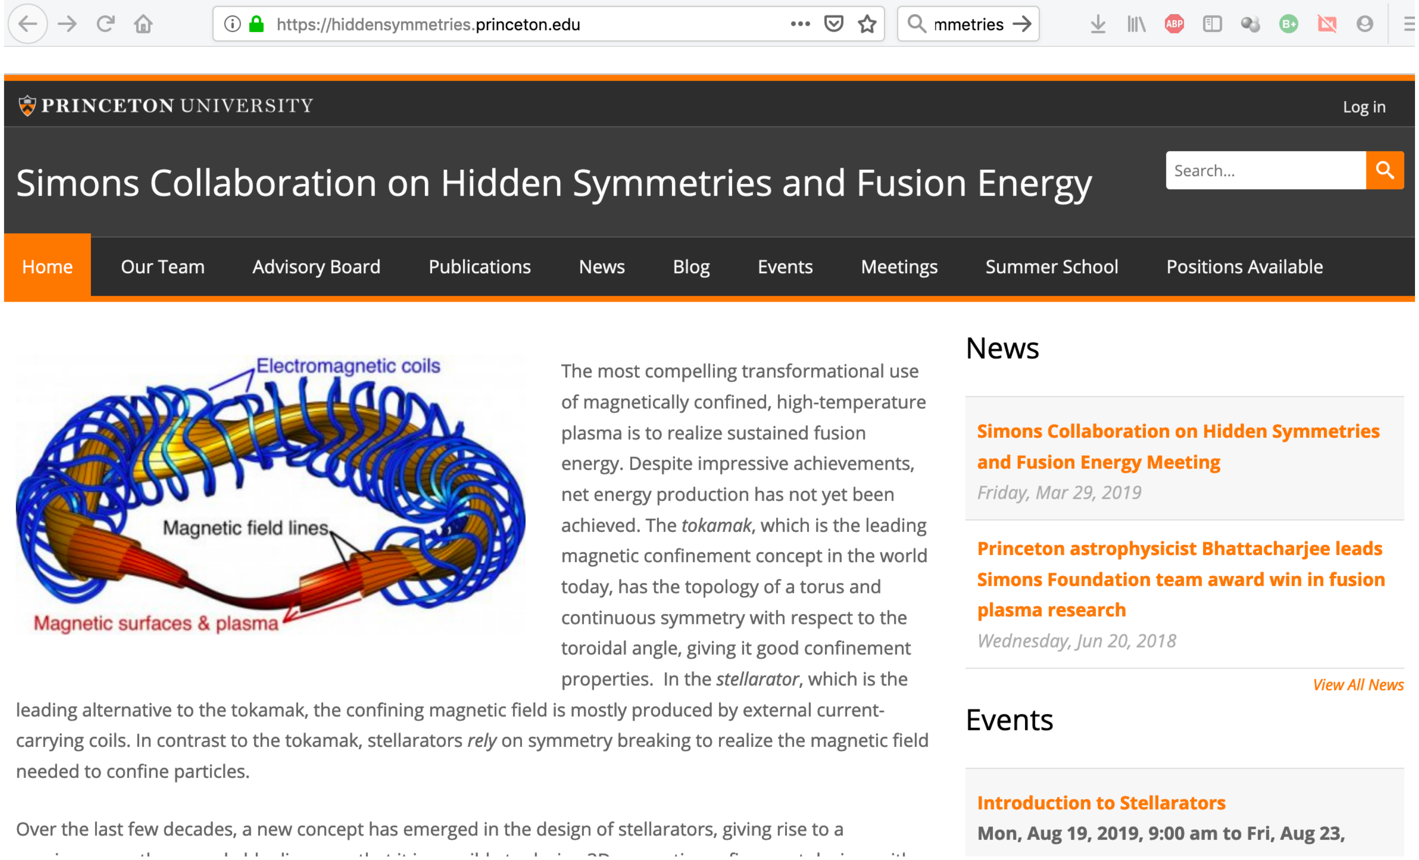

PRINCETON UNIVERSITY

Log in

Simons Collaboration on Hidden Symmetries and Fusion Energy

Search...

Home | Our Team | Advisory Board | Publications | News | Blog | Events | Meetings | Summer School | Positions Available

The most compelling transformational use of magnetically confined, high-temperature plasma is to realize sustained fusion energy. Despite impressive achievements, net energy production has not yet been achieved. The *tokamak*, which is the leading magnetic confinement concept in the world today, has the topology of a torus and continuous symmetry with respect to the toroidal angle, giving it good confinement properties. In the *stellarator*, which is the leading alternative to the tokamak, the confining magnetic field is mostly produced by external current-carrying coils. In contrast to the tokamak, stellarators *rely* on symmetry breaking to realize the magnetic field needed to confine particles.

Over the last few decades, a new concept has emerged in the design of stellarators, giving rise to a

News

Simons Collaboration on Hidden Symmetries and Fusion Energy Meeting

Friday, Mar 29, 2019

Princeton astrophysicist Bhattacharjee leads Simons Foundation team award win in fusion plasma research

Wednesday, Jun 20, 2018

View All News

Events

Introduction to Stellarators

Mon, Aug 19, 2019, 9:00 am to Fri, Aug 23,