DSMC Wind Tunnel

The DSWT Program User's Guide

Version 1.0
November 1999

G.A.B. Consulting Pty Ltd
5 Fiddens Wharf Rd., Killara, N.S.W. 2071, Australia
Phone: Country code 61 City code 2 9498 7809
FAX: 61 2 9499 6372
gabird@compuserve.com
http://ourworld.compuserve.com/homepages/gabird
http://gab.com.au

Table of Contents

2 The Scope of the DSWT Program

1 The DSWT Implementation of the DSMC Method

The direct simulation Monte Carlo method is a technique for the computer modelling of a real gas by some thousands or millions of simulated molecules. The velocity components and position co-ordinates of these molecules are stored in the computer and are modified with time as the molecules are concurrently followed through representative collisions and boundary interactions in simulated physical space. This direct simulation of the physical processes contrasts with the general philosophy of computational fluid dynamics which is to obtain solutions of the mathematical equations that model the processes. The computational task associated with the direct physical simulation becomes feasible when the gas density is sufficiently low. It also becomes necessary under these conditions because the Navier-Stokes equations do not provide a valid model for rarefied gases, and conventional CFD methods are unable to cope with the large number of independent variables that are involved in applications of the Boltzmann equation to realistic multi-dimensional problems.

The degree of rarefaction of a gas flow is generally expressed through the overall Knudsen number, which is the ratio of the mean free path in the undisturbed stream to a typical dimension of the flowfield. It is desirable to also define the local Knudsen number as the ratio of the local Knudsen number to the scale length of the gradient of a specified macroscopic (or continuum) flow property. The Navier-Stokes equations are valid only when the local Knudsen numbers based on flow speed, density, temperature and pressure are small in comparison with unity. The limit as the Knudsen numbers tend to zero may be identified with the inviscid limit that is modelled by the Euler equations. The opposite limit as the Knudsen number tends to infinity is the collisionless or free-molecule flow limit in which intermolecular collisions may be neglected. The flow regime between free-molecule and the limit of validity of the Navier-Stokes equations is generally referred to as the transition flow regime. A Knudsen number of 0.1 has traditionally been quoted as the boundary between the continuum and transition regimes, but the characteristic dimension of complex flow fields may be specified in many different ways and the use of the overall Knudsen number may be misleading.

The conservation equations of fluid mechanics are valid for all flow regimes, but the Navier-Stokes equations depend also on the Chapman-Enskog theory for the shear stresses, heat fluxes and diffusion velocities as linear functions of the velocity, temperature and concentration gradients. The Chapman-Enskog theory assumes that the velocity distribution is a small perturbation of the equilibrium or Maxwellian distribution. (In an isentropic flow for which the Euler equations are valid, the distribution function conforms everywhere to the Maxwellian). The formulation of the Chapman-Enskog distribution involves the local Knudsen numbers based on the velocity and temperature gradients. It has been found that errors become significant when these local Knudsen numbers exceed 0.1 and the continuum theory is hardly useable when they exceed 0.2. The transport property terms become zero in an isentropic flow and it might be thought that the Euler equation yield correct inviscid flow results at all Knudsen numbers. However, as the density decreases, the collision rate in the gas eventually becomes too low to maintain the isotropy of the pressure tensor and the continuum model breaks down.

Although it was shown many years ago that the Chapman-Enskog expansion for the distribution function is not uniformly valid, attempts are still being made to extend the range of validity of the Navier-Stokes equations to lower densities. However, the low density effects such as the development of an anisotropic pressure tensor are of a very basic nature and it is unlikely that much progress will be made for other than one-dimensional steady flows. In addition, effects such as thermal and pressure diffusion become more prominent at low densities and these are not generally included in the Navier-Stokes formulations. It is certain that the necessary extensions (in the event that adequate ones can be developed) will add greatly to the difficulty of the continuum approach. On the other hand, once the density becomes sufficiently low for the DSMC solution to be computationally feasible, it is a much easier method to apply. The main reasons for this are:-

(i) The calculation is always unsteady with physical time as a variable in the simulation. If the flow becomes steady, it is obtained as the large time state of the unsteady flow. The method does not require an initial approximation to the flow field and there is no iterative procedure or convergence to the final solution. (In the case of a time averaged steady flow or an ensemble averaged unsteady flow, there will be a gradual decline in the statistical scatter as the sample increases, but "convergence" is not the appropriate description of this process.)

(ii) Additional effects, such as non-equilibrium chemistry, may be included simply by adding to the complexity of the molecular model and the fact that these may change the basic nature of the mathematical models is of no consequence.

(iii) Most importantly, there are no numerical instabilities!

Personal computers now readily permit DSMC calculations of two-dimensional and axially symmetric flows at overall Knudsen numbers in the range 0.01 to 0.001. This is well into the continuum regime. Note that consideration of the Knudsen number does not introduce an additional independent parameter because the Reynolds number (Re) in air is related to the Knudsen number (Kn) and the Mach number M by

$$(Re)=1.21\, M/(Kn).$$

The total number of simulated molecules is limited and, in a three-dimensional calculation at normal and low densities, each simulated molecule must be regarded as representing an extremely large number of real molecules. The macroscopic or continuum flow properties are established through averages over the microscopic or molecular properties and are therefore subject to an unrealistic level of statistical scatter. However, the number of molecules in a cubic mean free path is inversely proportional to the square of the number density. The scatter in a typical DSMC calculation is, in fact, similar to the real scatter in a gas flow at a density of about five times normal or sea-level density. A flow with extremely small physical dimensions may therefore be calculated with a 1:1 correspondence between real and simulated molecules. The DSMC method may be modified to take account of the size of the molecules and the fluctuations in the simulation have then been shown to have the same physical characteristics as the real fluctuations. Fluctuations are neglected in the formulation of the Boltzmann equation, so that, in this respect, the DSMC solution has a higher degree of physical realism than a (hypothetical) solution of the Boltzmann equation. Fluctuations are included in gas theory only at the level of the Fokker-Planck equations and are, of course, disregarded in the Navier-Stokes equations.

The physical realism of the fluctuations has important consequences for DSMC calculations with two-dimensional or axial symmetry, as in the DSWT program. This is because the calculation can be regarded as a simulation within a slice of the flow that is sufficiently thin for the number of simulated molecules to be realistic. The statistical scatter generally decreases as the square root of the sample size and, in order to attain a sufficiently small standard deviation in the results, DSMC programs employ either time averaging for steady flows or ensemble averaging for unsteady flows. There is a serious statistical problem if a significant effect in the real gas is a consequence of the few molecules towards the extremities of the velocity distribution, but this is a very rare occurrence. The fluctuations are therefore a consequence (albeit an undesirable consequence) of the physical model rather than an unrealistic computational approximation.

The uncoupling of the molecular motion and collisions over small time steps and the division of the flow field into small cells are the key computational assumptions associated with the DSMC method. The time step should be much less than the mean collision time and a typical cell dimension should be much less than the local mean free path. The cell dimension should also be small compared with the distance over which there is a significant change in the flow properties. The DSMC method uses the cell system only for the sampling of the macroscopic properties and for the selection of possible collision partners, although the sampled density is used in the procedures for establishing the collision rate. This means that the cell geometry should be chosen to minimise the changes in the macroscopic properties across an individual cell.

A detailed exposition of the method is available in the reference that will be referred to as Bird (1994). However, the DSWT program owes its unprecedented ease of use to more recent developments in DSMC methodology.

Primitive implementations of the DSMC method choose the collision partners from any location in the same cell. Later implementations, including those in Bird (1994), employed fixed sub-cells to reduce the spacing of collision partners. The condition relating to the cell size relative to the mean free path then applies to the sub-cell size, rather than to the cell size. The most recent programs, including the DSWT program, generate a transient rectangular background grid to one cell at a time within the collision routine. The molecules are then indexed to this transient grid. The size of the background grid is set such there is approximately one simulated molecule within each grid element. The variables for the transient grid are allocated when the size is decided and are deallocated when the collisions have been calculated for the cell. The first molecule for a potential collision is chosen at random from the cell. The transient sub-cell in which it lies is determined from the indexing and, if there is more than one molecule in the sub-cell, the collision partner is chosen from these. If there are is only one molecule in the sub-cell, the collision partner is chosen from an adjacent sub-cell. The access to the adjacent sub cells is random in order to avoid artificial correlations in the location of collision partners. This yields nearest-neighbour collisions and is efficient with regard to both computation time and storage requirements. The ratio of the mean separation between collision partners to the local mean free path should be well under unity over the flowfield. The transient adaptive sub-cells ensure that the lowest possible values of this ratio are achieved. Further reductions can be obtained only by an increase in the total number of simulated molecules.

The transient adaptive sub-cells mean that expertise on the part of the user is no longer required in the setting of the cell structure. The cell structure of the DSWT program is set entirely by the program logic and no input is required from the user.

The other area in which expertise has been required is in the setting of the time step over which the molecular motion and collisions is uncoupled. In addition, the traditional specification of a single value of this time step over the whole flowfield is computationally inefficient, particularly for supersonic and hypersonic flows. This is because a value that is small in comparison with the mean collision time in the stagnation region is extremely small in comparison with the collision time in the outer regions of the flow. This problem is avoided by new procedures that make the time step vary with every molecule as well as with every cell. To this end, separate time variables are associated with every molecule and every cell. The procedures keep all these variables concurrent with the overall time variable.

The overall time variable is advanced in very small time steps, but only a small fraction of the molecules are moved and collisions are calculated in only a small fraction of the cells at any one time step. The mean collision time is kept for each cell and collisions in a cell appropriate to a time interval of one sixth[*] the local mean collision time are calculated whenever the cell time variable falls one twelfth the local mean collision time behind the overall flow time. Similarly, a molecule will generally be moved through a time interval of one sixth the mean collision time of its current cell when its time variable falls one twelfth this local mean collision time behind the overall flow time. However, the logic can often be simplified if there is a molecule velocity dependent restriction that limits the distance that a molecule can move in a single step. The step set by the local collision time is then the maximum value of the time step.

Because of the very small overall time step and the small number of cells for which collisions are calculated at each step, the implementation of the traditional indexing routine at each step would involve a prohibitive computational effort. This problem has been solved through the introduction of “bi-directional” indexing. The traditional index array is comprised of the molecule numbers in order of the cells. If the address of the corresponding entry in this array is stored for each molecule, the indexing may be kept up to date continuously as molecules move from cell to cell. The index array must also be extended with sufficient free space between cells to cope with the fluctuations in the number of molecules in individual cells. Some general rearrangement takes place from time to time to allow for particularly large fluctuations or systematic changes in the distribution of molecules. In many other procedures, the ability to assign different times to molecules leads to a simplification of the program.

---

[*] This is the fraction employed in DSWT. It may vary with the implementation.

The new procedures are computationally faster than the traditional procedures but, because of the data structures necessitated by the continuous indexing, the number of simulated molecules that can be employed for a given amount of memory is reduced. Because of the automatic setting of the time step, it is no longer possible for this to be incorrectly specified. The transient-adaptive sub-cells effectively prevent the incorrect setting of cell size and a poor calculation can be made only if the total number of simulated molecule is too small. This can cause the mean separation of collision partners to exceed the local free path but, since this is continuously checked, the program can be made to stop if an obviously bad calculation is attempted.

These new procedures effectively enable DSMC programs to be made foolproof. They have been implemented in a new general program, called the DS2A program, for two-dimensional and axially symmetric flows. The DSWT program is a restricted derivative of the DS2A program that incorporates a “user-friendly” graphical interface. The restrictions permit the automatic generation of the computational variables that must be set by the user of the DS2A program**. The user of the DSWT program has only to specify the approximate number of megabytes that are available to the program, the size of the simulated wind tunnel working section, the shape of the model and the stream properties.**

Reference

BIRD, G.A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Oxford University Press, 1994.

2 The Scope of the DSWT Program

2.1 Geometry

The DSWT program restricts the DS2A geometry to that corresponding to the working section of either a rectangular cross-section wind tunnel with a two-dimensional model or a circular cross-section tunnel with an axially symmetric model. The working section size is specified by the minimum and maximum values of the x and y co-ordinates. The x axis ($y = 0$) is the axis of an axially symmetric flow and the y co-ordinate is then a radius. If a two-dimensional model is symmetric about the x axis, the minimum y co-ordinate ($y = 0$) may be specified as a plane of symmetry. Otherwise, the boundaries are “stream boundaries” with incoming molecules characteristic of the undisturbed stream and outgoing molecules are removed from the flow. These boundary conditions are obviously different from those of a real tunnel, but cause less serious “wind tunnel interference” effects. Should the user wish to study more realistic interference effects, a solid surface may be set parallel to the flow direction. If the flow is two-dimension without a flow symmetry, the stream may have a component in the y direction as well as the x direction.

The “wind tunnel model” may comprise one or more separate surfaces. Each of these must be either a closed surface or an open surface that starts and finishes on a boundary of the working section. Each surface is specified either by a combination of straight line and circular arc segments or by an externally generated polyline file.

The flowfield grid is generated automatically by the program. It consists of a background rectangular grid that is uniformly spaced in each direction plus a number of cell layers that are wrapped around the surfaces. The surface cells mean that there are no difficulties associated with thin surfaces that terminate within a background cell. There is also the opportunity, within the program options, to have very thin cells adjacent to the surface. A fan shaped set of surface cells radiates from any surface vertex where there is a large change in surface slope. Cells for the sampling of surface properties are also generated automatically. Their size is unrelated to the size of the surface definition segments and, while their size is similar to that of the flow cells, they are not directly related to the cells.

The DSWT program calculates the intersections of the molecules with axially symmetric bodies as three-dimensional events. Most of the earlier programs employed two-dimensional logic with reduced time steps near the axis.

2.2 Gas Specification and Model

The gas is either ideal or real air and can comprise oxygen, nitrogen, atomic oxygen, atomic nitrogen, and nitric oxide. The default initial state is 20% oxygen and 80% nitrogen although these proportions and a fraction of atomic oxygen (to model upper atmosphere conditions) may be reset in the program options. The program considers the stream temperature and velocity and chooses an ideal gas, a non-reacting gas with vibrational excitation, or a real reacting gas.

The program employs the physical gas models that have been described and validated in Bird (1994). The gas is a mixture of the VHS models and the cross-sections, the viscosity-temperature index (which determines the way in which the cross-section changes with the relative velocity), are set separately for every molecular species. A classical Larsen-Borgnakke model is employed for the rotational degrees of freedom, while a quantum model is used for the vibrational modes. The chemical reaction model calculates reactive cross-sections that are consistent with the measured rate constants.

2.3 Gas-Surface Interactions

The classical diffuse reflection model with complete accommodation of the gas to the surface temperature is appropriate to "engineering surfaces" that have not been exposed for a long period to ultra-high vacuum. A uniform temperature may be specified for the surface or it may be specified as an adiabatic surface with zero heat transfer. In the latter case, the temperature distribution on the surface is one of the output quantities. Alternatively, the other "classical" model of specular reflection may be specified for the artificial case of complete slip and no energy transfer at the surface. This option can be applied to a surface along a boundary in order to convert the boundary to a plane of symmetry.

3 Computer Requirements

(i) The program requires Windows 95, 98 or NT 4.0.

(ii) A 100 MHz Pentium™ processor or higher is recommended.

(iii) 32MB RAM is required for the program to run efficiently.

(iv) The program auto-detects either 800×600 or 1024×768 screen resolutions and custom full-screen displays are generated for each of these cases. For higher resolutions, the 1024×768 display appears in a window. The program is not usable at 600×480. **1024×768 should be selected whenever possible.**

(v) High color (16 bit) is strongly recommended. (Should only 256 colors be available, the number 3 color map should be chosen in the DSWT window.)

(vi) The "Small Fonts" (normal) option must be chosen in Windows. (This is selected under Display Properties / Settings / Advanced)

(vii) A mouse or other pointing device is required.

4 Program Installation and Operation

The program is delivered as the compressed executable **DSWTZ.EXE**. This is run to produce the uncompressed executable **DSWT.EXE**.

Running DSWT produces the main DSWT window with the logo and title **DSMC Wind Tunnel (for two-dimensional plane and axially-symmetric flows)** and the pop-up **Run Type** window. The Run Type widow contains six push button style radio buttons and an **OK** button to move on when the desired button has been selected.

This is the pre-selected button, but is valid only when a calculation has been stopped after it had reached the stage at which the flow and surface property displays are refreshed. Restart files are written every time this "print-out" stage is reached. The selection of this option causes the last set of restart files to be read and re-starts the run from that flow time.

This is similar to the preceding option, but causes all the variables for flow and surface property sampling to be re-initialized. It will be chosen only when it is desired to override the automatic choice of the flow time at which the flow is assumed to become steady.

This option requires a valid data file **DSWT.DAT** to be present. Its selection causes the flow to be re-initialized and the calculation to start again from zero time. Non-standard program options are reviewed automatically at this stage.

Set New or Modified Data

This option allows a new calculation to be specified. If there is no valid data file on the computer, it must be selected before a calculation can be made. The standard program options are set up when the program was installed but, because they may have been changed, they are reviewed automatically at when this option is chosen. The user may select to input completely new data, or may choose to modify an existing **DSWT.DAT** file. The details of the input menu windows are described in Section 5.

Reset the Program Options

The program options specify the input units, the gas composition, and the factor for the reduction in cell height adjacent to the surface. The input window entitled **Program Options** appears when this selection is made.

The default input units are metric (SI), but the "US engineering" variant of British units may be chosen. The default gas composition is 20% oxygen and 80% nitrogen. These fractions may be changed and a fraction of atomic oxygen may be included. The cells that are wrapped around the surfaces normally have a height approximately equal to the size of the background rectangular cells. Smaller cells may be required if the flow is such that the density at the surface is very large in comparison with the stream density. The default value of the surface cell height ratio is unity. The default value may be reduced by a factor of up to twenty. The cell height distribution across the surface layer is an arithmetic progression and the number of cells is increased to keep the height of the outer cells approximately equal to the size of the rectangular cells.

About Program and Method

This option brings up a menu of four screens with information about the DSMC method and the DSWT program.

5 Data Input Screen and Menus

CREATE OR MODIFY DATA Menu

Generate Completely New Data

This option must be selected if a valid **DSWT.DAT** file does not exist. All data items must be specified, but most of the menus contain default values which may correspond to the desired values.

Modify the Existing Data Set

This is the default option, but a **DSWT.DAT** file must exist and its values are used to initialize the menus. Note that the input reverts to the first option if there is any change to the number of surfaces or the number of segments along any of the surfaces.

Note that the **DSWT.DAT** file contains a description of the data items and these mat be edited directly as long as the number of surfaces and the number of segments along any of the surfaces are not altered.

SET THE COMPUTER MEMORY Menu

This requires the user to specify the approximate amount of computer memory that is to be used by the program. This is the only data item that is not directly related to the problem specification. All other computational variables are set automatically by the program. As with the previous menu item, the choice is recorded on the right half of the screen.

Two-Dimensional Plane Flow Menu

Two-Dimensional Plane Flow

This is the default selection and sets up a flow in the *x-y* plane on the assumption that there are no gradients in the *z* direction.

Axially Symmetric Flow

This assumes that the flow has cylindrical symmetry about the axis *y=0*. The *y* co-ordinate becomes a radius.

Flow Symmetry Menu

This appears only when for two-dimensional plane flows (axially symmetric flows are always symmetric about y=0).

Symmetry is Not Assumed

Flow is Symmetric About y=0.

The first is the default choice and allows the definition of asymmetric bodies. It also permits the specification of a stream velocity component in the *y* direction. The second option sets the lower boundary of the flowfield to *y=0* and this is automatically defined as a plane of symmetry.

Working Section Size Menu

The working section of the simulated wind tunnel is rectangular and this menu requires values for the minimum and maximum *x* co-ordinates, the maximum *y* co-ordinate and, for the asymmetric case, the minimum *y* co-ordinate. An image of the working section appears as a black rectangle in the lower right corner of the DSWT window, and the range of co-ordinates is printed for each axis.

Number of Surfaces Menu

This requires an integer for the number of separate surfaces. There can be any number of surfaces, but every surface must be either a “closed” surface with identical initial and final points, or an “open” surface which must start and terminate on one of the boundaries of the working section.

The following menus are in a loop over the number of surfaces.

(Headings on the menus indicate the code numbers of the surfaces and segments.)

Surface Definition **Menu**

An integer is required for the number of segments (each segment is either a straight line or a circular arc) along the line.

In addition, the user must select between the **Input of the Individual Segments** via additional menus or **Read from a Polyline File**.

If segments are input individually, the following menus appear:

Set Initial Point **Menu**

This sets the *x* and *y* co-ordinates of the initial or starting point of the surface. In choosing the starting point, it is essential to keep in mind that the flow is on only one side of the surface and, if a thin surface is required, both sides must be defined. The flow is on the **right hand side** of the line looking from a point to the next point that is, or is to be, defined.

The following menus are in a loop over the number of segments:

Segment Type **Menu**

The segment is defined as a **Straight Line**, an **Arc Clockwise from Initial Point**, or an **ARC Clockwise from Initial Point**.

If it is a straight line segment, the following menu appears:

Final Point *(of Segment)* **Menu**

The *x* and *y* co-ordinates of the end point of the segment are input to this menu. As soon as these are accepted by pressing the “OK” button, the segment is drawn on the working section plot. An arrow pointing to the flow side of the surface is also drawn. This allows the user to verify that the points are being input in the correct order.

Else, if an arc segment is selected:

Arc Data **Menu**

The menu shows the initial point of the segment and the direction of the arc. It requires the input of the co-ordinates of both the final point of the segment and the center of the circle that defines the arc. These values must be consistent and, if inconsistent values are set, a warning window appears and the *Arc Data* Menu reappears. The arc and an arrow appears on the working section plot when the value is accepted.

End of loop over the segments.

Else, if the segment information is input from a polyline:

The standard Windows file selection menu appears. The selected file must be a text file with separate lines, each containing the co-ordinates of successive points along the line. The co-ordinates must be separated by spaces or by spaces and a comma. The number of lines must be equal to the number of segments plus one.

The polyline file may have been produced by a separate program that did not follow the convention that the flow must be on the right hand side as one proceeds from point to point along the line. Therefore, after the polyline has been read, the ***Forward or Reverse Order of Points*** **Menu** appear. This provides the option to either **Retain Point Order** or **Reverse Point Order**. This allows the flow to be set on the desired side of the line and, following this selection, the whole surface and the arrows on the segments appear on the working section plot.

End of segment information.

Surface Properties **Menu**

This specifies the way in which the simulated molecules are reflected from the surface.

Diffuse Reflection at a Set Temperature

This model is appropriate to “engineering” surfaces and will be the model that is chosen for most problems. This program assumes that each surface is at a uniform temperature and this is set as the single numerical input item on the menu.

Diffuse with Zero Heat Transfer

The diffuse reflection model is again employed, but the program adjusts the surface temperature automatically to the value at which the energy carried away by the reflected molecules is equal to the energy transferred to the surface by the incident molecules. It is therefore a zero heat transfer or adiabatic surface temperature. The adiabatic surface temperature distribution is then one of the output items.

Specular Reflection

This is an idealized limiting case with perfectly elastic reflection from the surface. The velocity component normal to the surface is reversed, while that parallel to the surface is retained. There is therefore zero shear stress and heat transfer and equal contributions to the pressure by the incident and reflected molecules. While it is an artificial model, it is very useful for studying a flow with no boundary layer. It is also useful for setting planes of symmetry within a flow.

End of loop over the surfaces.

***Stream Properties* Menu**

The number density, temperature and velocity of the uniform stream are defined in this menu.

The number density (number of real molecules per unit volume) may be an unfamiliar property to many uses, but is simply the familiar gas density divided by the average mass of a single molecule. It, together with the gas temperature (which affects the cross-section of the molecules) sets the mean free path in the gas.

The magnitude of the mean free path is a critical consideration in setting the data because a good DSMC calculation is not possible unless there are at least about five simulated molecules in each square mean free path of the flowfield. The input windows are therefore dynamically linked to an output window in the menu that displays the approximate value of the mean free path.

The input windows are also dynamically linked to output windows that shows the Mach number corresponding to the current stream definition and also the Reynolds number per unit length that corresponds to the Mach number and mean free path.

The stream definition input values may therefore be adjusted until the desired mean free path and Mach number are obtained and is then selected by clicking on the OK button. (Because of the dynamic items in this menu, it is a different class of window to the other menus and the alternative selection mechanism through pressing the enter key is not available.)

***The Data Set is Complete* Menu**

This gives the user an opportunity to **Review and/or Correct the Data** before the calculation is commenced. This is the only opportunity to go back through the menus and alter the data. Therefore, should an error be made in a data item, the data input should be continued until this option can be exercised.

The **Proceed with Flow Initialization** option causes the calculation to commence.

The computer compares the number of simulated molecules that are allowed by the specified computer memory with the dimensions of the working section (measured in multiples of mean free paths in the undisturbed stream). The likely number of simulated molecules in a square mean free path is calculated and the program displays an information, warning or stop screen based on this value. The program does not allow the continuation of a calculation that is clearly beyond the available computer resources.

The next step is to set and initialize all the quantities associated with the geometry and cell structure. The cells are displayed on the working section plot and, when the initial gas has been generated, the **Initialization Complete** window appears in the center of the data input screen. The clicking of the OK button, or the pressing of the enter key, causes the flow simulation to start from zero time.

6 The Interactive Flow Display Screen

This screen is visible throughout the running of the calculation and contains three graphics windows, about thirty numeric output windows, and nearly fifty push button controls. In addition, the graphical flow display accepts left button mouse clicks to display the flow properties at a point, and right button clicks to generate flow streamlines through the point. The screen is largely self-explanatory but some comment and clarification is useful for some of the contents.

(i) The **Flow time** is the time in seconds since the surfaces were instantaneously inserted into the uniform stream. The flow time advances in steps equal to the one sixth the smallest value of the sampled mean collision time in any of the cells. The flow properties are sampled at set intervals and, until and unless the flow becomes steady, the samples are reset each time the restart files are written. The flow and surface property displays are therefore time-averaged over time intervals that are small in comparison with the steady flow time and give a realistic picture of the unsteady flow development process.

(ii) The **Molecule Number History** window plots the number of simulated molecules over a time interval that immediately precedes the current flow time. There is usually an increase in the number of simulated molecules during the unsteady phase of a flow but, when a flow has become steady, the fluctuations should be characteristic of the statistical fluctuations. These have a standard deviation of approximately the inverse square root of the number of simulated molecules. The magnitude of the standard deviation is indicated by the length of the double-ended yellow arrow at the center of the window. The scatter in a steady two-dimensional flow shows the expected characteristics. Axially symmetric flows employ radial weighting factors that cause molecules to be discarded and duplicated as they move towards and away from the axis. This does not affect the magnitude of the average number, but it may affect the nature of the fluctuations. This is because the discarding and duplication process means that mass momentum and energy are not conserved exactly, but also only on the average. The lack of exact conservation can lead to the superimposition of a random walk on the physically real scatter. Since the molecules in the working section are being constantly regenerated, the random walk should not build to the level where it significantly affects the results. If problems do occur, they will be most serious at low Mach numbers. Care sould therefore be exercised when calculating low speed axially symmetric flows.

The program makes an initial estimate of the time to steady flow based on the transit time of the stream past the body. As this time is approached, the number of simulated molecules is compared with the maximum allowable number and is increased by random molecule duplication to more nearly fill the available memory. This step increase in the molecule number is visible in this screen. Should the number of molecules depart by more than about seven standard deviations from the value at the start of the sampling interval, the sample is reset. Should a flow become periodic rather than steady, it should be clear from the behavior of the total molecule number in this screen.

(iii) The log of the computation in the top left hand corner of the screen shows the number of **molecule round-off substitutions** that have been made during the calculation. A molecule is removed from the flow and is replaced by randomly selected molecule from the same cell if the computer is unable to fix its position within the flow. This problem arises from the six or seven significant digit limitation of 32 bit arithmetic. The internal calculations are made with metric units and will be no substitutions in flows with low to moderate flow speeds and typical lengths of order unity. The number of removals may exceed one in a thousand of the entering molecules if the physical dimensions of the flow are very small and/or the flow speed is extremely high. This level of removals borders on the unacceptable and the errors can be noticeable if they are concentrated near the axis of an axially symmetric flow.

In the case of small physical dimensions, a similar geometry can be chosen with much larger size and the number density adjusted to maintain the Knudsen number. The non-dimensional output should then be unaffected, but adjustments would have to be applied to the dimensioned quantities and the flow time. The problem with the extremely high velocities is that the time step becomes extremely small and this can lead to round-off errors even when the lengths are of order unity.

There should be no problem with round-off on a 64 bit computer.

(iv) The units are quoted for dimensioned output and the forces are based on a flow of unit width in two dimensional flow or the full 360° azimuth in an axially symmetric flow. The **non-dimensional** output option requires specific definitions for some of the output quantities.

The flow quantities are expressed as their ratios to selected undisturbed stream quantities. The non-dimensional positions are the x and y co-ordinates divided, respectively, by either the length or height of the working section. The speeds and velocity components are normalized by the speed U_∞ of the undisturbed flow. The overall temperatures and the separate components of the temperatures are normalized by the overall temperature of the stream. The pressure is normalized by the pressure p_∞ in the stream. The densities and number densities are normalized by the stream density r_∞ and number density n_∞.

A number flux coefficient is $\dot{N}/(n_\infty U_\infty)$.

A surface pressure coefficient is $(p - p_\infty)/(\frac{1}{2} r_\infty U_\infty^2)$.

A shear stress coefficient is $t/(\frac{1}{2} r_\infty U_\infty^2)$.

A heat flux coefficient is $h/(\frac{1}{2} r_\infty U_\infty^3)$.

A characteristic length or area is required for the normalization of the surface forces and moments. The input of this parameter is through a pop-up menu that appears the first time that non-dimensional output option is selected. The force coefficients are the force divided by the product of $\frac{1}{2} r_\infty U_\infty^2$ and a characteristic area. For two-dimensional cases, the force is that applying to a surface of unit width and the characteristic area is the characteristic length times unity. For symmetric two-dimensional bodies, the force is that applying to only one side of the body. The force and the characteristic length are doubled if the option to display both sides of the flow is chosen. For axially symmetric flows the force is for a surface that extends over the full 360° azimuth and the characteristic area must be consistent with this. The moment coefficient appears only for two-dimensional cases without symmetry and the denominator includes the square of the characteristic length.

(v) In flows with stagnation temperatures that are sufficiently high for the vibrational modes to be included, there is an apparent anomaly in that the vibrational temperature is usually output as zero in a low temperature stream. This is a result of the quantum model that is employed in the program. The minimum vibrational energy of any molecule is the Boltzman constant times the characteristic vibrational temperature. The characteristic vibrational temperatures in air are of the order of two or three thousand degrees Kelvin and, at temperatures of the order of one or two hundred degrees, only about one in a million molecules would be in the first vibrational level. This means that there is likely to be no vibrational energy in the number of molecules that are involved in a typical DSMC sample. The amount of vibrational energy at the low temperatures is so small that the zero values are of no consequence. To balance the zero vibrational temperatures in most samples in a low temperature gas, a small sample that happened to include one molecule at the first vibrational level would be assigned an extremely high vibrational temperature. These would appear even more anomalous to the casual user and they have been suppressed in the output. The calculation contains the correct average amount of vibrational energy despite the exaggerated scatter in the vibrational temperature at low overall temperatures.

(vi) The ratio of the local mean separation of the collision partners to the local mean free path is the key verification parameter for the calculation. The largest, and therefore the worst, value in any cell is output in the flow log. There could be some error in the calculation unless this is small compared with unity. On the other hand, higher values might be tolerated in a small region of the flowfield and the largest value might be an extreme fluctuation in one of the very small cells at the boundary of the surface cell region and the rectangular cells. Therefore, if the largest value is undesirably large, the **Coll. Sep. / m.f.p** option for the property in the flow display. The distribution of the ratio can then be studied and a judgement made about the quality of the calculation. The ratio can be reduced only by the use of a larger number of simulated molecules in the calculation.

(vii) The average number of simulated molecules in the cells may also be selected for the flow property display. This is a computational parameter and, unlike the flow properties, aspects of the cell structure are apparent in the display. The smallest numbers will generally be associated with the small cells that are remnants of the rectangular cells that are not quite cut out when the surface cells are superimposed on the background grid. Small cells are also occur adjacent to sharp corners in the surface cell regions, while the outermost cells in the corner regions are larger than normal. The total sample in the cell is obtained by multiplying the average number of molecules by the total number of samples that is shown in the flow log.

(viii) Axially symmetric calculations employ radial weighting factors that, in a uniform stream, equalise the number of molecules in the evenly spaced rows of cells in the *y* direction. This means that molecules may have to be duplicated or discarded as their distance from the axis changes. The number of real molecules represented by each simulated molecule is proportional to the radius. This effectively increases the number of simulated molecules by a factor that is inversely proportional to the radius. This factor cannot go to infinity at the axis and the radius of the innermost row of cells has been employed as a cut-off radius for the weighting factors. The leads to poor statistics very near to the axis and there can be some anomalies in the flow values very near to the axis.

Appendix A Using TECPLOT output

The built-in flow displays should be sufficient for most applications and either the whole screen or the separate flow and surface property windows can be output as atandard Windows (**.BMP)** bitmap files. Should post-processing be desired in order to obtain higher quality output and the TECPLOT program is available, TECPLOT compatible output files **DS2ASU.TEC** and **DS2AFF.TEC** may be output. This output is an optional selection on the 1024×768 screen and is always activated with the 600×800 display mode is being used.

DS2ASU.TEC is for the surface properties and the descriptive column headings are:

S the distance along a surface from the initial point
NF the number flux to the surface
PR the pressure
SSP the shear stress
SSC not used
EI incident energy flux
ER reflected energy flux
ET net energy flux
TS surface temperature
N1 the sample of species 1
N2 the sample of species 2 etc....

DS2AFF.TEC is for the flowfield properties and the descriptive column headings are:

X the x coordinate
Y the y coordinate
ND number density
D density
U u velocity component
V v velocity component
W not used
TTR translational temperature
TRT rotational temperature
TVB vibrational temperature
TOV temperature
MA Mach number
MC molecules per cell
MCT mean collision time
MFP mean free path
SOF ratio (mean collisional separation) / (mean free path)
F1 the fraction of species 1
F2 the fraction of species 2 etc...

The surface file is quite straightforward with a separate zone for each surface, and a line for each surface element. No additional instructions are needed for the surface file.

Zone 1 of the flowfield file is for the values in the interior of the flow. These are at the center of the cells with one line for each cell. Because TECPLOT assumes that flow values are at grid points rather than at the cell centers, additional zones have been added for extrapolated values along the flow boundaries. There is one zone for each closed line that defines a flow boundary. For example, if there is one isolated body in the working section, there will be one zone for the values on the body and one zone for the boundaries of the working section. Because of the unstructured grid, TECPLOT must triangulate the flowfield. To do this:-

Apply TECPLOT's PREPLOT program to DS2AFF.TEC to produce DS2AFF.PLT.

Apply TECPLOT to DS2AFF.PLT.

In the **Tecplot** window,

(i) Select **2D**

(ii) Select **Data / Triangulate**

In the **Triangulate** window:

(1) Select **ZONE 001** as the **Source Zone**
(2) Check **Use Boundary Zone(s)**
(3) Select **ZONE 002** and higher numbered zones as **Boundary Zone(s)**
(4**)** Press the **Compute** button

In the **Information** window:

Press the **OK** button.

(5) Press the **Close** button.

(iii) Select **Data / Delete Zone**

In the **Delete Zone** window:

(1) Select all zones other than the **Triangulation** zone.
(2) Press the **Delete** button.

In the **Information** window:

Press the **OK** button.

(3) Press the **Close** button

TECPLOT may now be used to draw contours, show streamlines etc.