GYRO Performance on MPP Systems

J. Candy
General Atomics, San Diego, CA

M.R. Fahey
Oak Ridge National Laboratory, Oak Ridge, TN

Cray User Group 2005 Meeting
Petroglyphs to Petaflops
Albuquerque, NM
16–19 May 2005

GENERAL ATOMICS

The Most Promising Concept for Power Production by Fusion Reactions is the Tokamak

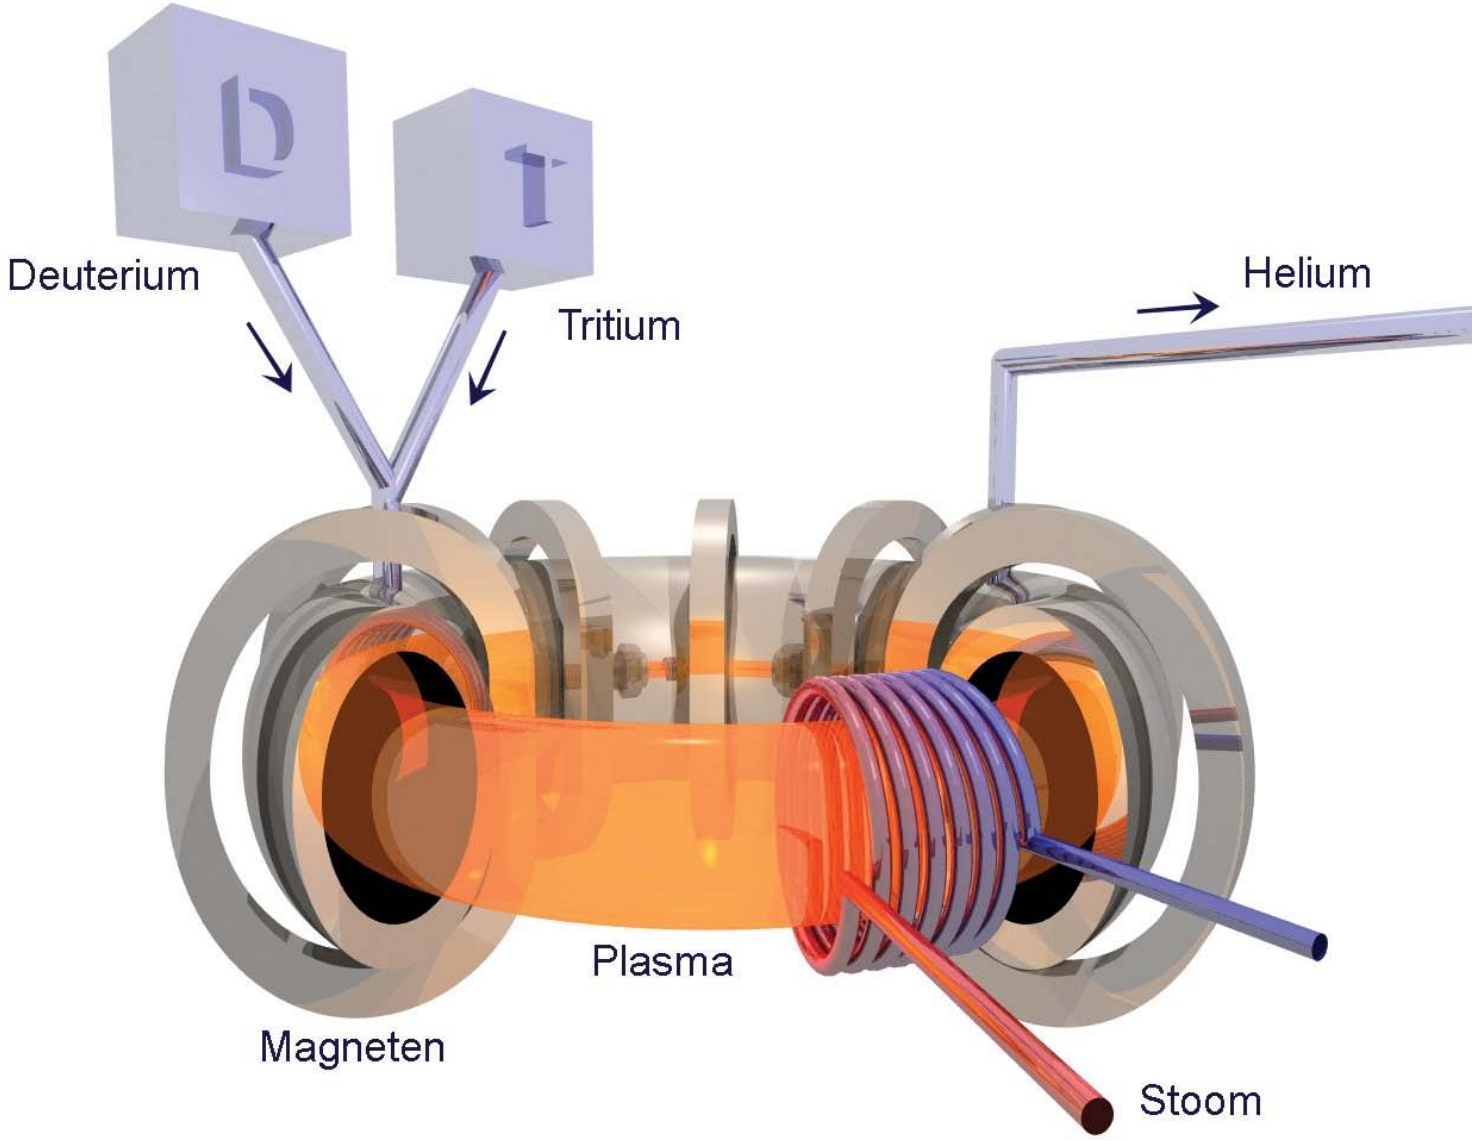

GENERAL ATOMICS

The Construction of a Tokamak Burning Plasma Facility is a Prudent Scientific and Humanitarian Undertaking

GENERAL ATOMICS

The Gyrokinetic-Maxwell Equations Provide the Foundation for Direct Numerical Simulation of Plasma Turbulence

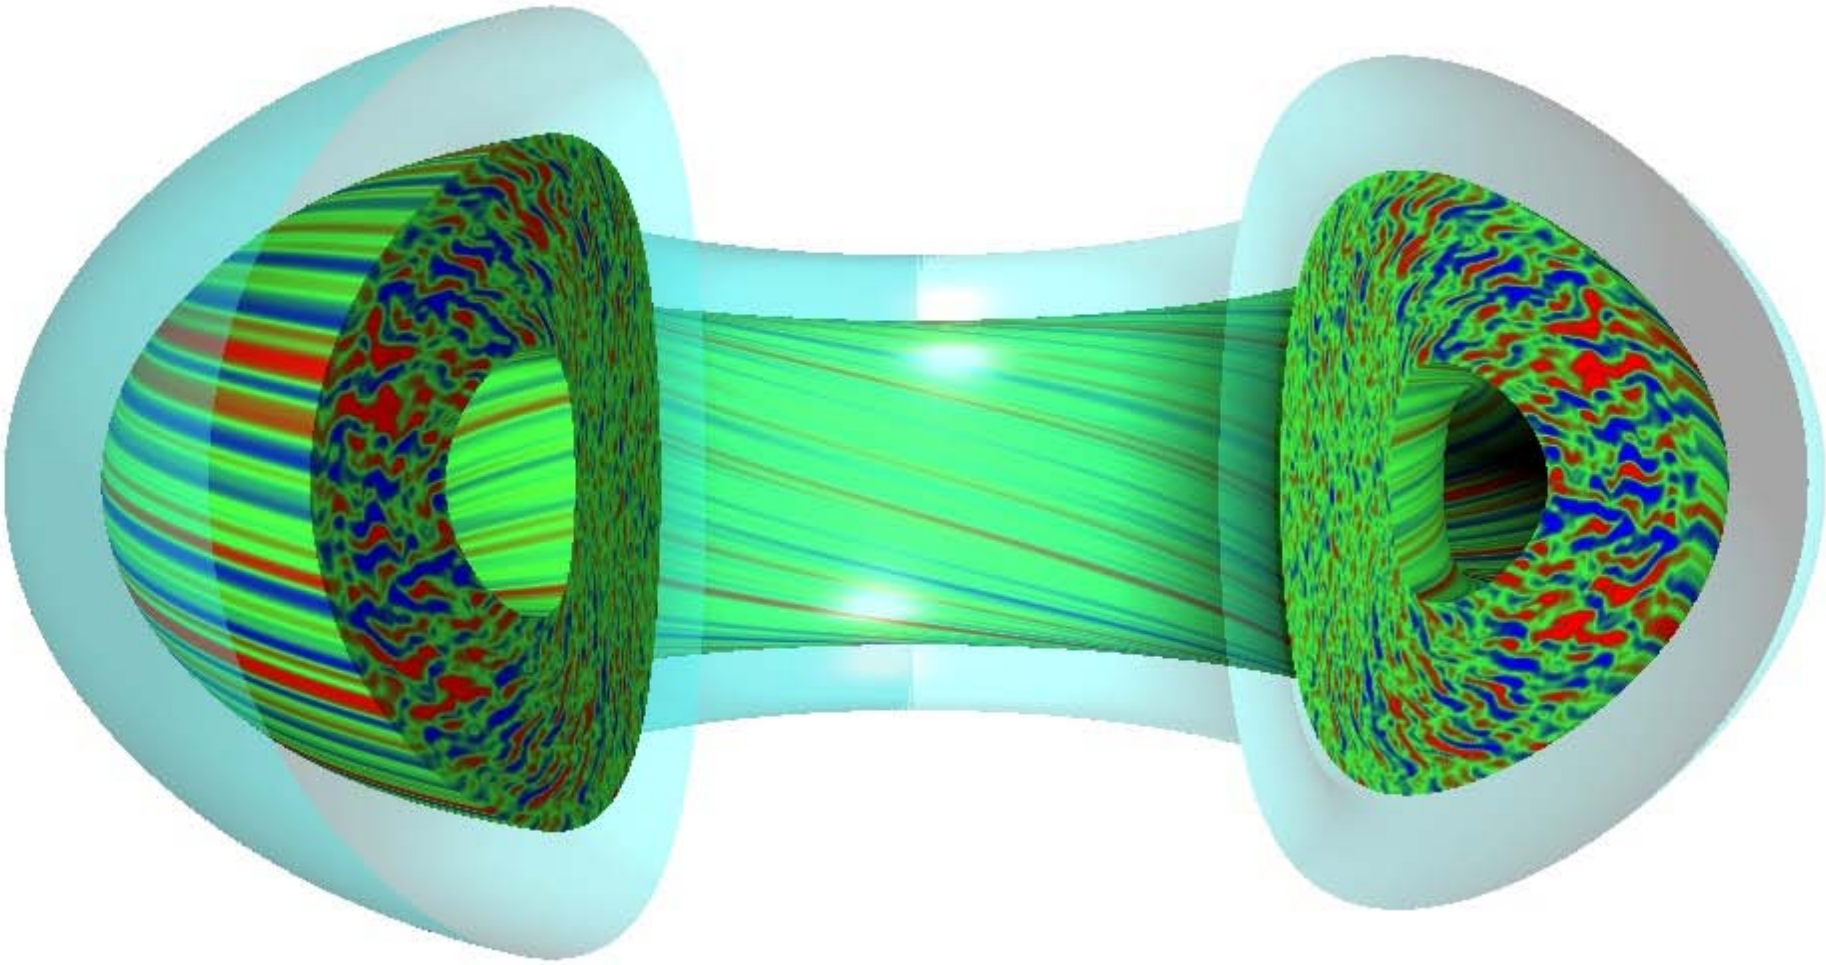

The Gyrokinetic Equations Replace the Older, Simpler Gyrofluid Model (below)

$$\frac{dn}{dt}+[\tfrac{1}{2}\hat{\nabla}_\perp^2\mathbf{v}_\Psi]\cdot\nabla T_\perp+B\nabla_\parallel\frac{u_\parallel}{B}-\left(1+\frac{\eta_\perp}{2}\hat{\nabla}_\perp^2\right)i\omega_*\Psi+(2+\tfrac{1}{2}\hat{\nabla}_\perp^2)i\omega_d\Psi+i\omega_d(p_\parallel+p_\perp)=0,$$

$$\frac{du_\parallel}{dt}+[\tfrac{1}{2}\hat{\nabla}_\perp^2\mathbf{v}_\Psi]\cdot\nabla q_\perp+B\nabla_\parallel\frac{p_\parallel}{B}+\nabla_\parallel\Psi+\left(p_\perp+\frac{1}{2}\hat{\nabla}_\perp^2\Psi\right)\nabla_\parallel\ln B+i\omega_d(q_\parallel+q_\perp+4u_\parallel)=0,$$

$$\frac{dp_\parallel}{dt}+[\tfrac{1}{2}\hat{\nabla}_\perp^2\mathbf{v}_\Psi]\cdot\nabla T_\perp+B\nabla_\parallel\frac{q_\parallel+3u_\parallel}{B}+2(q_\perp+u_\parallel)\nabla_\parallel\ln B-\left(1+\eta_\parallel+\frac{\eta_\perp}{2}\hat{\nabla}_\perp^2\right)i\omega_*\Psi+\left(4+\frac{1}{2}\hat{\nabla}_\perp^2\right)i\omega_d\Psi$$
$$+i\omega_d(7p_\parallel+p_\perp-4n)+2|\omega_d|(\nu_1T_\parallel+\nu_2T_\perp)=-\frac{2}{3}\nu_{ii}(p_\parallel-p_\perp),$$

$$\frac{dp_\perp}{dt}+[\tfrac{1}{2}\hat{\nabla}_\perp^2\mathbf{v}_\Psi]\cdot\nabla p_\perp+[\hat{\bar{\nabla}}_\perp^2\mathbf{v}_\Psi]\cdot\nabla T_\perp+B^2\nabla_\parallel\frac{q_\perp+u_\parallel}{B^2}-\left[1+\frac{1}{2}\hat{\nabla}_\perp^2+\eta_\perp\left(1+\frac{1}{2}\hat{\nabla}_\perp^2+\hat{\bar{\nabla}}_\perp^2\right)\right]i\omega_*\Psi$$
$$+\left(3+\frac{3}{2}\hat{\nabla}_\perp^2+\hat{\bar{\nabla}}_\perp^2\right)i\omega_d\Psi+i\omega_d(5p_\perp+p_\parallel-3n)+2|\omega_d|(\nu_3T_\parallel+\nu_4T_\perp)=\frac{1}{3}\nu_{ii}(p_\parallel-p_\perp),$$

$$\frac{dq_\parallel}{dt}+(3+\beta_\parallel)\nabla_\parallel T_\parallel+\sqrt{2}D_\parallel|k_\parallel|q_\parallel+i\omega_d(-3q_\parallel-3q_\perp+6u_\parallel)+|\omega_d|(\nu_5u_\parallel+\nu_6q_\parallel+\nu_7q_\perp)=-\nu_{ii}q_\parallel,$$

$$\frac{dq_\perp}{dt}+[\tfrac{1}{2}\hat{\nabla}_\perp^2\mathbf{v}_\Psi]\cdot\nabla u_\parallel+[\hat{\bar{\nabla}}_\perp^2\mathbf{v}_\Psi]\cdot\nabla q_\perp+\nabla_\parallel\left(T_\perp+\frac{1}{2}\hat{\nabla}_\perp^2\Psi\right)+\sqrt{2}D_\perp|k_\parallel|q_\perp+\left(p_\perp-p_\parallel+\hat{\bar{\nabla}}_\perp^2\Psi-\frac{1}{2}\hat{\nabla}_\perp^2\Psi\right)\nabla_\parallel\ln B$$
$$+i\omega_d(-q_\parallel-q_\perp+u_\parallel)+|\omega_d|(\nu_8u_\parallel+\nu_9q_\parallel+\nu_{10}q_\perp)=-\nu_{ii}q_\perp.$$

Gyrokinetic Equations Look Deceptively Simple

$$\frac{\partial f}{\partial t} = \mathcal{L}_a f + \mathcal{L}_b \langle \Phi \rangle + \{ f, \langle \Phi \rangle \}$$

$$\mathcal{F}\Phi = \int \int dv_1 \, dv_2 \, \langle f \rangle$$

- f is the **gyrocenter** distribution (measures the deviation from a Maxwellian), and $\Phi(\mathbf{r}) = [\phi, A_{\|}]$ are EM fields.
- $\mathcal{L}_b$, $\mathcal{L}_b$ and $\mathcal{F}$ are linear operators
- $\langle \cdot \rangle$ is a **gyroaveraging operator**
- The function $f(\mathbf{r}, v_1, v_2)$ is discretized over a 5-dimensional grid

Eulerian Schemes Solve the GKM Equations on a Fixed Grid

$$f(r, \tau, n_{\mathrm{tor}}, \lambda, E) \longrightarrow f(i, j, n, k, e)$$

$$i = 1, 2, \ldots, N_i$$

$$j = 1, 2, \ldots, N_j$$

$$n = 1, 2, \ldots, N_n$$

$$k = 1, 2, \ldots, N_k$$

$$e = 1, 2, \ldots, N_e$$

BASE DISTRIBUTION: $f([n], \{e, k\}, i, j)$ (1)

Distribution requirements for different code stages (i.e., evaluation of different operators)

- The distribution of an index across processors is incompatible with the evaluation of operators on that index
- For example, a derivative in r requires that all i should be on a processor

Stage	On-processor indices
Linear with field solve	i, j
Pitch-angle scattering	j, k
Energy diffusion	e
Nonlinear	i, n

Base Distribution: Velocity-Space over Columns and Toroidal Modes over Rows.

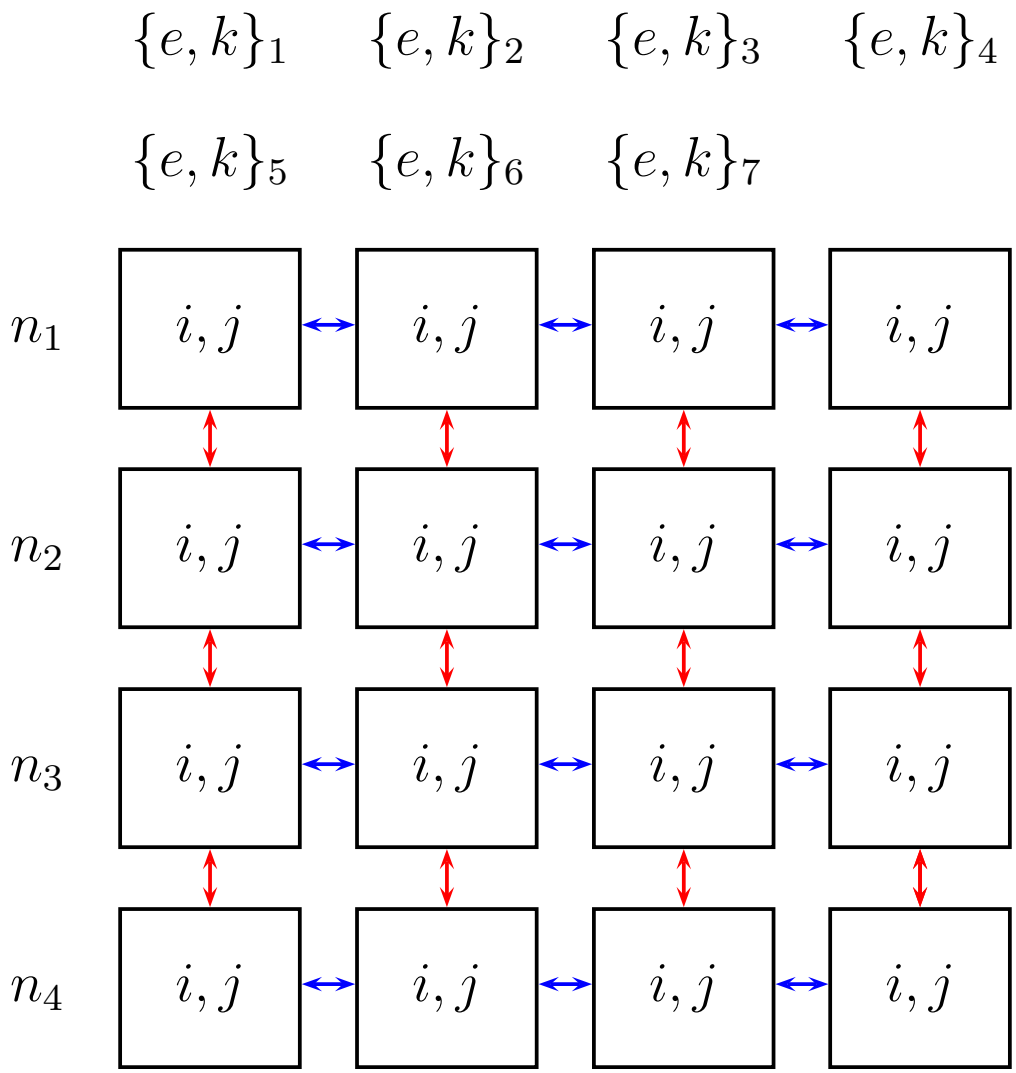

3-index Row Transpose: Symbolic Notation

We can define a generalized 3-index transpose operator, R, which acts individually on processor rows

$$R : \{e, k\}, i \longrightarrow \{i, e\}, k$$

The omitted index, j, is left on-processor. Because there are three indices, three applications of the operator R yields the identity:

$$\begin{aligned} R^3 : \{e, k\}, i &= R^2 : \{i, e\}, k \\ &= R : \{k, i\}, e \\ &= \{e, k\}, i \end{aligned}$$

3-index Row Transpose: Schematic Description

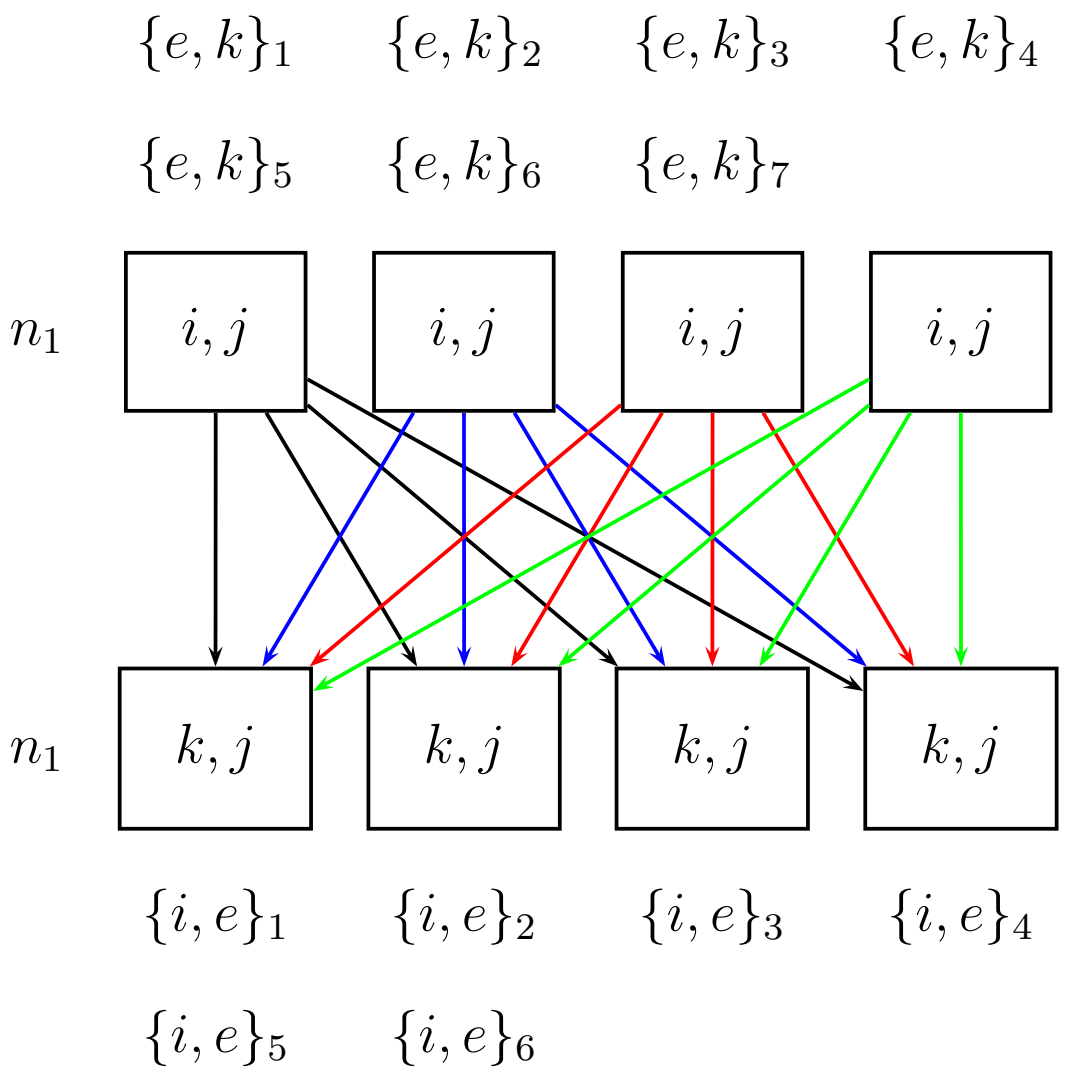

2-index Column Transpose: Schematic Description

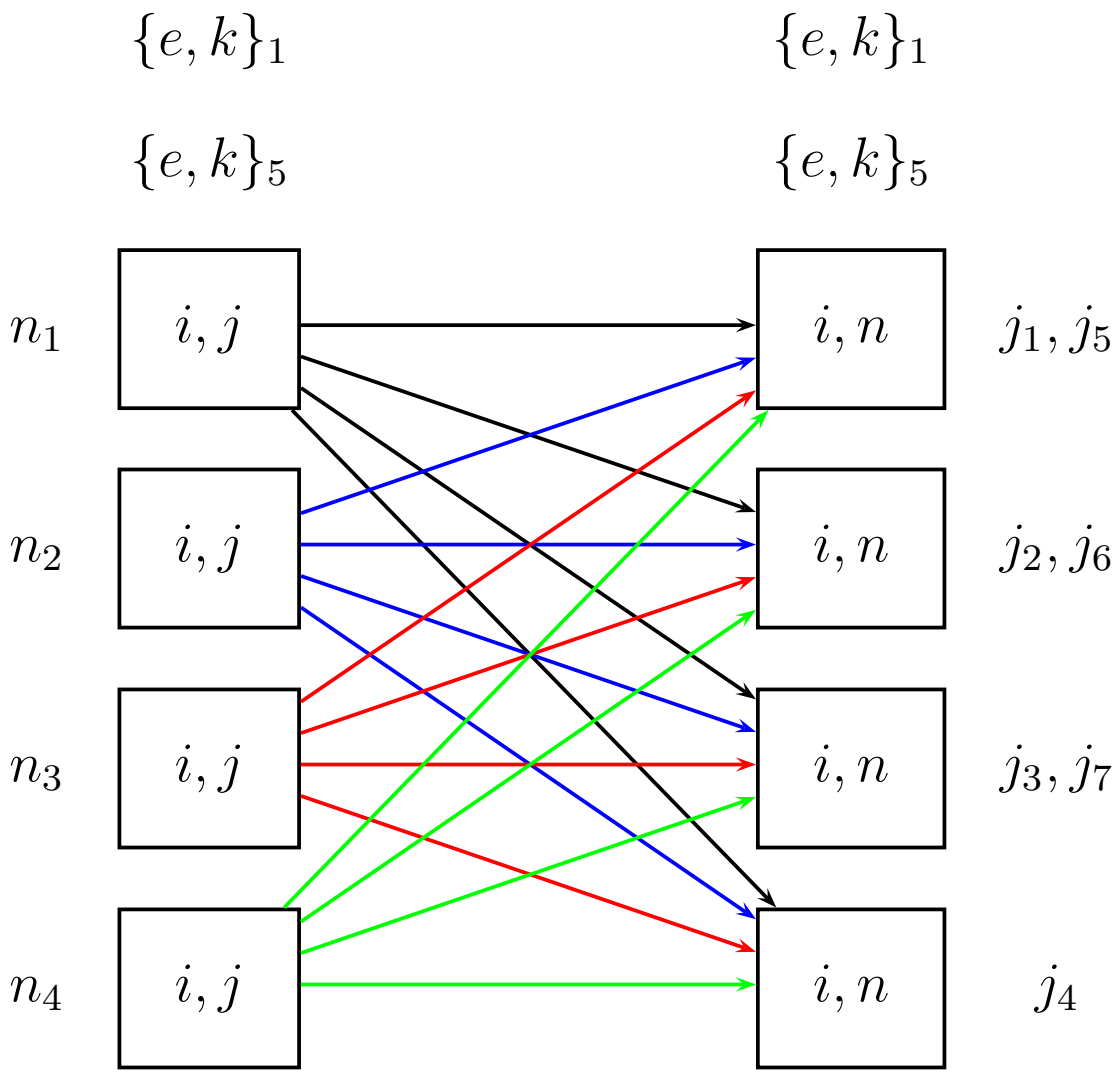

Distributions of Indices in Each Stage

Stage	Distribution
Linear terms with field solve	$f\left([n],\{e,k\},i,j\right)$
Pitch-angle scattering	$f\left([n],\{i,e\},k,j\right)$
Energy diffusion	$f\left([n],\{k,i\},e,j\right)$
Nonlinear	$f\left([j],\{e,k\},i,n\right)$

Typical sizes:

$$N_n = 16 \quad N_e = 8 \quad N_k = 8 \quad N_i = 128 \quad N_j = 28$$

GYRO: Overall Performance Comparison on 5 MPP Systems using the Waltz Standard Case Parameters (B1-std)

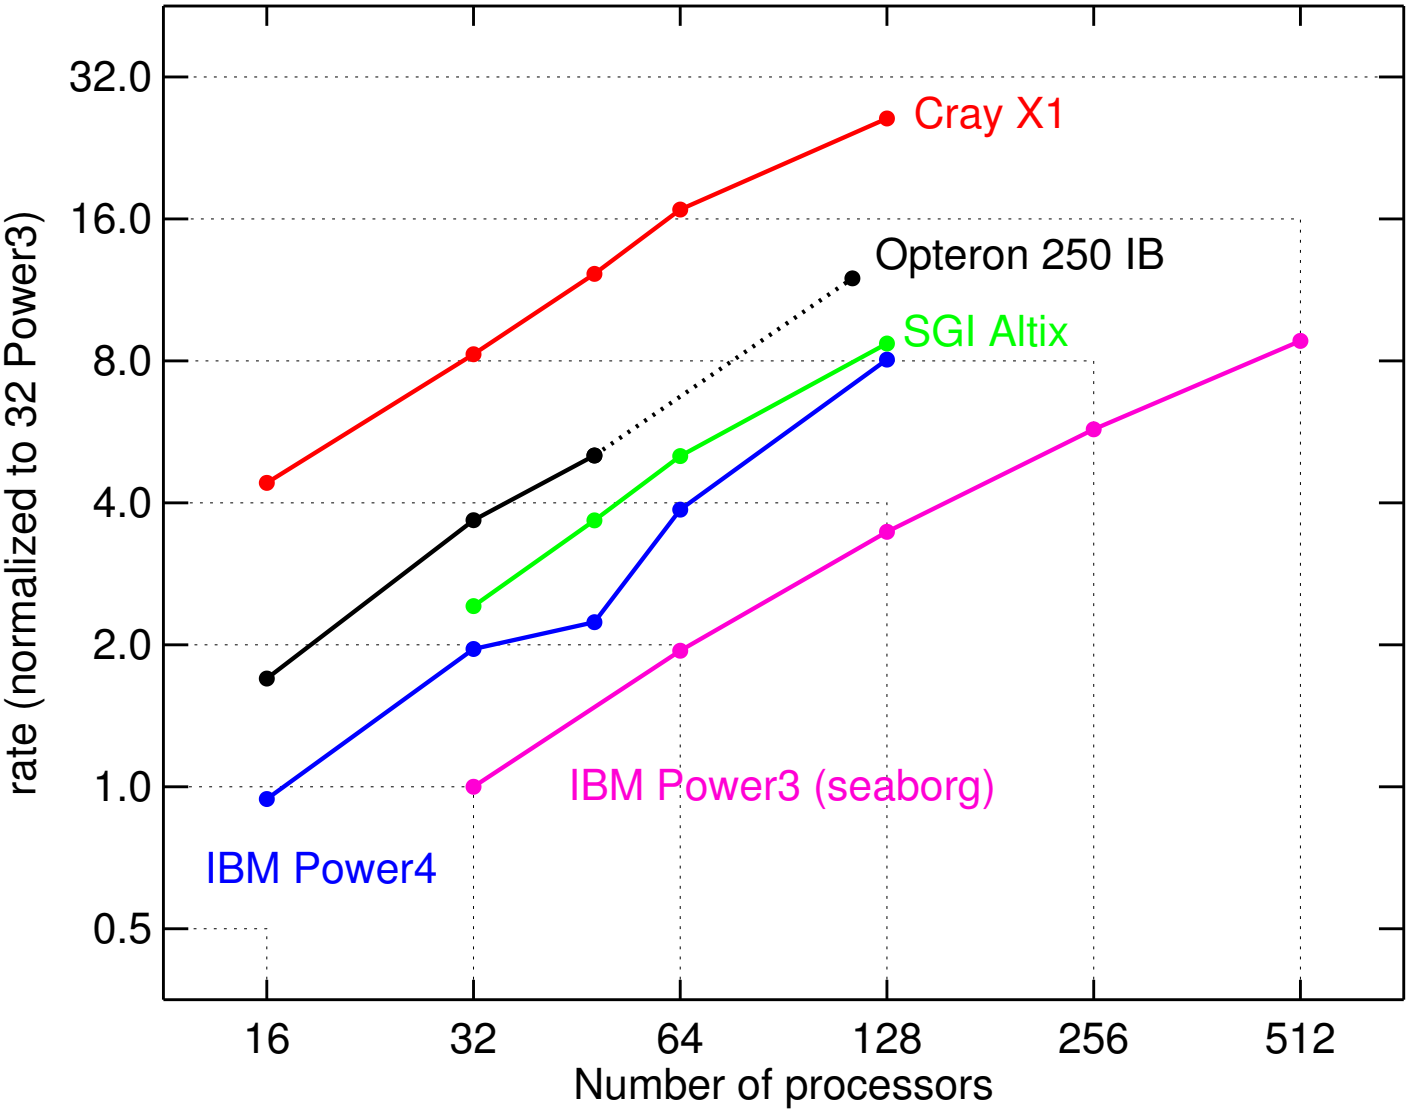

GENERAL ATOMICS

Summary of Overall GYRO Performance

- All systems **scale well** up to and past 128 processors.
- The **Cray X1** is the hands-down winner in per-processor performance:
  - $8\times$ the Power3
  - $4\times$ the Power4
  - $2\times$ the Opteron-IB
- The **IBM Power 4** is twice as fast as the IBM Power 3
- The **Opteron** cluster is four times as fast as the IBM Power 3

Absolute Communication Time For Forward+Reverse Column Transpose (fixed problem size)

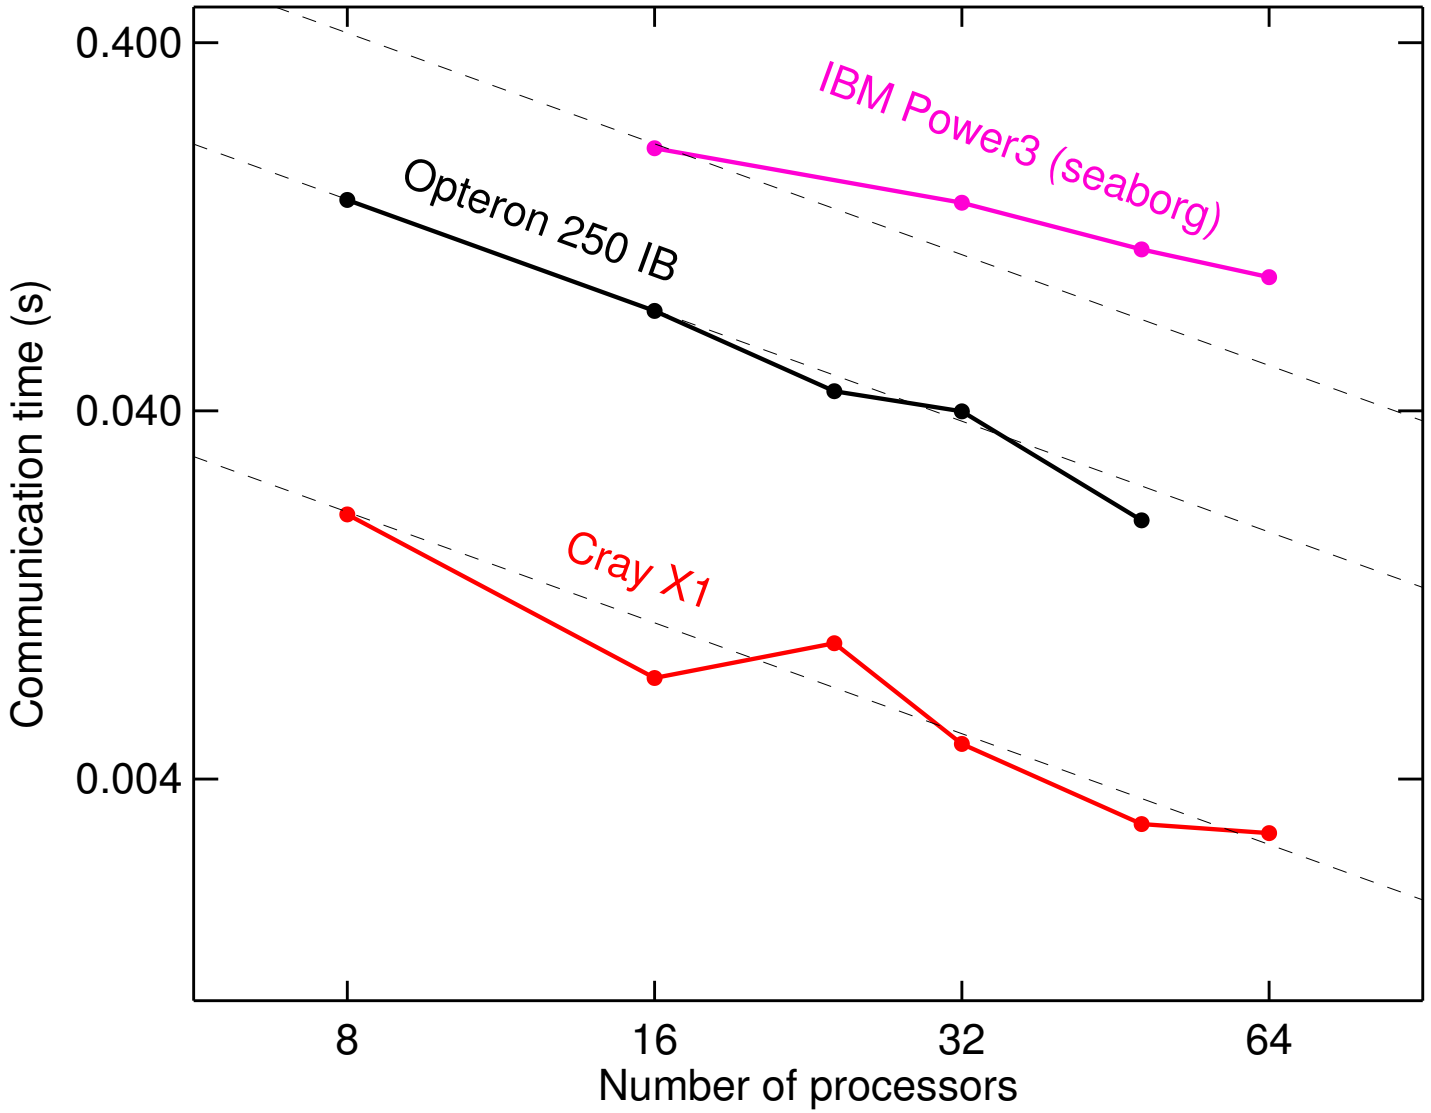

Ratio of Communication Time to Computation Time for Evaluation of Nonlinear Terms (fixed problem size)

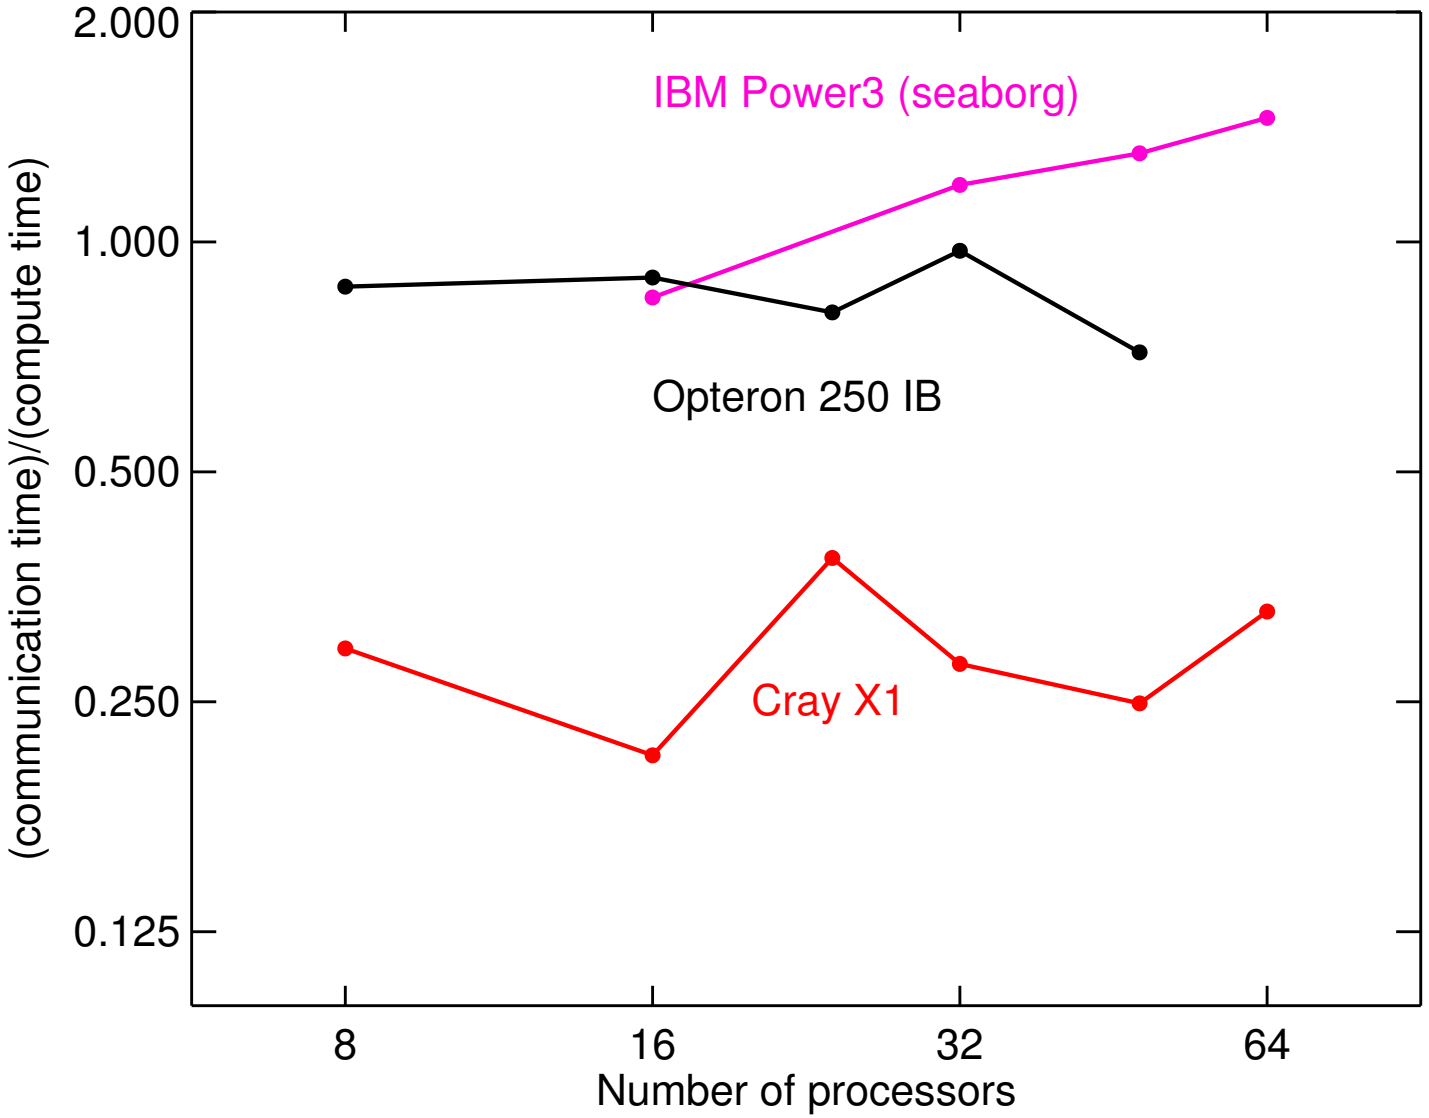

Summary of Essential Results for Column Transpose Timing

- Communication **scales perfectly** on the Cray and Opteron systems
- Communication **scales reasonably well** on the IBM Power3
- The **communication-to-computation** ratio is near unity on the Opteron and Power3 systems, and about 0.25 on the Cray
- The Cray X1 is the only system for which GYRO is not significantly communication-bound.

Preliminary Consideration for the FFT Algorithm: Libraries and Transform Length

- The fields to be transformed are **complex**
- The real-space products need to be **dealiased** for **conservation** of **density**, as well as (generalized) **energy** and **enstrophy** (fluid vorticity in NS turbulence).
- We compute $6N_i$ FFTs of length $3N_n$ during each call
- The FFT libraries are different on each machine:
  - **FFTW 2.1.5** on the AMD
  - **ESSL** on the IBM
  - **LibSci** on the Cray

Preliminary Consideration for the FFT Algorithm: Algebraic Structure

- The discretization uses the **Arakawa symmetrization** to enforce conservation laws

$$
\begin{aligned}
\{F, G\} &= \frac{\partial F}{\partial \alpha}\frac{\partial G}{\partial r} - \frac{\partial G}{\partial \alpha}\frac{\partial F}{\partial r} \\
&= \frac{1}{3}\frac{\partial}{\partial \alpha}\left(F\frac{\partial G}{\partial r} - G\frac{\partial F}{\partial r}\right) \\
&\quad + \frac{1}{3}\frac{\partial}{\partial r}\left(G\frac{\partial F}{\partial \alpha} - F\frac{\partial G}{\partial \alpha}\right) \\
&\quad + \frac{1}{3}\left(\frac{\partial F}{\partial \alpha}\frac{\partial G}{\partial r} - \frac{\partial G}{\partial \alpha}\frac{\partial F}{\partial r}\right) \quad .
\end{aligned}
$$

Comparison of Direct vs. FFT Method for Poisson Bracket Evaluation: Dotted Line is FFT

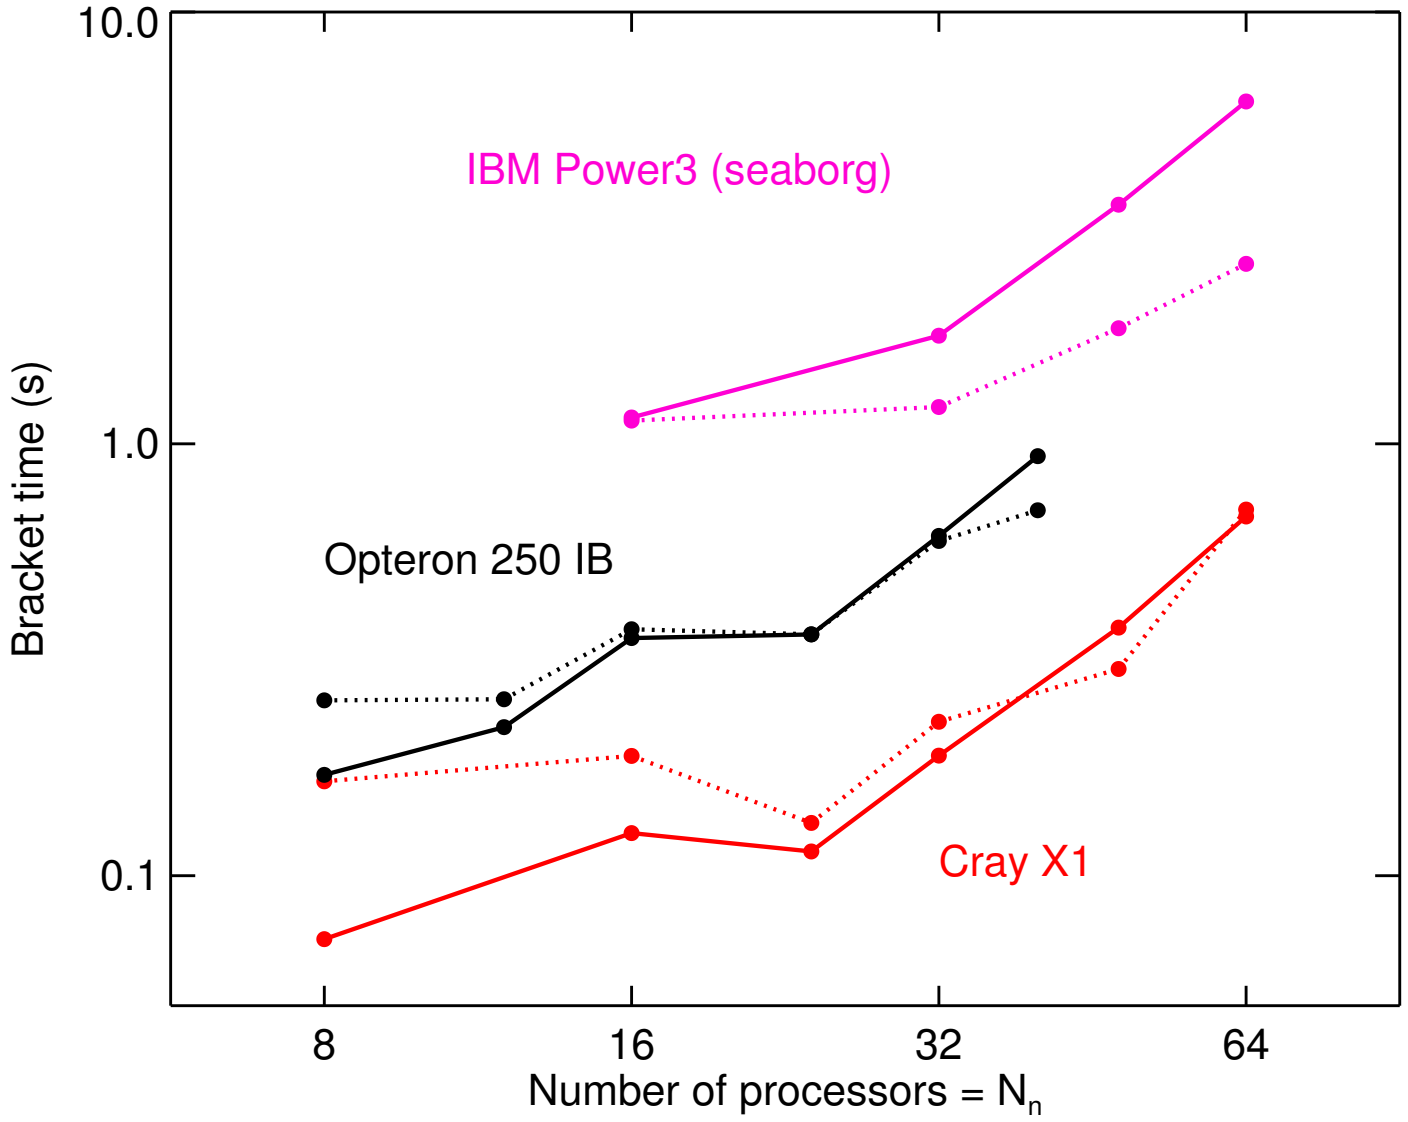

GENERAL ATOMICS

Conclusions Regarding Use of FFT Method (in place of direct method) for GYRO Simulations

- The FFT method (dotted curve) is preferred over the direct method (solid curve) for:
  - $N_n \geq 16$ modes on the IBM Power3
  - $N_n \geq 32$ modes on the AMD Opteron cluster
  - $N_n \geq 48$ modes on the Cray X1

Comparison of Direct vs. FFT Method Using Large Poloidal Grid: Dotted Line is FFT

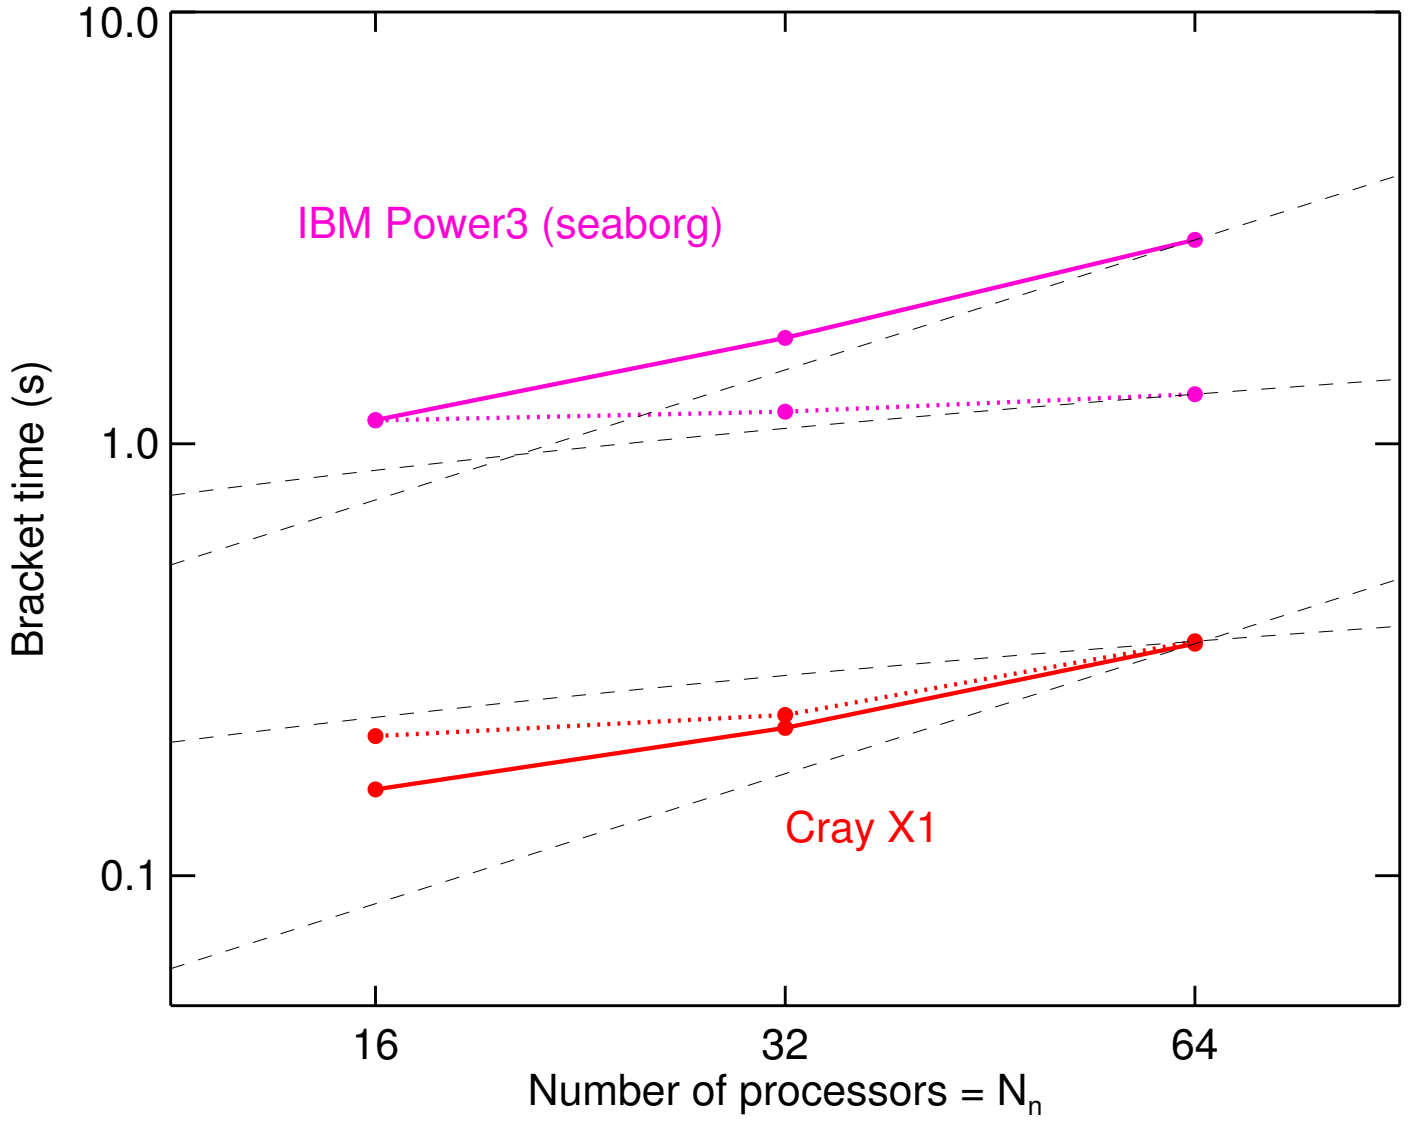

Conclusions Regarding Use of FFT Method for (perfectly load balanced) GYRO Simulations

The FFT behaviours on the IBM and Cray differ.

- The **IBM FFT cost** is, surprisingly, linear in N_n over the range $16 \leq N_n \leq 64$.
- The **Cray FFT cost** is **comparable** to the direct cost over the range $N_n \leq 16 \leq 64$
- We are never in a truly asymptotic $\mathcal{O}\left(N_n \log N_n\right)$ regime for which the FFT is "spectacular"