Computing and Storage Requirements for FES

J. Candy
General Atomics, San Diego, CA

Presented at

DOE Technical Program Review
Hilton Washington DC/Rockville
Rockville, MD
19-20 March 2013

GENERAL ATOMICS

Drift waves and tokamak plasma turbulence

Role in the context of fusion research

- **Plasma performance**:

  In tokamak plasmas, performance is limited by turbulent radial transport of both energy and particles.

- **Gradient-driven**:

  This turbulent transport is caused by drift-wave instabilities, driven by free energy in plasma temperature and density gradients.

- **Unavoidable**:

  These instabilities will persist in a reactor.

- **Various types (asymptotic theory)**:

  ITG, TIM, TEM, ETG . . . + Electromagnetic variants (AITG, etc).

Fokker-Planck Theory of Plasma Transport

Basic equation still unsolved in tokamak geometry

The Fokker-Planck (FP) equation provides the **fundamental theory** for plasma evolution:

$$\left[\frac{\partial}{\partial t} + \mathbf{v}\cdot\nabla + \frac{e_a}{m_a}\left(\mathbf{E} + \frac{\mathbf{v}}{c}\times\mathbf{B}\right)\cdot\frac{\partial}{\partial \mathbf{v}}\right] f_a = \sum_b C_{ab}(f_a, f_b) + S_a$$

where $\mathbf{E}$ and $\mathbf{B}$ satisfy the Maxwell equations.

Fokker-Planck Theory of Plasma Transport

Comprehensive ordering in series of papers by Sugama and coworkers

Systematic ordering for **plasma equilibrium**, **fluctuations**, and **transport**:

$$\left[\frac{\partial}{\partial t} + \mathbf{v}\cdot\nabla + \frac{e_a}{m_a}\left((\mathbf{E}+\hat{\mathbf{E}}) + \frac{\mathbf{v}}{c}\times(\mathbf{B}+\hat{\mathbf{B}})\right)\cdot\frac{\partial}{\partial \mathbf{v}}\right](f_a + \hat{f}_a)$$
$$= C_a(f_a+\hat{f}_a) + S_a$$

$f_a \longrightarrow$ **ensemble-averaged distribution**

$\hat{f}_a \longrightarrow$ **fluctuating distribution**

$S_a \longrightarrow$ **sources (beams, RF, etc)**

$C_a = \sum_b C_{ab}(f_a+\hat{f}_a, f_b+\hat{f}_b) \longrightarrow$ **nonlinear collision operator**

GENERAL ATOMICS

Fokker-Planck theory

Comprehensive, consistent framework for equilibrium profile evolution

The general approach is to separate the FP equation into **ensemble-averaged**, $\mathcal{A}$, and **fluctuating**, $\mathcal{F}$, components:

$$\mathcal{A} = \left.\frac{d}{dt}\right|_{\rm ens} f_a - \langle C_a \rangle_{\rm ens} - D_a - S_a \ ,$$

$$\mathcal{F} = \left.\frac{d}{dt}\right|_{\rm ens} \hat{f}_a + \frac{e_a}{m_a}\left(\hat{\mathbf{E}} + \frac{\mathbf{v}}{c} \times \hat{\mathbf{B}}\right) \cdot \frac{\partial}{\partial \mathbf{v}}(f_a + \hat{f}_a) - C_a + \langle C_a \rangle_{\rm ens} + D_a \ ,$$

where

$$\left.\frac{d}{dt}\right|_{\rm ens} \doteq \frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla + \frac{e_a}{m_a}\left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B}\right) \cdot \frac{\partial}{\partial \mathbf{v}} \ ,$$

$$D_a \doteq -\frac{e_a}{m_a}\left\langle \left(\hat{\mathbf{E}} + \frac{\mathbf{v}}{c} \times \hat{\mathbf{B}}\right) \cdot \frac{\partial \hat{f}_a}{\partial \mathbf{v}} \right\rangle_{\rm ens} \ .$$

▷ D_a is the **fluctuation-particle interaction operator**.

Fokker-Planck theory

Space- and time-scale expansion in powers of $\rho_* = \rho_s/a$

Ensemble averages are expanded in powers of ρ_* as

$$
\begin{aligned}
f_a &= f_{a0} + f_{a1} + f_{a2} + \dots \ , \\
S_a &= \qquad\qquad\qquad\quad S_{a2} + \dots \ \text{(transport ordering)}, \\
\mathbf{E} &= \mathbf{E}_0 + \mathbf{E}_1 + \mathbf{E}_2 + \dots \ , \\
\mathbf{B} &= \mathbf{B}_0 \ .
\end{aligned}
$$

Fluctuations are also expanded in powers of ρ_* as

$$
\begin{aligned}
\hat{f}_a &= \hat{f}_{a1} + \hat{f}_{a2} + \dots \ , \\
\hat{\mathbf{E}} &= \hat{\mathbf{E}}_1 + \hat{\mathbf{E}}_2 + \dots \ , \\
\hat{\mathbf{B}} &= \hat{\mathbf{B}}_1 + \hat{\mathbf{B}}_2 + \dots \ .
\end{aligned}
$$

Built-in assumption about scale separation **hard to escape**.

GENERAL ATOMICS

Fokker-Planck theory

Lowest-order conditions for flow and gyroangle independence

Lowest-order Constraints

The lowest-order ensemble-averaged equation gives the **constraints**

$$\mathcal{A}_{-1} = 0 : \qquad \mathbf{E}_0 + \frac{1}{c}\mathbf{V}_0 \times \mathbf{B} = 0 \qquad \text{and} \qquad \frac{\partial f_{a0}}{\partial \xi} = 0$$

where ξ is the gyroangle.

Large mean flow

The only **equilibrium flow** that persists on the fluctuation timescale is

$$\mathbf{V}_0 = R\,\omega_0(\psi)\mathbf{e}_\varphi \quad \text{where} \quad \omega_0 \doteq -c\frac{\partial \Phi_0}{\partial \psi} \; .$$

[F.L. Hinton and S.K. Wong, Phys. Fluids **28** (1985) 3082].

Fokker-Planck theory

Equilibrium equation is a formidable nonlinear PDE

Equilibrium equation

The gyrophase average of the zeroth order ensemble-averaged equation gives the **collisional equilibrium** equation:

$$\int_0^{2\pi} \frac{d\xi}{2\pi} \mathcal{A}_0 = 0 : \qquad \left(\mathbf{V}_0 + v_\parallel' \mathbf{b} \right) \cdot \nabla f_{a0} = C_a(f_{a0})$$

where $\mathbf{v}' = \mathbf{v} - \mathbf{V}_0$ is the velocity in the rotating frame.

Equilibrium distribution function

The exact solution for f_{a0} is a **Maxwellian in the rotating frame**, such that the centrifugal force causes the density to vary on the flux surface:

$$f_{a0} = n_a(\psi, \theta) \left(\frac{m_a}{2\pi T_a} \right)^{3/2} e^{-m_a (v')^2 / 2T_a} \ .$$

GENERAL ATOMICS

Fokker-Planck theory

Equations for neoclassical transport and turbulence at $\mathcal{O}(\rho_*)$

Drift-kinetic equation

Gyroaverage of first-order $\mathcal{A}_1$ gives expressions for **gyroangle-dependent ($\tilde{f}_{a1}$)** and **gyroangle-independent ($\bar{f}_{a1}$)** distributions:

$$\int_0^{2\pi} \frac{d\xi}{2\pi} \mathcal{A}_1 = 0 : \qquad f_{a1} = \tilde{f}_{a1} + \bar{f}_{a1} \; , \quad \tilde{f}_{a1} = \frac{1}{\Omega_a} \int^{\xi} d\xi \, \widetilde{\mathcal{L} f_{a0}}$$

▷ Ensemble-averaged $\bar{f}_{a1}$ is determined by the **drift kinetic equation (NEO)**.

Gyrokinetic equation

Gyroaverage of first-order $\mathcal{F}_1$ gives an expression for **first-order fluctuating distribution ($\hat{f}_{a1}$)** in terms of the distribution of the gyrocenters, $h_a(\mathbf{R})$:

$$\int_0^{2\pi} \frac{d\xi}{2\pi} \mathcal{F}_1 = 0 : \qquad \hat{f}_{a1}(\mathbf{x}) = -\frac{e_a \hat{\phi}(\mathbf{x})}{T_a} + h_a(\mathbf{x} - \rho)$$

▷ Fluctuating $\hat{f}_{a1}$ is determined by the **gyrokinetic equation (GYRO)**.

Drift-Kinetic Equation for Neoclassical Transport

NEO gives complete solution with full kinetic e-i-impurity coupling

$$v_\parallel' \mathbf{b} \cdot \nabla \bar{g}_a - C_a^L(\bar{g}_a) = \frac{f_{a0}}{T_a} \left[-\frac{1}{N_a} \frac{\partial N_a T_a}{\partial \psi} W_{a1} - \frac{\partial T_a}{\partial \psi} W_{a2} + c \frac{\partial^2 \Phi_0}{\partial \psi^2} W_{aV} + \frac{\langle B E_\parallel^A \rangle}{\langle B^2 \rangle^{1/2}} W_{aE} \right]$$

$$\bar{g}_a \doteq \bar{f}_{a1} - f_{a0} \frac{e_a}{T_a} \int^\ell \frac{dl}{B} \left(B E_\parallel - \frac{B^2}{\langle B^2 \rangle} \langle B E_\parallel \rangle \right) \;,$$

$$W_{a1} \doteq \frac{m_a c}{e_a} v_\parallel' \mathbf{b} \cdot \nabla \left(\omega_0 R + \frac{I}{B} v_\parallel' \right) \;,$$

$$W_{a2} \doteq W_{a1} \left(\frac{\varepsilon}{T_a} - \frac{5}{2} \right) \;,$$

$$W_{aV} \doteq \frac{m_a c}{2 e_a} v_\parallel' \mathbf{b} \cdot \nabla \left[m_a \left(\omega_0 R + \frac{I}{B} v_\parallel' \right)^2 + \mu \frac{R^2 B_p^2}{B} \right] \;,$$

$$W_{aE} \doteq \frac{e_a v_\parallel' B}{\langle B \rangle^{1/2}} \;.$$

GENERAL ATOMICS

Gyro-Kinetic Equation for Turbulent Transport

GYRO gives complete solution with full $(\phi, A_\parallel, B_\parallel)$ electromagnetic physics.

$$\frac{\partial h_a(\mathbf{R})}{\partial t} + \left(\mathbf{V}_0 + v_\parallel' \mathbf{b} + \mathbf{v}_{da} - \frac{c}{B}\nabla\hat{\Psi}_a \times \mathbf{b}\right) \cdot \nabla h_a(\mathbf{R}) - C_a^{GL}\left(\hat{f}_{a1}\right)$$

$$= f_{a0}\left[-\frac{\partial \mathrm{ln}(N_a T_a)}{\partial \psi}\hat{W}_{a1} - \frac{\partial \mathrm{ln} T_a}{\partial \psi}\hat{W}_{a2} + \frac{c}{T_a}\frac{\partial^2 \Phi_0}{\partial \psi^2}\hat{W}_{aV} + \frac{1}{T_a}\hat{W}_{aT}\right]$$

$$\hat{W}_{a1}(\mathbf{R}) \doteq -\frac{c}{B}\nabla\hat{\Psi}_a \times \mathbf{b} \cdot \nabla\psi \ ,$$

$$\hat{W}_{a2}(\mathbf{R}) \doteq \hat{W}_{a1}\left(\frac{\varepsilon}{T_a} - \frac{5}{2}\right) \ ,$$

$$\hat{W}_{aV}(\mathbf{R}) \doteq -\frac{m_a R c}{B}\left\langle (\mathbf{V}_0 + \mathbf{v}') \cdot \mathbf{e}_\varphi \nabla\left(\hat{\phi} - \frac{1}{c}(\mathbf{V}_0 + \mathbf{v}') \cdot \hat{\mathbf{A}}\right) \times \mathbf{b} \cdot \nabla\psi\right\rangle_\xi \ ,$$

$$\hat{W}_{aT}(\mathbf{R}) \doteq e_a\left\langle \left(\frac{\partial}{\partial t} + \mathbf{V}_0 \cdot \nabla\right)\left(\hat{\phi} - \frac{1}{c}(\mathbf{V}_0 + \mathbf{v}') \cdot \hat{\mathbf{A}}\right)\right\rangle_\xi \ .$$

$$\hat{\Psi}_a(\mathbf{R}) \doteq \left\langle \hat{\phi}(\mathbf{R} + \boldsymbol{\rho}) - \frac{1}{c}(\mathbf{V}_0 + \mathbf{v}') \cdot \hat{\mathbf{A}}(\mathbf{R} + \boldsymbol{\rho})\right\rangle_\xi$$

$$\to J_0\left(\frac{k_\perp v_\perp'}{\Omega_a}\right)\left(\hat{\phi}(\mathbf{k}_\perp) - \frac{\mathbf{V}_0}{c} \cdot \hat{\mathbf{A}}(\mathbf{k}_\perp) - \frac{v_\parallel'}{c}\hat{A}_\parallel(\mathbf{k}_\perp)\right) + J_1\left(\frac{k_\perp v_\perp'}{\Omega_a}\right)\frac{v_\perp'}{c}\frac{\hat{B}_\parallel(\mathbf{k}_\perp)}{k_\perp} \ .$$

Gyro-Kinetic Equation for Turbulent Transport

GYRO gives complete solution with full $(\phi, A_\parallel, B_\parallel)$ electromagnetic physics.

Must also solve the electromagnetic field equations on the **fluctuation scale**:

$$\frac{1}{\lambda_D^2}\left(\hat{\phi}(\mathbf{x}) - \frac{\mathbf{V}_0}{c}\cdot\hat{\mathbf{A}}\right) = 4\pi\sum_a e_a \int d^3v\,\hat{h}_a(\mathbf{x}-\boldsymbol{\rho}) \ ,$$

$$-\nabla_\perp^2 \hat{A}_\parallel(\mathbf{x}) = \frac{4\pi}{c}\sum_a e_a \int d^3v\,\hat{h}_a(\mathbf{x}-\boldsymbol{\rho})v_\parallel' \ ,$$

$$\nabla \hat{B}_\parallel(\mathbf{x}) \times \mathbf{b} = \frac{4\pi}{c}\sum_a e_a \int d^3v\,\hat{h}_a(\mathbf{x}-\boldsymbol{\rho})\mathbf{v}_\perp' \ .$$

▷ Can one compute equilibrium-scale potential Φ_0 from the Poisson equation?

▷ Practically, no; need higher-order theory and extreme numerical precision.

▷ All codes must take care to avoid **nonphysical potential** at long wavelength

▷ TGYRO gets $\omega_0(\psi) = -c\partial_\psi \Phi_0$ from the **momentum transport equation**.

Transport Equations

Flux-surface-averaged moments of Fokker-Planck equation

$$\left\langle \int d^3v \, \mathcal{A} \right\rangle_\theta$$ **density**

$$\left\langle \int d^3v \, \varepsilon \mathcal{A} \right\rangle_\theta$$ **energy**

$$\sum_a \left\langle \int d^3v \, m_a v'_\varphi \mathcal{A} \right\rangle_\theta$$ **toroidal momentum**

Only terms of order ρ_*^2 survive these averages

$$\boxed{\rho_*^{-1} = 10^3 \quad \rho_*^0 = 1 \quad \rho_*^1 = 10^{-3} \quad \rho_*^2 = 10^{-6}}$$

Transport Equations

Flux-surface-averaged moments of Fokker-Planck equation to $\mathcal{O}(\rho_*^2)$

$$n_a(r): \quad \frac{\partial \langle n_a \rangle}{\partial t} + \frac{1}{V'} \frac{\partial}{\partial r} (V' \Gamma_a) = S_{n,a}$$

$$T_a(r): \quad \frac{3}{2} \frac{\partial \langle n_a T_a \rangle}{\partial t} + \frac{1}{V'} \frac{\partial}{\partial r} (V' Q_a) + \Pi_a \frac{\partial \omega_0}{\partial \psi} = S_{W,a}$$

$$\omega_0(r): \quad \frac{\partial}{\partial t} (\omega_0 \langle R^2 \rangle \sum_a m_a n_a) + \frac{1}{V'} \frac{\partial}{\partial r} (V' \sum_a \Pi_a) = \sum_a S_{\omega,a}$$

$$S_{n,a} = S_{n,a}^{\mathrm{beam}} + S_{n,a}^{\mathrm{wall}} \quad \text{and} \quad \Gamma_a = \Gamma_a^{\mathrm{GV}} + \Gamma_a^{\mathrm{neo}} + \Gamma_a^{\mathrm{tur}}$$

$$S_{W,a} = S_{W,a}^{\mathrm{aux}} + S_{W,a}^{\mathrm{rad}} + S_{W,a}^{\alpha} + S_{W,a}^{\mathrm{tur}} + S_{W,a}^{\mathrm{col}} \quad \text{and} \quad Q_a = Q_a^{\mathrm{GV}} + Q_a^{\mathrm{neo}} + Q_a^{\mathrm{tur}}$$

$$\Pi_a = \Pi_a^{\mathrm{GV}} + \Pi_a^{\mathrm{neo}} + \Pi_a^{\mathrm{tur}}$$

RED: TGYRO   GREEN: NEO   BLUE: GYRO (TGLF)

Acknowledgments

Thanks for input, assistance and labour from

Yang Chen, Univ. Colorado

Stephane Ethier, PPPL

Chris Holland, UCSD

Scott Parker, Univ. Colorado

Weixing Wang, PPPL

GENERAL ATOMICS

Project Description

Overview and Context

The overall objective of this research is to better understand the fundamental physics of transport (collisonal and turbulent) in tokamaks using a theoretical framework that approximates the solution of the 6D Fokker-Planck-Landau equation. An approximate separation of these equations into collisional (neoclassical) and turbulent (gyrokinetic) components forms the basis of the current approach. Most of the computer time is required by the turbulent component via massively parallel gyrokinetic simulations. The goal is to translate this level of understanding into a predictive modeling capability – which includes design, optimization, and interpretation of future experiments and reactors.

Project Description

Overview and Context

In this report, we cover work based on the codes GTS (PPPL) and TGYRO (General Atomics), but in the final tabulation we also include data for GEM (Univ. of Colorado).

Project Description

Scientific Objectives for 2017

What are your projects scientific goals for 2017? Do not limit your answer to the computational aspect of the project.

GTS: Experimental validation for NSTX and DIII-D data is our main objective so that we can predict the transport levels in upcoming NSTX-U (upgrade of the NSTX experiment currently underway) and ITER experiments.

Project Description

Scientific Objectives for 2017

What are your projects scientific goals for 2017? Do not limit your answer to the computational aspect of the project.

Project Description

Scientific Objectives for 2017

TGYRO: Having completed years of validation exercises, and identified regimes where gyrokinetic theory succeeds and fails (so-called L-mode shortfall), we wish to continue development of “sufficiently accurate” models of turbulent transport that can profitably be used in integrated whole-device modeling frameworks. These in turn will be used to design and optimize future experiments. An important point is that agreement isnt perfect, but good enough to provide actual guidance, especially with respect to uncertainty quantification. Specifically, we will continue with detailed validation of gyrokinetic (and reduced gyrofluid) model predictions against experimental observations in US tokamaks (DIII-D, C-Mod, NSTX) to better qualify where current models perform well, and identify parameter regimes where they must be improved. Some attepts to understand aspects of the (near-marginal) core turbulence in ITER are also expected.

Computational Strategies (now and in 2017)

Approach

Give a short, high-level description of your computational problem and your strategies for solving it.

GTS: An important goal right now is the development of a robust algorithm for electromagnetic, finite-beta physics in the GTS code. The more complex field equations greatly increases the time spent in the solver, which is currently implanted with the PETSc library routines, which is currently not multi-threaded. Since the rest of GTS is multi-threaded, we are looking at other possible numerical solvers.

Computational Strategies (now and in 2017)

Approach

Give a short, high-level description of your computational problem and your strategies for solving it.

TGYRO: We have all the tools for solving the full model hierarchy described previously: first-prinicples nonlinear electromagnetic gyrokinetic calcualtions (GYRO), neoclassical transport (NEO), and a reduced quasilinear model (TGLF) of the turbulent fluxes. This suite of capabilities is managed by TGYRO to supply the coupling and feedback for production predictive modeling.

Computational Strategies (now and in 2017)

Codes and Algorithms

Please briefly describe your codes and the algorithms that characterize them.

GTS: The Gyrokinetic Tokamak Simulation code, GTS, is a global, gyrokinetic particle-in-cell application in general toroidal geometry. There is also an associated neoclassical component, GTC-NEO, the solves a time-dependent form of the neoclassical kinetic equations using a δf particle method.

Computational Strategies (now and in 2017)

Codes and Algorithms

Please briefly describe your codes and the algorithms that characterize them.

TGYRO: This application combines GYRO (to solve the nonlinear electromagnetic gyrokinetic equations using Eulerian spectral and finite-difference methods), NEO (to solve for the neoclassical distribution and fluxes/flows with exact linearized Landau collison operator using a spectral expansion scheme in velocity space), and TGLF (as a proxy to GYRO using a fast quaslinear model that approximates the turbulent transport coefficients). TGYRO itself is a transport manager that couples the above modules to give steady-state profile prediction, as a function of input heating power, for existing devices and future devices.

HPC Resources Used Today

Computational Hours

NERSC will enter the hours your project used at NERSC in 2012. If you have significant allocations and usage at other sites please describe them here.

None at the moment

Data and I/O

How much storage space do you typically use today at NERSC for the following three categories?

NOTE: a single TGYRO ensemble run for m1574 created 55GB data. When carrying out many ensemble runs, not all data from each ensemble must be archived (some users care only about transport coefficients). This means that, realistically, less than 1TB of storage will be required per user per year.

GENERAL ATOMICS

HPC Resources Used Today

Data and I/O

Scratch (temporary) space:

5 TB/user (GTS), $<$ 500 GB/user (GYRO)

Permanent (can be shared, NERSC Global Filesystem /project):

20 TB / 4 TB (GTS), $<$ 1 TB/user (GYRO)

HPSS permanent archival storage:

60 TB (GTS), $<$ 1 TB (GYRO)

HPC Resources Used Today

Data and I/O

Please briefly describe your usage of these three types of storage and their importance to your project.

Nothing out of the ordinary: the scratch filesystem is where we run all of our simulations. We share codes and data with the other members of our project through the project space. We archive important simulation data in HPSS.

Between which NERSC systems do you need to share data?

Hopper, Carver, Edison, dtn.

HPC Resources Used Today

Data and I/O

If you have experienced problems with data sharing, please explain.

On the contrary, the project area has greatly facilitated data sharing.

How do your codes perform I/O?

GTS: Some parallel I/O with ADIOS library calls for large data sets and checkpoint-restart, as well as single-core FORTRAN ASCII I/O.

TGYRO: Large restart files are written via MPI-IO, simulation data written in ASCII, option for HDF5 output, but not generally favoured by users.

If you have experienced problems or constraints due to I/O, please explain.

HPC Resources Used Today

Parallelism

How many (conventional) compute cores do you typically use for production runs at NERSC today? What is the maximum number of cores that your codes could use for production runs today?

GTS: Typically use 16,512 cores (2,752 MPI tasks with 6 OpenMP threads/task) for a single run, which allows us to get the most out of our allocation. To code can scale to a larger number of cores but the efficiency of the solver goes down due to the lack of multithreading in PETSc. Weak scaling is more important in order to simulate ITER plasmas.

HPC Resources Used Today

Parallelism

How many (conventional) compute cores do you typically use for production runs at NERSC today? What is the maximum number of cores that your codes could use for production runs today?

TGYRO: Typical TGYRO ensemble run with synthetic diagnostics uses 42,240 cores (11 GYRO instances each using 1280 MPI tasks and 3 OpenMP threads, run for 12 hours, or 506,880 MPP-hours).

HPC Requirements in 2017

TGYRO note: Identifying range of scales needed for modeling burning plasma conditions. A hugely important question: is ITG-scale physics enough, or is multiscale required? This and other key questions to be answered in near term will impact our future HPC requirements. Also, will we want to move towards production level sims and ensemble UQ; how big will each ensemble element need to be? Plausible 2017 predictive workflow: improved version of TGLF or some other model as workhorse, with GYRO “spot checks” at periodic and critical points.

HPC Requirements in 2017

Computational Hours Needed

How many compute hours (Hopper core-hour equivalent) will your project need for CY 2017? Include all hours your project will need to reach the scientific goals you listed in 1.2 above. If you expect to receive significant allocations from sources other than NERSC, please list them here. What is the primary factor driving the need for more hours?

GTS: We expect our project to need between 200M to 300M core-hours during CY 2017 for ITER simulations with finite-beta physics.

TGYRO: Probably about 100M core hours.

HPC Requirements in 2017

Data and I/O

For each of the three categories of data storage described in 3.2, above, list the storage capacity and I/O rates (bandwidth in GB/sec, if known) that you will need in 2017. What are the primary factors driving any storage capacity and/or bandwidth increases?

GTS: ITER data sets will be much larger so the I/O numbers should increase by about an order of magnitude.

TGYRO: We do not expect data needs to increase significantly per user – perhaps by a factor of 2 or 3. However, in the future we envision an increasing number of TGYRO users.

HPC Requirements in 2017

Scientific Achievements with 32X Current Resources

Historically, NERSC computing resources have approximately doubled each year. If your requirements for computing and/or data requirements listed in 4.1 and 4.2 are less than 32X your 2012 usage, briefly describe what you could achieve scientifically by 2017 with access to 2X more resources each year.

HPC Requirements in 2017

Scientific Achievements with 32X Current Resources

One could probably achieve as much by reducing bureaucracy and barriers to resource usage as could be by adding more computational resources. The ease-of-access to NERSC resources is valuable and should be maintained. Moreover, true scientific advances require improvements on 2 fronts: (1) better coupling and workflow for predictive simulations of the type achievable today, (2) an improved theoretical formulation that can treat the pedestal and separatrix correctly – using perhaps the original 6D FPL equation. This is a new area which goes beyond traditional gyrokinetic. There is no accepted approach and no codes to solve this problem. By 2017, we suspect there will be. To reiterate, progress in (2) is a matter of progress in theoretical physics much apart from access to resources.

HPC Requirements in 2017

Parallelism

How many conventional (e.g., Hopper) compute cores will your code(s) be able to use effectively in 2017? (If you expect to be using accelerators or GPUs in 2017, please discuss this in 4.7 below.) Will you need more than one job running concurrently? If yes, how many?

GTS: We have already demonstrated with another gyrokinetic PIC code, GTCP, that we can achieve very high scalability ($>$ 700,000 cores) by adding a radial domain decomposition along with optimized multi-threading and data layout. The same optimizations can be implemented in GTS in order to improve efficiency and scalability.

HPC Requirements in 2017

Parallelism

How many conventional (e.g., Hopper) compute cores will your code(s) be able to use effectively in 2017? (If you expect to be using accelerators or GPUs in 2017, please discuss this in 4.7 below.) Will you need more than one job running concurrently? If yes, how many?

TGYRO: We expect to be using GPUs in 2017, and effort is underway now to develop this capability on titan at ORNL. We believe that the GYRO/TGYRO framework is now maximally parallelized using MPI and OpenMP, with scalability beyond 100K cores if required. By 2017 we expect to also solving more complex 6D equations, and we cannot make projections about that for 2017.

GENERAL ATOMICS

HPC Requirements in 2017

Memory

For future NERSC systems we need your memory requirements. Please describe your memory requirements for 2017 in terms of the minimum amount needed per node and aggregate memory required for you to run.

GTS: 16 Gbytes per node is a reasonable number although more is better. 320 Tbytes of memory in total will be needed for large ITER simulations.

TGYRO: 16 Gb/node is acceptable. Probably less than 100TB aggregate.

HPC Requirements in 2017

Many-Core and/or GPU Architectures

It is expected that in 2017 and beyond systems will contain a significant number of "lightweight" cores and/or hardware accelerators (e.g., GPUs). Are your codes ready for this? If yes, please explain your strategy for exploiting these technologies. If not, what are your plans for dealing with such systems and what do you need from NERSC to help you successfully transition to them?

Our codes use OpenMP multi-threading and we are actively working on a GPU port of key routines to take advantage of this specialized architecture for both GTS and GYRO.

HPC Requirements in 2017

Software Applications and Tools

What HPC software (applications/libraries/tools/compilers/languages) will you need from NERSC in 2017? Make sure to include analytics applications and/or I/O software.

GTS: Unless we find something better we will still use the PETSc library to implement our solvers, as well as a random number generator (SPRNG) and spline routines (PSPLINE). Our I/O is implemented with ADIOS.

TGYRO: netCDF, multithreaded fftw3, BLAS, LAPACK, mumps/superlu, HDF5.

HPC Requirements in 2017

HPC Services

What NERSC services will you require in 2017? Possibilities include consulting or account support, data analytics and visualization, training, collaboration tools, web interfaces, federated authentication services, gateways, etc.

All of the above.

Requirements Summary Worksheet

Table 1: Present and Future Requirements (GTS)

	Used at NERSC in 2012	Used at NERSC in 2017
Computational Hours (Hopper core-hour equivalent)	14.5 Million	300 Million
Scratch storage	5 TB	100TB
Scratch bandwidth		
Shared global storage (project)	4 TB	100TB
Shared global bandwidth		
Archival storage (HPSS)	60 TB	600TB
Archival bandwidth		
Number of cores used for prod. runs	16512	160512
Memory per node	32 GB	32 GB
Aggregate memory	22 TB	220 TB

Requirements Summary Worksheet

Table 2: Present and Future Requirements (TGYRO)

	Used at NERSC in 2012	Used at NERSC in 2017
Computational Hours (Hopper core-hour equivalent)	14.7 Million (m1574)	150 Million
Scratch storage	< 1 TB	10 TB
Scratch bandwidth		
Shared global storage (project)	< 1 TB	10 TB
Shared global bandwidth		
Archival storage (HPSS)	< 1 TB	10TB
Archival bandwidth		
Number of cores used for prod. runs	20K	100K
Memory per node	16 GB	16 GB
Aggregate memory	10 TB	100 TB

Requirements Summary Worksheet

Table 3: Present and Future Requirements (GEM)

	Used at NERSC in 2012	Used at NERSC in 2017
Computational Hours (Hopper core-hour equivalent)	10M	100 Million
Scratch storage		
Scratch bandwidth		
Shared global storage (project)	2 TB	50 TB
Shared global bandwidth		
Archival storage (HPSS)		
Archival bandwidth	0.1-1 Gbit/s	1-10 Gbit/s
Number of cores used for prod. runs	10K	100K
Memory per node	16GB	16GB
Aggregate memory	10 TB	100 TB