Data Parallel Simulation and 3-D Visualization of Tokamak Gyro-Landau Fluid Turbulence [1]

G.D. KERBEL, R.E. WALTZ[1], and J.L. MILOVICH

National Energy Research Supercomputer Center
Lawrence Livermore National Laboratory
Livermore, California 94550 USA
gdk@kerbel.nersc.gov

ABSTRACT

High performance data parallel algorithms for simulating and visualizing gyro-Landau fluid (GLF) tokamak turbulence in three dimensions with the TMC CM5^{TM} Connection Machine are used in predictive transport scaling studies.
The time advancement algorithm is composed of data parallel block structured matrix solvers and multiple instance fast Fourier transforms used in operator splitting and pseudospectral convolution evaluation. The overall performance of the algorithm is optimized by choosing a sequence of data arrangements such that the computation in each arrangement is local and there is a fast data rearrangement algorithm setting the next stage of the sequence. The existence of fast communication kernels connecting certain classes of arrangements helps to determine how to structure the algorithm.
The complexity of 3D tokamak turbulence in toroidal geometry requires that advanced visualization techniques be used to understand and communicate results effectively. We have developed an interactive distributed visualization system to examine, compare and display physics results based on the same high performance computing and communications algorithms used in the physics simulation.
The visualization system provides a wide range of flexibility for image composition and animation, making it easy to select, navigate and frame large data sets for further study and presentation. The system also enables the visual experience of concurrent viewing by collaborators, which is helpful to debug algorithms and suggest meaningful diagnostics.
The distributed visualization system uses a CM5 as a remote compute and storage server coordinated through CMAVS/AVS^{TM} with options for local display ranging from a standard X-server to a 3-D multi-screen stereoscopic projection system using multiple SGI Reality Engine^{TM} processors (CAVE[2]).

1. Introduction

The character of computational science is undergoing revolutionary change. Computational physicists have realized a windfall in the amount of computation that can be accessed regularly, and this amount is growing appreciably on the code development timescale.

[1]This work was supported by the USDOE at the Lawrence Livermore National Laboratory under contract number W-7405-ENG-48.

[1]Currently at General Atomics, San Diego, California 92138-5608, USA, waltz@gav.gat.com
[2]Currently at National Energy Research Supercomputer Center, jlm@nersc.gov

Predictive scaling studies of 3-D toroidal GLF model turbulence in tokamaks are useful to characterize plasma transport in current and proposed tokamak experiments. The computing power these studies require varies with the grid size N as N^4, thus in the recent past rather coarse grids have been typical of computations using conventional supercomputers. Massively parallel (MP) computers are now becoming available which can extend these limits.

In response to this growth in resources, computational models are becoming more realistic. Qualifying assumptions which had been retained to make computations more tractable are being relaxed. Three dimensional models are being developed and used today where two dimensional models had to suffice just a few years ago. But as our models have become more realistic, the solutions have as well. The amount of data to be digested has increased enormously. The models and their solutions are so complex that results are typically presented as global statistics. Model and code comparison routinely involve only broad measures of computed physical quantities such as global time averages or simply measures of code or algorithm performance. While global time averages may be the most relevant *bottom line* check, they are often less than helpful in diagnosing systematic differences in scaling due to model or algorithm variation. Better measures are needed.

Physics issues which could not be adequately addressed before are now closer at hand. But harnessing the full power of MP computers often implies redesigning algorithms, including those used to extract useful information for study.

We have created high performance versions of a tokamak gyrofluid turbulence simulation code, **Trb**, for the TMC CM5 and the Cray C90 computers designed to study 3-D ion temperature gradient (ITG) turbulence in a novel non-linear ballooning mode representation (BMR). This representation resolves toroidal modes with $k_\perp \rho_s \sim 1$ more efficiently than the conventional multiple helicity Fourier basis and diagonalizes the toroidal GLF moment closure described by Waltz, Dominguez and Hammett[3] making it much more convenient to apply numerically. The principal non-linearities in the model, arising from the $\mathbf{E} \times \mathbf{B}$ drift convection, are computed by evaluating convolutions pseudospectrally. Time evolution is done by operator splitting between the dual transform spaces along the magnetic field using multiple instance parallel fast Fourier transforms and block diagonal inversion.

These high performance physics codes generate complex non-linearly saturated 3D turbulent plasma flows in toroidal geometry in a few hours of computing on a dedicated resource.

The development of sophisticated graphical user interfaces, specialized hardware for high performance visual displays and extensible software systems to access these features from high level languages[1] have made it possible to augment the high performance data processing capabilities of large supercomputers with the visual display hardware capabilities of high performance graphics workstations at the scientific program development level.

Taken together, these developments have enabled a revolutionary advance in the way we use supercomputers to do computational physics that could not have come at a more opportune time. Never before have our simulation models generated data more rapidly. Yet visualization and communication systems exist which can *match* these data generation rates. Rendering algorithms connected through high speed networks to fast framebuffer displays can allow us to see the emergence and development of complex features in computed data with our eyes. Proximity to the supercomputing resource can enable this even as the simulation is running.

It is only natural that this kind of combined resource should stimulate the development of ever better ways to visualize the results of computations. Although it is difficult to be entirely objective as to what might be meant here by *better*, one can imagine discussing objective criteria for what visual elements might be termed *helpful* or evocative in the sense that perceptual visual processing in itself might help build an understanding of the behavior of the model system.

In this article we will compare several visualization strategies applied to a gyro-Landau fluid (GLF) model of tokamak turbulence. A description of two specific visualization algorithms we have developed follows after a rudimentary synopsis of the physical and computational models which we include for completeness. We conclude with a brief discussion of some visual elements incorporated in the images and animations produced with respect to the criterion that they help lead to recognition and understanding of emergent features in the model simulations.

2. The Physical Model

Turbulence among ion temperature gradient (ITG) modes is thought to be a useful paradigm for understanding transport in tokamaks. While this turbulence is primarily two dimensional in character with the dominant nonlinear motions directed perpendicular to the magnetic field, 2-D simulations do not adequately address the role of magnetic field shear and curvature in plasma transport scaling. Studies of shear and curvature require fundamentally three dimensional representations of the waves or modes. Shear creates linear radial coupling and curvature couples the poloidal harmonics. This causes a nonlinear *ballooning* motion of the plasma across the equilibrium magnetic field.

Hammett and Perkins[4] showed that a collisionless heat diffusivity $\chi = 2v_{th}/\sqrt{\pi}|k_{||}|$ added to the parallel pressure (highest) moment represents a three-pole approximation to the kinetic plasma dispersion $Z(\omega/k_{||}v_{th})$. **Trb** codifies a generalization of this model due to Waltz, et al[3] which includes in addition the toroidal curvature drift motion of gyro-centers. The principal result follows from taking velocity moments of the (Fourier transformed) nonlinear ion gyro-kinetic equation to get fluid equations for perturbed density $\widetilde{\nu}_{\mathbf{k}}$, parallel motion $\widetilde{u}_{\mathbf{k}}$, parallel pressure $\widetilde{p}_{||\mathbf{k}}$ and perpendicular pressure $\widetilde{p}_{\perp\mathbf{k}}$ of gyro-centers. The higher order perturbed moments are closed as fixed linear combinations of the lower moments by first assuming a perturbed Maxwellian and then adding *Maxwellian deviations* fitted to the ion kinetic response function. A charge quasi-neutrality relation based on adiabatic electron response is then used to obtain the electrostatic potential.

2.1. Ballooning Mode Representation

Trb uses a novel non-linear *ballooning mode representation* (BMR) which is different from the conventional multiple helicity (m-poloidal,n-toroidal,r-radial) Fourier representation.[5] This basis is aligned with the magnetic field and is thus well suited for resolving fine structure ITG modes satisfying $k_\perp \rho_s \sim 1$.

The formulation and numerical implementation of this technique has undergone intense recent development by the present authors as well as Beer et al[6] who developed an efficient method[6] for enforcing poloidal periodicity.

Using the *twisted eddy basis* of Cowley et al[7] we have formulated a simple and practical nonlinear BMR; the coordinates of this system (x', y', z') are related to the lab fixed basis (x, y, z) by

$$x' = x; \qquad y' = y - xz\widehat{s}/Rq; \qquad z' = z \tag{1}$$

with

$$\frac{\partial}{\partial x'} = \frac{\partial}{\partial x} + z\widehat{s}/Rq\frac{\partial}{\partial y}; \qquad \frac{\partial}{\partial y'} = \frac{\partial}{\partial y}; \qquad \frac{\partial}{\partial z'} = \frac{\partial}{\partial z} + x\widehat{s}/Rq\frac{\partial}{\partial y} \tag{2}$$

where $\widehat{s}$ is the normalized magnetic shear length, $\widehat{s} = Rq/L_s$ and $L_s = \partial \ln q/\partial r$. Here z' stays along the sheared field as radial variable x' changes with x', y' perpendicular to the sheared field. The corresponding Fourier space (k'_x, k'_y, z') is then

$$k'_x = k_x + z\widehat{s}/Rqk_y; \qquad k'_y = k_y; \qquad z' = z \tag{3}$$

An advantage of using BMR coordinates for this problem is that they diagonalize the operator involving the magnetic curvature drift in the dissipative closure[3]

$$|\omega_D| = \frac{2{v_{th}}^2}{\Omega R}|k'_y\left(\cos\theta' + \widehat{s}(\theta' - \widehat{\theta})\sin\theta'\right)| \tag{4}$$

where $\theta' \equiv z'/Rq$ and $\widehat{\theta} = k'_x/k'_y\widehat{s}$ and Ω is the ion cyclotron frequency. Noting $k'_y = nq/r$ one can identify the usual BMR where n is the toroidal mode label and $\widehat{\theta}$ is the ballooning angle or phase between poloidal harmonics.

2.2. Charge Quasi-Neutrality

The gyro-averaged potential $\varphi_{\mathbf{k}}$ can be related to the dynamical fields through the quasi-neutrality relation,

$$\widetilde{\nu}_{\mathbf{k}} = \widetilde{n}_{\mathbf{k}} + (1 - \Gamma_o)\varphi_{\mathbf{k}} \tag{5}$$

where $\tilde{n}_{\mathbf{k}}$ is the perturbed physical ion density, $\Gamma_o(b) = e^{-b} I_n(b)$ and $b = k_\perp^2 \rho_s^2$ transforms as

$$b \to b' = {k_\perp}'^2 \rho_s^2 = k_y'^2 (1 + \hat{s}^2 (\theta' - \widehat{\theta})^2) \rho_s^2. \tag{6}$$

To model ITG or η_i-mode turbulence, we take the adiabatic response to be

$$\tilde{\nu}_{\mathbf{k}} = \varphi_{\mathbf{k}} - \langle \varphi_{\mathbf{k}} \rangle \tag{7}$$

where $\langle \varphi_{\mathbf{k}} \rangle$ is the flux surface average $\varphi_{\mathbf{k}}$. For finite-n ballooning modes $\langle \varphi_{\mathbf{k}} \rangle = 0$, but for modes with $n = 0$ it is important to realize that electron density can have a Boltzmann response only to deviations from the average potential along a field line.[8] The simple adiabatic electron response represented by (7) is a convenience. The replacement of this approximation with more accurate electron dynamics is the subject of current research.

3. The Computational Model

3.1. Time Evolution Algorithm

The **Trb** time advancement algorithm uses fast Fourier transforms and operator splitting. The technique is a four step Crank-Nicholson predictor-corrector scheme with each of the two pairs of steps, (a predictor pair and a corrector pair) split into two parts; one split advances the part of the operator which diagonalizes in θ'-space, the other advances the part of the operator which diagonalizes in the dual transform space, ${k_\parallel}'$-space. Fast Fourier transforms are used between splits.

The time advancement uses an MNJ-fold multiple instance parallel (4X4) matrix inversion scheme in three out of the four steps of the algorithm. Here M and N are the discretization levels of the k_y' and k_x' meshes, respectively; J is the co-discretization level for the θ' and ${k_\parallel}'$ meshes. The (4X4) elemental block structure arises from the fourth degree moment hierarchy truncation closure.[5]

The algorithm used for time advancement for each field $f_{\mathbf{k}} = (\tilde{\nu}_{\mathbf{k}}, \tilde{u}_{\mathbf{k}}, \tilde{p}_{\parallel \mathbf{k}}, \tilde{p}_{\perp \mathbf{k}})$ can be cast in the form

$$\frac{\partial}{\partial t} f_{\mathbf{k}i} = \{\varphi_{\mathbf{k}}, f_{\mathbf{k}i}\} + \sum_l \widehat{A}_{il} f_{\mathbf{k}l} + \widehat{B}_i \varphi_{\mathbf{k}} + \mathcal{F}_\parallel^{-1} (\sum_l \widehat{C}_{il} \mathcal{F}_\parallel (f_{\mathbf{k}l}) + \widehat{D}_i \mathcal{F}_\parallel (\Gamma_o \varphi_{\mathbf{k}})) \tag{8}$$

where $\widehat{A}$, $\widehat{B}$ are functions of $(\mathbf{k}_\perp', \theta')$, $\widehat{C}$, $\widehat{D}$ are functions of $(\mathbf{k}_\perp', {k_\parallel}')$, and $\mathcal{F}_\parallel, \mathcal{F}_\parallel^{-1}$ are transforms between these dual spaces θ' and ${k_\parallel}'$ parallel to the magnetic field. The convective non-linearity appearing as the convolution $\{\varphi_{\mathbf{k}}, f_{\mathbf{k}}\}$ is evaluated pseudospectrally using fast Fourier transforms in the perpendicular directions.

Supressing the subscipts and $\mathbf{k}_\perp$ for the matrices, the four step algorithm can be expressed as

$$\begin{aligned}
\mathcal{F}_\parallel({f_{\mathbf{k}}}'^{n+\frac{1}{2}} - {f_{\mathbf{k}}}^n)/\tfrac{1}{2}\Delta t &= \widehat{C}\mathcal{F}_\parallel(\vartheta {f_{\mathbf{k}}}'^{n+\frac{1}{2}} + (1-\vartheta){f_{\mathbf{k}}}^n) + \widehat{D}\mathcal{F}_\parallel(\Gamma_o {\varphi_{\mathbf{k}}}^n) \\
({f_{\mathbf{k}}}^{n+\frac{1}{2}} - {f_{\mathbf{k}}}'^{n+\frac{1}{2}})/\tfrac{1}{2}\Delta t &= \widehat{A}(\vartheta {f_{\mathbf{k}}}^{n+\frac{1}{2}} + (1-\vartheta){f_{\mathbf{k}}}^n) + \widehat{B}{\varphi_{\mathbf{k}}}^n + \{{\varphi_{\mathbf{k}}}^n, {f_{\mathbf{k}}}^n\} \\
\mathcal{F}_\parallel({f_{\mathbf{k}}}'^{n+1} - {f_{\mathbf{k}}}^n)/\Delta t &= \widehat{C}\mathcal{F}_\parallel({f_{\mathbf{k}}}^{n+\frac{1}{2}}) + \widehat{D}\mathcal{F}_\parallel(\Gamma_o {\varphi_{\mathbf{k}}}^{n+\frac{1}{2}}) \\
({f_{\mathbf{k}}}^{n+1} - {f_{\mathbf{k}}}'^{n+1})/\Delta t &= \widehat{A}(\vartheta {f_{\mathbf{k}}}^{n+1} + (1-\vartheta){f_{\mathbf{k}}}^n) + \widehat{B}{\varphi_{\mathbf{k}}}^{n+\frac{1}{2}} + \{{\varphi_{\mathbf{k}}}^{n+\frac{1}{2}}, {f_{\mathbf{k}}}^{n+\frac{1}{2}}\}
\end{aligned} \tag{9}$$

Prior to each step, the fields are transformed appropriately in $\theta' \leftrightarrow {k_\parallel}'$ with fast Fourier transforms.

4. Data Parallel Algorithms

Creating the data parallel CM5 version of **Trb** required special consideration of data distribution not necessary for the Cray C90 version. Massively parallel computers routinely

access some data in memory local to the processing unit which computes on it, while other necessary data reside in memory associated with remote processors and must be communicated to be processed. Access to remote data is often significantly slower than access to local data, thus the efficiency of an algorithm can depend strongly on data distribution in memory.

In addition, there are certain patterns of data motion which compile into *fast moves*. Using data distributions whose access patterns with respect to the computation are either local or which coincide with these *fast moves* tends to optimize overall code performance. It is often the case, however, that different parts of a computation accessing the same data may generate very different access patterns. Here the choice can be a compromise between a suboptimal data distribution on the one hand or a costly transpose between two optimal distributions on the other.

5. Visualization Algorithms

The geometry natural to the BMR, our computational domain, is mathematically convenient but difficult to visualize directly. The presumed (multiple) periodicity in the representation admits an implicit *tiling* of the torus with 3-D helical tiles. One visualization technique we have developed uses data parallel ray casting to render 3-D projections of computed data based on this tiling. A second and complementary technique, designed specifically for flow visualization, uses data parallel fluid particle advection in the computational (BMR) domain and fast parallel transformations to real geometry for viewing.

The parallel ray casting procedure has been developed to represent 3-D data on a set of surfaces corresponding to conventional projections such as flux surfaces and poloidal crossections. The parallel flow visualization scheme includes a graphical user interface for a variety of tracer particle initialization templates. The 3-D GLF flow field is animated by following these tracer particles and drawing streamers through their subsequent positions.

We have developed a geometry previewer using advanced features of AVS which allows low resolution animations to be run in real time on the local AVS server concurrently with full scale remote automated animation production. This sort of multi-resolution approach is extremely useful for overcoming some of the network load balance difficulties presented by limited bandwidth between a local AVS-X-server and a remote AVS-compute-server. It is especially useful in production mode where interactivity is necessary for setup, but must be severely curtailed for image sequence generation.

The visualization system can be run as a set of cooperative remote CMAVS modules on the CM5 as part of an AVS network. There is a command line interface (CLI) and a graphical user interface (GUI) to control the system. The variety of control parameters allows the user a great deal of latitude in data exploration and image creation.

Most of the control features are in fact precursors of a virtual reality capability in that they are operated by manipulation of on-screen graphical widgets, radio buttons, dials, sliders, spaceballs and the like, rather than more automatic peripheral equipment. To that end, in collaboration with collegues at the UIChicago Electronics Visualization Lab[2] we have joined a scientific visualization virtual reality experiment called the CAVE with our CMAVS/AVS tokamak visualization system. This has made possible visualizations which run simultaneously on the CM5 and the CAVE with a scientifically sophisticated graphical user interface based on AVS. The results give new meaning to the concept of data presentation.

6. Explanatory Visual Experience

It may be quite impossible to articulate objective criteria that perceptual visual processing itself can help lead to recognition and understanding of emergent features in model simulations. It is easier to recount subjective experience and try one's best to explain it.

Increasing the dimensionality of simulations in computational plasma physics from 2D to 3D comes at the cost of added complexity in algorithms used for both physics computation and visualization. 2D data is aptly presented in the form of contour maps and one can always use this same technique for 3D data by showing only part of the data at a time, ie. as a set of contour maps. This method is simply the analog of showing the 2D data as a set of line plots corresponding to slices through the 2D data set. The price paid for reducing the complexity of the rendering algorithm is the loss of relational information. Data variation with respect to only a subset of the independent spacial variables is presented at a time.

Procedures which increase the dimensionality of the presentation without increasing the

dimensionality of the data typically add little more information to the rendering. Using the function value $f(x, y)$ to determine surface elevation, for example, can tend to obscure spacial information. Determining what goes where often becomes rather more an intellectual exercise in projective geometry than the simpler perceptual experience involved in reading a contour map.

The underlying geometry for a contour map dimensionally matches the information contained in the rendering. Our visual experience of position on the map coincides directly with the objects we are trying to place and relate to one another, namely the level curves of the function. The interpretation is direct and objective.

The surface rendering algorithm can be applied to more powerful effect by contouring the (elevated) surface with a second function, $g(x, y)$ say, whose covariation relative to $f(x, y)$ is of interest. Here although the underlying spacial information may be obscured, new information relating how variations in f and g relate to one another is created *at the perceptual level.* Independent and parallel modes of presentation have been used to generate increased direct objective visual information.

We have tried to incorporate generalizations of these simple ideas in the visualization algorithms we have used to render 3D data. The ray casting technique generates spacial information by creating a simple geometrical embedding or *context* which is naturally suited to the physical model. The context is a set of surfaces on which 3D field data will be contoured. Choices for these surfaces include combinations of magnetic flux surfaces and surfaces of constant toroidal or poloidal angle. Displaying field data defined on these surfaces creates direct visual spacial information about geometric effects in the physical system. For example, seeing η_i *modes ballooning most strongly in the region of worst curvature* displayed in toroidal geometry lends meaning to the statement, especially to an observer who is familiar with the physical model, though not necessarily an expert.

Animation as an independent and parallel mode of presentation can in some cases work to generate additional objective direct visual information. This depends strongly on the ease with which the primary visual system can connect the sequence of images comprising the animation. It is difficult to determine a priori at what point an animation becomes a sequence of independent images connected by an intellectual process, but it is clear that the degree to which each image differs from its predecessor in the sequence and the frame rate are important factors, all else held constant. Although as a general rule, the more that changes from frame to frame, the more taxing is the visual experience, clearly some changes can act constructively to ease the burden. For example, consistent perspective, depth cueing and hidden surface removal from frame to frame seem to work together to ease the perception of spacial relationships in animations relative to the fixed image counterparts. A smooth animation appears to move so easily one can follow the motion and attend to changing internal details simultaneously. An obvious trivial example of this is the case of a moving viewpoint observing a static scene.

Our flow visualization algorithm is a rather radical departure from the embedding ray casting technique. Our intent is to capture a sense of the 3D character of the GLF flow. The geometric context here is merely suggested by the magnetic field lines on a flux surface: the context *coexists* with the flow as opposed to defining the boundaries of its representation as in the embedding. The flow visualization presentation happens via the display of GLF fluid element trajectories determined by parallel solution of the system

$$\frac{d}{dt}\mathbf{x}_i(t) = \mathbf{v}_i(\mathbf{x}_i(t), t)$$

with $\mathbf{x}_i(0) = \mathbf{x}_{i0}$.

Our experimentation with this new technique is just beginning, but there are a few things which seem more or less apparent.

Interactive involvement is important in creating meaningful flow visualizations. It is essential to preserving perceptual continuity to be able to easily navigate the flow field, ie. move the focus of the visualization. It is particularly helpful also to be able to set $\mathbf{x}_{i0}$ interactively. We have found that for GLF micro-turbulence, small scale, dense initializations or large scale highly structured initializations give interesting results as long as they are advanced on time scales slow relative to the time for the initial configuration to dissasemble. We have developed schemes to prolong the impression of high density and to highlight the connectivity along

trajectories. The loss of either of these properties causes a rapid deterioration in the direct visual impact of the visualization as a recognizable flow. At this writing the authors have not yet experienced the effects of autostereopsis in these visualizations to be demonstrated in the CAVE at SIGGRAPH 1994 (Orlando, FL).

7. Conclusions

The matching of high performance computing data generation rates with display technology has led to the development of better ways to visualize the results of computations. Objective measures of the explanatory value of visualization techniques are difficult to elucidate, but some general themes can be identified. In particular, it seems the more perceptual and direct the visual experience, the wider is the range of feature recognition that can be elicited. Contextual information providing orientation through geometric or primary perceptual cues is helpful.

In this article we have compared several visualization strategies applied to a gyro-Landau fluid (GLF) model of tokamak turbulence. A description of two specific visualization algorithms we have developed followed a rudimentary synopsis of the physical and computational models. A brief discussion of specific aspects of images and animations produced with these algorithms was given with regard to recognition and understanding of emergent features in the model simulations.

8. Acknowledgement

This work is supported in part by the National High Performance Computing and Communications Initiative Grand Challenge Numerical Tokamak Project. The authors are grateful to the staff of the Advanced Computing Laboratory at Los Alamos National Laboratory and to the staff of the Electronics Visualization Laboratory at the University of Illinois at Chicago for the use of their facilities and for their help and cooperation in bringing this project to its current status.

9. Figures

Figure 1a,b Electrostatic potential for ion temperature gradient driven turbulence. Figure (b) shows the turbulent field in figure (a) after an imposed sheared rotational flow field has nearly quenched the turbulence. The results are displayed with a parallel ray tracing technique developed by G. Kerbel[9]

References

1. AVS^{TM}, Explorer^{TM}, Khoros^{TM}, PVWave^{TM} and IDL^{TM} are examples of commercial visualization software packages.
2. CAVE is a virtual reality theatre created in the Electronic Visualization Laboratory of the University of Illinois at Chicago.
3. R.E. Waltz, R.R. Dominguez and G.W. Hammett, Phys. Fluids **B4**, 3138 (1992).
4. G.W. Hammett and F.W. Perkins, Phys. Rev. Lett. **64** (1990) 3019.
5. R.E. Waltz, G.D. Kerbel and J.L. Milovich, *Toroidal Gyro-Landau Fluid Model Turbulence Simulations In A Nonlinear Ballooning Mode Representation With Radial Modes*, submitted to Phys. Fluids.
6. M. Beer, G.W. Hammett, W. Dorland and S.C. Cowley, Bull. Am. Phys. Soc., **37**, 1478 (1992) 5P23.
7. S.C. Cowley, R.M. Kulsrud and R. Sudan, Phys. Fluids, **B3**, 2767 (1991).
8. W. Dorland and G.W. Hammett *Gyrofluid Turbulence Models with Kinetic Effects*, Phys. Fluids **B5**, 812 (1993) [see also W.Dorland, Princeton Ph.D. Thesis.
9. G.D. Kerbel, R.E. Waltz and J.L. Milovich, *Vortex Structure In Tokamak Gyro-Landau Fluid Turbulence Studies*, International Sherwood Fusion Theory Conference, (Dallas, March 1994).