

KINETIC PLASMA TURBULENCE SIMULATIONS ON TOP SUPERCOMPUTERS WORLDWIDE

William M. Tang*

**Princeton Institute for Computational Science & Engineering (PICSciE)
and Princeton Plasma Physics Laboratory
Princeton University, Princeton, New Jersey**

**10th West Lake International Symposium on Magnetic Fusion
& 12th Asia Pacific Plasma Theory Conference (WLIS-APPTC)**

Hangzhou, China

May 9-13 , 2016

***Collaborators:** *Bei Wang (PU), S. Ethier (PPPL), Bruce Scott (IPP-Garching), K. Ibrahim (LBNL), K. Madduri (Penn State U), S. Williams (LBNL), L. Olier (LBNL), T. Williams (ANL), C. Rosales (TACC), T. Hoefler (ETH-Zurich), G. Kwasniewski (ETH-Zurich), Y. Lu (NUDT)*

INTRODUCTION

I. FOCUS: HPC Performance Scalability and Portability with FES as appropriate example application domain

→ *Illustration of domain application that delivers discovery science with good performance scaling, while also helping provide viable metrics on top supercomputing systems such as “portability,” “time to solution,” & associated “energy to solution”*


Reference: “*Scientific Discovery in Fusion Plasma Turbulence Simulations @ Extreme Scale*,” W. Tang, B. Wang, S. Ethier, Computing in Science and Engineering (CiSE), vol. 16. Issue 5, pp.44-52, 2014

II. PHYSICS & COMPUTATIONAL PROGRESS: *Enabled by deployment of innovative algorithms – e.g., using MPI & OpenMP within modern code that delivers new scientific insights on world-class systems → currently: Mira; Sequoia; K-Computer; Titan; Piz Daint; Blue Waters; Stampede; TH-2*

& in near future on: Summit (via CAAR), Cori, Stampede-II, Tsubame 3.0, -----

III. FUTURE: *Ability to utilize computing at exascale & beyond will require algorithmic & solver advances enabled by Applied Mathematics – in an interdisciplinary “Co-Design” type environment together with Computer Science & Extreme-Scale HPC Domain Applications*

Performance Development of HPC over the Last 22 Years from the Top 500 (J. Dongarra)



Applications Impact → Actual value of extreme Scale HPC to scientific domain applications & industry

Context: new US Govt. announcement of NATIONAL STRATEGIC COMPUTING INITIATIVE

- Practical Considerations: “Better Buy-in” from Science & Industry requires:

- Moving beyond “voracious” (more of same - just bigger & faster) to “transformational” (achievement of major new levels of scientific understanding)
- Improving experimental validation and verification to enhance realistic predictive capability of both hypothesis-driven and big-data-driven statistical approaches
- Deliver software engineering tools to improve “time to solution” and “energy to solution”
- David Keyes (KAUST/Columbia U) → “Billions of \$ of scientific software worldwide hangs in the balance until better algorithms arrive to span the architecture-applications gap.”

- Associated Challenges:

- Hardware complexity: Heterogeneous multicore; gpu+cpu → Summit; mic+cpu → Aurora
- Software challenges: Rewriting code focused on data locality

- Applications Imperative: “Accountability” aspect

→ Need to provide specific examples of impactful scientific and mission advances enabled by progress from terascale to petascale to today’s multi-petascale HPC capabilities

Demonstration of GTC Productivity & Impact → Delivery of Scientific Advances
with use of Increasingly Powerful Supercomputing Systems

GTC Simulation	Computer name	PE # used	Speed (TF)	# Particles used	T i m e steps	Physics Discovery (Publication)
1998	Cray T3E NERSC	10^2	10^{-1}	10^8	10^4	Ion turbulence zonal flow (<i>Science</i> , 1998)
2002	IBM SP NERSC	10^3	10^0	10^9	10^4	Ion transport size scaling (<i>PRL</i> , 2002)
2007	Cray XT3/4 ORNL	10^4	10^2	10^{10}	10^5	Electron turbulence (<i>PRL</i> , 2007); EP transport (<i>PRL</i> , 2008)
2009	Jaguar/Cray XT5 ORNL	10^5	10^3	10^{10}	10^5	Electron transport scaling (<i>PRL</i> , 2009); EP-driven MHD modes
2012 to present	Cray XT5→Titan ORNL Tianhe-1A (China)	10^5	10^4	10^{11}	10^5	Kinetic-MHD (<i>PRL</i> , 2012); Turbulence + EP + MHD TAE Modes (<i>PRL</i> , 2013)
2018 (future)	Path to Exascale HPC Resources	TBD	10^6	10^{12}	10^6	Turbulence + EP + MHD + RF

*** GTC is first FES code to deliver production run simulations @ TF in 2002 and PF in 2009

DOE SciDAC Success Story from Fusion Energy Science: PIC Code Simulations of Confinement Loss from Turbulent Transport

- **Mission Importance:** Fusion reactor size & cost determined by balance between loss processes due to turbulent transport & self-heating rates from fusion reactions


- **“Scientific Discovery”** - Transition to favorable scaling of confinement for ITER-size plasmas [Good News for ITER!]

- $a/\rho_i = 400$ (JET, largest present lab experiment)
- $a/\rho_i = 1000$ (ITER, ignition experiment)

from Multi-TF simulations using 3D PIC code [Z. Lin, et al, **2002** → 1st ITER-scale simulation with ion gyroradius resolution with GTC-Code

- **“Co-Design Enabled Advances @multi PF”** → Excellent scalability of 3D PIC Codes on top 7 multi-PF supercomputers worldwide

→ enables unprecedented resolution/physics fidelity needed for better understanding of large fusion systems




Accelerated progress enabled by SciDAC-based interdisciplinary approach with Computer Science (CS), Applied Math (AM), & Domain Applications

→ Recent achievements enabled by SciDAC approach plus “CoDesign” inclusion of Hardware Design

Particle Resolution (ppc) Convergence Study


[Results from C-version GTC-Princeton Code for ITER (D-size) Case on IBM BG-Q]

Time History of Thermal Diffusivity from ITG Instability @ Different Resolutions



High Resolution Ion Transport Scaling Results enabled by “Mira” at ALCF


[vertical axis represents transport level and horizontal axis the plasma size with ITER at 1000]
two weights scheme + remapping



New Trends: “rollover” significantly more gradual than established earlier in much lower resolution, shorter duration studies with magnitude of transport now reduced by ~ 2

Picture of Particle-in-Cell Method


- Charged particles sample distribution function
- Interactions occur on a grid with the forces determined by gradient of electrostatic potential (calculated from deposited charges)
- *Grid resolution dictated by Debye length (“finite-sized” particles) up to gyro-radius scale*



Specific PIC Operations:

- “**SCATTER**”, or deposit, charges as “nearest neighbors” on the grid
- Solve Poisson Equation for potential
- “**GATHER**” forces (gradient of potential) on each particle
- Move particles (**PUSH**)
- Repeat...

GTC-P: six major subroutines



- **Charge:** particle to grid interpolation (**SCATTER**)
- **Smooth/Poisson/Field:** grid work (local stencil)
- **Push:**
 - grid to particle interpolation (**GATHER**)
 - update position and velocity
- **Shift:** in distributed memory environment, exchange particles among processors

KEY ROLE OF OPEN-MP in ADDRESSING MODERN HPC CHALLENGES


Open-MP-enabled scalable scientific software for extreme scale applications: FES as illustrative application domain

- **Extreme concurrency**: → *Adopting OpenMP is one of most efficient algorithmic approaches to facilitate efficient multi-threading methods*
- **Portability**: → *except for GPU hardware, OpenMP works with all multicore processors*
- **Ease of Deployment**: → *OpenMP is now a mature implementation relatively easy to use*
 - easiest approach to deploy OpenMP is at loop level;
 - OpenMP worked best at loop level beginning in late 1990s/early 2000s and has remained best approach since then;
 - Example: deployed this way in all prominent Global FES PIC codes – GTC-P, GTC, GTS, and XGC;

KEY ROLE OF OPEN-MP (continued)


- **NEED FOR OPEN-MP CAPABILITY IN MULTI-GRID SOLVERS SUCH AS LLNL'S "HYPRE"**
- **Significant challenge/goal:** *Incorporation of multi-grid Poisson solvers with OpenMP to efficiently deal with extreme concurrency, multi-threading issues characteristic of near-future systems (e.g., 100 PF systems such as Summit and Aurora)*
- **FES Application:** GTC (UC Irvine) & GTC-P Project in Fusion Energy Science were selected for current portfolio of OLCF CAAR Early Science Program for Summit → will require multi-grid electromagnetic field-solver with OpenMP such as HYPRE
- *Choosing a portable and threaded solver (e.g., HYPRE) is critically important for GTC and GTC-P*

ILLUSTRATION OF GTC-P CODE PORTABILITY



- Broad range of leading multi-PF supercomputers worldwide
- Percentage indicates fraction of overall nodes currently utilized for GTC-P experiments
- NOTE: Results in this figure are only for CPU nodes on Stampede and TH-2

Weak Scaling of GTC-P (GPU-version) on Heterogenous (GPU/CPU) “Titan” and “Piz Daint”




of nodes: 64 256 1024 4096

- The number of particles per cell is 100
- GTC-P GPU obtains 1.7x speed up

Same code for all cases → Performance difference solely due to
hardware/system software

*Aries
Network on
Piz Daint

● Mira ▲ Titan ◆ Piz Daint ⊗ Stampede (OFLD)
 □ TH-2 (CPU) △ Titan (CPU) ◇ Piz Daint (CPU) * Stampede (CPU)



GTC-P (kinetic electron) weak scaling performance using a fixed problem size per node across all systems allows comparisons of node performance.

“ENERGY TO SOLUTION” ESTIMATES
(for Mira, Titan, and Piz Daint)


	CPU-Only			CPU+GPU		
	Mira	Titan	Piz Daint	Titan	Piz Daint	
Nodes	4096	4096	4096	4096	4096	
Power/node (W)	69.7	254.1	204.9	269.4	246.5	
Time/step (s)	13.77	15.46	10.00	10.11	6.56	
Energy (KWh)	1.09	4.47	2.33	3.10	1.84	

- **Energy per ion time step (KWh) by each system/platform for the weak-scaling, kinetic electron studies using 4K nodes.**

$(\text{Watts/node}) * (\# \text{nodes}) * (\text{seconds per step}) * (1\text{KW}/1000\text{W}) * (1\text{hr}/3600\text{s})$


- **Power/Energy estimates obtained from system instrumentation including compute nodes, network, blades, AC to DC conversion, etc.**

GTC-P SIMULATION OF MICROTURBULENCE IN ITER-SIZE PLASMA



Density fluctuations during the non-linear phase of an ITER-size GTC-P simulation of plasma microturbulence



ILLUSTRATION OF GTC-P CODE CAPABILITY FOR INCREASING PROBLEM SIZE



New Physics Results: Fusion system size-scaling study of “trapped-electron-mode” turbulence showing the “plateauing” of the radial electron heat flux as size of tokamak increases.

- **Verification** – via (1) *GTC Simulation Results* (up to C) – Y. Xiao, Z. Lin (PRL – 2009); and (2) *Analytic Foundations* – **Liu Chen** & Z. Lin hybrid kinetic model for electrons

Kolmogorov-like Rollover Spectral Results from GTC-P on multi-PF BG-Q“Mira” (collaboration with B. Scott, IPP-Garching)



Associated Findings

- Advanced phase space remapping and Krook-type collision models deployed for long temporal duration global PIC simulations with GTC-P (*Bei Wang and Bruce Scott*)

→ Current findings of realistic Kolmogorov-type spectral roll-over indicate:

- (i) numerical dissipation in some form needed to ensure proper “steady state” behavior for “collisionless” systems;
- (ii) while further improvements are needed, present results from remapping and Krook-type collision models effectively reduce noise levels for long-duration simulations where the amplitudes are small; and
- (iii) Inherent turbulence-driven dissipation sustainable in meaningful way – without invoking “ad-hoc” artificial dissipation models

APPLIED MATH LOCALITY CHALLENGE: GEOMETRIC HAMILTONIAN APPROACH TO SOLVING GENERALIZED VLASOV-MAXWELL EQUATIONS

Hamiltonian → Lagrangian → Action → Variational Optimization → Discretized Symplectic Orbits for Particle Motion

I. “Ultrahigh Performance 3-Dimensional Electromagnetic Relativistic Kinetic Plasma Simulation

Kevin J. Bowers, et al., *Phys. Plasmas* 15, 055703 (2008)

- Basic foundation for symplectic integration of particle orbits in electromagnetic fields without frequency ordering constraints
- Foundational approach for present-day simulations of laser-plasma interactions on modern supercomputing systems
- Limited applicability with respect to size of simulation region and geometric complexity

II. “Geometric Gyrokinetic Theory for Edge Plasmas”

Hong Qin, et al., *Phys. Plasmas* 14, 056110 (2007)

- Basic foundation for symplectic integration of particle orbits in electromagnetic low-frequency plasma following GK ordering
- Still outstanding challenge: Address reformulation of non-local Poisson Equations structure for electromagnetic field solve

Concluding Comments

- Modern FES HPC domain application code capable of scientific discovery while providing good performance scaling and portability on top supercomputing systems worldwide – together with illustrating the key metrics of “time to solution” and associated “energy to solution”

Reference: *“Scientific Discovery in Fusion Plasma Turbulence Simulations @ Extreme Scale;”* W. Tang, B. Wang, S. Ethier, Computing in Science and Engineering (CiSE), vol. 16. Issue 5, pp.44-52, 2014

- Current physics and computational progress enabled by deployment of innovative algorithms within a modern application code (GTC-P) that delivers new scientific insights on world-class systems → currently: *Mira; Sequoia; K-Computer; Titan; Piz Daint; Blue Waters; Stampede; TH-2*

with future targets: *Summit (via CAAR), Cori, Aurora, Stampede-II, Tsubame 3.0, -----*

- Future progress will require algorithmic & solver advances enabled by Applied Mathematics – *in an interdisciplinary “Co-Design” type environment together with Computer Science* (e.g., OpenMP4.5, OpenACC, etc.) *& Extreme-Scale HPC Domain Applications*