GYRO: Analyzing new physics in record time

M. Fahey and J. Candy
ORNL, Oak Ridge, TN
General Atomics, San Diego, CA

20 May 2004

Cray User Group
Knoxville, TN

CUG 2004

Acknowledgment

- Research was 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.

- Oak Ridge National Laboratory is managed by UT-Battelle, LLC for the United States Department of Energy under Contract No. DE-AC05-00OR22725.

Outline

- GYRO
- Test platforms
- Performance results
  - GTC.n64.500a
  - Waltz standard case benchmark
  - Exploratory Plasma Edge simulation
- Physics Results
- Recent and Future work
- Conclusions

GYRO

- is an Eulerian gyrokinetic-Maxwell (GKM) solver developed by Jeff Candy and Ron Waltz at General Atomics
- computes the turbulent radial transport of particles and energy in tokamak plasmas
- uses a 5-D grid and advances the system in time using a second-order, implicit-explicit Runga-Kutta integrator
- is the only GKM code worldwide that has both *global* and *electromagnetic* operational capabilities
- is partially funded by the DOE SciDAC Plasma Microturbulence Project
- has been ported to a wide variety of machines including commodity clusters

GYRO on the X1 - history

- Port (mid '03) required no source-code changes
- Functional tests did identify a few bugs in GYRO
- First set of X1-related optimizations accepted back into GYRO release in late '03
  - 14 routines modified ($< 10\%$)
  - Mostly directives added
  - Pushed 1 loop down into subroutine call
  - Few instances of rank promotion/demotion
  - A few optimizations rejected

Platforms

Cray X1 at ORNL

- 256 Multistreaming Processors
- 1024 GB total memoory
- 3.2 GF/s peak performance

Other platforms

- **AMD cluster at PPPL (Princeton):** 48 2-way Athlon MP2000+ (1.667 GHz) with gigE interconnect
- **IBM p690 cluster at ORNL:** 27 32-way p690 SMP nodes (1.3 GHz Power4) and the Federation Switch[a]
- **IBM Nighthawk II cluster at NERSC:** 416 16-way SMP nodes (375 MHz Power3) and SP2 Switch
- **SGI Altix at ORNL:** 256-way single-system image with a NUMAflex fat-tree interconnect

[a]Striping does not work properly for adapters with 2 links. So the current settings are to use only 1 communication paths for the network protocol, i.e. no striping.

GYRO performance

Three real problems, problem size fixed in each case (strong scaling)

- GTC.n64.500a
  - 64-toroidal-mode adiabatic, 64x400x8x8x20x1 grid
  - extremely high resolution
  - electron physics ignored allowing large timestep
- Waltz Standard Case Benchmark (WSCk)
  - 16-toroidal-mode electrostatic, 16x140x8x8x20x2 grid
  - domain is relatively small
  - electromagnetics off, electron collisions on
- Exploratory Plasma Edge
  - prototype simulation, new for the parameter regime it addresses
  - 28 modes

Caveat

Note that because of

- Sporadic benchmarking on evolving system software and hardware configurations
- Continued evolution of OS and compilers and libraries
- Evolution of GYRO

performance results are transient and performance characteristics are slightly changing over time.

GYRO performance - GTC.n64.500

Comparing overall performance

- X1 is faster
  - about 4× faster than Altix
  - about 7× faster than IBM Power4

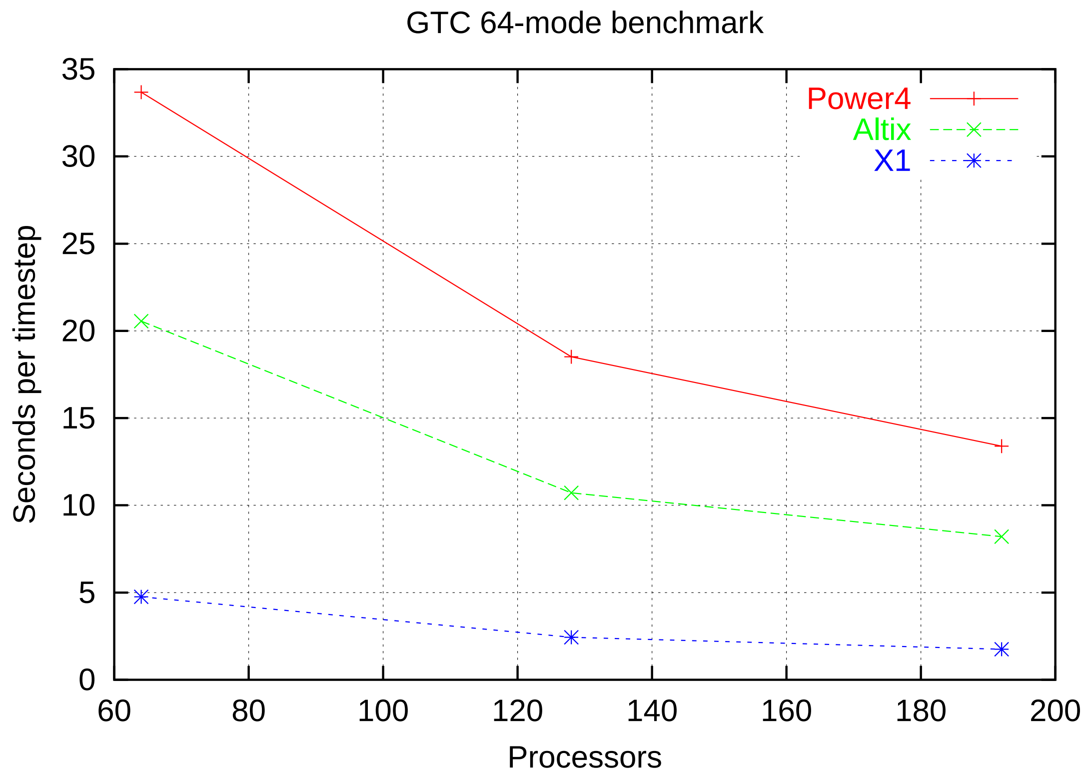

GYRO performance - GTC.n64.500 (cont.)

Comparing communication time

- IBM and SGI performance is limited by communication overhead
- X1 communication ratio is at least 5× better

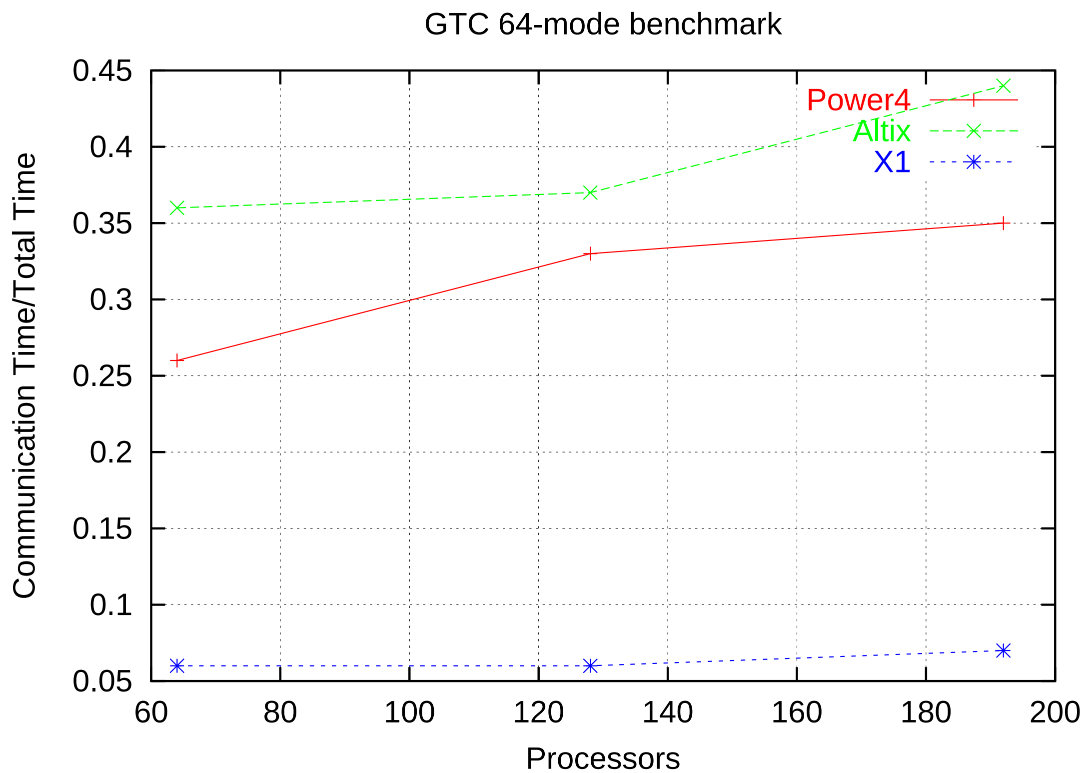

GYRO performance - Waltz standard case

- X1 (only) 2× as fast
- Why?

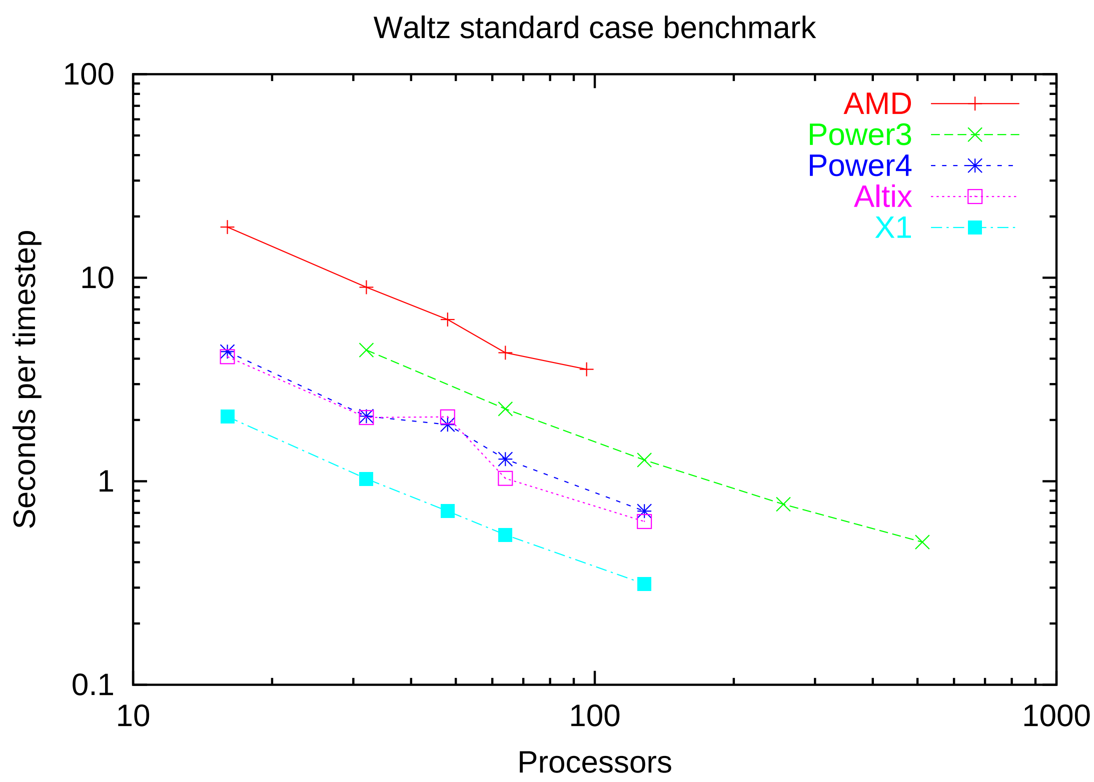

GYRO performance - Waltz standard case (cont.)

- X1 provides much more bandwidth
- Again, why?

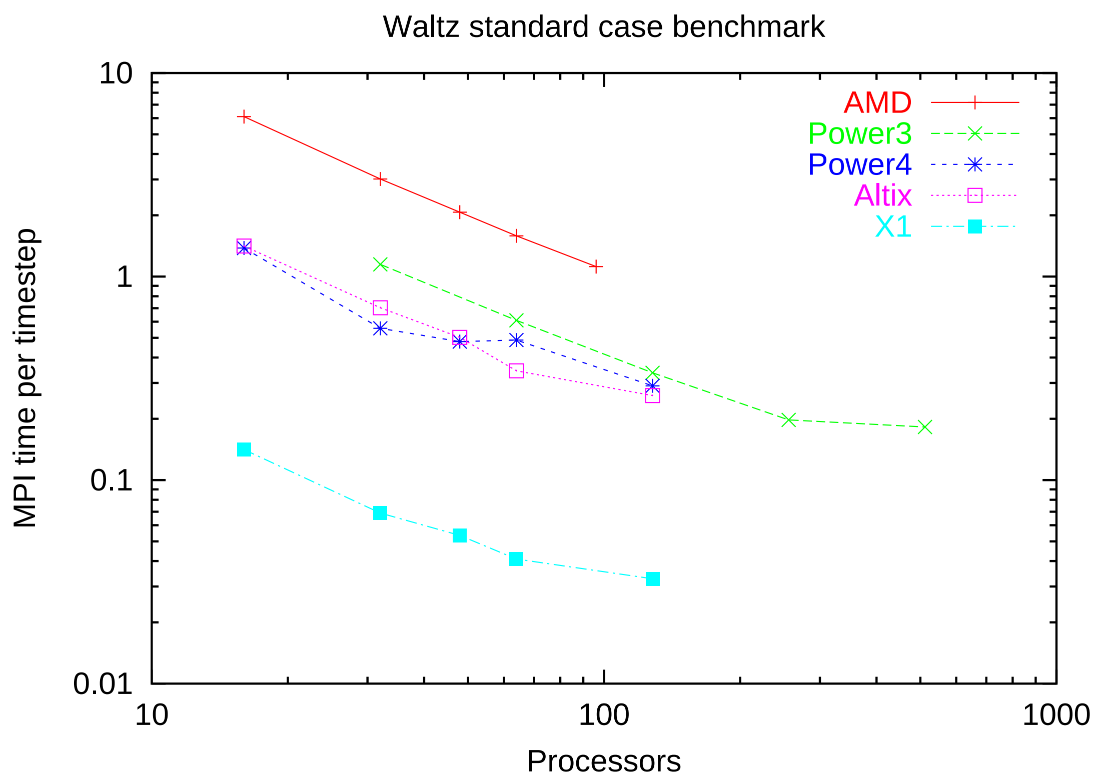

GYRO performance - Waltz standard case (cont.)

timings for the **collision step**

- X1 is several times slower than the other architectures
- Q: why is the X1 slower?
  A: the collision routine has a significant amount of scalar operations
- If collisions ignored, then X1 is at least 5× faster

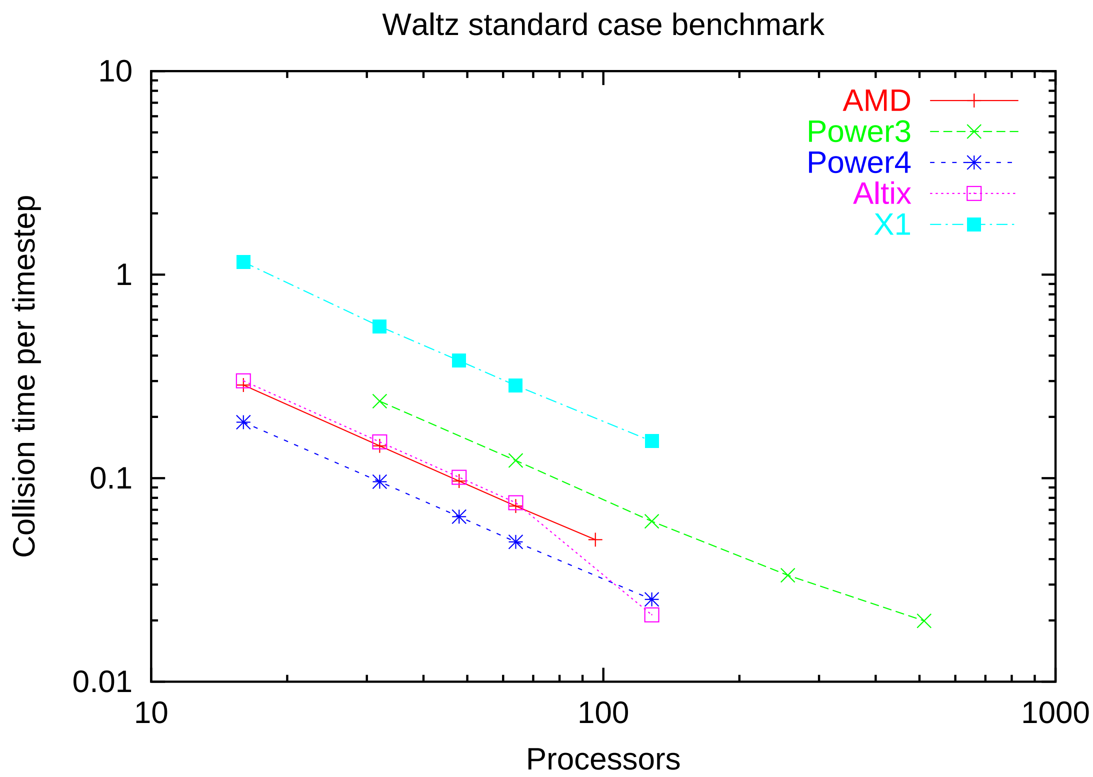

GYRO performance - Exploratory Plasma Edge

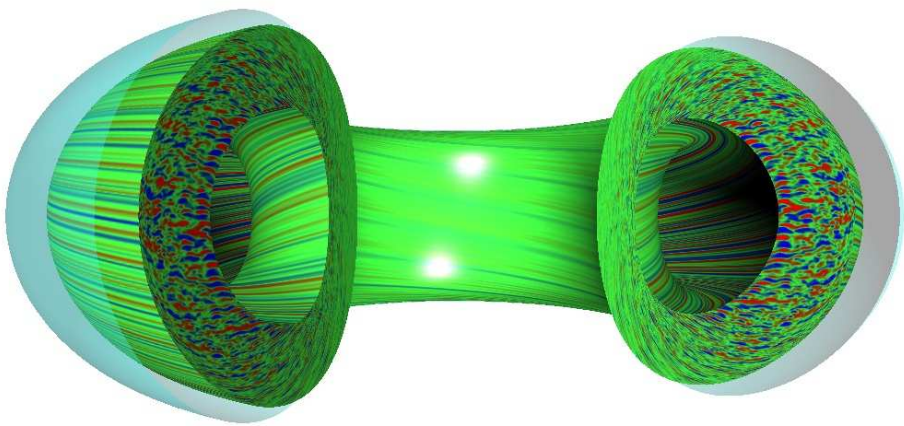

GYRO performance - Exploratory Plasma Edge (cont.)

Machine	processors	time(s)/step	MPI-time(s)/step
IBM Power3 cluster	896	0.602450	0.103694
	1344	0.544581	0.081436
	1792	0.405187	0.067532
	2240	0.431481	0.073186
	2688	0.422913	0.066386
Cray X1	504 MSP	0.072615	0.005889

Using the inverse of column two:

- The X1 can do 13.8 steps per second (maybe more with more MSPs)
- The IBM Power3 can do at best 2.5 steps per second

GYRO accomplishments on the X1

- **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
  - calculations matched experimental results for electron and ion energy transport [1] within experimental error bounds
  - Bohm-scaled diffusivity of the experiments was also reproduced
  - the most physically comprehensive tokamak turbulence simulations ever undertaken
- **Evaluation of minimum-q theory of transport barrier formation:**
  - shown that a minimum-q surface (where $s = 0$) in a tokamak plasma does not act as the catalyst for ion transport barrier formation [3]
  - it was clearly shown that transport is smooth across an $s = 0$ surface due to the appearance of *gap modes*

- **Resolving the local limit of global GK simulations:**
  - an existing transport scaling study [5] overestimated the *Cyclone base case* [4] benchmark value
  - contradicts the *local hypothesis* which states that global and flux-tube simulations should agree at sufficiently small ρ_*
  - GYRO found an ion diffusivity χ_i that closely agrees with the Cyclone value at small ρ_* [2]
  - GYRO further showed for these large-system-size simulations, there is a very long transient period for which χ_i exceeds the statistical average
- **Particle and impurity transport:**
  - first systematic gyrokinetic study of particle transport, including impurity transport and isotope effects
  - 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
  - found to be independent of temperature gradient and electron collision frequency

GYRO recent issues

In Dec '03, results were found to agree to only 9 decimal digits compared to the IBM and AMD clusters

- just after the setup phase; which machine was (more) right?

Primary contributor was found to be *catastrophic cancellation* in two routines

- $f = \sqrt{(1 - x)}$ where $x \approx 1$
- implemented exceptional cases; if $x \approx 1$ then $f = 0$
- improved agreement between all architectures
- accuracy loss was roughly equivalent to adding a stochastic source term with amplitude 1e-9
- Can be shown to make little difference in "time-averaged" turbulent diffusivity
- thus previous results were valid, and now GYRO more robust

GYRO recent issues (cont.)

Optimize the collision step

- inlined LAPACK tridiagonal solve and eliminated pivoting
- vectorized across tridiagonal solves
  - ignoring matrix setup and assuming each solve of the same order, then a 20x speedup could be attained
  - BUT matrix order not uniform and matrix setup not negligible
  - final result: 40% speedup
- matrix setup in collision routine now the largest cost
- have to rewrite routine to vectorize better
  - recent attempts look promising
  - a test (last week) showed 5x speedup on X1 and slightly faster on Power3

GYRO future work

1. Continue optimizations to collision step
2. Fully parallelize field solves, rather than replicate work
3. Improve the nonlinear step by evaluating the transformation of the toroidal angle in real space, will involve FFTs
4. Possibly replace sparse solver

Conclusions

- X1 has provided a platform where new physics scenarios have been quickly designed and analyzed just in the last year
- the performance of GYRO on nonvector machines is constrained by communication bandwidth, not true on X1
- For collisionless scenarios, the X1 provides performance many times faster than other modern machines, up to 20× on the exploratory edge simulation
- Collisions perform poorly on the X1, and are being evaluated as to how it can be optimized for the X1 without negatively affecting other platforms

Acknowledgment

Wish to thank Pat Worley for the benchmark data he provided on the GTC problem.

References

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

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

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

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

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