GYRO: A 5-D gyrokinetic-Maxwell solver

Mark R. Fahey
Oak Ridge National Laboratory
P.O. Box 2008 MS6008
Oak Ridge, TN 37831-6008
faheymr@ornl.gov

Jeff Candy
General Atomics
P.O. Box 85608
San Diego, CA 92186-5608
candy@fusion.gat.com

ABSTRACT

GYRO solves the 5-dimensional gyrokinetic-Maxwell equations in shaped plasma geometry, using either a local (*flux-tube*) or global radial domain. It has been ported to a variety of modern MPP platforms including a number of commodity clusters, IBM SPs and the Cray X1. We have been able to quickly design and analyze new physics scenarios in record time using the Cray X1: (i) transport barrier studies (Phys. Plasmas **11** (2004) 1879), (ii) the local limit of global simulations (Phys. Plasmas **11** (2004) L25), (iii) kinetic electron and finite-beta generalizations of a community-wide benchmark case, and (iv) impurity transport with application to fuel separation in burning D-T plasmas (to be submitted to Nuclear Fusion). We report on recent physics progress and studies. Further, we discuss GYRO performance across several architectures.

Categories and Subject Descriptors

J.2 [**Computer Applications**]: Physical Sciences and Engineering—*Physics*; D.2.8 [**Software Engineering**]: Metrics—*performance measures*; G.4 [**Mathematics of Computing**]: Mathematical Software—*Parallel and vector implementations,portability*

General Terms

Performance, Theory

Keywords

GYRO, gyrokinetic, Eulerian, turbulence, parallel performance

1. INTRODUCTION

The most promising and aggressively studied concept for power production by fusion reactions is the tokamak. Advances in the understanding and control of tokamak plasmas are continuously being realized, although uncertainties remain in predicting confinement properties and performance of larger reactor-scale devices. The coupled gyrokinetic-Maxwell (GKM) equations provide a foundation for the first-principles calculation of turbulent tokamak transport. For years, the numerical solution of the nonlinear GKM equations has been a computational physics "Grand Challenge".

GYRO is a code that simulates tokamak turbulence by solving the time-dependent, nonlinear gyrokinetic-Maxwell equations for both ions and electrons. It uses a five-dimensional grid and advances the system in time using a second-order, implicit-explicit (IMEX) Runge-Kutta (RK) integrator. GYRO is the only GKM code that has both *global* and *electromagnetic* operational capabilities. Development has been partially funded by the *Plasma Microturbulence Project*, a fusion Scientific Discovery through Advanced Computing (SciDAC) project [1]. GYRO has been ported to a variety of modern MPP platforms including a number of commodity clusters, IBM SPs and the Cray X1. It has shown good scalability on all these platforms; the extent to which this scalability is enhanced by greater per-processor efficiency will be reported. Recently, the GYRO developers have been able to quickly formulate and analyze new physics scenarios in record time using the Cray X1 at Oak Ridge National Laboratory (ORNL).

Section 2 reviews the GKM equations and outlines the numerical methods used for their solution. Section 3 discusses recent physics accomplishments. Section 4 compares the performance of GYRO across several platforms. Section 5 presents conclusions and future work.

2. GYRO OVERVIEW

2.1 Heuristic picture

The GKM equations couple the *gyrocenter* distribution, f, to the electromagnetic fields, Φ:

$$\frac{\partial f}{\partial t} = \mathcal{L}_a f + \mathcal{L}_b \Phi + \{f, \Phi\} \quad \text{where} \quad \mathcal{F}\Phi = \int\int dv_1\, dv_2\, f. \tag{1}$$

$\mathcal{L}_b$, $\mathcal{L}_b$ and $\mathcal{F}$ are linear operators. Strictly speaking, f measures the deviation from a Maxwellian background. For global simulations, a linear adaptive source technique [14] is used to inhibit the evolution of f and Φ on equilibrium scales. Within the standard gyrokinetic ordering, the sole nonlinearity has a Poisson bracket structure. The function $f(\mathbf{r}, v_1, v_2)$ is discretized over a 5-dimensional grid (three spatial and two velocity coordinates), while the 3-dimensional electromagnetic fields $\Phi(\mathbf{r}) = [\phi, A_\parallel]$ are independent of velocity. Here ϕ and $A_\parallel$ are the *electrostatic* and *electromag-*

Proceedings of the ACM/IEEE SC2004 Conference November 6-12, 2004
0-7695-2153-3/04 $20.00 ©2004 IEEE.

netic potentials, respectively. One averages over the fast orbital motion (gyro-orbit) to eliminate the third velocity-space dimension (gyro-angle). However, this so-called *gyro-averaging* operation introduces *nonlocal* spatial operators ($\mathcal{F}$ above, for example) perpendicular to the magnetic field.

The first attempts to solve the GKM equations made use of the particle-in-cell (PIC) method [13, 7] as it provided a relatively straightforward algorithm to numerically evolve a PDE with complicated characteristics (particle orbits). However, PIC codes stagnated for many years in the depth of physics that they were able to treat: only cases without electron dynamics (so-called adiabatic electron) were tractable, and probably not the best way to treat real electrons. In contrast, Eulerian solvers [8, 9] like GYRO have proved their effectiveness for treating the complexities of electromagnetic turbulence.

2.2 Basic equations

First, a fully spectral decomposition of the fluctuating quantities $(f, \phi, A_\parallel)$ is made,

$$\phi = \sum_n \phi_n(r,\theta)\, e^{-\imath n[\varphi - q(r)\theta]} \quad . \tag{2}$$

The integer n labels *toroidal eigenmodes*. Linear studies, still very important for research problems, can be carried out using a single value of n. In Eq. (2), φ is the *toroidal angle*, θ is the *poloidal angle* and q is the *safety factor* (which measures the helicity of the equilibrium magnetic field). Although the physical field, ϕ, is 2π-periodic in θ, the Fourier representation has the implication that the Fourier coefficients, ϕ_n, are nonperiodic, and satisfy the phase condition

$$\phi_n(r,\pi) = e^{-2\pi\imath n q(r)}\phi_n(r,-\pi). \tag{3}$$

Since ϕ is real, the coefficients satisfy $\phi_n^* = \phi_{-n}$.

In the following, we work with a transformation of the normalized gyrocenter distribution function $f_{\rm s}$. Note $f_{{\rm s}n}$ denotes the nth toroidal harmonic of the normalized gyrocenter distribution function $f_{\rm s}$. We use a transformation of $f_{{\rm s}n}$ to eliminate the time derivative of $A_\parallel$ given by $h_{{\rm s}n} = f_{{\rm s}n} + z_{\rm s}\alpha v_{\parallel{\rm s}}(\mathcal{G}A_\parallel)_n$, where $z_{\rm s}$ is the *charge*, $\alpha_{\rm s} \doteq {\rm n_s}/T_{\rm s}$, ${\rm n_s}(r)$ is the *equilibrium density*, $T_{\rm s}$ is the *equilibrium temperature*, and $\mathcal{G}$ is the *gyro-averaging operator*. And, in the this paper, we restrict attention to circular $s-\alpha$ (simple) geometry for brevity. Thus, the spectral form of the gyrokinetic equation, including collisions, in terms of $h_{{\rm s}n}$ is

$$\frac{\partial h_{{\rm s}n}}{\partial t} - \mathcal{C}[h_{{\rm s}n} - z_{\rm s}\alpha_{\rm s} v_{\parallel{\rm s}}(\mathcal{G}A_\parallel)_n] = {\rm RHS}_n(r,\theta;\lambda,\epsilon) \tag{4}$$

$$\begin{aligned} {\rm RHS}_n \;=\; & \left[-\frac{v_{\parallel{\rm s}}(r,\theta)}{R_0 q}\frac{\partial}{\partial\theta} + \imath\omega_{\rm ds}^{(1)} + \imath\omega_{\rm ds}^{(r)}\frac{\partial}{\partial r}\right] \\ & (h_{{\rm s}n} + z_{\rm s}\alpha_{\rm s}(\mathcal{G}{\rm u_s})_n) \\ & -\imath{\rm n_s}\,\omega_{*{\rm s}}\,(\mathcal{G}{\rm u_s})_n + \imath(q/r)\rho_s\,\{\mathcal{G}{\rm u_s}, h\}_n \quad , \end{aligned} \tag{5}$$

where ${\rm u_s} \doteq \phi - v_{\parallel{\rm s}}A_\parallel$, and

$$v_{\parallel{\rm s}}(r,\theta) \doteq \sigma\mu_{\rm s}\sqrt{2\epsilon T_{\rm s}(r)\left(1-\lambda\hat{B}(r,\theta)\right)} \quad . \tag{6}$$

Above, $\rho_s = c_s/\Omega_{ci}$ is the *ion-sound gyroradius*, $c_s = \sqrt{T_e/m_i}$ is the *sound speed*, $\Omega_{ci} = eB/m_i$ is the *ion cyclotron frequency*, R_0 is the *plasma major radius*, $\hat{B}(r,\theta)$ is the *effective magnetic field strength*, and σ is the *sign of the parallel velocity*. The two velocity coordinates are (ϵ,λ). It is important to note that $v_{\parallel{\rm s}}$ depends on the mass through $\mu_{\rm s} = \sqrt{m_i/m_{\rm s}}$ (the equations are normalized to the primary ion mass m_i). For electrons, $\mu_e \simeq 60$, making the parallel term stiff and difficult to treat numerically. The bracket $\{\cdot,\cdot\}$ describes nonlinear $\mathbf{E}\times\mathbf{B}$ and magnetic flutter dynamics. The *curvature drift* coefficients, $\omega_{\rm d}^{(1)}$ and $\omega_{\rm d}^{(r)}$, as well as the *diamagnetic frequency* ω_*, are given by

$$\omega_{*{\rm s}} \;\doteq\; k_\theta\rho_s\left[\frac{a}{L_{n{\rm s}}} + (\epsilon - 3/2)\frac{a}{L_{T{\rm s}}}\right] \quad , \tag{7}$$

$$\begin{aligned} \omega_{\rm ds}^{(1)} \;\doteq\; & k_\theta\rho_s\,\frac{2T_{\rm s}}{z_{\rm s}\hat{B}R_0}\,\epsilon\left(1-\frac{\lambda\hat{B}}{2}\right) \\ & [\cos\theta + (s\theta - \alpha_{\rm MHD}\sin\theta)\sin\theta] \quad , \end{aligned} \tag{8}$$

$$\omega_{\rm ds}^{(r)} \;\doteq\; -i\rho_s\,\frac{2T_{\rm s}}{z_{\rm s}\hat{B}R_0}\,\epsilon\left(1-\frac{\lambda\hat{B}}{2}\right)\sin\theta \quad . \tag{9}$$

Above, $k_\theta \doteq nq/r$ is the *poloidal wavenumber*, s is the *magnetic shear* and $\alpha_{\rm MHD}$ is the *MHD pressure gradient*. The source of free energy to drive instabilities is embodied in the ω_* term, where $L_{n{\rm s}}(r)/a$ and $L_{T{\rm s}}(r)/a$ are the normalized density and temperature gradient scale lengths. For typical parameters, turbulent fluxes tend to peak at $k_\theta\rho_s \sim 0.2$.

The Maxwell equations are written in terms of the $h_{{\rm s}n}$ as

$$\sum_{\rm s=1}^{n_{\rm ion}} \alpha_{\rm s} z_{\rm s}^2\,(1-\mathcal{R}_{\rm s})\,\phi_n = \sum_{\rm s=1}^{n_{\rm ion}+1} z_{\rm s}{\rm V}[\mathcal{G}h_{{\rm s}n}] \quad , \tag{10}$$

$$-\frac{2\rho_{\rm s}^2}{\beta_e}\nabla_\perp^2 A_{\parallel\,n} + \sum_{\rm s=1}^{n_{\rm ion}+1}\alpha_{\rm s} z_{\rm s}^2\,{\rm V}[v_{\parallel{\rm s}}^2 A_{\parallel\,n}] = \sum_{\rm s=1}^{n_{\rm ion}+1} z_{\rm s}{\rm V}[v_{\parallel{\rm s}}\mathcal{G}h_{{\rm s}n}] \quad , \tag{11}$$

where $n_{\rm ion}$ is the number of ion species (electrons are denoted by ${\rm s} = n_{\rm ion}+1$), and $\mathcal{R}_{\rm s} = {\rm V}[\mathcal{G}\mathcal{G}]$ is a velocity-space-integrated double gyroaverage. We avoid the disastrous *Ampère cancellation problem* by rewriting the ${\rm V}[v_{\parallel{\rm s}}^2 A_{\parallel\,n}]$ term in an equivalent form that lends itself to a more robust numerical calculation. To do electrostatic simulations, one normally sets $A_\parallel = 0$, which is consistent with the limit $\beta_e \to 0$.

The object V above is the 2-dimensional *velocity-integration operator*, defined as

$$\begin{aligned} {\rm V}[f_{\rm s}] \;\doteq\; & \sum_\sigma \frac{1}{2\sqrt{\pi}}\int_0^\infty d\epsilon\, e^{-\epsilon}\sqrt{\epsilon} \\ & \int_0^1 \frac{d(\lambda\hat{B})}{\sqrt{1-\lambda\hat{B}}}\, f_{\rm s}(r,\theta,\varphi;\epsilon,\lambda,\sigma;t) \quad . \end{aligned} \tag{12}$$

For concreteness, we note that ${\rm V}[1] = 1$, ${\rm V}[\epsilon] = 3/2$ and ${\rm V}[v_\parallel] = 0$.

Finally, we remark that nonlinear simulations normally reach a statistical steady state on a timescale much shorter ($100\times$) than an energy confinement time. To be precise, for typical tokamak parameters, we normally run simulations out to $t \simeq 1000(a/c_s)$.

2.3 Discretization schemes

We briefly sketch the type of discretization scheme used in each dimension. A detailed treatment is beyond the scope of the present paper.

- **toroidal angle:** fully spectral [see Eq. (2)].
- **radius:** linear advective derivatives on f are treated with upwind differences, whereas derivatives on fields are treated with centered differences. The gyrooperators $\mathcal{G}$ and $\mathcal{R}$ are approximated using a (banded) pseudo-spectral technique. The order of all discretizations is adjustable at run-time.
- **poloidal angle:** for f, there is no fixed grid in θ. Instead, the transformation

$$\frac{v_{\|\mathrm{s}}(r,\lambda,\theta)}{R_0 q(r)}\frac{\partial}{\partial\theta} \rightarrow \Omega(r,\lambda)\frac{\partial}{\partial\tau} \qquad (13)$$

  is used to eliminate the singularity at bounce points, $v_{\|\mathrm{s}} = 0$. Then, an upwind scheme in τ is used to discretize $\partial f/\partial\tau$. Centered differences only are used on τ-derivatives of fields, otherwise numerical instability will result. The use of a τ-grid (leading to a different set of points in θ for every value of λ) for the GK equation dictates that the Maxwell equations are solved by expansion of fields in complex finite-elements: $\phi(r_i,\theta) = \sum_m c_m^i F_m^i(\theta)$. The F_m^i satisfy the phase condition in Eq. (3).
- **velocity-space:** A transformation property of Eq. (12) under integration over θ is used to recast the velocity-space integration. Then, in both ϵ and λ, an exact Gauss-Legendre quadrature scheme is numerically generated (by nonlinear root-finding) at run-time. This is different at each radius and for different plasma equilibria.
- **nonlinearity:** The nonlinear Poisson bracket is evaluated with a *conservative* difference-spectral analogue of the Arakawa method. This scheme ensures exact conservation of density and generalized entropy at vanishing time step (independent of grid resolution).
- **collisions:** Collisions are represented by a second-order diffusive-type operator in λ. This operator is split from the collisionless problem and an irregular-grid generalization of the Crank-Nicholson method is used.
- **time-advance:** A 2nd-order IMEX RK scheme is used, with the electron parallel motion ($\partial/\partial\theta$) in Eq. (6) treated implicitly. This is exceptionally complicated due to the use of a τ-grid, as well as the presence of the field quantity $\mathrm{u_s}$ (see Eq. 6) in the advection. However, the implicitness is crucial for the elimination of a numerical instability connected with pathological *electrostatic Alfvén waves*. Indeed, our experience shows that the numerical difficulties associated with kinetic electrons become rapidly more severe for radial boxes larger than about $80\rho_s$.

2.4 Comparison with other GK codes

Table 1 shows a brief comparison of gyrokinetic codes. In Table 1, we have the following abbreviations: UM=Univ. of Maryland, GA=General Atomics, PPPL=Princeton, UCI=Univ. of California at Irvine, UCol=Univ. of Colorado at Boulder, LLNL=Lawrence Livermore National Lab, PIC=particle-in-cell, Euler=Eulerian, Loc=Local, Glb=Global, and Elc=Electromagnetic, and Shp=Shape.

Table 1: Gyrokinetic codes comparison.

Code	Lab	Type	Loc	Glb	Elc	Shp
GS2	UM	Euler	•		•	•
GYRO	GA	Euler	•	•	•	•
GTC	PPPL/UCI	PIC		•		
TUBE	UCol	PIC	•		•	
PG3EQ	LLNL	PIC	•			

Throughout the development of GYRO, various benchmarking efforts have been undertaken. Linear comparisons with GS2, and nonlinear comparisons with PG3EQ, are reported in [3]. More recently, global comparisons with GTC have been reported in [4]. In that work, GYRO results are shown to agree with those from GS2 and PG3EQ in the limit where the latter flux-tube simulations are valid. The results of Ref. [4] are discussed in more detail in the next section.

3. RECENT PHYSICS ACCOMPLISHMENTS

Development of the GYRO core solver began in 1999, and all physics capabilities in the original design specification were realized in 2002. Refinements to the numerical schemes were added after that and concluded by early 2003. For the last year, the developmental focus has been on usability and support for a growing user base. Currently, there are four users at GA, two at PPPL and one at the Univ. of Texas.

The following subsections detail some of the more important findings made in the last year:

1. **Comparison with DIII-D L-mode ρ_* experiments:** An exhaustive series of global, full-physics GYRO simulations of DIII-D L-mode ρ_*-similarity discharges was made. In these famous experiments, only a single dimensionless parameter was varied: $\rho_* \doteq \rho_s/a$. The GYRO calculations matched experimental results for electron and ion energy transport [2] within experimental error bounds. The Bohm-scaled diffusivity of the experiments was also reproduced, for which inclusion of electron collisions together with equilibrium sheared rotation was found to be crucial. These were the most physically comprehensive tokamak turbulence simulations ever undertaken.
2. **Evaluation of minimum-q theory of transport barrier formation:** It was shown that a minimum-q surface (where $s = 0$) in a tokamak plasma does not act as the catalyst for ion transport barrier formation [5]. Although there are theories of this process in the literature which argue that a barrier should form near $s = 0$, the reality of this process had never been convincingly demonstrated with simulation. Using the X1, it was clearly shown that transport is smooth across an $s = 0$ surface due to the appearance of *gap modes*.
3. **The local limit of global GK simulations:** A controversial transport scaling study by Lin *et al.* [10] substantially overestimated the *Cyclone base case* [6]

benchmark value as ever-larger global simulations (at successively smaller ρ_*) were run. This contradicted the *local hypothesis*, which states that global and flux-tube simulations should agree at sufficiently small ρ_*. Lin repeated the same study with minor revisions in 2002 [12] and yet again in 2004 [11] obtaining an ion diffusivity, χ_i, still 36% higher than the Cyclone value. Lin's global scenario was revisited with GYRO [4], with the finding that at small ρ_*, χ_i closely matches the Cyclone value. Further, it was shown that for these large-system-size simulations, there is a very long transient period for which χ_i exceeds the statistical average. By running simulations for twice as long as the Lin simulations (900 a/c_s versus 400 a/c_s), a true statistical turbulent steady state was recovered.

4. **Particle and impurity transport:** The first systematic gyrokinetic study of particle transport was made, including impurity transport and isotope effects. We found that in a burning D-T plasma, the tritium is better confined than deuterium, with the implication that the D-T fuel will separate as tritium is retained. This conclusion was found to be independent of temperature gradient and electron collision frequency.

4. GYRO PERFORMANCE

In this section we describe the performance of GYRO on five separate platforms. As mentioned previously, GYRO has been ported to a variety of modern MPP platforms including a number of commodity clusters, IBM SPs, an SGI Altix, and a Cray X1. Since code portability and flexibility are considered crucial, only a single source is maintained. Thus the ports to new architectures often involve nothing more than creating a new makefile. This is mentioned now because, although the single-source philosophy has made GYRO a very portable code, the single-source requirement has consequence of preventing some machine-specific code optimizations which adversely affect performance on other architectures. However, this requirement does not preclude certain machine-specific optimizations. For example, GYRO has had some vectorization for the X1. The changes were minimal and in some cases simultaneously beneficial on other architectures. These changes were mostly the addition of directives, but there were selected instances of rank promotion/demotion, and an instance of "pushing" a loop down into a subroutine call.

We now compare the performance of GYRO on the following five platforms:

- **AMD Athlon cluster**
  The AMD Athlon cluster at Princeton Plasma Physics Laboratory (PPPL) has 64 2-way Athlon MP2000+ (1.667 GHz) nodes with gigabit Ethernet interconnect. Each processor has a peak performance of $1.667 \times 2 = 3.334$ GFlops/s. GYRO was compiled with Lahey/FujitsuFortran 95 L6.10c.
- **IBM Nighthawk II (Power3) cluster**
  The IBM Nighthawk II cluster with Colony interconnect at the National Energy Research Scientific Computing Center (NERSC) has 416 16-way Power3 (375 MHz) nodes. The peak performance of a Power3 processor is $375 \times 4 = 1.5$ GFlops/s. GYRO was compiled using ESSL 3.3 and XL Fortran 8.1. The machine has parallel environment 3.2 and AIX 5.1.
- **IBM p690 (Power4) cluster**
  The IBM p690 cluster with Federation interconnect[1] at ORNL has 27 32-way Power4 (1.3 GHz) nodes. The peak rate of each processor is $1.3 \times 4 = 5.2$ GFlops/s. GYRO was built using ESSL 4.1 and XL Fortran 8.1. The machine has parallel environment 4.1 and AIX 5.2.
- **SGI Altix**
  The SGI Altix at ORNL is a single-system image running Linux with 256 Itanium 2 processors (1.5 GHz). The peak rate of each processor is $1.5 \times 4 = 6$ GFlops/s. GYRO was built using the Intel Fortran 8.0 compiler and SGI's SCSL library.
- **Cray X1**
  The Cray X1 at ORNL is also a single-system image of 504 multistreaming processors. Each processor has a peak rate of 12.8 GFlops/s. GYRO was built using Programming Environment 5.2.0.2 and MPT 2.3.0.4. The OS level was UNICOS/mp 2.4.17. For our benchmarks, the Multistreaming processor (MSP) is considered the basic processing unit rather than the Singlestreaming processor (SSP). Note that an MSP is composed of four SSPs.

The current performance is described by comparing timings across architectures, as shown in the next two subsections.

4.1 Waltz standard case benchmark

This test case is 500 timesteps of the *Waltz Standard Case* [3] with kinetic electrons and electron collisions, but no electromagnetic effects. This is a flux-tube 16-toroidal mode electrostatic (electrons and ions, 1 field) case on a $16 \times 140 \times 8 \times 8 \times 20 \times 2$ grid. The numerical grid resolution used for this benchmark is the same as that used in recent production runs. It represents, roughly, the *minimum grid size* required to obtain physically accurate results. We could have improved the scalability on all platforms by moving to a finer grid, but felt that results for a grid of typical production size are of greater practical value .

Figures 1–3 show the results of running the Waltz standard case benchmark on the various platforms mentioned above. Figure 1 shows the seconds per timestep – the ultimate measure of how well this code does on any machine. Figure 2 shows the average MPI time per timestep. Figure 3 shows the average collision time per timestep. We chose the MPI communication time and the collision time in addition to the entire run-time because they demonstrate interesting performance characteristics across the test platforms.

A few observations can be made from Figures 1–3. First, Figure 1 shows for this benchmark that the X1 is four times as fast as the nearest competitor. Second, Figure 2 clearly shows that the communication time on the X1 is an order of magnitude less than the next best. This is due to the fact that GYRO consumes a lot of bandwidth, and the X1 is designed to provide more bandwidth than the other platforms. Collisions on the X1 were several times slower than

[1] Striping across Federation adapters currently inhibits MPI performance. IBM recommends not striping until a later microcode release, thus striping is not currently enabled on this machine.

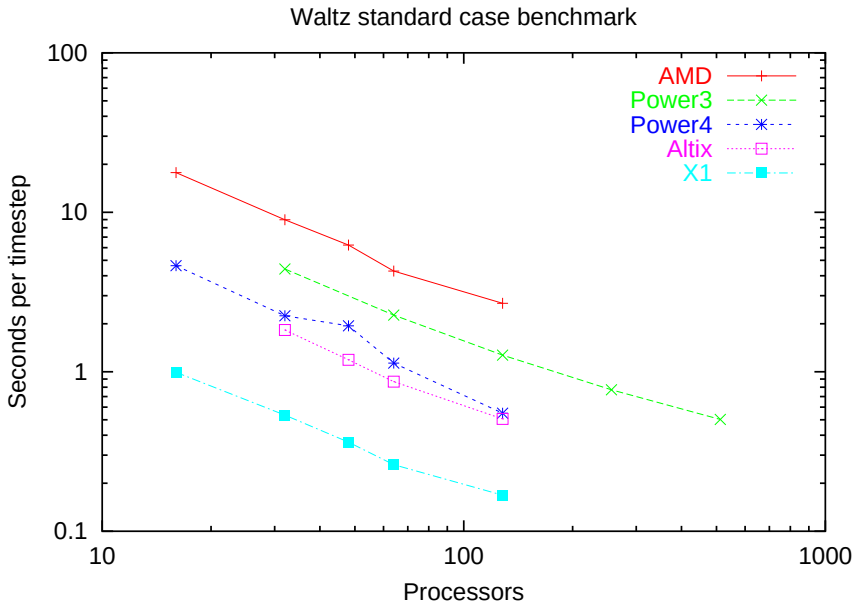

Figure 1: Seconds per timestep on the Waltz standard case benchmark.

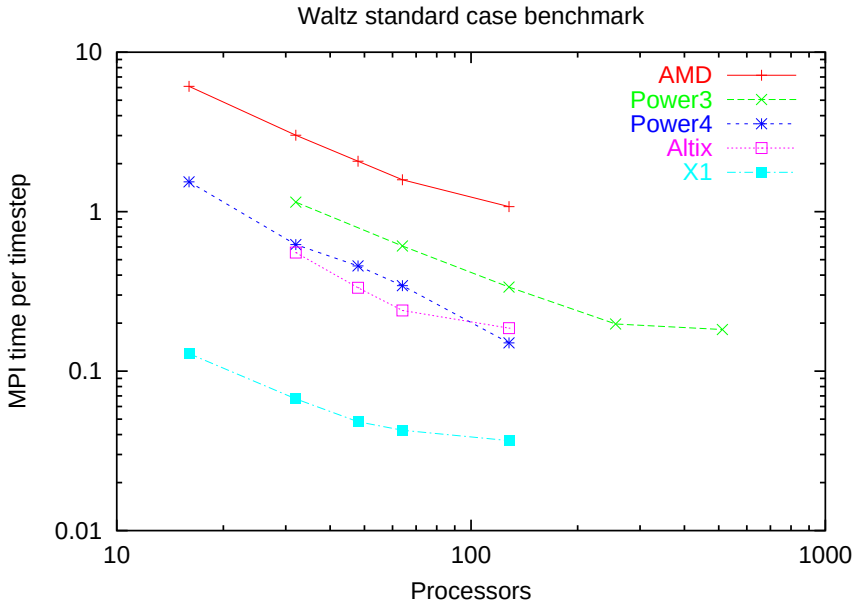

Figure 2: MPI time per timestep on the Waltz standard case benchmark.

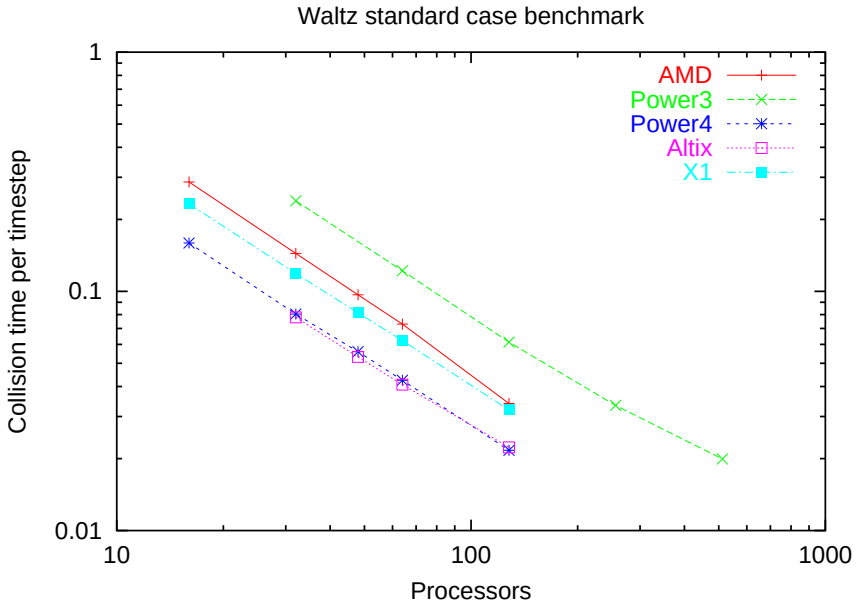

Figure 3: Collision time per timestep on the Waltz standard case benchmark.

competing platforms until recent optimizations dramatically improved its performance. Figure 3 shows that the collision calculations are slightly slower on the X1 than the Altix and Power 4 platforms. However, in terms of seconds per timestep, the X1 is still a factor of 4 better than the next best machine. We also see from Figures 1 and 2 that the IBM Power4 cluster with Federation interconnect compares favorably with the other machines and that it scales fairly well even with its current handicap (see footnote 1.)

4.2 Exploratory plasma edge simulation

This test case is a gyrokinetic simulation meant to explore the parameter space characteristic of the plasma edge. Magnetic shear and safety factor are very high at the outer plasma boundary, and equilibrium gradients are steep, making the simulations more challenging than for core parameters. This is a global 28-toroidal-mode electrostatic case with adiabatic electrons on a $200 \times 10 \times 10 \times 28 \times 28 \times 1$ grid.

A visualization of the results from a 504 MSP simulation on the Cray X1 is shown in Figure 4. This plots contours of the turbulent electrostatic potential mapped back onto an (elongated) torus. Note the extreme elongation of the turbulent eddies in the toroidal direction.

Table 2 compares the performance of running this exploratory simulation on several machines at relatively large processor counts (in some cases as large as the machine allowed). The "time/step" column shows the average time per step in seconds for 400 time steps of this simulation, and the "MPI/step" column shows the average time per step spent in MPI. This was a collisionless simulation, therefore collision time is not reported.

This study shows, using the inverse of column 3 in Table 2, that the Cray X1 at 504 MSPs can take almost 16 steps per second. The SGI Altix does 1.8 steps per second with 224 processors, the IBM Power4 cluster about 1.8 steps per second with 448 processors and the IBM Power3 machine

Table 2: Exploratory plasma edge simulation timings in milliseconds.

Machine	procs	time/step	MPI/step
IBM Power3 cluster	896	602.45	103.69
	1344	544.58	81.44
	1792	405.19	67.53
	2240	431.48	73.19
	2688	422.91	66.39
IBM Power4 cluster	224	994.94	209.60
	448	561.21	133.90
SGI Altix	224	554.12	188.40
Cray X1	224	121.98	15.03
	448	72.25	6.94
	504	62.35	6.07

reaches a maximum rate of 2.5 steps per second at 1792 PEs and beyond. Based on the peak speeds of each machine and the data in Table 2, it is clear that the X1's vector processors deliver greater per-processor efficiency, which leads to greater scalability.

5. CONCLUSIONS

The GYRO development effort has reached all its original capability objectives, although refinements are continuously being implemented. Currently, developmental focus is on usability and support for a user base that is growing. Through the use of GYRO several important research results have been obtained over the last year. These include detailed comparisons with DIII-D L-mode ρ_* experiments, a critical assessment of the minimum-q theory of transport barrier formation, a demonstration that local and global simulations are confluent in the $\rho_* \rightarrow 0$ limit, and a systematic study of particle and impurity transport.

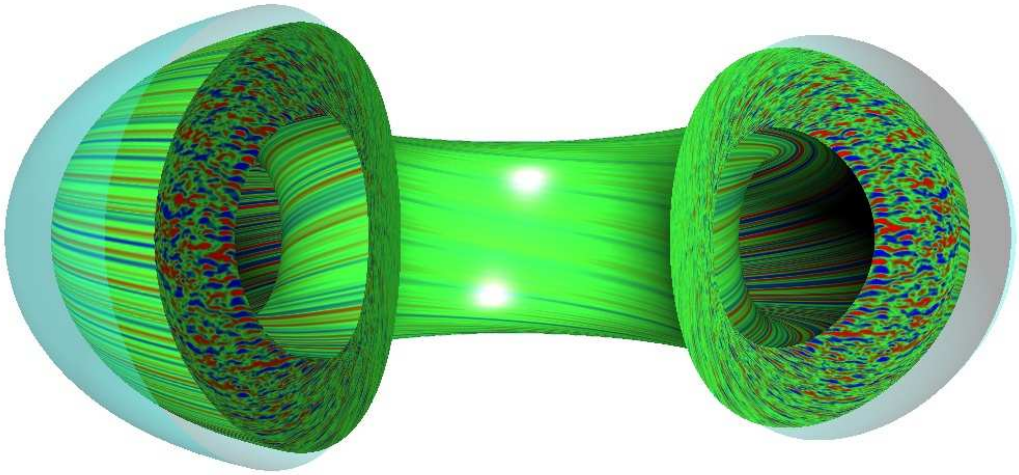

Figure 4: Turbulent potential fluctuations in shaped toroidal geometry for plasma edge simulation.

The performance of GYRO on nonvector systems is constrained by communication bandwidth, which is not true on the X1. The X1 has exceptionally high-bandwidth, low-latency communication hardware, so communication times for various expensive ALL TO ALL operations dropped to the lowest levels ever for GYRO. Moreover, the powerful vector processing units each give a performance that is at least a factor of 4 greater than the closest microprocessor-based supercomputers.

Note that the collision step is somewhat slower on the X1 than on competing architectures. This is true despite the several optimizations which were made to improve vectorization to its current state. We believe the collision step as currently implemented has probably been optimized as much as possible; and therefore the underlying algorithm must be changed to effect significant improvements. In addition, we plan to review alternatives to the current sparse solver, to improve the nonlinear step by evaluating the Poisson bracket fully in real space, and to implement a new parallel algorithm for the field solves, which currently replicates work.

The X1 is an impressive supercomputer platform on which new physics scenarios have been designed and analyzed in record time. The fast communication coupled with the intrinsically lower processor count define a system which, for a given production problem size, can in many cases provide a time-to-solution which is smaller than a traditional supercomputer of any size.

6. ACKNOWLEDGMENTS

This research was sponsored by the U.S. Department of Energy under Grant No. DE-FG03-95ER54309. It was also sponsored by the Office of Mathematical, Information, and Computational Sciences, Office of Science, U.S. Department of Energy under Contract No. DE-AC05-00OR22725 with UT-Batelle, LLC. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

7. REFERENCES

[1] Scientific Discovery through Advanced Computing. http://www.osti.gov/scidac/.

[2] J. Candy and R. Waltz. Anomalous transport in the DIII-D tokamak matched by supercomputer simulation. *Phys. Rev. Lett.*, 91:045001–1, 2003.

[3] J. Candy and R. Waltz. An Eulerian gyrokinetic-Maxwell solver. *J. Comput. Phys.*, 186:545, 2003.

[4] J. Candy, R. Waltz, and W. Dorland. The local limit of global gyrokinetic simulations. *Phys. Plasmas*, 11:L25, 2004.

[5] J. Candy, R. Waltz, and M. Rosenbluth. Smooothness of turbulent transport across a minimum-q surface. *Phys. Plasmas*, 11:1879, 2004.

[6] A. Dimits, G. Bateman, M. Beer, B. Cohen, W. Dorland, G. Hammett, C. Kim, J. Kinsey, M. Kotschenreuther, A. Kritz, L. Lao, J. Mandrekas, W. Nevins, S. Parker, A. Redd, D. Shumaker, R. Sydora, and J. Weiland. Comparisons and physics basis of tokamak transport models and turbulence simulations. *Phys. Plasmas*, 7:969, 2000.

[7] A. Dimits, T. Williams, J. Byers, and B. Cohen. Scalings of ion-temperature-gradient-driven anomalous transport in tokamaks. *Phys. Rev. Lett.*, 77:71, 1996.

[8] W. Dorland, F. Jenko, M. Kotschenreuther, and B. Rogers. Electron temperature gradient turbulence. *Phys. Rev. Lett.*, 85:5579, 2000.

[9] F. Jenko, W. Dorland, M. Kotschenreuther, and B. Rogers. Electron temperature gradient driven turbulence. *Phys. Plasmas*, 7:1904, 2000.

[10] Z. Lin, S. Ethier, T. Hahm, and W. Tang. Size scaling of turbulent transport in magnetically confined plasmas. *Phys. Rev. Lett.*, 88:195004, 2002.

[11] Z. Lin and T. Hahm. Turbulence spreading and transport scaling in global gyrokinetic particle simulations. *Phys. Plasmas*, 11:1099, 2004.

[12] Z. Lin, T. Hahm, S. Ethier, W. Lee, J. Lewandowski, G. Rewoldt, W. Tang, W. Wang, L. Chen, and P. Diamond. Size scaling of turbulent transport in

tokamak plasmas. In *Plasma Physics and Controlled Nuclear Fusion Research (Proc. 19th International Conference on Fusion Energy, Lyon, 2002)*. IAEA, Vienna, 2003.

[13] R. Sydora, V. Decyk, and J. Dawson. Fluctuation-induced heat transport results from a large global 3d toroidal particle simulation model. *Plasma Phys. Controlled Fusion*, 38:A281, 1996.

[14] R. Waltz, J. Candy, and M. Rosenbluth. Gyrokinetic turbulence simulation of profile shear stabilization and broken gyrobohm scaling. *Phys. Plasmas*, 9:1938, 2002.