NIMROD Team Meeting Minutes
July 29-31, 2009
Madison, Wisconsin

Coding Topics

Carl Sovinec discussed three topics related to NIMROD's algorithm. The first concerns the boundary conditions for the advance of magnetic field. The electrostatic ∇T_e term in Ohm's law is eliminated analytically to reduce numerical error in two-fluid computations. The surface term from integration by parts therefore needs to use $\mathbf{E}+\nabla T_e$ to be consistent. Tests show that errors arise when this term is not included and that there is an implied electrostatic profile with $T_i>0$ unless ion diamagnetic drift is included. The second topic is the data reordering for parallel efficiency. Most of the changes that will affect developers are in the field-computation and integrand routines. The third topic is the implementation of a 3D semi-implicit operator. It is shown to improve accuracy during the resistive nonlinear saturation of an ideal-MHD 1/1 kink. The cylindrical test case shows converged evolution at $\Delta t=46\ \tau_A$ or $\gamma\Delta t=0.1$. Further testing will consider multi-helicity conditions, resistive instability, and two-fluid computations.

Jacob King presented tests of the new incomplete-LU (ILU) factorization that is available in version 4.0 of sequential SuperLU. The goal is to improve parallel scaling when a parallel version is available; ILU tends to require much less parallel communication and to use much less memory than full LU. The Fortran-to-C bridge routines are similar to those for full LU, but King had to reinstall an option to refactor without reallocating memory. His tests scan the ILU drop tolerance and show that while ILU is slower at tight tolerances, it can reduce NIMROD's overall memory use by ~25%. In a nonlinear two-fluid test, the time consumed with ILU is only slightly greater than with LU, but the LU matrix storage is reduced by 75%. This is an encouraging result, and King will report them to Sherry Li to encourage parallel implementation.

Scott Kruger gave an update on the status of the NIMDEVEL repository at revision 632. Most improvements for the hot-particle closure are incorporated; though, Charlson Kim indicated that there are a couple more bug fixes. NIMDEVEL is now using C-preprocessing to help manage whether different closure modules are loaded. This avoids the use of dummy libraries. Many Fortran compilers will execute C-preprocessing automatically, but a separate PYTHON script version would help debugging. Regression tests are now running well. Also, the nimrodteam.org website will be moved to a new server. [Later in the meeting, Kruger stated that he would like a volunteer to take over management of mp21.]

Model Development

Jeong-Young Ji reported on progress with the general moment method for developing closures. All of the work discussed this time is for the collisional limit in the sense that $\lambda<L$ and aims to improve the accuracy of our closure relations for $\mathbf{q}$, $\underline{\Pi}$, Q, and $\mathbf{R}$ in terms of n, $\mathbf{V}$, and T for each species. He briefly discussed work on nonlinear terms in the collision operator. However, he focused on the transition from small to large magnetization (measured by $x=\Omega\tau_{coll}$). At large magnetization, the moment approach better approximates the asymptotic result when increasing the number of moments but eventually transitions to different limiting behavior. This is attributed to the electron-ion collision matrix, which is not diagonally dominant. His recommendation is to use fits for each transport coefficient that transition to the asymptotic behavior found from the Lorentz-gas model. The new rational function

approximations (with terms of fractional powers in x) capture the information from the moment approach at low- to medium x-values and the asymptotic behavior at large-x. They each have less than 1% error for all x and for arbitrary ionization.

Eric Held discussed topics related to continuum solutions of kinetic equations in NIMROD. He first discussed performance of the toroidal preconditioning scheme with the mixed method for collisional heat conduction. It appears to help in some cases but not all when taking large steps such that the transient is not well resolved. The mixed method may help reduce resolution requirements for stellarator cases, where the $|B|$-variation in the periodic direction is large. Regarding the main topic, tests of solving velocity-moments of the kinetic distortion as functions of position are well underway. In the current implementation, the speed dependence does not appear in the numerically implicit coupling, which allows greater parallelization. His tests consider heat conduction across a 2/1 magnetic island, and computations performed in the collisional regime reproduce Braginskii modeling. A variation of the algorithm staggers the solution of the kinetic distortions and temperature. A separate study investigates the use of finite elements for representing velocity dimensions. Working with a new student (Andrew Spencer), Held has 0-spatial D tests running that show kinetic evolution to a Maxwellian distribution.

Charlson Kim gave a comprehensive review of the hot-particle module in NIMROD. An important aspect of the implementation is that the evolved center-of-mass velocity includes the hot particles; though, the formulation assumes that the hot particles do not contribute appreciable momentum or mass. The equilibrium hot-particle distribution is a 'slowing down' distribution that is coded directly into the weight-evolution equations. It would be possible to replace it with a function call for more generality, but optimization would suffer. Kim covered the F90 data structures and the input required when using the module. The parameters affect the type of kinetic evolution (full vs. drift), quantitative information for the equilibrium, and its spatial distribution. There was also discussion of how to set the normalization constant `psi0`. Particle parameters include the number of particles (Kim recommends $\sim 10^6$ for linear computations) and the subcycling needed to resolve toroidal transits ($v\Delta t \sim 0.1R$*`orbit`). Spatial localization, which was meant to improve particle statistics (by concentrating the particle load in the region of interest, e.g. within q=1 surface) and avoid particle loss due to finite orbits that intersect the boundary, does NOT work because of the finite trajectories. Some particles loaded outside the region of interest will have orbits that pass inside. These particles are the most energetic and are important to include. The slowest part of the algorithm is the Newton search to find a particle's logical coordinates after an advance. An important check when running the module is the Vpar diagnostic. It shows an average parallel speed, which is assumed to be 0, and significant errors appear if it approaches the Alfven speed. Improvements to the use of block parallelization allow efficient 2D decomposition, and computations with 10 M particles have been run efficiently on 100 processors. The group discussed the electric-field computation for combined two-fluid/hot-particle simulations, and the Milroy-Barnes PFB paper from '91 may be relevant. Kim is now applying the hot-particle module to TAEs and EPMs.

<u>Applications</u>

Fatima Ebrahimi presented a study of pressure-driven modes in the RFP, motivated by the 26% β obtained in MST through a combination of pellet injection and PPCD. In a cylinder, one would expect instability in the MHD model. Here, Ebrahimi is considering two-fluid and FLR effects in toroidal geometry. She has generated the equilibria using NIMEQ (for accuracy) with the current and pressure profiles from MSTFIT. Linear computations show that the m=1 modes have tearing parity, while the $m>1$ modes have interchange parity. In general, two-fluid Ohm's law without gyroviscosity leads to small reductions in growth rate, while two-fluid Ohm's with gyroviscosity significantly reduces growth rate but does not provide complete stability. Ion diamagnetic drift is not included in the profiles, so an electrostatic potential profile is implied. The omega-star drift does not exceed twice the MHD growth rate. Nonlinear single-helicity computations have been performed in a cylinder and are dominated by $m=0$ modes that are driven by pressure. Multihelicity zero-beta MHD computations have also been performed to confirm the possibility of running at $S=2\times10^5$ and Pm=20. A new project to model the plasma dynamo experiment in Physics is being funded by NSF.

Eric Howell presented a linear study for decaying spheromak plasmas. It has been noted that when modeling high-performance SSPX discharges, NIMROD MHD simulations tend to under-predict electron temperature by ~40%. The Hooper '08 PoP paper concludes that MHD modes tend to be more unstable in the simulations than in reality. This study focuses on two-fluid and FLR stabilization effects as a possible explanation (there are others). SSPX-relevant equilibrium profiles are generated with NIMEQ, and the computations have $S=10^6$. The two-fluid effects are more significant for large-n harmonics and tend to be stabilizing; though, complete stabilization is not observed. Howell has considered profiles with uniform density and (motivated by the recent Zhu PRL) with $P\sim\rho^\gamma$, and the latter tend to show larger changes in growth rate at medium-n.

Highlights from a recent seminar by Bonita Burke on the ELM benchmark study were presented by Sovinec, with some additions regarding modes near the stability threshold. The focus has been on two circular cross-section toroidal equilibria that have been developed by Scott Kruger and Phil Snyder to emphasize ballooning. Burke's study shows that it is important to separately test for ideal internal behavior and for external vacuum-like behavior when applying nonideal codes to ideal external modes. She models a 'halo' region with a sharp transition to very large resistivity, with an outer S-value approaching unity, and a significant reduction in mass density. For the more unstable equilibrium, there is excellent quantitative agreement between NIMROD and the other codes when the halo region is used. There is also good agreement on the less unstable case, but the cases near marginality are problematic if the halo region is not used. Sovinec has tested a number of aspects regarding the equilibrium (regenerating it with NIMEQ) and the mesh and has found that even at $S=4\times10^8$, the near-marginal modes are resistive and challenging to resolve given that they exhibit mm-size features at many nearby rational surfaces. Difficulties arise if the mesh packing does not cover all active regions, given NIMROD's tendency to converge from the unstable side on these cases. Including a halo limits this response. Burke is returning to cases with a significant parallel current spike to check peeling-type behavior, and checks for ideal behavior without a halo are complete. In preparation for nonlinear studies, she is also investigating results with a realistic temperature-dependent resistivity profile.

Ping Zhu gave an update on his efforts to apply NIMROD to two-fluid ballooning. Since Sherwood, he has completed stability computations for a circular cross-section equilibrium for toroidal harmonics up to $n=62$. The equilibrium has a pressure length scale of 5% of the minor radius, and the tests consider ion gyroradii to be 0.5%. Computations with two-fluid Ohm's law and gyroviscosity show growth rates reduced by as much as 50% relative to MHD but no complete stabilization. He has considered possible explanations including coupling of the drift-Alfven mode to drift-acoustic behavior (described in a 2003 Hastie paper). Chris Hegna suggested varying compressibility to check this. Zhu also showed plots of the eigenmodes and how they shift outward with the two-fluid effects. They appear to move into a region with smaller diamagnetic drift, an effect considered in a 2000 Hastie PoP paper. An initial nonlinear computation appears to need more spatial resolution.

Mark Schlutt noted high-β results from stellar experiments that motivate his study of 3D island evolution, retaining finite parallel thermal conductivity. His analytical computations evaluate the island-width equation with finite $\chi_{||}$ and now include nonresonant contributions. When running cylindrical stellarator computations with NIMROD, he finds a dependence on $\chi_{||}$ in steady conditions, which has been puzzling. He has checked poloidal resolution, and Kruger noted that these computations may be very sensitive to toroidal resolution. Pfirsch-Schluter currents are particularly important to the theory, and results from the simulations appear to be in at least rough agreement with analytical estimates. Schlutt hopes to be able to create 2D equilibria that are above an MHD stability threshold, perturb them, and observe nonlinear evolution to a 3D saturation.

Andrea Montgomery presented some initial work for resistive wall modes in a cylinder. Her first steps are to reproduce linear behavior for a 2/1 mode and compare with the '95 Finn paper. She is using the analytical solution for vacuum fields beyond the resistive wall, which is modeled with the thin-wall equation. The corresponding tangential electric field is used in the surface integrand for the update of magnetic field. Sovinec recommended setting *B*-normal as an extra condition, as required for the divergence cleaner. The most recent cases use a large change in resistivity at the plasma surface. A growing mode that depends on the wall resistivity is found, but the growth rates at this point are significantly different than what is expected.

Ryoji Takahashi gave on update on his study of hot-particle effects on tearing modes and their nonlinear evolution. Previously, he found a significant stabilizing effect on linear modes in the experimentally relevant parameter regime in terms of S and β_N/l_i. His nonlinear base-case uses resistive MHD with n up to 10 and saturates with a small 2/1 island. This resolution may be too costly with particles. When cases with n up to 2 are reset with energetic particles, the n=0 and n=2 kinetic energies suddenly increase by several orders of magnitude, and the larger the hot-particle β, the larger the jump. With β_{hot}=1%, the nonlinear simulation runs to saturation. With two-fluid effects and no hot particles in linear computations, the rotation is in the opposite direction, but noise develops near the separatrix.

Dylan Brennan described a study of the effects of flow on nonlinearly generated tearing modes in tokamaks to address island-size and torque scaling. He summarized previous linear and nonlinear results of flow, differential flow between rational surfaces, and flow shear at rational surfaces. Initial-value computations use significant toroidal flows (~sound speed) that affect the equilibrium profile to some extent. The 1/1 surface rotates past the 3/2 surface several times during a 1/1 growth time. The nonlinear simulation saturates, and Brennan is analyzing torques with a computation that is based on phases at the original rational surfaces. However, the phase diagnostics show flips of pi in close proximity to the rational surfaces, which are due to updown asymmetry and flow, and necessitate high accuracy in tracking the movement of the actual rational surfaces. From amplitudes, it is clear that the 1/1+2/1 channel is the dominant coupling, but in nonlinear saturation, 1/0 is also large, so the 2/2+1/0 channel may become important. The core 1/1 mode also drags other modes forward in the saturated stage. The 3/2 island remains small through the 1/1 crash phase but becomes large after the crash during recovery; the crash itself may have little to do with island generation.

Jacob King gave an update on his study of two-fluid magnetic tearing in a pinch profile with uniform pressure. With just a two-fluid Ohm's law, the linear growth rate increases with ρ_s. When gryoviscosity is added, the linear growth rates are reduced over a range of small ρ_s-values and loses the stabilization effect at large ρ_s, where ions are poorly coupled. In nonlinear single-helicity computations, the saturated island size decreases somewhat for parameters where GV is linearly stabilizing. King presented a diagnostic of forces from the gyroviscosity and showed that they oppose the tearing flow pattern. Dan Barnes suggested checking other phases to verify that the GV is not damping.There is progress on a multi-helicity case with two-fluid Ohm's law and GV, but it has yet to run through saturation with Pm=1.

Tom Jenkins presented his study on island evolution with a simple model for RF current drive. He is comparing and contrasting the D_{ECCD} and long- and short-term Δ' effects. To investigate Δ' effects, he applies the current-drive model before evolving the nonsymmetric perturbation. Movement of the rational surface at short and long times is important for the stabilization effect. The short-time response moves the rational surface away from the point of RF deposition, but over the resistive diffusion time-scale, the response is an outward drift. Thus, an application of RF can be destabilizing at short time and stabilizing later. In addition, different levels of applied RF power can better target D_{ECCD} and can stabilize through that effect even when the short-time Δ' effect is destabilizing. Jenkins also gave on update on a more realistic model that couples propagation information from GENRAY through the Integrated Plasma Simulation software from SWIM.

Dylan Brennan described a tearing-mode study being performed by his student, Matt Behlmann. They would like to see the approach to ideal instability with constant current gradient in slab geometry. This will add J_0' effects to the analysis of Cole & Fitzpatrick to address penetration of magnetic perturbations with respect to varying stability. They have performed analytical work for the viscous regime for initial comparisons with NIMROD, and will compare computations in the inertial regime with the extensions to Cole and Fitzpatrick to be applicable to D3D parameters. Behlmann is running the slab computations with NIMROD and is reproducing layer-width and growth-rate scalings.

Val Izzo gave an update on her study of runaway-electron generation during disruption. The modeling advances equations of motion for a group of test electrons during each NIMROD time-step, but there is no feedback on the evolving fields. The parallel implementation started from what Held uses for characteristics in the nonlocal closure computations, but the master-slave load balancing is not needed for the relatively small number of test electrons. The equations of motion include curvature, ∇B, and $\mathbf{E}\times\mathbf{B}$ drift effects in the perpendicular directions and $\mathbf{E}$-acceleration, collisions, synchrotron, and Bremsstrahlung radiation for the parallel speed. Computational results show that for relativistic-γ, the curvature drift smoothes over magnetic stochasticity, and the fast electrons are better confined than what one would expect from following field-lines. Izzo has also performed simulations that are relevant to disruption induced by argon particle injection with resonant magnetic perturbation (RMP); though, the particle effect is spread throughout the volume. The simulated quench time of 1 ms may be too fast, and many electrons escape when the n=3 mode breaks magnetic surfaces. Izzo also described a proposal for ITER to explore whether RMP can be used to shed electrons and what amplification is needed. She has an implementation of boundary conditions to apply **E** and *B*-normal for time-dependent surface fields of general shape; it may be useful for other applications.

Dalton Schnack summarized the status of the giant sawtooth project. Since the Sherwood meeting, the necessary coding modules have been assembled in NIMDEVEL, and growth rates from NIMDEVEL and NIMPSI agree. Next steps include modifying the particle equilibrium to model RF-generated particles that are critical for preventing smaller-amplitude sawteeth. Modifications to Ohm's law may be necessary for combining two-fluid (& gyroviscosity?) and hot-particle effects; reviews of the formulation are needed. Schnack suggested a separated meeting to focus on the giant sawtooth topic.

Chris Carey presented nonlinear simulation results on the evolution of kink-unstable jets. The modeling follows the magnetic-tower configuration that may be produced by a dynamo within the disk such that there is no net flux. The general morphology is that a short-length helix forms near the propagation front at early times, generates other harmonics and saturates, stretches, and leaves a long-wavelength helix. Collimation is influenced by the decreasing Alfven-speed profile away from the disk. The nonlinear saturation of the m=1 mode relaxes the mean profile. Rotation influences the mode but does not stabilize at the parameters considered. Line-tying is thought to be important in the region immediately above the disk, providing a starting point for collimation. Similar to spheromaks, the relaxation process leads to amplification of the initial poloidal flux, despite the absence of a close-fitting conducting shell.

Charlson Kim gave a summary of different PSI-Center applications with NIMROD. He first discussed simulations of the co-planar jet experiment at Caltech. Results do not appear to show as much compression as expected, but it is possible that a significant amount of power is being dissipated in the boundary layer above the injector. Simulations of LDX use heating, drive interchange modes unstable, and generate fluctuations during saturation. The pressure appears to be evolving toward the marginally unstable profile. Computations with energetic particles using full kinetics apply a correction for the equilibrium distribution. This is needed for FRC applications and may be important for high-energy particles in tokamaks. Computations of tearing modes with hot particles for RFPs show broadening of the eigenfunctions. So far, this does not include off-diagonal pressure-tensor components to avoid errors due to noise, but their effects may be important. Simulations of rotating magnetic field (RMF) current drive for FRCs are progressing. The simulations model even-parity RMF, which is applied to boundary conditions on **E** and **B**, and make use of toroidal preconditioning. Kim showed simulation results on FRC formation from RMF alone. Initial computations used periodic geometries. When end walls are used, kink distortions develop. Future improvements will include anisotropic transport that can handle the field null, better resolution, less dissipation, and neutral modeling.

Cihan Akcay presented simulations of the HIT-SI experiment. His computations use a resistive layer of 1 mm thickness along the inner surface of the wall, and resistivity increases by a factor of 10^5 relative to the interior. The injector fields are modeled as a force-free distribution over the patches where the injector loops intersect the flux conserver, and they are applied through *B*-normal and *E*-tangential conditions. Akcay is first trying to reproduce Izzo's results, but the implementation of the resistive layer and the use of viscosity near the convex corners of the wall are different. There are discrepancies during the drive phases, but the decay phases look similar. Checks of both poloidal and toroidal resolution are underway.

Tom Bird presented his study of magnetic relaxation of a current filament from a miniature plasma gun in Pegasus. As known from the experiment, the important criteria for the onset of relaxation is the development of reversed flux inboard of the gun. In these conditions, the self-induced field is at least comparable to the vacuum field and can have a large influence on the current path. He is simulating one gun located in the bottom of the chamber. His simulations use a helicity source that is localized toroidally and poloidally. Simulations with a simple zero-β model and small dissipation show oscillations that are attributed to the current helix collecting and redistributing while the current builds. When approaching relaxation conditions, the current filament bunches near the injector. After reversal, there is mixing, and the helical structure is largely lost. Computations with temperature-dependent coefficients evolve more quiescently through the initial phase.

John O'Bryan is modeling more recent experiments of current injection on Pegasus, where the injectors are located near the outboard side of the chamber and poloidal-field (PF) induction is applied to maintain force-balance. He implemented the magnetization effect in 3D temperature-dependent thermal conductivity, according to the Braginskii formulation, to keep perpendicular conductivity from exceeding parallel conductivity at low temperature. [Ji's recent and more accurate formulas could be used instead.] Two-dimensional zero-β tests of PF induction modeling show current multiplication of approximately 2 that becomes independent of S for $S>10^4$. Simulations with temperature evolution are performed in two ways: 1) profiles from a 2D source are allowed to equilibrate before compression, and 2) estimates from the experiment are loaded as initial conditions, and PF is applied immediately. For the former, the evolution of toroidal current and its multiplication are similar to the zero-β tests. In the latter, there is a large initial transient, and the compression adds to the heating.

Dylan Brennan presented a study of spheromak evolution based on ideal stability; some of the slides were taken from his Sherwood poster on the topic. The stability computations use a model 2D equilibrium with separate values of parallel current for open- and close-flux regions. With this he is able to provide a stability map for the parallel-current parameter space. He is examining results from the SPHEX experiment with a geometrically centered rod and a modest amount of I_{TF}. He would like to show that the relaxed profile remains near the ideal stability limit as the equilibrium evolves.

Chris Hegna presented analysis on two-fluid stabilization of ballooning in the presence of a gradient in the drift parameter ω_*. He reviewed basics of the ballooning formalism for the large-*n* limit, which simplify the equations to ODEs along field-lines. For large-*n*, the WKB analysis leads to a quantization rule, where $\omega^2=\lambda$ is the eigenvalue. For uniform ω_*, the Roberts-Taylor stabilization criterion results. Hegna redid the computation with ω_* varying linearly in poloidal flux. With a very large gradient, the result is essentially MHD (no stabilization) as the mode shifts outward to a location of lower ω_*, similar to what was observed in Zhu's computations. This behavior does not go away with increasing *n*, so drift stabilization may be lost in some cases. Hegna noted that this result is essentially the same as what is in Hastie's 2000 paper.