SANDIA REPORT
SAND2009-0963
Unlimited Release
Printed February 2009

Hybrid Plasma Modeling

Matthew M. Hopkins, Edward S. Piekos, Lawrence J. Dechant, Timothy D. Pointon

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2009-0963
Unlimited Release
Printed February 2009

Hybrid Plasma Modeling

Matthew M. Hopkins[1], Edward S. Piekos[2], Lawrence J. Dechant[3], Timothy D. Pointon[4]

[1]Nanoscale and Reactive Processes, [2]Microscale Science and Technology, [3]Aerosciences,
[4]Electromagnetic and Plasma Physics Analysis
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-0836

Abstract

This report summarizes the work completed during FY2007 and FY2008 for the LDRD project "Hybrid Plasma Modeling". The goal of this project was to develop hybrid methods to model plasmas across the non-continuum-to-continuum collisionality spectrum. The primary methodology to span these regimes was to couple a kinetic method (e.g., Particle-In-Cell) in the non-continuum regions to a continuum PDE-based method (e.g., finite differences) in continuum regions. The interface between the two would be adjusted dynamically based on statistical sampling of the kinetic results. Although originally a three-year project, it became clear during the second year (FY2008) that there were not sufficient resources to complete the project and it was terminated mid-year.

CONTENTS

FIGURES

TABLES

NOMENCLATURE

L	characteristic length (m)
m	molecular mass (kg)
N	number of samples / particles
T	temperature (K)
p	pressure (Pa)
R	gas constant (J/kg/K)
x	spatial coordinate (m)
v	particle velocity (m/s)
V	fluid velocity (m/s)
λ	mean free path (m)
μ	viscosity (Pa•s)
ρ	density (g/cm^3)
σ	standard deviation of the velocity distribution (m/s)
Kn	Knudsen number ($Kn = \lambda/L$)

1. INTRODUCTION

There are numerous plasma applications that include both non-continuum and continuum phenomena that must be resolved. For example, vacuum arc initiation evolves from a high vacuum non-continuum state to one in which there are mixed areas of high and low collisionality, potentially varying by neutral and charged species within the plasma. The goal of this project was to develop a methodology for spanning the non-continuum-to-continuum spectrum as it varies over time, space, and species.

The following sections will describe the proposed methodology (section 2), some intermediate results (section 3), and progress in developing a needed statistical test (section 4).

We assume the reader has some familiarity with general Particle-In-Cell (PIC) [1], [2], Direct Simulation Monte Carlo (DSMC) [3], and PDE solution methods. Although we make explicit reference to finite difference methods (FDMs), we believe there are benefits to using finite element or finite volume methods (FEMs or FVMs) of various pedigrees.

2. SUMMARY OF HYBRID METHODOLOGY

We refer to "hybrid" methods as those that couple a non-continuum kinetic description to a continuum bulk one. Typically, particle-based methods (PIC, DSMC) are used for the former and mesh-based PDE methods (FDM, FEM, FVM) for the latter. Although there are some coupled particle-PDE methods used for plasma simulation [4], [5], the particle-PDE methods in the neutral gas community are more mature. Current hybrid methods in the neutral gas community [6] assume a default continuum description and employ methods to identify regions that should be treated as non-continuum. In contrast, our methodology assumes a non-continuum description and employs methods to identify regions that should be treated as continuum. The former methods are appropriate for systems that contain small regions of non-continuum physics which, ideally, do not influence the macroscopic continuum region. For example, the non-continuum region may be limited to a boundary layer, with negligible influence on the macroscopic continuum flow outside the boundary layer. In contrast, the mental picture of the applications we are interested in, primarily vacuum arc discharge, are initially non-continuum (vacuum) everywhere, and grow regions of higher and higher collisionality over time. These regions of higher collisionality become continuum; however, they are surrounded by regions of non-continuum (or boundary) and the coupling between the two is strong in both directions. The goals of our methodology are to:

1. Identify regions modeled with a non-continuum description that are collisional enough to be modeled with a continuum method,
2. Develop methods to switch from a continuum description to a non-continuum description,
3. Develop a methodology to account for circumstances where a region with a continuum description can revert back to non-continuum, and
4. Couple the mass, momentum, energy, and charge across the continuum-to-non-continuum interface.

The overall algorithm must of course be stable, and ideally, efficient. In particular, it should be more efficient than performing a purely particle-based simulation.

Although we initially wanted to use discontinuous Galerkin finite element methods, we did not achieve this goal. In the work presented, and the intermediate methodology we implemented for the LDRD, we used finite difference methods. Whether an actual "element" in a finite element method, or the rectangles/cubes within a finite difference method, we refer to the enclosed area containing particles as a "cell".

We developed almost all the methodology in the context of neutral gas flow. A primary difficulty in developing these models for plasmas is to handle the charge density distribution. We planned ways to handle this, but did not progress far enough to implement and test any of them.

Finally, we specifically do not employ methods that try to identify regions of the domain with a continuum model that should be modeled with a non-continuum description – an approach employed by the neutral gas hybrid method community. Instead, our methodology makes this assessment unnecessary.

2.1 What is "collisional enough"?

To address goal (1) we need a method to assess when a set of particles within a cell is "collisional enough" to be considered continuum. We equate this to identifying when a set of particles' discrete velocities appear to come from a drifted Maxwellian distribution. That is, there is an isotropic Gaussian-distributed fluctuation in velocities about some mean flow rate. We perform this identification by applying a statistical test. The remaining details of the statistical testing are presented in section 4 – it is one of the primary tasks we addressed during the LDRD.

2.2 Switching from continuum to non-continuum

Switching a cell's description from continuum to non-continuum is straightforward. Given a macroscopic cell species density, ρ, species mean velocity, $\mathbf{v}$, and species temperature, T, we sample n particles from the drifted Maxwellian distribution,

$$\mathbf{v}_i = \mathbf{v} + \left(N\left(0, \sqrt{\frac{kT}{m}}\right), N\left(0, \sqrt{\frac{kT}{m}}\right), N\left(0, \sqrt{\frac{kT}{m}}\right) \right)$$

where $\mathbf{v}_i$ is the velocity of the i^{th} particle, $i = 1, \ldots, n$, k is the Boltzmann constant, m is the mass of the particle being introduced, and $N(\mu, \sigma)$ represents a normal distribution with mean μ and standard deviation σ. Particle locations are chosen to be uniformly distributed within the cell. The number of particles created, n, is chosen to conserve mass,

$$n = \lfloor \rho\Omega \rfloor + \begin{cases} 1 & \text{if } U[0,1] < \lfloor \rho\Omega \rfloor - \rho\Omega \\ 0 & \text{otherwise} \end{cases}$$

where Ω is the volume of the cell in question, $\lfloor x \rfloor$ is the largest integer smaller than x (the "floor" operator), and $U[a,b]$ represents a random variable from the uniform distribution with endpoints a and b.

2.3 Switching from non-continuum to continuum

Switching a cell's description from non-continuum to continuum is even simpler than the reverse. The macroscopic values of density, velocity, and temperature are found by taking ensemble averages over all particles within the cell,

$$\rho = \frac{nm}{\Omega}$$

$$\mathbf{v} = \frac{1}{n}\sum \mathbf{v}_i$$

$$T = \sqrt{\frac{\sum m\left(\mathbf{v}_i - \mathbf{v}\right)}{3nk}}$$

where the sums are over all particles within the cell.

2.4 Cell switching algorithm

As stated earlier, we wish to simulate a gas or plasma that transitions from a vacuum condition to one that is collisional enough in one or more species in one or more spatial regions to warrant switching the model description to a continuum one (for those species in those locations). For this discussion, we presume we have a working statistical test (as described in section 4).

Each cell has a designation, by species, for two pieces of information: (a) the current model description type (continuum or non-continuum), and (b) whether the velocity distribution can be described as Maxwellian (i.e., is, or is convertible to, a continuum description). Because we assume all cells described with a continuum model can in fact be described with one (e.g., we are not assessing when a continuum description should be switched to a non-continuum one), we have a total of three possibilities for each cell and species:

NC-NC:	A non-continuum description and non-continuum collisionality.
NC-C:	A non-continuum description and continuum collisionality.
C:	A continuum description and continuum collisionality.

Our methodology will enclose regions of type C with regions of type NC-C. Surrounding a continuum region with cells that have continuum characteristics, but have non-continuum descriptions, will allow us to match quantities at the boundary in a more stable way. We disallow, by construction, the possibility of a cell with a continuum description being adjacent to a cell with non-continuum collisionality. The cost incurred is managing a potentially high number of NC-C cells; these are cells that are being modeled with a non-continuum description (e.g., DSMC particles) but are in fact collisional enough to be modeled with continuum methods if considered in isolation. However, it is expected that this region scales as one dimension

smaller than the general problem description (i.e., it will be a surface for a three-dimensional simulation), and so the cost should not be prohibitive.

Given an additional constraint that a cell will be considered NC-C instead of NC-NC only if it satisfies our statistical test over some test period of time builds confidence that we can take time-averaged ensemble averages of particle moments and consider them bona fide macroscopic descriptions of that cell's state. For example, we only consider a mean particle velocity useful to a macroscopic continuum boundary matching process if that cell has been considered continuum for some time.

We only allow a cell to convert from type NC-C to type C if all of its neighbors are NC-C or C, and possibly all of its neighbors' neighbors (and so on). The connectivity distance is expected to be a parameter within the method: requiring larger patches around the candidate cell to be NC-C or C should lead to a more robust, yet more expensive, overall algorithm.

An NC-C cell can easily become NC-NC (a failure of the statistical test, say). We rely on this to drive cells of type C to cells of type NC-C via a neighbor distance requirement similar to the one for switching an NC-C cell to a C cell. That is, if some neighbor within a prescribed distance switches from NC-C to NC-NC, then this C cell becomes an NC-C cell (and the continuum model description is replaced with a non-continuum one). In reality, this latter neighbor distance should be smaller than the prior one to avoid "cell model thrashing", where some set of cells oscillate rapidly between types NC-C and C due to stochastic "noise" at the NC-NC/NC-C interface further away.

To summarize the switching possibilities:

	NC-NC	NC-C	C
NC-NC to	no-op	pass statistical test sufficiently many times	disallowed by construction
NC-C to	fail statistical test	no-op	all neighbor cells within prescribed connectivity are NC-C or C
C to	disallowed by construction	some neighbor cell within prescribed distance becomes NC-NC	no-op

Table 1: Cell type transitions.

2.5 Boundary conditions

We provide dual descriptions to at least the first layer of C cells at the NC-C/C interface. These are re-seeded each timestep with a non-continuum description according to the continuum macroscopic quantities, to act as a boundary to the adjacent NC-C cells. This allows particles to pass between non-continuum cells where we care about the non-continuum description (NC-C cells); we do not need to implement specialized boundary conditions for the non-continuum

description cells. If a particle leaves this dual non-continuum cell into an adjacent continuum cell, we simply throw it away. Particles from the NC-C cell should not transit the entire adjacent C cell (with non-continuum dual description) because of kinetic method time step constraints (e.g., for DSMC, a particle should only be capable of moving approximately 1/3 of a cell in one timestep). Larger patches of dual cells can be used to relax the constraint of re-seeding every timestep.

Although this dual-description method satisfies boundary condition requirements for the non-continuum method, we have not yet described how to supply boundary conditions for the continuum one. Because we require some region of NC-C cells around the C cells, we derive a macroscopic boundary condition from the NC-C cells. The underlying mass, velocity, and temperature calculations may be quite noisy and require time-averaging and possibly spatial smoothing as well. However, as previously discussed, we can control the quality of the NC-C data by placing stronger requirements on the NC-C designation vs. the NC-NC one (e.g., must pass the statistical test for multiple timesteps).

2.6 Full algorithm description

To summarize the algorithm:

```
for each timestep:
    for each NC-* cell:
        compute collisionality statistical test
        set the cell type to be NC-NC or NC-C according to the test
    for each NC-C cell:
        compute the minimum connectivity distance to an NC-NC cell
        if distance > neighbor connectivity distance for NC-C to C switching:
            switch cell from non-continuum to continuum description
    for each C cell:
        compute the minimum connectivity distance to an NC-NC cell
        if distance < neighbor connectivity distance for C to NC-C switching:
            switch cell from continuum to non-continuum description
    re-seed set of boundary C cells with non-continuum description
    advance non-continuum method
    derive boundary conditions for the C region from the adjacent NC-C cells
    advance continuum method
```

Note that we have avoided explicitly stating what non-continuum or continuum methods are employed. It is expected that the above algorithm would be compatible with all, or most, combinations. Some simplifications can be gained by using one or another method. For example, a finite element method can use the derived boundary conditions at the nodes and surfaces at the C/NC-C interface. Similarly, a finite difference method may use ghost cells that are in the adjacent NC-C region.

3. RESULTS

3.1 One-dimensional static boundary results

We constructed a one-dimensional test case to couple a non-continuum DSMC description to a continuum FDM description. The one-dimensional domain was static – we were unable to achieve enough progress to implement the dynamic switching methods, as described in section 2.4, although we did use the interface coupling strategies described in section 2.5.

The DSMC solution technique was the basic unmodified DSMC method. The only additional step in its solution method was to re-seed the adjacent continuum cells with particles. The FDM method was a standard MacCormack scheme with optional artificial viscosity.

Two test problems were investigated. The first was to open one end of a one-dimensional domain containing air (at 1 atm and 300 K) to vacuum and track the resulting rarefaction wave as it travelled through both model descriptions (non-continuum DSMC and continuum FDM). The other was to track a Mach 4 impulse through a one-dimensional domain containing air at the same conditions.

Results for the rarefaction wave are shown in Figure 1-4 below. In all of these figures, the blue left half of the domain ($0 \leq x \leq 1$) is modeled with a non-continuum (DSMC) description, and the black right half of the domain ($1 \leq x \leq 2$) is modeled with a continuum (FDM) one. Starting at the upper left and moving clockwise, the graphs depict the density, x-velocity, temperature, results of a statistical test to assess how non-Maxwellian the particles x-velocity distribution is within a cell, the number of computational particles in each cell, and the particle phase space (x-velocity vs. x-position). The bottom row of graphs shows information that is only available within the solution that is non-continuum (DSMC). In addition, the statistical test results in the bottom half are not included in the rest of the document. For more information on what exactly they measure, please contact the authors.

It is important to determine what kind of continuum boundary condition to apply at the DSMC/FDM interface. In particular, when deriving the FDM boundary condition from the DSMC data at the interface, the particle moments must be sampled and time-averaged. Some massaging beyond that is sometimes also necessary. This is due to one (or more) of the solution characteristics alternating from pointing into the domain to pointing out of the domain, and vice versa, in a very noisy fashion. When the DSMC-derived velocity switches from subsonic to supersonic in and/or supersonic out, one or more boundary conditions in the FDM cannot be enforced. This was an active area of investigation when the project terminated. We were applying a stochastic boundary condition to a regular FDM problem to help us understand the issue.

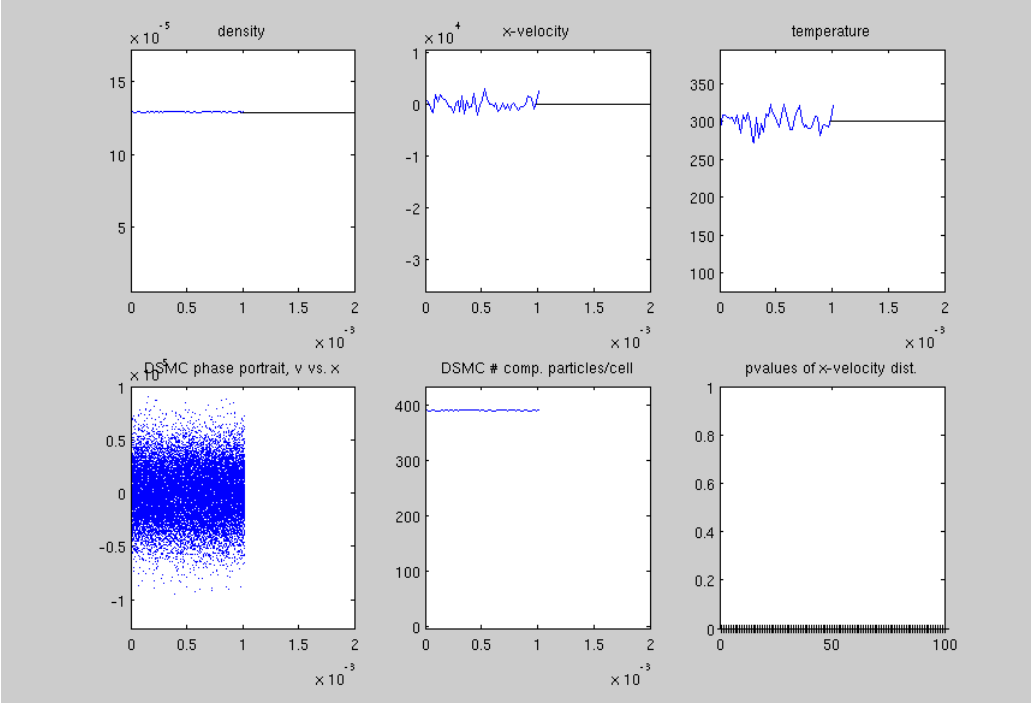

Figure 1: Initial conditions for hybrid rarefaction simulation.

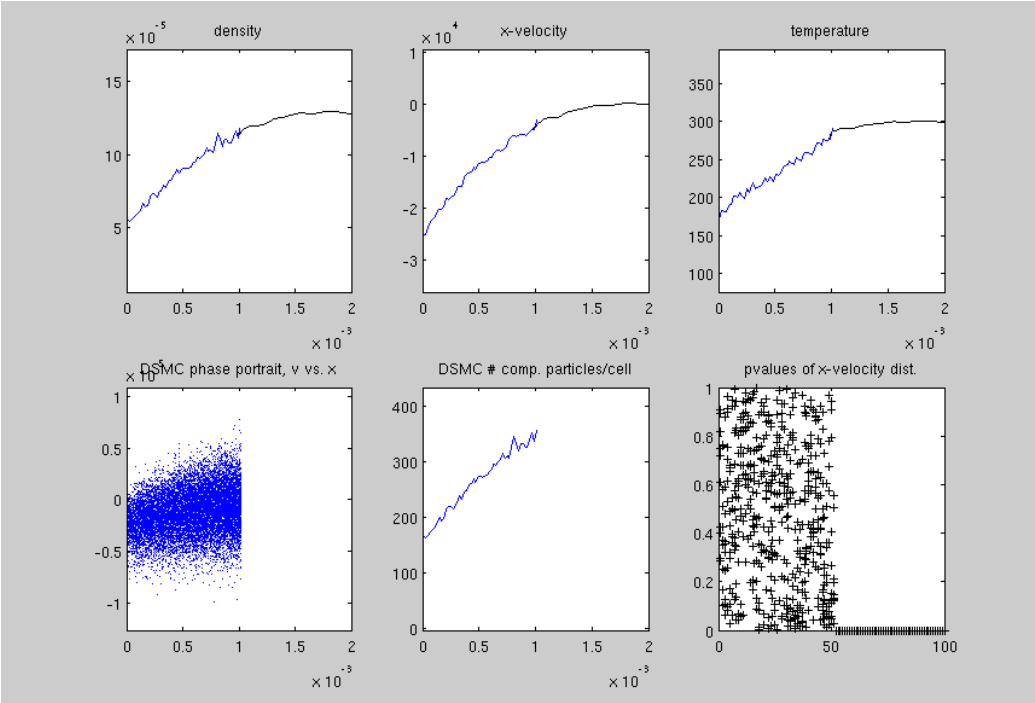

Figure 2: Rarefaction wave began at the left (*x*=0) and has just moved through the non-continuum DSMC region into the continuum FDM region.

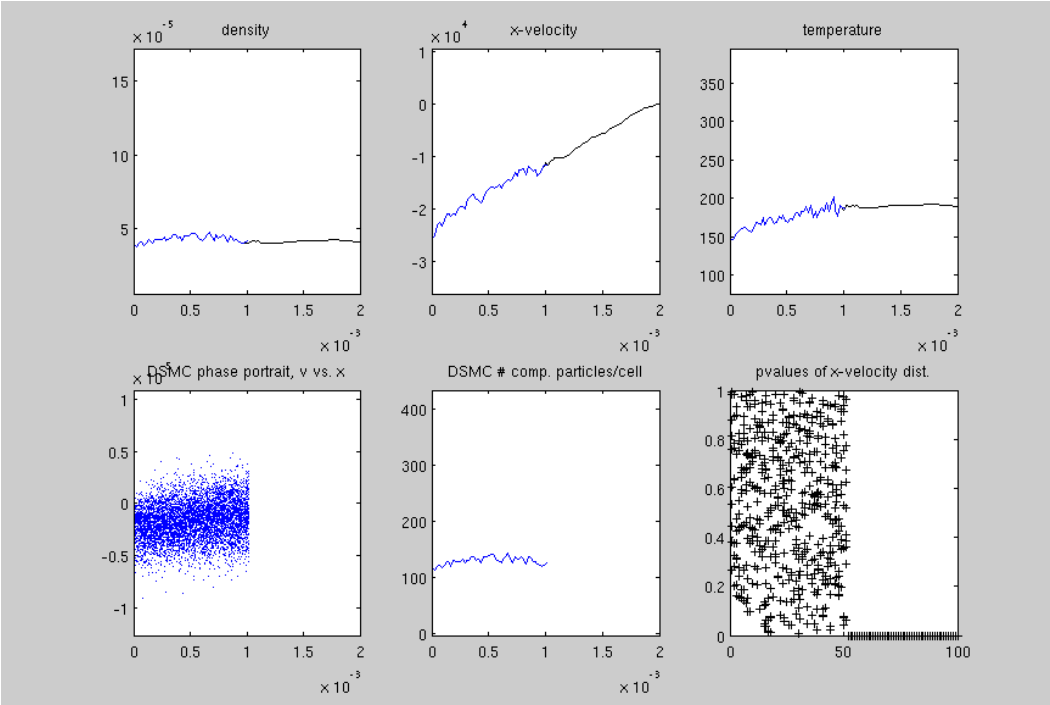

Figure 3: The rarefaction wave has reached the right boundary of the continuum region. Note the overall drop in pressure (density), and linear *x*-velocity profile.

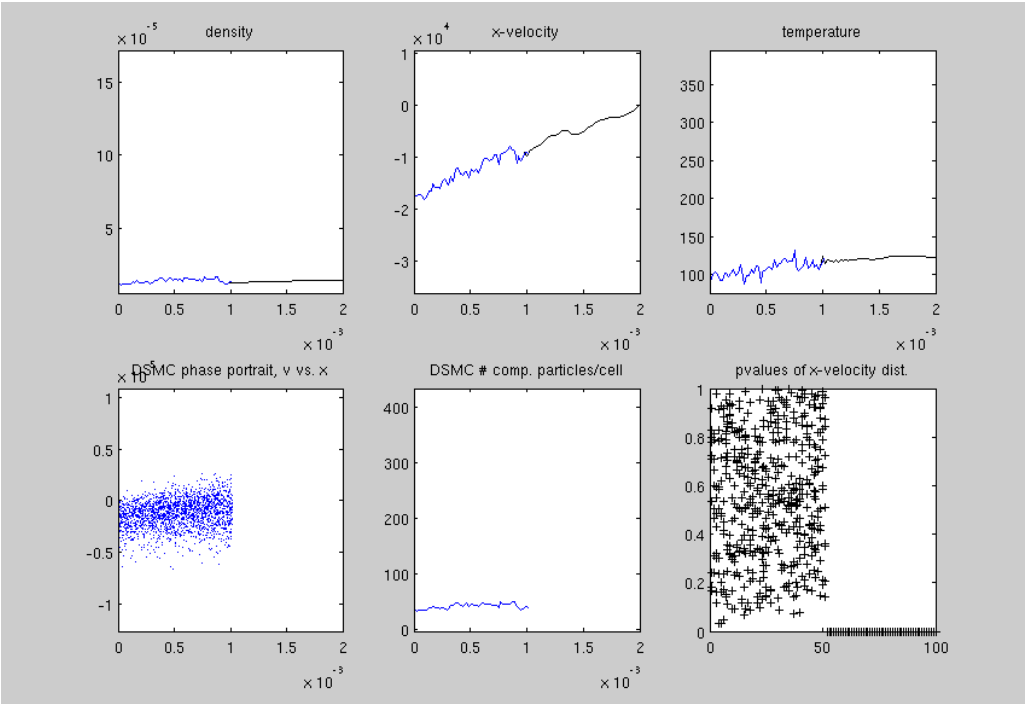

Figure 4: Near the end of the simulation, the density, velocity, and temperature moments are all converging to "vacuum".

3.2 Implementation within Quicksilver

We also implemented a Van Leer FDM within Quicksilver. We performed a short study of MacCormack methods vs. Van Leer ones, paying particular attention to the monotonicity characteristics. We were proceeding to assess how they would behave in the presence of noisy boundary conditions when the project was terminated.

4. STATISTICAL TESTING FOR EQUILIBRIUM

A mechanism for determining where to switch between particle and continuum representations is a key part of any hybrid scheme. Switching from a continuum to a particle description is reasonably straightforward. In these cases, switching is usually based on the Knudsen number:

$$Kn = \frac{\lambda}{L}$$

where λ is the mean free path and L is a characteristic length. The mean free path can be calculated as a function of macroscopic variables available in a continuum solver, such as pressure and temperature, and the characteristic length is often based on geometry or gradients that are also available. It is generally considered that $Kn < 0.01$ places the flow is comfortably in the continuum regime [7]. As Kn increases, the validity of the continuum model begins to degrade, so a non-continuum approach is adopted, generally between $Kn = 0.01$ and $Kn = 0.1$.

For vacuum-initiated problems, switching in the opposite direction (from particle to continuum description) is paramount. This is a much more difficult problem because recovering quantities from a stochastic field is complicated by statistical noise. In the current context, the quantity of interest is a distribution function because we are looking for an approach to the equilibrium, or Maxwellian, distribution. Unfortunately, it is difficult to construct a distribution function from a group of particles in the best of circumstances because each particle does not provide a point on the curve; it provides a single contribution to a sum that provides a single point on the curve. A great many particles are therefore necessary to construct a distribution of even marginal fidelity.

This difficulty is demonstrated in Figure 5, where a distribution with 201 points (-1000 to 1000 m/s in 10 m/s increments) is constructed from a varying number of samples from a Maxwellian velocity distribution for Argon at 273 K. It can be seen from this figure that the basic shape of the distribution is visible at 1000 samples, but it would be difficult to discern quantitative information from the resulting curve. At 10,000 samples, the curve is still very rough, but the structure is visible. Discerning departures from equilibrium, which are often subtle changes in shape, in this manner is therefore very difficult.

Fortunately, the problem of determining if a certain data sample follows a certain distribution occurs in diverse statistical applications. It is known generally as a “goodness-of-fit” (GOF) problem and numerous methods for treating it have been proposed in the literature. In the current case, we are aided greatly by the fact that we know which distribution we are testing against: the Maxwellian, or Normal distribution.

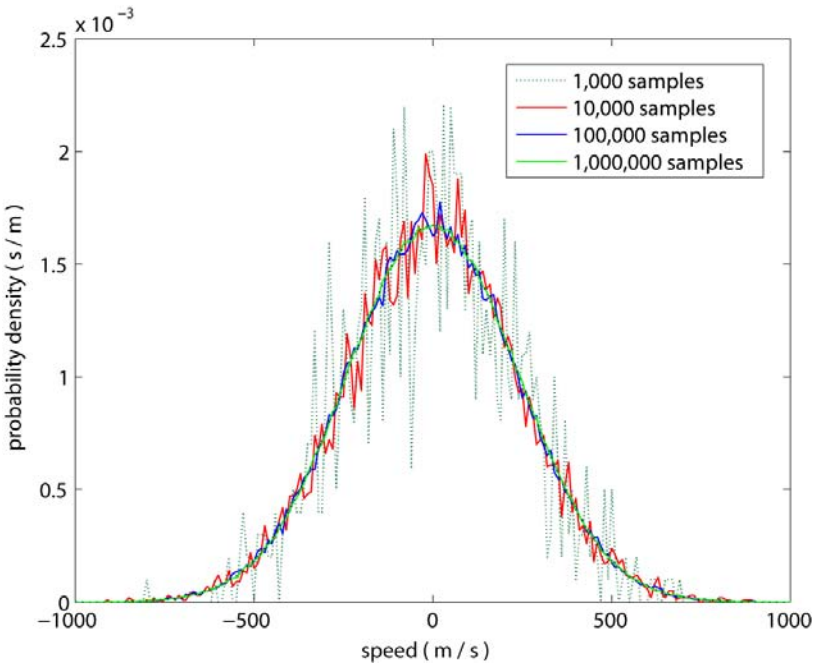

Figure 5: Equilibrium Maxwellian distribution recovered using a varying number of samples.

Several GOF tests were examined, including two Kolmogorov-Smirnov variants, Cramér-von Mises, Kuiper, Watson, and Anderson-Darling tests. A further test was constructed from a measure proposed for the purpose examined herein: looking for near-equilibrium conditions in particle-based fluid simulations.

In the literature, the power of a GOF test to identify conformity to the target distribution is generally tested by feeding it populations selected from other well-known distributions. For example, a test intended to detect the Normal distribution is given particles from an Exponential, Lognormal, or Uniform distribution. In a fluid simulation, observed distributions are almost always Maxwellian in general character, so this is an unrealistic method of assessment. Several alternatives are therefore explored in the current work.

As a first means of evaluation, the rejection rate (or the percentage of tests that concluded the distribution was nonequilibrium) was computed for an equilibrium distribution, as well as distributions with varying levels of nonequilibrium. Nonequilibrium in this case was provided by sampling half the trials for a given set from a distribution with the same mean but with a standard deviation multiplied by some factor. In a fluid, this is equivalent to a bimodal distribution with two underlying temperatures. All operations were performed in MATLAB.

The results of this evaluation for several tests detailed in [8] are shown in Figure 6. In this figure, the rate at which the test flags a nonequilibrium condition is normalized by the rate at which it does so for an equilibrium input. That is, some proportion of samples taken from the correct distribution will naturally fail the test due to occasionally generating an exceptional sample (i.e., outlier). We can set the sensitivity of our test to produce a prescribed proportion of natural failures. This proportion, termed a level of significance, is set to 5% in our work. With a nonequilibrium input, we would like to see a result significantly greater than unity, signifying

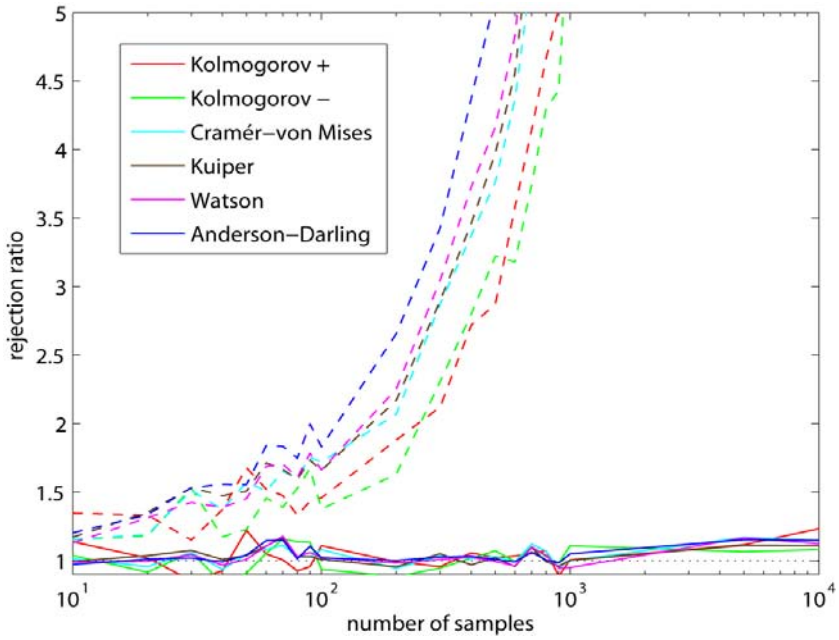

Figure 6: Comparison of sensitivities of several tests to nonequilibrium simulated by bimodal distributions with standard deviations that differ by 10% (solid lines) and 50% (dashed lines).

that the test can distinguish between equilibrium and nonequilibrium conditions. For the case where the standard deviations in the bimodal distribution differ by 10%, none of the tests show a rejection ratio reliably greater than unity for sets with fewer than a thousand samples, which translates to one thousand particles per cell in the current context if testing is to be done on an instantaneous basis. For the case where the standard deviations in the bimodal distribution differ by 50%, all the tests reliably flagged the departure from equilibrium, with the Anderson-Darling test appearing to be the most sensitive.

A literature search for a more sensitive measure of departure from equilibrium yielded several potential alternatives, the most promising of which was presented by Tiwari and Rjasanow [9] and is based on a Sobolev norm. Their method uses a more direct measure of distance in velocity space from an equilibrium distribution than the statistical tests evaluated previously. A usual Sobolev norm depends on the size of the function as well as the size(s) of its derivative(s). The measure suggested by Tiwari and Rjasanow can be expressed as:

$$S = \left[\left\| \partial_{w_N} \right\|_{H^{-2}}^2 - 2 \left\langle f_M, \partial_{\omega_N} \right\rangle_{H^{-2}} \right]^{1/2}$$

where

$$\left\| \partial_{w_N} \right\|_{H^{-2}}^2 = \frac{\pi^2}{N^2} \sum_{j=1}^{N} \sum_{l=1}^{N} e^{-|v_j - v_l|}$$

and

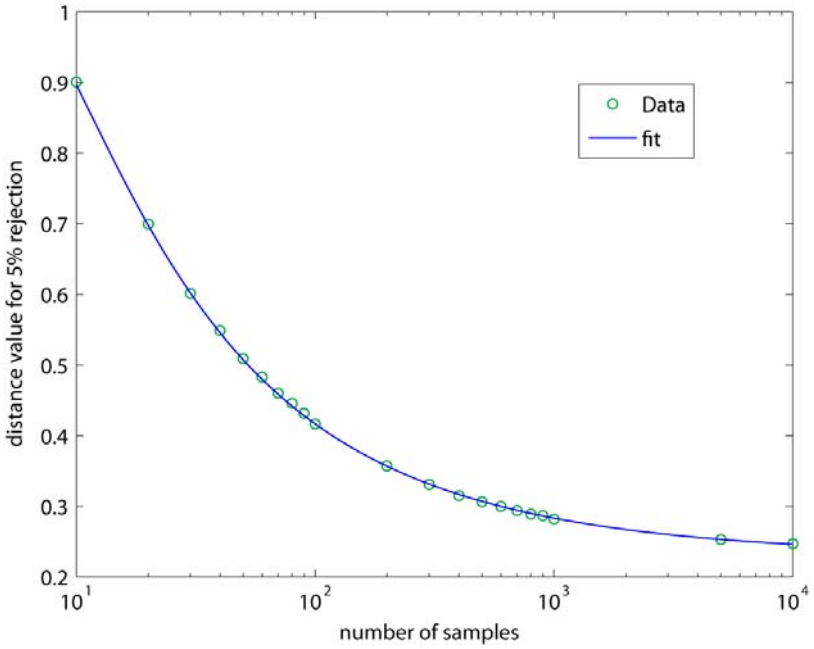

Figure 7: Sobolev norm-based distance measure for 5% false positive as a function of number of samples

$$\left\langle f_M, \partial_{\omega_N} \right\rangle_{H^{-2}} = \sqrt{e}\,\frac{\pi^2}{2N}\sum_{j=1}^{N}\left[e^{-|v_j|}\left(1-\frac{1}{|v_j|}\right)erfc\left(\frac{1-|v_j|}{\sqrt{2}}\right) + e^{|v_j|}\left(1+\frac{1}{|v_j|}\right)erfc\left(\frac{1+|v_j|}{\sqrt{2}}\right)\right].$$

where the v_i's are the discrete samples (velocities), $i = 1, ..., N$. This derivation applies directly to one-dimensional velocity space; extensions to higher dimensions are straightforward.

In order to perform a head-to-head comparison with the methods evaluated previously, this measure must be converted into a statistical test. To find the threshold in S representing the 5% level of significance used for the previous tests, ten thousand ensembles were sampled from the equilibrium distribution with a given number of particles and sorted in ascending order by their corresponding distance measure. The value at the index 5% below the maximum (sample number 9500) then represents a good estimate of the S distance value corresponding to a 5% rejection rate at that sample size, N. This process was repeated for a varying number of particles and a fit was performed to provide a continuous function:

$$S_{crit_{p=0.05}} = 0.23 + 2.8/N + 1.6/\sqrt{N} - 11.9/N^2.$$

The data and the fit are presented in Figure 7.

The behavior of this measure, acting as a test, is compared to the best-performing of the previous tests in Figure 8. From this figure, it is clear that the new measure is more sensitive. For the 10% nonequilibrium case, a clear deviation from unity is visible at 100 samples, whereas the Anderson Darling test did not show clear results until the sample size was greater than 1000. For

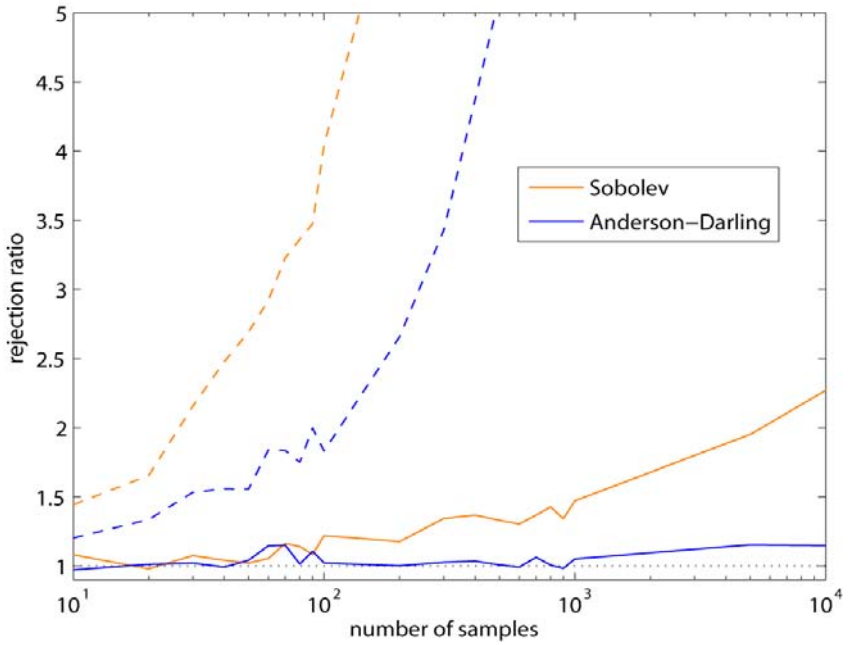

Figure 8: Comparison of Sobolev norm-based test to results from Anderson-Darling test shown in Figure 6.

the 50% nonequilibrium case, the difference is clear at all sample counts, with the Sobolev norm rejecting twice as many cases at only ten particles per cell than the Anderson-Darling test.

In order to assess these tests on a distribution from a real fluid problem, direct-simulation Monte Carlo simulations were performed on the Fourier problem using the DSMC1 code of Bird [3]. The Fourier problem, whose convergence to theory for DSMC calculations was shown in [10], consists of a quiescent gas between two plates held at differing temperatures. In this work, the walls were 1mm apart and were held at 323 K and 223 K. The intervening space was filled with Argon and the Knudsen number was controlled by varying the pressure to change the mean free path, which was computed via the hard-sphere expression [3]:

$$\lambda = \frac{16}{5}\left(\frac{m}{2\pi kT}\right)^{1/2}\frac{\mu}{p}$$

where m is the molecular mass, k is the Boltzmann constant, μ is the viscosity, and p is the pressure. The variable hard sphere (VHS) molecular model, with parameters taken from Appendix A of Bird [3], was employed in all DSMC calculations presented in this statistical test creation work.

The distribution in each sub-mean free path cell is expected to approach equilibrium at the local temperature as the Knudsen number tends to zero. As noted previously, the delineation between continuum and non-continuum flow is commonly set between $0.01 < Kn < 0.1$, so we would like to see the tests begin to fail in this vicinity. For the Knudsen number in this problem, we will define the characteristic length, L, based on the temperature gradient according to:

$$L = \frac{T}{dT/dx}.$$

To construct data for the test, the velocities of all particles in the cell at the domain center were written to a file at each timestep, following attainment of a steady state, until a total of five million entries were recorded. The distribution functions recovered from this data are presented in Figure 9. It can be seen that *Kn* = 0.01, the distribution closely resembles the Maxwellian distribution at the mean temperature (it is actually the Chapman-Enskog distribution [11], but the difference between these distributions is subtle). At *Kn* = 10, however, the positive velocities approximate the Maxwellian distribution at the cold wall temperature while the negative

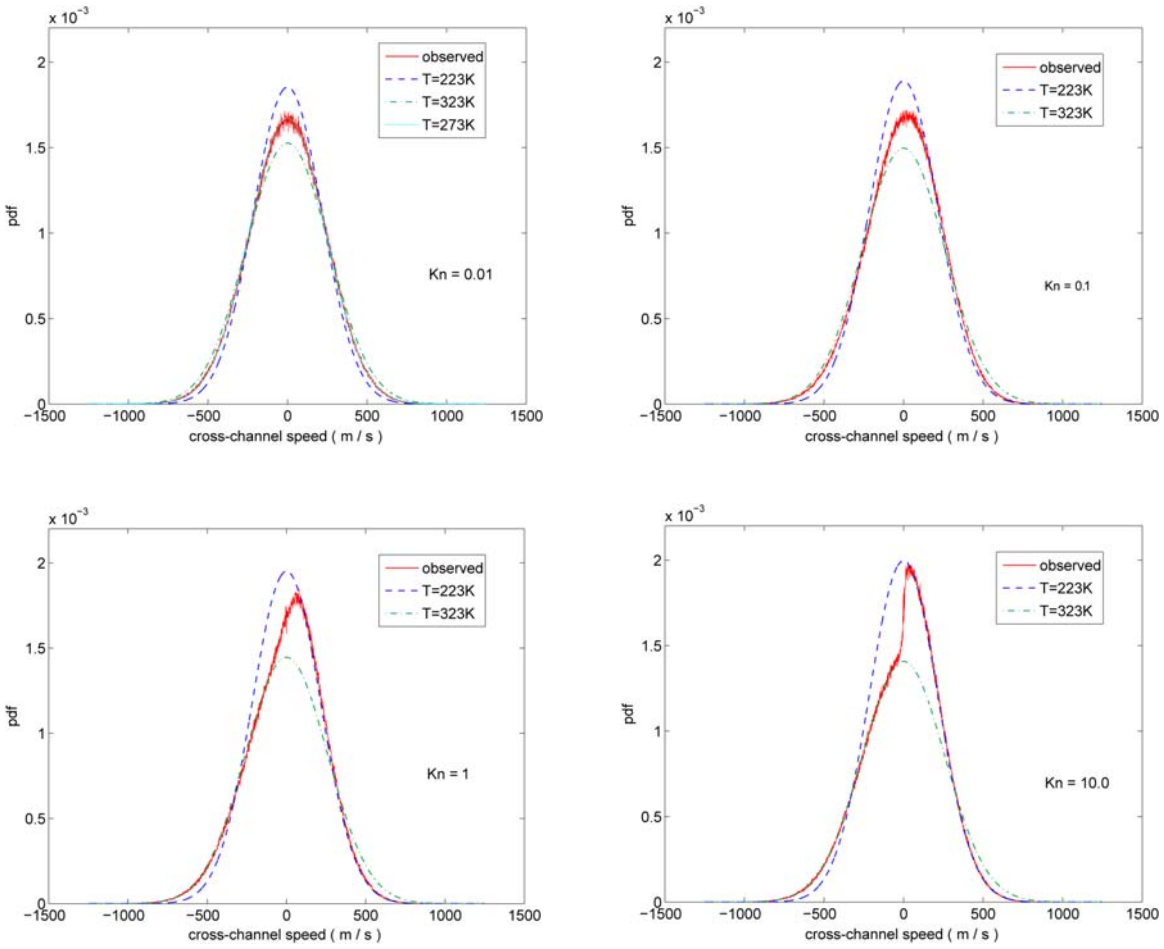

Figure 9: Cross-channel velocity distribution functions at various Knudsen numbers for the Fourier problem.

velocities approximate the Maxwellian distribution at the hot wall temperature. At intermediate Knudsen numbers, the flow transitions between these two extremes.

Running the Anderson-Darling and the Sobolev norm tests on these datasets, subdividing the particles to create 5,000 trials with 1,000 particles each, the rejection rate versus Knudsen

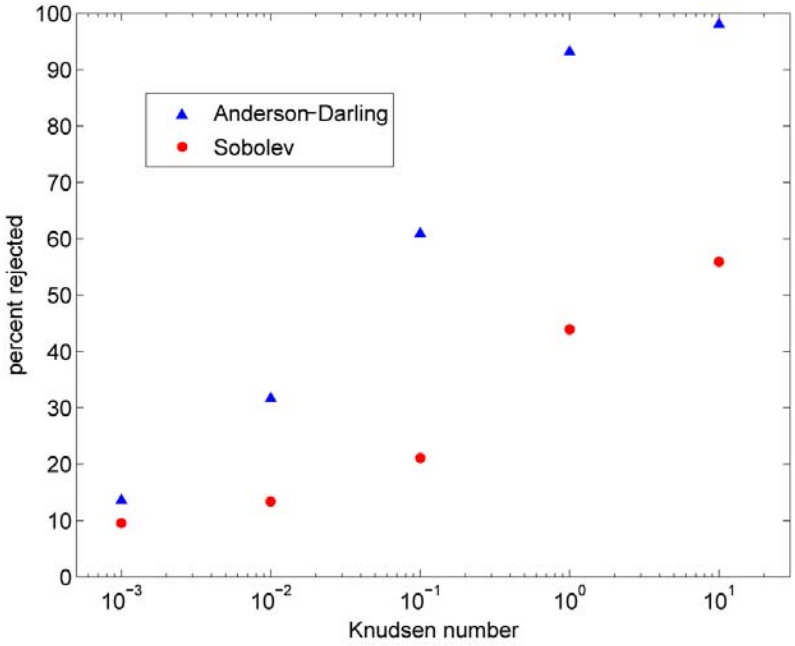

Figure 10: Rejection rate versus Knudsen number for Anderson-Darling and Sobolev-norm tests on Fourier problem.

number is computed and the results are presented in Figure 10. Several features are immediately evident from this figure. First, the general trend follows our expectation, with a low rejection rate at low *Kn* and a higher rejection rate at high *Kn*. Also, as we had hoped, a large increase in rejection rate occurs for both tests between *Kn* = 0.01 and *Kn* = 0.1. Two unexpected features are evident as well. First the Anderson-Darling test outperforms the Sobolev test, contradicting the results obtained on the manufactured nonequilibrium case. Second, neither test reaches its 5% rejection rate floor, even at *Kn* = 0.001.

To investigate these anomalies, DSMC1 was run on an equilibrium case (entire domain at a fixed temperature) and the rejection rate versus the number of samples in an ensemble was computed. This arrangement should yield a reasonably constant 5% regardless of the number of particles. Instead, an increasing rejection rate with particle count was observed. Further investigation showed that this behavior was due to duplicate velocities in the particle list. In the purely MATLAB investigation, this was not a problem because the distribution was sampled on a point-by-point basis and each entry was almost certainly unique. In the DSMC1 case, however, the points on the distributions are obtained by sampling particles passing through the cell. Slow-moving particles can therefore be sampled more than once. This issue is particularly serious in one-dimensional simulations, like those used in this evaluation, because the residence time of a particle depends only on one velocity component. In this case, it is very easy to sample the same particle multiple times.

For constructing a distribution, multiple samples are not problematic because the results are binned. The tests evaluated in this work function on a different level, however, and are apparently sensitive to the duplication. Removing the duplicates is not a resolution to this problem because the distribution would then be underpopulated at low speeds, which is also visible to the tests. The best solution is to make several timesteps between samples to allow slow

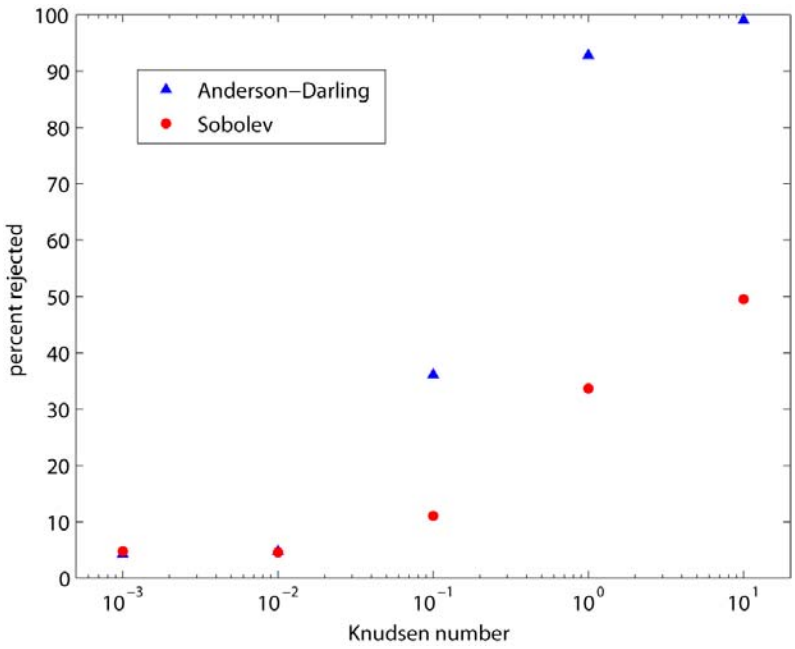

Figure 11: Repeat of rejection rate evaluation on Fourier problem with 100 steps between samples to avoid particle duplication.

particles to leave their cells. Earlier work [12] showed that approximately 25 steps between samples ensures that they are statistically independent. This consideration increases the desirability of finding a test that requires few particles to detect equilibrium because these samples may have to be accumulated over a great number of timesteps.

To examine the test results without the confounding effects of duplicate particles, the simulations were repeated with 100 timesteps between samples. The results of this effort are presented in Figure 11. The expected behavior is now seen at low Knudsen numbers, with the rejection rate of both tests approaching 5% at $Kn = 0.01$. Interestingly, the shape of the Sobolev norm-based test does not change significantly with the elimination of duplicates; the entire curve merely shifts downward. Conversely, the two highest points on the Anderson-Darling test are not affected by the change, but the points at and below $Kn = 0.1$ are strongly affected. This makes the test appear even more attractive, showing a very large jump in rejection rate between $Kn = 0.01$ and $Kn = 0.1$, precisely where non-continuum effects become important. The jump in the Sobolev norm-based test is considerably smaller, making the Anderson-Darling test clearly preferable in this case, still contradicting the ordering observed in the manufactured nonequilibrium test case.

The Couette problem is the final test case examined in this work. Similar in geometry to the Fourier problem, the Couette problem consists of parallel plates at the same temperature, but moving in opposite directions at (generally) the same speed. In this case, nonequilibrium at high Knudsen numbers is expected to be found in the velocity distribution in the direction parallel to the plate motion.

Once again, a length scale based on the gradient will be employed in the definition of Kn. In this case, however, the value of the forced quantity at the domain center cannot be used as a

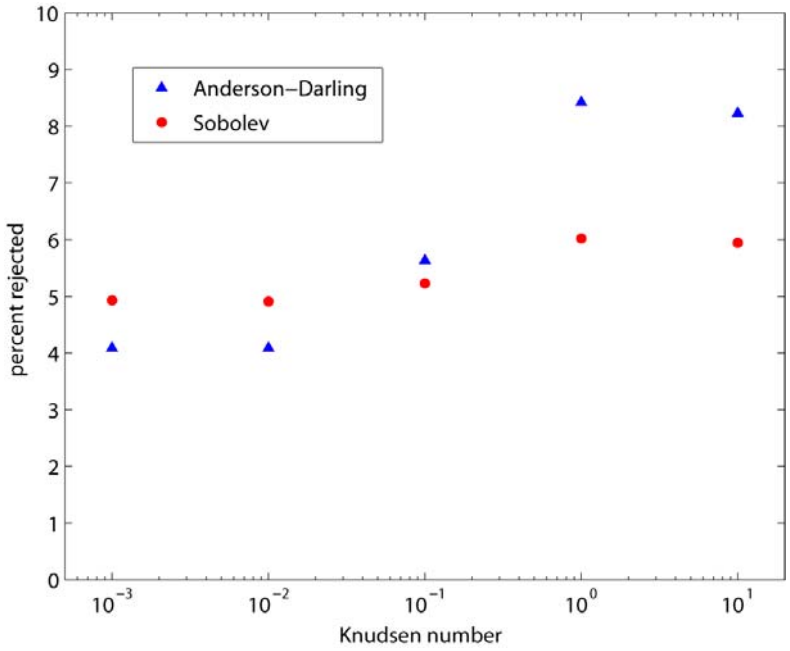

Figure 12: Rejection rate versus Knudsen number for the Couette problem.

normalizing factor because the mean velocity is zero. We therefore choose to use the most probable molecular speed [**Error! Bookmark not defined.**] at the reference temperature for this purpose, thus:

$$L = \frac{\sqrt{2RT}}{dV/dx}$$

where R is the gas constant, which is 208.24 J/kg/K for Argon. It may be noted that this definition leads to a Knudsen number that differs by an order unity constant from the shear stress Knudsen number often used in related problems.

In order to match L, and thus maintain the pressure-Kn relationship between the Fourier and Couette cases, the wall velocity was set to ±61.8 m/s. The results of these calculations are shown in Figure 12.

The first important feature visible in Figure 12 is the relative insensitivity of both tests to rarefaction in this problem, so much so that it was necessary to reduce the vertical scale by an order of magnitude to make the trends visible. The reduction in rejection rate between Kn = 1 and Kn = 10 provides a clue regarding the underlying issue. While it is expected that the strongly collisional (low Kn) cases would have near-equilibrium distributions in each cell, this particular problem also produces near-equilibrium velocity distributions in the free-molecular (high Kn) limit *in each direction*. The nonequilibrium nature of the flow is reflected in the fact that the velocity distributions in the directions parallel and perpendicular to the wall motion reflect different temperatures, defined as:

$$T = \sigma^2 / R$$

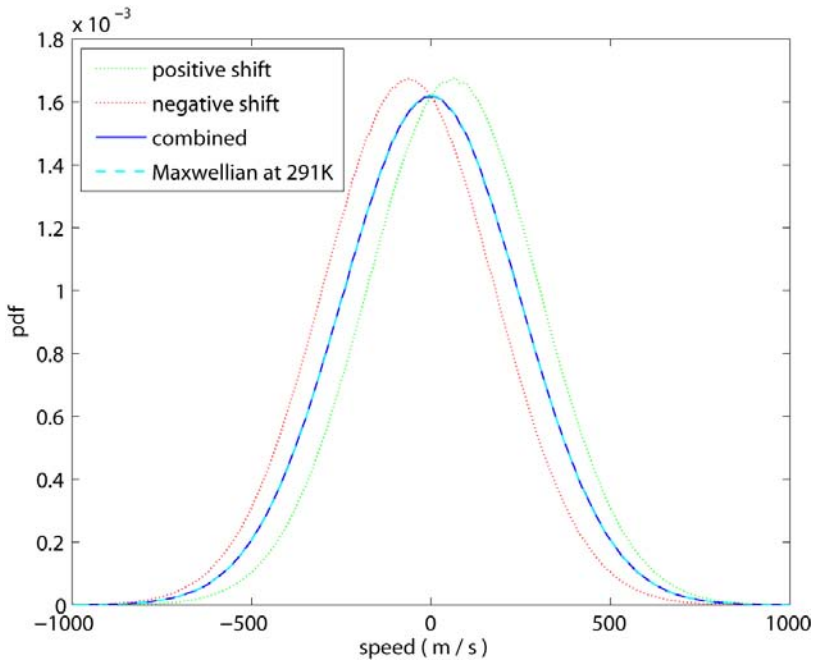

Figure 13: Construction of a free-molecular Couette flow distribution function in the direction parallel to wall movement at 273 K and its comparison to equilibrium at 291 K.

where σ is the standard deviation of the distribution. This behavior is demonstrated in Figure 13, where Maxwellian distributions at 273 K are shifted left and right by 61.8 m/s and then combined into a single distribution. The result is a Maxwellian distribution with temperature 291 K.

Because the two endpoints are indistinguishable from equilibrium, the distributions at midrange Knudsen numbers never depart far from equilibrium, thus the tests show a low rejection rate. It may be noted, however, that the Anderson-Darling test again appears to be more sensitive in this case than the Sobolev norm-based test. The latter test expressed in terms of velocity magnitudes in three dimensions, instead of the one-dimensional implementation employed herein, may be able to discern nonequilibrium of this type. Investigation of this possibility is recommended for future work.

5. REFERENCES

[1] R.W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles*, Institute of Physics Publishing, Bristol and Philadelphia, 1988.

[2] C. K. Birdsall and A. B. Langdon, *Plasma Physics via Computer Simulation*, Institute of Physics Publishing, Bristol and Philadelphia, 1991.

[3] G.A. Bird, *Molecular Dynamics and the Direct Simulation of Gas Flows*, Oxford University Press, Oxford, 1994.

[4] J. Denavit, "Time-Implicit Simulation of Particle-Fluid Systems," Space Science Reviews, **42**(1-2), 1985.

[5] P. Chapelle, J. P. Bellot, H. Duval, A. Jardy, and D. Ablitzer, "Modelling of plasma generation and expansion in a vacuum arc: application to the vacuum arc remelting process," J. Phys. D: Appl. Phys., **35**, 2002.

[6] J.-S. Wu, Y.-Y. Lian, G. Cheng, R. P. Koomullil, K.-C. Tseng, "Development and verification of a coupled DSMC-NS scheme using unstructured mesh," Journal of Computational Physics, **219**(2), 2006.

[7] S.A. Schaff and P.L. Chambre, Flow of rarefied gases, in *Fundamentals of Gas Dynamics*, chapter H, Princeton University Press, New Jersey, 1958.

[8] M.A. Stephens, "EDF statistics for goodness of fit and some comparisions," Journal of the American Statistical Association, **69**(347), 1974.

[9] S. Tiwari and S. Rjasanow, "Sobolev norm as a criterion of local thermal equilibrium," European Journal of Mechanics, B / Fluids, **16**(6), 1997.

[10] D. J. Rader, M. A. Gallis, J. R. Torczynski, and W. Wagner, "DSMC convergence behavior of the hard-sphere-gas thermal conductivity for Fourier heat flow," Phys. Fluids **18**, 077102, 2006.

[11] S. Chapman and T.G. Cowling, *The Mathematical Theory of Non-Uniform Gases*, Cambridge University Press, Cambridge, 1970.

[12] E.S. Piekos and M.A. Gallis. "Accelerating DSMC Data Extraction," Sandia National Laboratories Report, SAND2006-6692, 2006.

6. DISTRIBUTION

1	MS0346	E. S. Piekos	01513	
1	MS0346	D. J. Rader	01513	
1	MS0825	L. J. DeChant	01515	
1	MS0825	J. L. Payne	01515	
1	MS0836	M. M. Hopkins	01516	
1	MS0836	T. P. Hughes	01516	
1	MS0836	J. S. Lash	01510	
1	MS1152	T. D. Pointon	01654	
1	MS1181	L. X. Schneider	01650	
2	MS0899	Technical Library	09536	(electronic copy)
1	MS0123	D.L. Chavez, LDRD Office	01011	


Sandia National Laboratories