

Variance Reduction for a Discrete Velocity Gas

A. B. Morris, P. L. Varghese, and D. B. Goldstein

ASE-EM Department University of Texas at Austin, 1 University Station, C0600, Austin, TX 78712, USA

Abstract. We extend a variance reduction technique developed by Baker and Hadjiconstantinou [1] to a discrete velocity gas. In our previous work, the collision integral was evaluated by importance sampling of collision partners [2]. Significant computational effort may be wasted by evaluating the collision integral in regions where the flow is in equilibrium. In the current approach, substantial computational savings are obtained by only solving for the deviations from equilibrium. In the near continuum regime, the deviations from equilibrium are small and low noise evaluation of the collision integral can be achieved with very coarse statistical sampling. Spatially homogenous relaxation of the Bobylev-Krook-Wu distribution [3,4], was used as a test case to verify that the method predicts the correct evolution of a highly non-equilibrium distribution to equilibrium. When variance reduction is not used, the noise causes the entropy to undershoot, but the method with variance reduction matches the analytic curve for the same number of collisions. We then extend the work to travelling shock waves and compare the accuracy and computational savings of the variance reduction method to DSMC over Mach numbers ranging from 1.2 to 10.

Keywords: Boltzmann equation, variance reduction, discrete velocity scheme, Monte Carlo method

PACS: 47.11.-j, 47.45.-n, 47.70.Nd

INTRODUCTION

Stochastic solution of the Boltzmann equation collision integral can solve non-equilibrium gas flows efficiently, but noise inherent to the scheme may be unacceptable for problems where ensemble averaging is not feasible. For many flow problems of interest, the deviations from equilibrium may be small and one can substantially reduce statistical scatter by operating on the deviations from equilibrium alone [1]. This variance reduction technique has been applied to particle methods such as LVDSMC and VRDSMC for low Mach number flows [5,6]. Here we apply this variance reduction technique to a stochastic discrete velocity gas. This approach allows for rapid evaluation of the collision integral when the deviations from equilibrium are small. Traditional DSMC has difficulty simulating flows that span the rarefied to continuum regimes because of the enormous computational cost in the continuum regime. The current method allows for efficient simulation of such problems because the flow is near equilibrium in the continuum regime and the collision integral can be evaluated rapidly. In our discrete velocity framework, the distribution function may be viewed as a set of delta functions whose magnitude represents the population of molecules with a specific velocity. To best evaluate the collision integral, collisions between all points in velocity space with all other points in velocity space should be performed, but for many problems this is not practical. Significant savings are obtained by only performing collisions with a set of randomly chosen velocity pairs, such that on the average the collision integral is computed exactly. To maximize efficiency, importance sampling is used to select velocity pairs that are more likely to participate in a collision. This is similar to how DSMC selects collision pairs. Unlike DSMC, the user can specify the number of collision partners to sample in the discrete velocity model. This flexibility allows for optimization of computer resources because expensive computations of the collision integral can be restricted to regions where it is necessary.

METHOD

This work is restricted to single species monatomic gas flows with no body forces but there is no implied loss of generality. Further, the differential cross section is assumed to be isotropic, so integration over all scattering angles yields the total cross section, σ_T . Under these simplifications, the non-dimensional Boltzmann equation reduces to

$$\frac{\partial \hat{\phi}}{\partial \hat{t}} + \hat{\eta}_i \frac{\partial \hat{\phi}}{\partial \hat{x}_j} = \frac{1}{Kn} \int_{\hat{\zeta}_i} \left[\hat{\phi}(\hat{\zeta}_i') \hat{\phi}(\hat{\eta}_i') - \hat{\phi}(\hat{\zeta}_i) \hat{\phi}(\hat{\eta}_i) \right] \hat{g} \hat{\sigma}_T d\hat{\zeta}_i, \quad (1)$$

where the Knudsen number, Kn , is the ratio between a reference mean free path, $(n_r \sigma_r)^{-1}$ and a characteristic length scale and $\hat{\phi}$ is the scaled density weighted distribution function. The velocities are scaled by a reference velocity $\sqrt{2kT_r/m}$ and are denoted by $\hat{\eta}$ and $\hat{\zeta}$, where the primes indicate post-collision velocities, and \hat{g} is the relative velocity. Time is scaled by the mean collision time and the space variable \hat{x} is scaled by the reference mean free path. The hats over the variables indicate they are scaled. If the distribution function is split, $\hat{\phi} \equiv \hat{\phi}^{MB} + \hat{\phi}^d$, into equilibrium, $\hat{\phi}^{MB}$, and deviation from equilibrium, $\hat{\phi}^d$, components, then the collision integral, I_c , becomes:

$$\begin{aligned} I_c &= \int_{\hat{\zeta}_i} 2 \left[\hat{\phi}^d(\hat{\zeta}_i') \hat{\phi}^{MB}(\hat{\eta}_i') - \hat{\phi}^d(\hat{\zeta}_i) \hat{\phi}^{MB}(\hat{\eta}_i) \right] \hat{g} \hat{\sigma}_T d\hat{\zeta}_i + \int_{\hat{\zeta}_i} \left[\hat{\phi}^d(\hat{\zeta}_i') \hat{\phi}^d(\hat{\eta}_i') - \hat{\phi}^d(\hat{\zeta}_i) \hat{\phi}^d(\hat{\eta}_i) \right] \hat{g} \hat{\sigma}_T d\hat{\zeta}_i \\ &+ \int_{\hat{\zeta}_i} \left[\hat{\phi}^{MB}(\hat{\zeta}_i') \hat{\phi}^{MB}(\hat{\eta}_i') - \hat{\phi}^{MB}(\hat{\zeta}_i) \hat{\phi}^{MB}(\hat{\eta}_i) \right] \hat{g} \hat{\sigma}_T d\hat{\zeta}_i \\ &= I_{c1} + I_{c2} + 0. \end{aligned} \quad (2)$$

The last integral with $\hat{\phi}^{MB} \hat{\phi}^{MB}$ is identically zero because collisions do not change the equilibrium distribution. The integral involving $\hat{\phi}^{MB} \hat{\phi}^d$ terms is the dominant contributor to the collision integral and will be referred to as integral I_{c1} throughout the remainder of this paper. The integral involving $\hat{\phi}^d \hat{\phi}^d$ typically contributes less to the collision integral, especially when the flow is near equilibrium, and will be referred to as integral I_{c2} . To solve the collision integral the velocity space is uniformly discretized by a scaled velocity spacing, β , and truncated where the distribution function becomes vanishingly small. The distribution function is constructed by a set of delta functions located at each node in velocity space and the integrations can be replaced by summations over delta functions.

Evaluating the Collision Integral

In the present numerical scheme, a variable hard sphere, VHS, model is used. The collision cross section is related to the relative velocity by $\hat{\sigma}_T = \hat{g}^{-\kappa}$, where κ is typically fractional. When $\kappa = 0$ the collision cross section reduces to a hard sphere model and when $\kappa = 1$ the collision cross section corresponds to Maxwell molecules. In the subsequent numerical cases Maxwell molecules are used. The collision integral is most accurately solved by considering collisions between all velocity pairs, but this procedure is very expensive and of order N_v^2 , where N_v is the number of points in velocity space. Alternatively, the collision integral can be evaluated statistically by randomly sampling a set of collision pairs. In [2], the collision integral was solved in a quasi-stochastic manner because collision pairs were created by combining every point in velocity space with randomly sampled partners. In the present work the collision integral is evaluated by using importance sampling to select both velocities as in [1]; DSMC effectively also does the same thing. Collision partners for I_{c1} are obtained by sampling one velocity from the equilibrium distribution and the other from the deviation distribution. For I_{c2} both collision partners are obtained by sampling from the deviation distribution. The probability, P , of a velocity being selected from a given distribution, in Eq. 3, is based on the magnitude of the value of the distribution function because parts of the deviation distribution are always negative:

$$P = \frac{|\hat{\phi}(\hat{\eta})|}{\sum_{\hat{\zeta}} |\hat{\phi}(\hat{\zeta})|}. \quad (3)$$

In Fig. 1a an arbitrary non-equilibrium distribution is shown as the solid line and the corresponding equilibrium distribution is the dotted line. The equilibrium distribution is obtained by projecting a continuous Maxwellian with given macroscopic properties onto the discrete set of velocities. The projected Maxwellian distribution onto a coarse grid, $\beta > 1$, may not exactly conserve mass, momentum, and energy. This occurs because the distribution is largely composed of a single spike for coarse grids. In our framework, a single spike has zero temperature, so on extremely coarse grids the temperature tends towards zero. In less extreme cases, this can be accounted for by iterating on the input macroscopic properties until the desired moments are obtained. However, extremely coarse velocity space grids are not needed to keep computations tractable because of the very efficient computation scheme for the collision integral.

Velocities near the center of mass velocity are more likely to be sampled from the equilibrium distribution than those velocities near the fringes of the domain. In Fig. 1b, the deviation distribution is plotted and the most probable

deviation velocity is located where the magnitude of the deviation is the largest. Figure 1c shows the case when velocity $\hat{\zeta} = -2$ is selected from the deviational distribution and $\hat{\eta} = 0$ is selected from the equilibrium distribution. The depletion occurs at the pre-collision velocities $\hat{\zeta}$ and $\hat{\eta}$ and the replenishment occurs at the post-collision velocities denoted by primes and a Δ . For elastic collisions, the post-collision velocities are obtained by a random rotation of the relative velocity vector about the center of mass velocity. For simplicity the collision circle is drawn in two dimensions, but in general the velocity space is three dimensional and the post-collision velocities can lie anywhere on a collision sphere. It is possible for the pre- and post-collision velocities to be identical, and if this occurs a null collision is performed because the replenishing velocities are the same as the depleting velocities. This scheme has the flexibility to choose an arbitrary number of replenishing pairs to sample, but only one replenishing pair is sampled to balance the noise associated with the depleting and replenishing contributions to the collision integral. A conservative interpolation scheme, discussed in [7], can be used to map the floating post-collision velocities onto nearby grid points. In the current work, interpolation is avoided by simply restricting post-collision velocities to grid points. This approach is similar to previous work such as the Δ - ε method [8] and IDSMC [9].




FIGURE 1. (a) An arbitrary non-equilibrium distribution (solid line) and the corresponding equilibrium distribution (dashed line). (b) The deviation distribution. Parts of the deviation distribution are negative, lightly colored, when not in equilibrium. (c) A sample collision between velocities ζ and η results in post-collision velocities that lie at random grid nodes on the dashed collision circle, the triangles.

For Maxwell molecules the selection is simplified because $\hat{\sigma}_T \hat{g}$ is constant. The net depletion rate is computed by integrating the depleting terms in equation 2 over all velocities $\hat{\eta}$ and $\hat{\zeta}$. To satisfy the net depletion rate, once a collision pair is determined for Maxwell molecules, the depletion at each partner is:

$$\Delta \hat{\phi}_\eta = \Delta \hat{t} \frac{\text{sgn}(\hat{\phi}_\eta^d) \hat{n}^{MB} \hat{n}^d}{N_c \beta^3} + \Delta \hat{t} \frac{\text{sgn}(\hat{\phi}_\eta^d) \text{sgn}(\hat{\phi}_\zeta^d) \hat{n}^d \hat{n}^d}{2N_c \beta^3}; \hat{n}^d \equiv \beta^3 \sum |\hat{\phi}^d|, \hat{n}^{MB} \equiv \beta^3 \sum |\hat{\phi}^{MB}|; \text{sgn}(x) = \begin{cases} +1; & x > 0 \\ -1; & x < 0 \end{cases}, \quad (4)$$

where N_c is the number of collision pairs. For I_{c1} a depleting collision always drives the distribution function towards the equilibrium value: $\hat{\phi}_\eta^d > 0 \Rightarrow \Delta \hat{\phi}_\eta, \Delta \hat{\phi}_\zeta < 0$; $\hat{\phi}_\eta^d < 0 \Rightarrow \Delta \hat{\phi}_\eta, \Delta \hat{\phi}_\zeta > 0$. In Fig. 1c, $\hat{\phi}^d > 0$ at $\hat{\zeta} = -2$ and the depleting collision reduces $\hat{\phi}$ at each collision partner location in velocity space.

In this method the statistical fluctuations can be adjusted by selecting how many collision pairs, N_c , are sampled. Further, since I_{c2} generally is much less than I_{c1} it is computationally efficient to sample fewer pairs for I_{c2} . In general the number of collision pairs to sample is arbitrary, but in the current work we balance the statistical noise in each integral by choosing the number of sampled collision pairs proportional to \hat{n}^d for I_{c1} and $\hat{n}^d \hat{n}^d$ for I_{c2} , Eq 5:

$$N_{c,1} = \Delta \hat{t} \hat{n}^d / r^2 \beta^3; \quad N_{c,2} = \Delta \hat{t} \hat{n}^d \hat{n}^d / r^2 2 \hat{n}^{MB} \beta^3. \quad (5)$$

The variable r is a user input that is proportional to the rms noise in the distribution function resulting from statistical evaluation of the collision integral. The timestep and velocity spacing are included in Eq. 5 such that the noise in the solution is independent of timestep and discretization. The ability to adjust how coarsely to evaluate the collision integral allows for additional levels of flexibility not available to DSMC. Since in general $\hat{n}^d \ll \hat{n}^{MB}$ the number of samples for I_{c2} is much less than I_{c1} and when the flow is near equilibrium very few collisions are computed.

NUMERICAL RESULTS

The above scheme has been implemented and applied to a homogeneous relaxation test case as well as a 1-D normal shock. Results are compared to a discrete velocity method where both collision partners are randomly sampled but variance reduction is not used, and to DSMC. The Bobylev-Krook-Wu (BKW) distribution has a highly non-equilibrium initial condition and a known analytic relaxation to equilibrium for Maxwell molecules [3,4]. In Fig. 2, relaxation of the BKW distribution is computed with and without variance reduction. In both simulations, the same number of simulated collisions were performed. The velocity grid was discretized with a velocity spacing $\beta = 0.5$ and extended to $\hat{\eta} = \pm 3.0$. The scaled time step was 0.2.

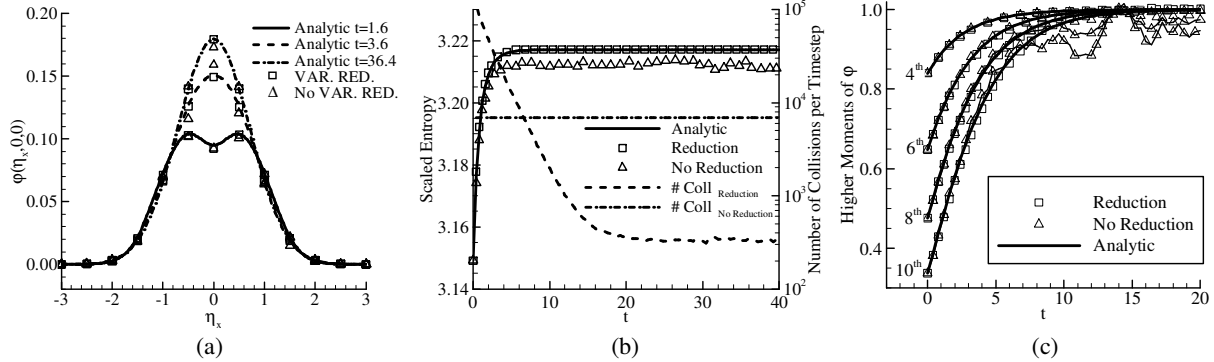


FIGURE 2. (a) Slices of the BKW distribution at different times. The square symbols are from the variance reduction method and the triangles are without variance reduction. (b) Time variation of entropy and number of collisions, note logarithmic scale for collisions. (c) Relaxation of the 4th, 6th, 8th, and 10th moments of the distribution function.

Figure 2a shows excellent agreement between the distribution function computed with the variance reduction method and the analytic solution. At early times, both discrete velocity methods have relatively little statistical scatter because the initial distribution is noise free. However, at later times the noise in the method without variance reduction becomes more noticeable. Figure 2b shows that statistical noise causes the entropy to undershoot when variance reduction is not used. As the solution approaches equilibrium, the number of simulated collisions per time step rapidly decreases because the number of collision pairs is proportional to the degree of non-equilibrium. After the distribution function equilibrates, the number of simulated collisions and RMS noise continue to decrease. This occurs because as the noise decreases the deviations from equilibrium are smaller and the magnitude of the collision integral is reduced and fewer collision pairs are required. The higher moments of the distribution are plotted for the two simulations with and without variance reduction in Fig. 2c. When variance reduction is used, the higher moments agree well with the analytic solution. However, for the same computational effort without variance reduction, the 8th and 10th moments do not track the analytical solution because of statistical noise.

Figure 3a shows the RMS in the equilibrium distribution plotted against the computational effort when a BKW distribution is relaxed with different input values of r . For $r > 0.1$, the RMS is larger than 10^{-2} and despite an increased number of simulated collisions, the noise does not decrease. This occurs because of an instability shown schematically in Fig. 3b. In this worst-case scenario, the entire depletion is distributed over very few points. If too much mass is depleted from certain points in velocity space the result is a distribution with greater non-equilibrium. In some cases large negative spikes appear in the distribution function. These large negative spikes are very likely to be selected to participate in further collisions, causing the distribution function to diverge further. If r is sufficiently small, e.g. $r < 0.2$, the fluctuations in the scheme damp very rapidly. As the computational load increases from 125 collisions per step to 200 collisions per step, the fluctuations about the equilibrium distribution caused by the stochastic computation of the collision integral decrease by approximately 3 orders of magnitude. This is a significant improvement to stochastic methods without variance reduction where the noise decays as $N_c^{-1/2}$. The noise decays very rapidly because as the solution approaches equilibrium, the deviations from equilibrium become smaller and the collision integral is reduced. Consequently, fewer collisions need to be sampled, as the noise continues to decrease.

All simulations with $r \geq 10^{-3}$ were computed on a velocity space that extended to $\hat{\eta} = \pm 3.0$. For this velocity space, when r is further decreased to $r = 2 \times 10^{-4}$ (the dotted line) the noise in the distribution is dominated by errors

associated with truncation of the high velocity tails of the distribution and velocity discretization. This occurs because the discrete equilibrium distribution is restricted to a set of truncated velocities while the continuous Maxwellian includes all velocities. The discrepancy between the discrete and continuous equilibrium distributions can be larger on coarse grids than the statistical noise, causing the RMS to asymptote despite increasing number of collisions. On velocity grids that extend further, the triangle, the errors between a discrete and continuous equilibrium distribution are lessened and the RMS continues to be dominated by statistical fluctuations.




FIGURE 3. (a) The noise about equilibrium for various numbers of sampled collisions. (b) Slices of the BKW distribution taken at three different times. This is an unstable case where too few simulated collisions are performed.

To verify the relaxation rate and accuracy, the scheme is now applied to 1-D normal shocks spanning a range of Mach numbers from 1.2 to 10. The advection step uses a 4th order finite difference scheme [8]. Figure 4 shows the *instantaneous* snapshots from our numerical scheme compared to *time-averaged* DSMC. In addition, the number of collisions is plotted to show the reduction in computational effort in areas near equilibrium.




FIGURE 4. Instantaneous snapshots of the instantaneous density and temperature profiles obtained by the variance reduction method (symbols) compared to time averaged DSMC solutions. The right axis shows the number of sampled collisions in the variance reduction method.

Excellent agreement is found between DSMC and the current scheme for Mach numbers ranging from 1.2 to 10, indicating the correct relaxation rate. For the Mach 1.2 shock wave, we used a scaled velocity discretization, $\beta = 0.5$, time step of $\Delta t = 0.05$, 100 physical cells with a spacing of $\Delta x = 1.0$, and $r = 4 \times 10^{-3}$. A Mach 2.0 shock wave spans a larger velocity space, so a larger $\beta = 0.7$ was used. Additionally, the number of spatial cells was doubled and the cell size was reduced to 0.5. For the Mach 10 shock wave, the velocity space is extremely large and in this calculation $\beta = 1.5$. The coarseness of the velocity grid on the cold region upstream of the Mach 10 shock wave causes slight noise in the temperature profile.

Statistical noise in the downstream density is plotted versus computational time in Fig. 5 in order to compare the performance between the current scheme and DSMC. For the Mach 1.2 shock wave, simulations were run to steady state with 100 cells, with a scaled physical spacing of $\Delta \hat{x} = 1$ and time step $\Delta \hat{t} = 0.05$. For the Mach 2.0 shock wave, 200 cells with a spacing of 0.5 and a time step of 0.05 were used. The symbols are obtained from the variance reduction results presented in Fig. 4.

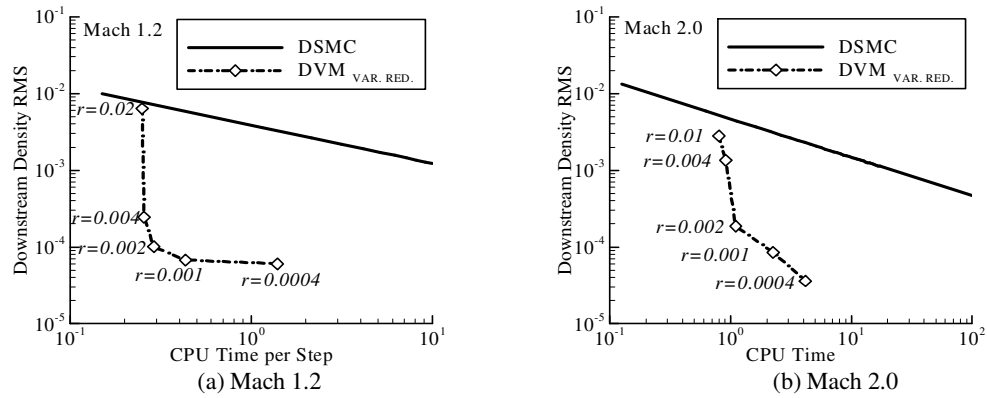


FIGURE 5. Noise in the downstream density versus computational time.

The RMS in the downstream density converges rapidly in the variance reduction method. The noise can be decreased by nearly two orders of magnitude for little additional computational work by decreasing r from 0.02 to 0.002. When $r = 0.002$ the present method is nearly 3 orders of magnitude faster than DSMC for the same statistical uncertainty in the Mach 1.2 simulation. When $r < 0.001$ the noise does not continue to decrease because of a truncated discrete velocity space. The velocity space extends to $\hat{\eta} = \pm 3.2$ for the Mach 1.2 shock wave where the reference temperature is the downstream temperature. For the Mach 2 shock wave the velocity space is larger and extends to $\hat{\eta} = \pm 4.0$ and the noise continues to decrease for very small r .

CONCLUSIONS

We have extended a variance reduction method developed by Baker and Hadjiconstantinou to a discrete velocity gas and have solved for the homogeneous relaxation of a BKW distribution and 1D normal shock wave. Excellent agreement in the entropy and reduction in computational cost is obtained for the BKW relaxation. As the distribution approaches equilibrium, the number of simulated collisions reduces by approximately 2 orders of magnitude. The rapid convergence towards equilibrium occurs because as the statistical fluctuations become smaller the non-equilibrium decreases and the collision integral is reduced. This scheme is vastly efficient for flows near equilibrium and is well suited for hybrid solvers where the interface between the continuum and DSMC regimes is often near equilibrium. Normal shock waves with strengths varying from Mach 1.2 to Mach 10.0 show excellent agreement with DSMC. For the same level of noise, the new scheme performed nearly 3 orders of magnitude faster than DSMC for the Mach 1.2 and Mach 2 shock waves.

ACKNOWLEDGMENTS

We gratefully acknowledge support from NASA under cooperative agreement NNX08AB27A.

REFERENCES

1. L. L. Baker, N. G. Hadjiconstantinou, *Physics of Fluids* **17**, 051703, 1-4, 2005.
2. A. B. Morris, P. L. Varghese, D. B. Goldstein, *Rarefied Gas Dynamics: Proc. of the 26th Intl. Symposium on Rarefied Gas Dynamics*, Kyoto, Japan, **1084**, pp. 458-463, 2008.
3. A. V. Bobylev, *Soviet Physics Doklady* **20**, 822-824 (1976).
4. M. Krook and T. T. Wu, *Physics of Fluids* **20**, 1589-1595 (1977).
5. T. M. M. Homolle, N.G. Hadjiconstantinou, *Journal of Computational Physics*, **226**, 2341-2358, 2007.
6. L. L. Baker, N. G. Hadjiconstantinou, *Journal of Theoretical and Computational Nanoscience* **5**, 165-174, 2008.
7. P. L. Varghese, *Rarefied Gas Dynamics: Proc. of the 25th International Symposium*, M. S. Ivanov and A. K. Rebrov, Editors, Novosibirsk, Russia, pp. 227-232, 2007.
8. Z. Tan and P. L. Varghese, *Journal of Computational Physics* **110**, pp. 327-340, 1994.
9. D. Goldstein, B. Sturtevant, J. E. Broadwell, *Rarefied Gas Dynamics: Proc. of the 16th International Symposium*, Pasadena, CA, pp. 100- 117, 1988.