JOURNAL OF COMPUTATIONAL PHYSICS 109, 30–36 (1993)

Reduction of Simulation Cost and Error for Particle Simulations of Rarefied Flows

M. A. FALLAVOLLITA,* D. BAGANOFF, AND J. D. MCDONALD†

Department of Aeronautics and Astronautics, Stanford University, Stanford, California 94305

Received August 3, 1992; revised February 25, 1993

The computational cost in simulating a steady, rarefied gas flow by use of a particle method is shown to be greatly reduced if the simulation employs an average number of particles per cell that is greater than a certain (approximate) minimum. The minimum value depends on the rms level of statistical fluctuations judged acceptable in the resultant data and is found to be significantly greater than values often used in practice when two- and three-dimensional flows are simulated. The computational cost is found to remain fixed for values greater than the minimum, showing that the ensemble- and time-averaging operations become interchangeable. For the problem studied, it is shown that a regime exists for which computational cost may be reduced by a factor of 10 by merely increasing the size of the simulation by a factor of five, when holding the rms level of statistical fluctuations fixed.

INTRODUCTION

When using a large collection of simulated particles to model a molecular flow on a computer, one is interested in knowing how many particles are needed in each small volume of space in order to properly model the relevant flow physics; and does the use of this number have a beneficial effect on computational cost? These important questions arise irrespective of the particular method used in the simulation, whether it is the method of molecular dynamics (Alder and Wainwright [1]), an approach used in simulating ionized gas motion (Hockney and Eastwood [2]), or the direct simulation Monte Carlo (DSMC) method employed in the study of rarefied gas flows (Bird [3]). In addition, it is clear that the number required to obtain a given level of simulation accuracy when dealing with a non-steady flow is far larger than that needed for a steady flow where time averaging may be permitted. Our discussion will focus on steady flows and the use of time averaging. In addition, Bird's DSMC method was selected for the study because the calculational effort grows roughly in proportion to the total number of particles N, for which the analysis to be presented is more straightforward, as opposed to N^2 or $N \log N$.

Beyond the requirement of steady flow, the use of time averaging to reduce the effect of statistical fluctuations, which are inherent in a particle method, is based on two assumptions: first, that a sufficient number of particles is in fact present in a computational cell to adequately model the physics of interest; and second, that a time average can be used to replace the cell average obtained if a still larger number of particles were used, or if a large number of repeated runs for the same conditions were carried out. If this replacement is permitted, then the concept of an ensemble average applies to this situation and its exchange with the time average leads to the assumption that the ergodic hypothesis holds. Use of the time average requires that random processes associated with a single cell in space at one time are statistically independent of those associated with the same cell at a different time. In other words, the time interval between samples is greater than the correlation time for the random quantity being averaged. If all these assumptions are valid, then one is able to employ a cell sample size given by $S_c = N_c T$, where N_c is the average number of particles in a single cell and T is the number of time steps used in the time averaging. Because the relative statistical error for an averaged quantity, defined by the ratio *rms/mean*, decreases as $S_c^{-1/2}$ for a statistically independent random process, one concludes that doubling T allows one to halve N_c, provided that N_c is initially large enough. Therefore, as long as the computational effort is simply proportional to S_c, which is true for Bird's DSMC method, then the simulation with the smaller N_c would be preferred because a smaller demand is placed on the amount of computer memory required, while the computational cost and modelling precision remain the same. Our objective is to fully explore these issues with regard to the DSMC method.

The act of doubling T and halving N_c is certainly limited because one would quickly arrive at a point where too few particles would be present in a cell to adequately model a physical gas flow; an obvious example is $N_c = 1$. Long

* Mailing address: NASA Ames Research Center, M/S 230-2, Moffett Field, CA 94035.

† Now at MasPar Computer Corp.

0021-9991/93 $5.00

before this point is reached, it is clear that the simulation would lose efficiency and longer and longer time averages would be required to attempt to obtain the same level of relative statistical error, thus driving up the computational cost. Because no theory exists to guide one in determining how many particles are needed to adequately model given fluid mechanical processes, to evaluate the corresponding computational cost, or to determine whether cost can be reduced by following a particular mode of operation, we conducted a series of numerical experiments using a modified version of the DSMC method to explore these questions. The basic approach followed was to repeatedly run the same simulation for a given problem while varying the duration of the time average and the total number of particles used in the simulation, and then collect the appropriate data to evaluate the level of statistical uncertainty present in the results.

SIMULATIONS

The representative problem chosen for study consisted of a two-dimensional rarefied flow past a flat plate placed normal to the oncoming stream as depicted in Figs. 1–3. The free stream Mach number was set at 8 and the Knudsen number, based on the plate height, was fixed at 0.1 to clearly place the flow in the transition regime. A unique characteristic of a rarefied gas flow is that the temperature field extends much further ahead of a blunt body than the density or pressure fields, which can be clearly seen by comparing Figs. 1–3. In order to limit the number of variables in this first study, the simulated gas chosen consisted of diatomic nitrogen with rotational nonequilibrium (collision number set to 5) but no vibrational nonequilibrium. The molecular collision cross section was modeled using Bird's variable hard-sphere model [4], where the value of the exponent in the inverse power force law was set to 10. The boundary condition on the flat plate was chosen as isothermal, with the plate temperature set equal to 7.5 times the free stream temperature, a value typical for high speed flight in the upper atmosphere. For particles contacting, the plate diffuse reflection was assumed. Our intention was to study a fairly straightforward rarefied flow so that the major effects of interest could be easily identified.

In order to properly explore the questions raised, one must have access to a method of simulation that has a very large dynamic range; otherwise the search for modelling limits would be thwarted by the limitations of the simulation itself. In addition to a large dynamic range, the simulation must be computationally efficient because very large simulations as well as small simulations must be fully explored. All our simulations were carried out on the Cray-YMP and made use of a highly vectorized code written by McDonald [5] which employs a specialized vectorization-compatible selection rule for modelling collisions (see

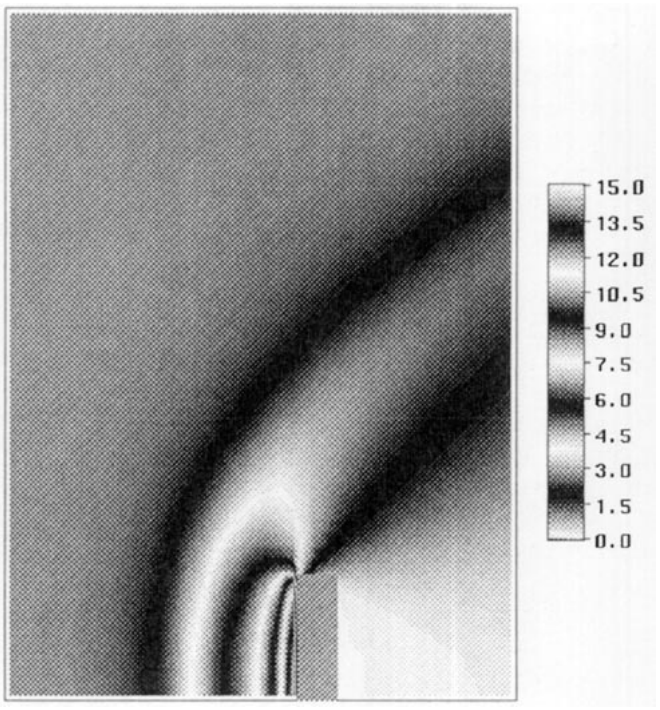

FIG. 1. Density distribution in a rarefied flow about a flat plate for $M = 8$ and $Kn = 0.1$.

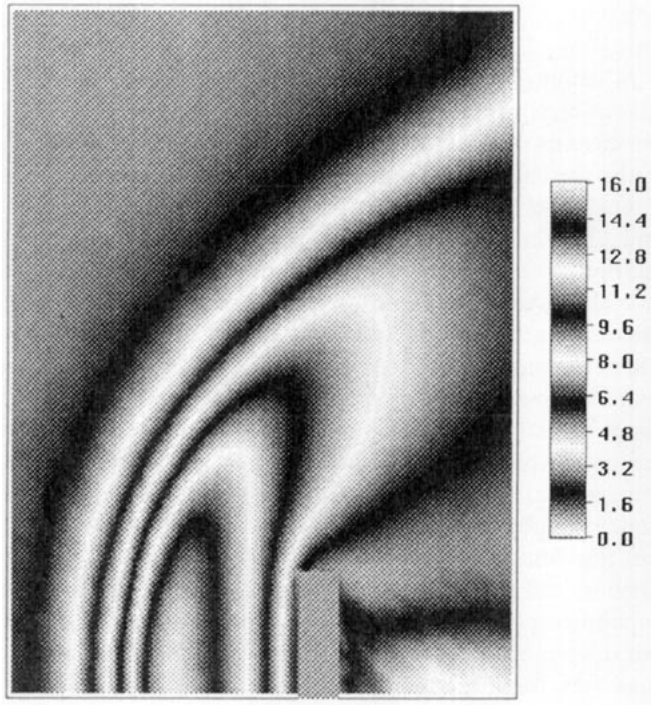

FIG. 2. Temperature distribution in a rarefied flow about a flat plate for $M = 8$ and $Kn = 0.1$.

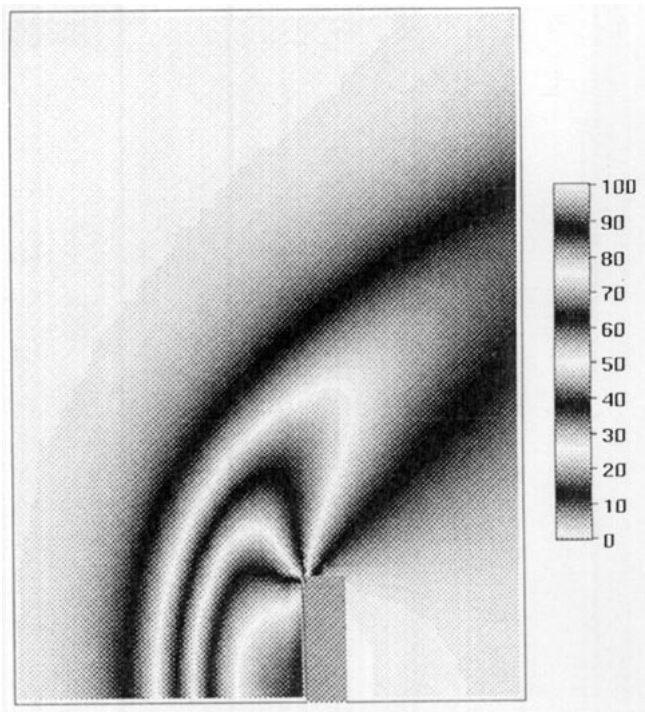

FIG. 3. Pressure distribution in a rarefied flow about a flat plate for $M = 8$ and $Kn = 0.1$.

Baganoff and McDonald [6]) and various programming steps taken to improve code efficiency, as discussed by McDonald [5] and Baganoff [7]. The resultant computational speed of the code was roughly 1.0 μs per particle per time step.

In defining the problem to be studied, performance considerations led to the selection of simulated wind tunnel dimensions of 40 cells in the streamwise directions, 55 cells in the vertical direction (half space), and 3 cells in depth for a total of 6600 cubical cells. The vertically oriented flat plate had a half-height of 10 cells and a streamwise thickness of 3; see Figs. 1–3. The upstream mean free path length was set at 2.0 cells, giving the Knudsen number of 0.1 quoted above. On varying the average particle number density (based on the entire simulation) from roughly 8 to 121 particles per cell, the overall size of the simulations thus varied from roughly 53,000 to 800,000 particles. The fairly large upper limit was the controlling factor in our selection of a two-dimensional problem for study, as opposed to a three-dimensional problem. Our use of a small three-unit depth for the simulated wind tunnel, while modelling a two-dimensional problem, was related to the three-dimensional capability of the code used. This particular selection of parameters resulted in a run time of approximately 1 s per time step for the larger simulations, and a corresponding total run time of roughly 0.5 h. The reference solution, see Eq. (1) below, employed 800,000 particles and 1689 time steps for time-averaging the data.

STATISTICAL ERROR

In order to determine the level of statistical fluctuations, or *rms* error, associated with a given simulation, one must consider two items: first, a reference solution is needed against which all others are compared; and second, a specific definition for the measure of rms error must be introduced. For the rarefied flow considered, an exact solution is simply not available to provide a reference. However, it will be shown that a procedure can be found for determining the absolute rms error for each run from an analysis of the entire group of runs, even without having the exact solution itself among the group. This apparent logical contradiction becomes more rational when one learns that at least one high quality solution must be present in the group to give reliable results and that the results of the full analysis are not much different from the straightforward approach of using the highest quality run, consisting of the largest number of particles and the longest time averaging, as the reference.

With regard to defining the rms level of statistical fluctuations, it is clear that the macroscopic fluid quantities for density, velocity, temperature, pressure, stress, and heat flux may exhibit different levels because they represent different moments of the velocity distribution function. Because density is the zeroth-order moment, it should exhibit the smallest ratio of rms error to mean, while pressure and temperature represent second-order moments and the corresponding ratios should be higher. Therefore, the analysis must distinguish between the different macroscopic variables. Most of the results given below will be presented for the temperature variable. In defining a single numerical measure of error for a particular macroscopic variable, one could consider a single point in the flow that corresponds to a particular position of interest or consider an average for the entire flow field. The definition to be applied will make use of an average over the flow field.

The appropriate concepts are most easily reviewed if the simplest approach is considered first, i.e., the case consisting of the largest number of particles and the longest time averaging is used as the reference, and all other runs are compared with it. In a simulation, a macroscopic fluid quantity is first determined from an appropriate average of data associated with the particles in a single cell and then it is further averaged as the simulation is advanced in time. Using the overbar notation to designate a time average, the symbol $\bar{q}_{\alpha,i}$ will be used to represent a time-averaged macroscopic fluid quantity q associated with run α and evaluated for cell i. If the corresponding reference quantity $\bar{q}_{r,i}$ is viewed as an exact mean value, then the square error can be defined by

$$\delta^2_{\alpha,i} = (\bar{q}_{\alpha,i} - \bar{q}_{r,i})^2. \tag{1}$$

A single dimensionless measure of relative error for the entire flow can then be introduced by writing

$$\hat{\mu}_\alpha^2 = \frac{1}{N_{\text{cells}}} \sum_{i=1}^{N_{\text{cells}}} \left(\frac{\bar{q}_{\alpha,i}}{\bar{q}_{r,i}} - 1 \right)^2, \tag{2}$$

where $\hat{\mu}_\alpha$ has the interpretation of a dimensionless rms value. An alternative to (2) that simplifies the computation somewhat makes use of a single reference mean value, such as the maximum, and the corresponding relation reads

$$\mu_\alpha^2 = \frac{1}{N_{\text{cells}}} \sum_{i=1}^{N_{\text{cells}}} \left(\frac{\bar{q}_{\alpha,i}}{\bar{q}_{r,\max}} - 1 \right)^2. \tag{3}$$

On comparing the value of a fluid variable at the stagnation point, or a point of maximum, to its free stream value, the ratios for density, temperature, and pressure are roughly 15, 15, 100, respectively, for the case studied; see Figs. 1–3. Because random fluctuations scale with the size of the local mean value, it is clear that the relative error μ_α is heavily weighted by the large values near the plate while $\hat{\mu}_\alpha$ is more evenly weighted. Figures 4 and 5 give the results from a series of tests for $\hat{\mu}_\alpha$ and μ_α, respectively, for the fluid temperature variable. The independent variable in the two figures is the average number of particles per cell defined by $N_c = N_{\text{total}}/N_{\text{cells}}$ and the parameter that was varied was the size of the sample for the entire run defined by $S = N_{\text{total}} T$, where T is the number of time steps used in the time averaging. In this analysis all time steps were used in the averaging; none was skipped. Generally, factors of two were used in varying the quantities N_{total} and T. Comparison of Figs. 4 and 5 shows that the respective curves look very similar except for their absolute numerical values, which are different because of the different normalization used in (2) and (3). The datum point that is missing in the two figures is the one for which the reference would be compared to itself. All the curves show the same trend, namely, that the relative rms error decreased monotonically with increasing N_c for fixed values of the total sample size S. Likewise, it also decreases monotonically with increasing T for fixed values of N_c. Because the computational effort grows in proportion to S, the data show that, for fixed computational cost and for most of the region studied, increasing N_c is clearly more effective than increasing T in producing a small *rms* error. Likewise, expression (3) is preferred over (2) because it is more convenient to evaluate and yet it predicts essentially the same results.

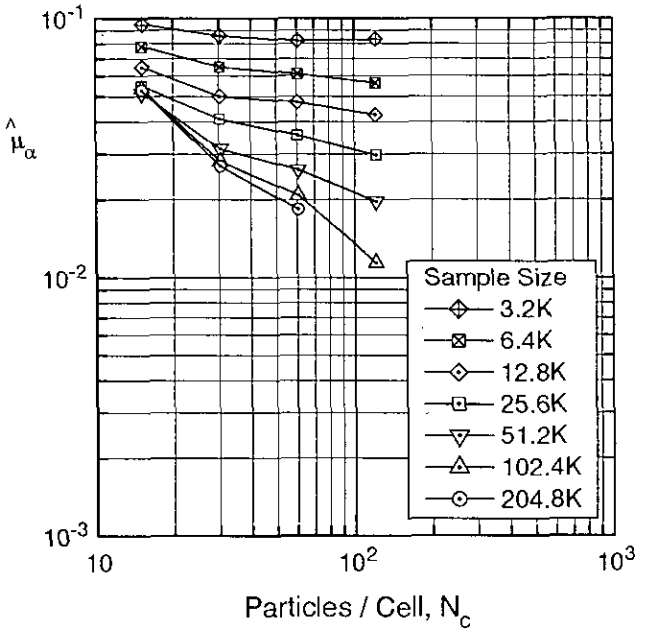

FIG. 4. Relative rms error $\hat{\mu}_\alpha$ for the temperature variable versus the average number of particles per cell N_c, holding the total sample size $S = N_{\text{total}} T$ fixed.

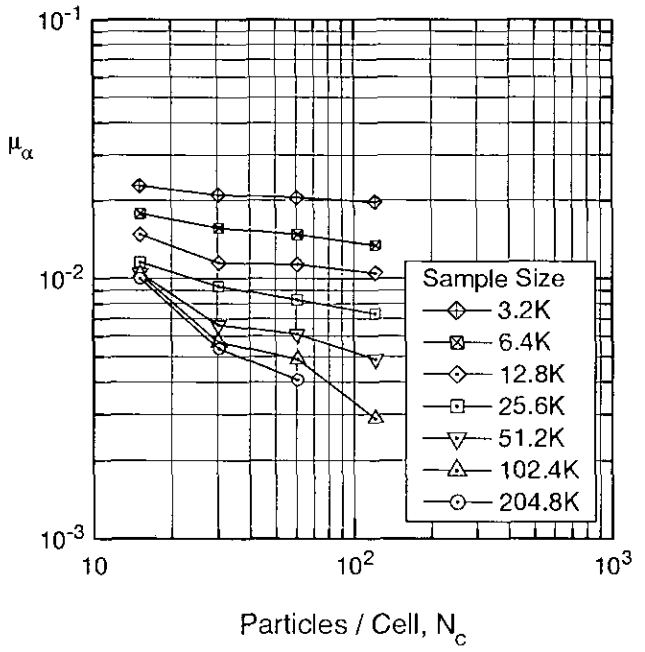

FIG. 5. Relative rms error μ_α for the temperature variable versus the average number of particles per cell N_c, holding the total sample size $S = N_{\text{total}} T$ fixed.

Alternatively, a theoretical determination of the absolute rms error, as opposed to the relative rms error, can be found by first considering two distinct runs (α, β), each carried out with a different number of particles and/or a different duration of time averaging. If $\bar{q}_{\alpha,i}$ and $\bar{q}_{\beta,i}$ represent time-averaged data for the same cell but for two different runs, then one is able to define a measure of their difference by

$$\delta_{\alpha\beta,i}^2 = (\bar{q}_{\alpha,i} - \bar{q}_{\beta,i})^2. \tag{4}$$

Now, both $\bar{q}_{\alpha,i}$ and $\bar{q}_{\beta,i}$ can be considered to be composed of the exact mean value plus an error; and therefore, their difference is merely the difference of the two absolute error terms alone. Consequently, we may write

$$\delta^2_{\alpha\beta,i} = (\bar{e}_{\alpha,i} - \bar{e}_{\beta,i})^2. \tag{5}$$

If an average is again taken over all cells in the flow, then (5) leads to the relation

$$\Delta^2_{\alpha,\beta} = \frac{1}{N_{\text{cells}}\bar{q}^2_{r,\max}} \sum_{i=1}^{N_{\text{cells}}} (\bar{e}^2_{\alpha,i} - 2\bar{e}_{\alpha,i}\bar{e}_{\beta,i} + \bar{e}^2_{\beta,i}), \tag{6}$$

where the constant $\bar{q}^2_{r,\max}$ is arbitrarily introduced to non-dimensionalize the equation. In view of the definition of the absolute error term $\bar{e}_{\alpha,i}$, it is reasonable to treat it as a random variable with respect to its subscript i; surely its mean value is zero. In addition, the two quantities $\bar{e}_{\alpha,i}$ and $\bar{e}_{\beta,i}$ ought to be statistically independent since they derive from two independent runs. On this basis the cross product term in (6) is expected to vanish when the sum is carried out over all cells ($N_{\text{cells}} = 6600$ for the example studied), thus reducing the relation to

$$\Delta^2_{\alpha,\beta} = \sigma^2_\alpha + \sigma^2_\beta, \tag{7}$$

where σ^2_α is a dimensionless spatially averaged absolute error term defined by

$$\sigma^2_\alpha = \frac{1}{N_{\text{cells}}\bar{q}^2_{r,\max}} \sum_{i=1}^{N_{\text{cells}}} \bar{e}^2_{\alpha,i}. \tag{8}$$

Equation (7) provides the means for obtaining the absolute rms error for each run, even without having the exact solution itself in hand. This follows from the fact that the left-hand side of (7) can be computed directly for the different combination of runs using the spatial average of definition (4), i.e.,

$$\Delta^2_{\alpha,\beta} = \frac{1}{N_{\text{cells}}\bar{q}^2_{r,\max}} \sum_{i=1}^{N_{\text{cells}}} (\bar{q}_{\alpha,i} - \bar{q}_{\beta,i})^2. \tag{9}$$

The quantities on the right-hand side of (7) are obtained from the solution of the resultant set of simultaneous equations. Clearly the system is over specified, because there are $r(r-1)/2$ distinct entries in the symmetric matrix $\Delta^2_{\alpha,\beta}$ and only r unknowns σ^2_α, where r is the number of different cases or runs. A discussion of (7) is most easily followed if the runs are first conceptually ordered with respect to their rms error, where the smallest is designated as run r. Using this ordering, it is clear that the entries near the diagonal in the symmetric matrix $\Delta^2_{\alpha,\beta}$ are not as useful (small relative error) as those further removed. Therefore, the system of equations can be conveniently reduced, by ignoring the less useful equations, to a properly specified set, without having to resort to a least square error method to solve the entire set. On retaining the subset of (7) for which $\alpha = 1, 2, \ldots (r-1)$ and $\beta = r$ and then arbitrarily including the equation $\alpha = 1$ and $\beta = (r-1)$, a closed set of equations is obtained and it is given by

$$\begin{pmatrix} \Delta^2_{1,r} \\ \Delta^2_{2,r} \\ \vdots \\ \Delta^2_{r-1,r} \\ \Delta^2_{1,r-1} \end{pmatrix} = \begin{pmatrix} 1 & 0 & \cdots & 0 & 1 \\ 0 & 1 & \cdots & 0 & 1 \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 1 \\ 1 & 0 & \cdots & 1 & 0 \end{pmatrix} \begin{pmatrix} \sigma^2_1 \\ \sigma^2_2 \\ \vdots \\ \sigma^2_{r-1} \\ \sigma^2_r \end{pmatrix}. \tag{10}$$

As expected, the solution of (10) varies slightly with changes in the selection of equations retained and the value of σ^2_r shows the greatest sensitivity to alterations in the selection. Nevertheless, Fig. 6 displays the resulting solution of (10) for our data, again for the temperature variable, and shows that agreement with Figs. 4 and 5 is quite good, demonstrating a firm consistency between the three approaches. The agreement between Figs. 4–6 therefore allows one to conclude that the approach defined by Eq. (3) is much preferred because of its ease in evaluation while still providing the desired information.

The logical procedure leading to (10) gives spatially averaged absolute rms error values for the flow. Equivalent quantities for individual cells could be found if one were

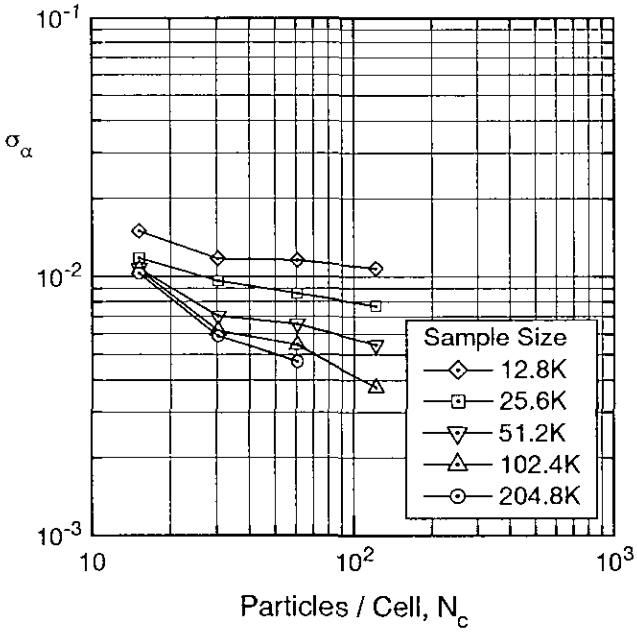

FIG. 6. Absolute rms error σ_α for the temperature variable versus the average number of particles per cell N_c, as obtained from the solution of Eq. (10).

willing to consider a greatly increased computational effort. In Eq. (6) the cross product term vanished because of the spatial averaging. The same term could be made to vanish, while retaining the index i, if an ensemble average were introduced instead. This would allow the same development leading to Eq. (10) except the index i would be preserved, thus giving values associated with individual cells. Clearly the large number of repeated simulations required to carry out ensemble averaging would be prohibitively expensive in practice. However, the concept that such data could be found in principle is important to our understanding of the method.

COMPUTATIONAL COST

The display of data in each of the Figs. 4–6 reflects the order in which the numerical simulations were conducted. For example, consider Fig. 5 and the sequence for which $N_c = 121$ particles/cell. In this case the total number of particles used in the simulation was set once steady state was reached, and the time averaging was carried out in steps, where the total time-averaging period for each step was double the previous period. If the conditions of a run happen to be compatible with the requirements of the ergodic hypothesis then the rms error should decrease as $T^{-1/2}$; and it can be seen from the data for $N_c = 121$ that a factor of 4 increase in T leads to a reduction of rms error by a factor of 2, which is consistent with the ergodic hypothesis. However, this rule clearly does not apply for all values of N_c displayed, but it is difficult to judge from the curves where the rule begins to fail. The same observation also applies if the data of Figs. 5 were displayed with T being the independent variable.

On the other hand, if the choice of variables is rearranged as shown in Fig. 7, then the judgement becomes much easier to make. The total sample size $S = N_{\text{total}} T$, which also represents the computational cost for the DSMC method, is shown as a function of the average cell particle density N_c, for fixed values of the dimensionless rms error μ_α. Because the simulations could not be conducted in this order, these results were obtained from suitable cross plots of the two graphical forms μ_α versus N_c and μ_α versus T. Focusing attention on the curve for a fixed 4% rms error, it is evident that two asymptotes exist. For N_c greater than approximately 100 particles/cell, the ergodic hypothesis clearly applies; i.e., the computational cost is constant and independent of the size of the simulation. This is because in this limit the rms error is proportional to $S^{-1/2}$ and therefore a fixed error implies a fixed S; and a fixed S results in a fixed computational cost because it is linearly related to S. Finally, from the definition $S = N_{\text{total}} T$, a fixed S allows a free choice of N_{total} (or T) and thus cost is independent of the size of the simulation N_{total}.

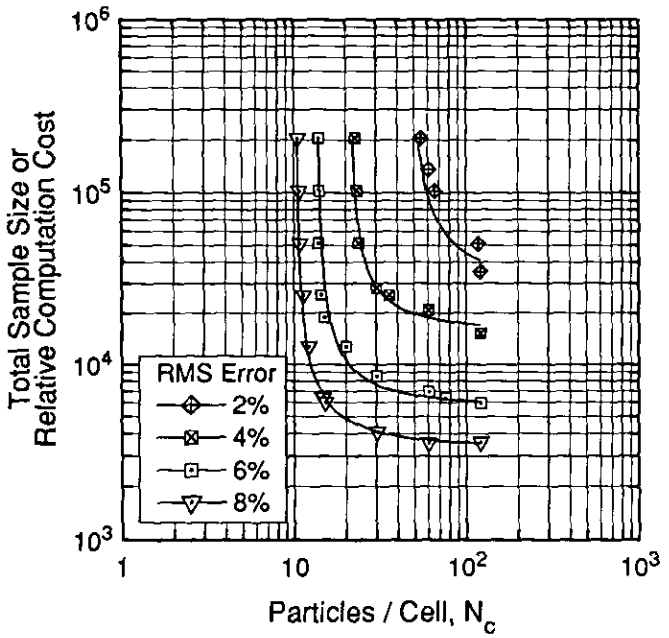

FIG. 7. Total sample size $S = N_{\text{total}} T$, or relative computational cost, versus the average number of particles per cell N_c for the temperature variable.

Following the same curve for 4% error, we find that for N_c less than approximately 30 particles/cell, the computational cost rises rapidly with just a small decrease in N_c. This is the region in which the simulation becomes very inefficient, because there are too few particles in a cell to adequately model the flow physics. Consequently, one attempts to make up the severe deficiency with a huge increase in the period of time averaging. This asymptotic limit obviously shows that, for a given level of rms error, there is a minimum N_c that is allowed, even if the period of time averaging were infinite. In retrospect, this is a conclusion that should be expected on physical grounds; however, the numerical simulations were needed to fix the actual numerical value at which this occurs. The division between efficient and inefficient simulations can be conveniently defined by the knee in the curve, which for the case of 4% error appears at roughly $N_c = 30$ particles/cell.

Continuing to review the curve for 4% error in Fig. 7, we see that a five-fold increase in N_c from 25 to 125 particles/cell leads to a ten-fold decrease in the computational cost. In other words, a large simulation is less costly than a small simulation! At first glance, this appears to be counter-intuitive, but in actual fact it is merely a reflection of the difference in simulation efficiency at the two extremes. This is a very important conclusion for this class of simulations, because it shows that access to greater computer memory can have a dramatic effect on reducing the computer run time. It also points out that for an extremely large simulation that makes use of all available computer memory and

still does not operate in an efficient mode and which would normally require several hours of run time, sufficient savings in time could be realized by switching to an efficient mode of operation to suggest the possible use of disk read/write to allow the necessary further increase in N_{total}. For this same 4% rms error level, we see that roughly 100 particles/cell are required for an efficient simulation and that for the two-dimensional example studied ($N_c = 6600$) this translates into a total of 660,000 total particles required. A similar three-dimensional problem would require over an order of magnitude greater number (exact ratio obtained from the 55-cell height to the 3-cell depth used). Because more than 10 words of data storage are needed for each simulated particle, especially if chemical reactions are modelled, it can be seen that roughly 120 Mwords of memory are needed for an efficient simulation, even when the geometrical resolution of the problem studied is fairly modest, as in our example problem.

DISCUSSION AND CONCLUSIONS

Many of the past applications of the DSMC method for two- and three-dimensional problems were conducted at average number densities of around 15 to 20 particles per cell. This was done for a number of closely coupled reasons relating to the size of available computer memory, code execution speed, total run time that could be committed, and the type of machine used. The clear conclusion drawn from Fig. 7 is that every effort should be made to employ an average particle number density four or five times greater, so that full advantage can be taken of the greater simulation efficiency. This is a result that is independent of machine architecture and depends solely on the physics of rarefied gas flow and its simulation. However, the ease with which the desired operating point can be reached is machine dependent and does require appropriate consideration.

The obvious questions left unanswered by this study relate to differences introduced by more complex flow geometries, the presence of multiple species and chemical reactions in the simulated gas, rms error specific to a particular cell as opposed to a single measure for an entire simulation, and the effect of varying cell size. The asymptotic limit suggested by each curve in Fig. 7 can be interpreted as the number of particles needed in a single cell to give the same accuracy in a single time step. However, the study was conducted for the case of a steady flow and it does not follow that this same number would necessarily be valid for time accurate results. This question would require a separate study dealing with transient flows. Likewise, in regions of flow where gradients are steep, as occur in regions close to solid boundaries where translational non-equilibrium becomes very important, one is also not able to conclude from this work that 100 particles per cell is sufficient to give the same 4% resolution, because the boundary layer was relatively thick in the example studied, owing to the fairly high Knudsen number chosen. What has been shown is that computational cost for the DSMC method can be reduced in a major way by conducting a simulation in a regime where the relevant physical processes are efficiently modelled, even though the modelling requires the use of significantly greater memory and/or data storage.

ACKNOWLEDGMENTS

This work was supported in part by the National Aeronautics and Space Administration under Grant NCA2-495 and in part by the Air Force Office of Scientific Research under Grant AFOSR 90-0232. All of the computational work carried out on the NAS Cray-YMP was supported by NASA-Ames Research Center.

REFERENCES

1. B. J. Alder and T. E. Wainwright, "Molecular Dynamics by Electronic Computers," in *Transport Processes in Statistical Mechanics*, edited by I. Prigogine (Interscience, New York, 1958).
2. R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles* (Adam Hilger, Bristol, 1988).
3. G. A. Bird, *Molecular Gas Dynamics* (Clarendon, Oxford, 1976).
4. G. A. Bird, "Monte-Carlo Simulation in an Engineering Context," in *Rarefied Gas Dynamics: Technical Papers from the 12th International Symposium, Progress in Astronautics and Aeronautics Series* (AIAA, New York, 1980), Vol. 74, Part 1, p. 239.
5. J. D. McDonald, Ph.D. thesis, Stanford University, December 1989.
6. D. Baganoff and J. D. McDonald, *Phys. Fluids A* **2** (7), 1248 (1990).
7. D. Baganoff, "Vectorization of a Particle Code Used in the Simulation of Rarefied Hypersonic Flow," in *Computational Technology for Flight Vehicles, Vol. I*, edited by A. K. Noor and S. L. Veneri (Pergamon, Oxford, 1990); also in *Comput. Systems Eng.* **1** (2-4), 437 (1990).