

THE NUMERICAL SIMULATION OF RADIATIVE SHOCKS. II. THERMAL INSTABILITIES IN TWO-DIMENSIONAL MODELS

RALPH S. SUTHERLAND, GEOFFREY V. BICKNELL, AND MICHAEL A. DOPITA

Research School of Astronomy and Astrophysics, Institute of Advanced Studies, The Australian National University,
 Cotter Road, Weston Creek, ACT 2611, Australia; ralph.sutherland@anu.edu.au

Received 2002 December 25; accepted 2003 March 11

ABSTRACT

We present the results of high-resolution hydrodynamic simulations of stable and overstable radiative shocks. A one-dimensional resolution study, incorporating both power-law and realistic astrophysical cooling functions, agrees well with analytical solutions both in the spatial structure of the shocked zone and in the frequencies of overstable oscillations. This builds upon previous work by Strickland & Blondin, evaluating the accuracy of our code and estimating the resolution required to construct credible multi-dimensional models of interstellar radiative shocks. These models show that accurate modeling of the spatial and temporal structure induced by cooling processes in a multidimensional hydrodynamic simulation requires high resolution. We then present inhomogeneous two-dimensional models with varying input density fluctuation spectra and show that the resulting postshock density and velocity structures are largely independent of the initial seed fluctuation spectrum and that small fluctuations can result in a dense filamentary structure in two dimensions being fully developed in a single cooling timescale. These inhomogeneous two-dimensional structures are described by a fractal dimension, which takes a characteristic value in these two-dimensional simulations. Cooling inhomogeneous shocks have enhanced cooling efficiency, due to their fractal structure, compared to homogeneous one- and two-dimensional models. The increased radiative efficiency is accompanied by a decrease in the conversion of kinetic to thermal energy as the additional degrees of freedom in the two-dimensional models allow kinetic energy to be redirected in other directions, resulting in two-dimensional turbulence.

Subject headings: hydrodynamics — instabilities — plasmas — shock waves

1. INTRODUCTION

A physical understanding of the role of thermal and dynamical instabilities is critical to our ability to model many classes of astrophysically interesting radiative shock waves, especially those with velocities in the range of 150 to $\sim 1500 \text{ km s}^{-1}$. These are encountered in such diverse environments as active galaxy narrow-line regions, high-redshift radio galaxies, supernova remnants, and the jets associated with newly born stars. While detailed spectral predictions from steady-flow shock models have been available for decades (Cox 1972; Daltabuit & Cox 1972; Dopita 1977; Raymond & Smith 1977; Raymond 1979; Dopita, Binette, & Tuohy 1984), it has been known for almost as long that steady, time-independent models may not be a good approximation for high-velocity shocks. Steady, plane-parallel shocks have been computed using the MAPPINGS III code. Although these models incorporate expanded ionization and temperature ranges, they neglect the important effects of local thermal instabilities (Sutherland 1993; Sutherland & Dopita 1993; Dopita & Sutherland 1995, 1996) as well as global shock oscillations induced by cooling. The latter occurs when a shock front initially collapses as a result of efficient post-shock cooling followed by reexpansion driven by the high pressure of freshly shocked gas.

1.1. Local Plasma Thermal Instability

The onset of thermal instability in cooling plasmas depends upon the characteristics of the cooling function and occurs whenever the cooling processes tend to enhance preexisting density fluctuations as the plasma cools. In particular, a hot gas is unstable to isobaric (constant pressure)

perturbations if

$$\left(\frac{\partial \dot{Q}}{\partial T}\right)_P > 0 \quad (1)$$

(Field 1965), while the gas will be thermally unstable to isochoric (constant density) perturbations if

$$\left(\frac{\partial \dot{Q}}{\partial T}\right)_V > 0 \quad (2)$$

(Parker 1953). In shocks, cooling occurs at (almost) constant pressure, so that the first condition is appropriate. Writing this explicitly in terms of the cooling function, Λ , and the temperature, T , the plasma is thermally unstable when

$$\frac{d \ln \Lambda}{d \ln T} < 2. \quad (3)$$

An inspection of the time-dependent, noncollisional equilibrium cooling functions given by Sutherland (1993) reveals that the slope of the cooling function, Λ , turns over and becomes negative between 2×10^5 and $\sim 3 \times 10^7 \text{ K}$ and that generally equation (3) is satisfied for $T \gtrsim 2 \times 10^5 \text{ K}$.

1.2. Shock Oscillation Instability

Chevalier & Imamura (1982) showed that when a shock occurs in a radiative plasma, the entire shock can oscillate as a result of thermal instability, in a cycle of collapse and reformation of the shock. This instability is related to the local thermal instability, and the criteria for

power-law cooling is similar to equation (3) above, with a more severe limit,

$$\frac{d \ln \Lambda}{d \ln T} < 1. \quad (4)$$

That is the shock will be unstable to large-scale oscillations of the shock-front location when the index of cooling is less than 1.0. This condition on the nonequilibrium interstellar medium (NEQ) cooling function is essentially the same as above, i.e., when $T \gtrsim 2 \times 10^5$ K, as this is the same point where the steep (and stable) cooling below 1×10^5 K turns over to much less steep cooling at higher temperatures.

For a fully ionized plasma of solar metallicity, the post-shock temperature, T_s , is related to shock velocity, v_s , by

$$T_s = 1.38 \times 10^5 (v_s / 100 \text{ km s}^{-1})^2 \text{ K} \quad (5)$$

(Dopita & Sutherland 2002). Thus any shock with a velocity greater than $\sim 120 \text{ km s}^{-1}$ can oscillate in an overstable fashion. This covers the entire range of autoionizing radiative shocks computed in Sutherland (1993) and Dopita & Sutherland (1995, 1996).

This was also seen in Innes et al. (1987a, 1987b), where one-dimensional hybrid hydrodynamic and spectrophotometric models in the regime between 150 and 200 km s^{-1} were shown to oscillate. These supported the expectation that normal shock diagnostics (i.e., the line ratios that are sensitive to the shock velocity) are invalid when the shock becomes thermally unstable. Later, Innes (1992) suggested that the oscillating process could result in limit cycle behavior—manifest as pulsations in the one-dimensional propagating shocks—that in a time-averaged sense could produce similar spectra to steady flow models.

1.3. Multidimensional Shocks

It remains an open question as to whether two- and, most importantly, three-dimensional thermally unstable shocks pulsate or whether the higher dimensionality allows other modes of to occur. Indeed, it is possible that pulsations may become incoherent along the face of an extended shock, reducing the shock region to a quasi-steady state. The effects of two- and three-dimensional local thermal instability, forming differentiations in temperature and density in the postshock gas (filaments, voids, and clumps) on the global shock oscillations needs to be investigated.

Another limitation of steady one-dimensional shock models, such as produced by the code MAPPINGS III is the assumption of an infinite plane-parallel geometry. This has direct implications for the radiative transfer models. When the homogeneous and steady state assumptions are removed, one of the key questions is that of radiative transfer throughout the volume; in particular how much of the downstream and upstream radiation fields are affected by more complex geometry?

Recently, detailed and spatially resolved observational studies of interstellar shocks have become available. These are ideal for comparison with theoretical models and include work by Danforth, Blair, & Raymond (2001) on the Cygnus supernova remnant (SNR) and by Ghavamian et al. (2000) on the Tycho remnant. These studies reveal complex multiple shock structures that are clearly neither plane-parallel nor steady. Moreover, the line diagnostics in the Cygnus SNR suggest shock velocities in the critical velocity

domain near and above 150 km s^{-1} , where thermal instabilities are important. The filamentary structures in the Danforth et al. (2001) observations are strongly suggestive of highly inhomogeneous preshock gas, which in turn can be reasonably expected to cool in an unstable fashion.

In this paper we describe our development and testing of the computational tools that will allow us to directly model thermally unstable, nonsteady shocks in time-dependent, three-dimensional simulations. In a previous paper (Sutherland, Bisset, & Bicknell 2003, hereafter Paper I) we described a method for eliminating a numerical shock instability that is a feature of many directionally split codes. In this paper we describe the construction of the code that is used in these studies and the results of one-dimensional simulations that are used to test fundamental aspects such as the grid resolution, which is necessary in order to properly describe radiative shocks.

We then present inhomogeneous two-dimensional models where the gas entering the shock front is given 5% density fluctuations with two different fluctuation spectra. We analyze the statistical properties and fractal structure of the solutions and compare these with steady or homogeneous models. These two-dimensional models provide us with insight into the physical properties of thermally unstable shocks. In future papers we shall extend this work to the computation of the full three-dimensional structure of radiative fast shocks.

2. DESCRIPTION OF THE CODE

For this, and for other purposes, we have developed a new hydrodynamic code based on the University of Virginia VH-1 *ppmlr* code (Blondin & Lukin 1993), which is in turn based on the seminal work of Colella & Woodward (1984). We have largely reorganized the code for improved efficiency, improved its Riemann solver, incorporated radiative cooling, and incorporated new options for eliminating numerical instabilities. In the context of radiative shocks, the most important of these is the striping instability that is seeded at the shock front, which is subsequently amplified in the postshock flow region perpendicular to the shock front, appearing as striated fluctuations in temperature, density, and pressure. This instability occurs in many directionally split codes. This has been eliminated by the use of an “oscillation filter,” which was described in detail in Paper I, where other approaches to eliminating this instability are also discussed.

We have not used an adaptive mesh because the advected structures introduced onto the grid in the two-dimensional models are fractal in nature. Consequently, any grid refinement would be forced to the maximum resolution over most of the grid in order to preserve the high spatial frequencies, which occur throughout the model domain. The one-dimensional tests performed here therefore also use a fixed-resolution grid so that the conclusions drawn from the one-dimensional models are then applicable to the final two-dimensional models.

Magnetic fields are neglected in the present simulations. While magnetic fields may play an important role in determining the full development of the filamentary structures in interstellar shocks, it is also true that in the hot postshock region the thermal pressure greatly exceeds the magnetic pressure even if the thermal and magnetic pressure are in equipartition in the precursor medium. Thus, the

development of filamentary structure in the cooling layer is initially controlled by the characteristics of thermally unstable cooling. Since the ratio of the magnetic to gas pressure increases in the postshock gas in proportion to the density, eventually the dense cooled layer is supported primarily by magnetic pressure, and the gas then continues to cool not only at constant pressure but also at constant density. In the simulations presented here we have eliminated cooling below 10^4 K so that the dense gas is artificially kept from becoming too dense. This mimics what occurs when the gas is supported by magnetic pressure. As long as this dense layer remains comparable in size to that computed for the shock tail in a model with a preshock equipartition magnetic field and steady flow cooling, then the dynamical behavior should be similar between the magnetically supported model and the field-free hydrodynamic simulation. To undertake magnetohydrodynamic simulations of sufficiently high resolution in multiple dimensions remains a challenging computational task, which we hope to undertake in the future. For the time being, however, we limit ourselves to consideration of the nonmagnetic case.

2.1. Treatment of Radiative Cooling

We have used the one-dimensional time-dependent plasma-modeling code MAPPINGS III to compute the nonequilibrium cooling function and ionization state of a plasma with the standard solar abundance set taken from Anders & Grevesse (1989). The MAPPINGS III code is described in a number of papers (Dopita 1977; Binette, Dopita, & Tuohy 1985; Sutherland 1993) and has been extensively tested against other codes (Ferland et al. 1994). MAPPINGS III can be used to compute the time-dependent ionization and cooling for up to 16 atomic species over a range 10^4 – 10^8 K, for both photoionization models and models dominated by collisional ionization.

We use an explicit scheme to compute the cooling in the *ppmlr* code. The time step is controlled by the cooling time in addition to the usual Courant condition. The cooling is computed in each cell after the hydrodynamical sweeps, introducing pressure gradients to which the gas flow reacts. The effect of the cooling after a time step Δt , in a given cell $C(i, j)$ is evaluated by reducing the initial pressure $P_0(i, j)$ using an explicit Runge-Kutta (RK) integration:

$$U_0 = P_0(i, j)/(\gamma - 1), \quad (6)$$

$$U_1 = U_0 + [\lambda_1 + 2(\lambda_2 + \lambda_3) + \lambda_4]/6, \quad (7)$$

$$P_1(i, j) = (\gamma - 1)U_1, \quad (8)$$

where

$$\lambda_1 = -\rho_0^2 \Lambda(U_0) \Delta t, \quad (9)$$

$$\lambda_2 = -\rho_0^2 \Lambda(U_0 + \lambda_1/2) \Delta t, \quad (10)$$

$$\lambda_3 = -\rho_0^2 \Lambda(U_0 + \lambda_2/2) \Delta t, \quad (11)$$

$$\lambda_4 = -\rho_0^2 \Lambda(U_0 + \lambda_3) \Delta t. \quad (12)$$

The NEQ cooling function $\Lambda(U)$, was evaluated from MAPPINGS III models of plasma cooling isobarically from high temperatures, and, together with its derivative, was tabulated over 1024 logarithmic temperature intervals from 10^3 to 10^{10} K; these tables are used for interpolation during a simulation.

The NEQ functions differ from the cooling function derived from a single-velocity steady shock model. The principal difference is the absence of a short-lived strongly cooling phase immediately after the shock front, where low-ionization preshock material suddenly becomes hot. This low-ionization gas rapidly ionizes in this region, strong collisional excitation also occurs, and a spike of efficient cooling results. However, the timescale of this phase is very short, being only $\sim 1\%$ of the steady flow in extent. The integrated losses in this region are therefore small, despite the high cooling rates. Since the shock front in the finite grid simulations is already subject to excessive cooling as a result of the interpolation of intermediate temperatures in the shock front, we prefer to use the smooth NEQ function in order to avoid exacerbating the cooling error in the shock front. For shocks above approximately 150 km s^{-1} , the NEQ function is a good approximation to the actual shock cooling function, apart from the initial spike; the postshock gas is hot enough to come close to equilibrium before cooling in shocks of this velocity. The NEQ function cools almost at equilibrium down to about 10^6 K, and all shocks above 200 km s^{-1} will cool similarly down to this temperature. The NEQ function may be a poorer approximation for shocks of 100 km s^{-1} or less and to solve this problem ionization calculations will need to be computed in the hydrodynamical models when determining cooling rates. Finally, the calculation of the maximum allowed time step, Δt during a simulation step includes the local cooling timescale, in addition to the timescales involved in the usual Courant condition. We limit the time step to be no longer than 10% of the shortest cooling time on the computational grid by taking the harmonic mean of the sound and flow crossing timescales and the cooling timescale.

The Runge-Kutta method for integrating the internal energy was compared with higher order and other implicit methods (such as reverse Eulerian) and was found to give comparable results with minimal computation. The nonequilibrium cooling function is smooth, and the Runge-Kutta method was found to be accurate to a level $\sim 10^{-6}$ when compared to analytical solutions for power-law cooling. This method of computation is preferable to a method in which the cooling is computed during the hydrodynamic sweeps because at the high Mach numbers here (15 and above) the internal energy of the gas is a small component of the total energy, and errors in the cooling estimates can at times produce unphysical negative pressures from the direct energy equation integrator. Energy conservation and the consistency of this method was confirmed (see § 5.4).

We also note that even with relatively low diffusion, as in high-order schemes such as the piecewise parabolic method (PPM), there is always the possibility of excessive cooling in zones separating hot and cool phases (that usually involve low and high densities). In regions of steep gradients (in density, temperature, etc.), the use of zone-averaged quantities to compute the cooling [i.e., $\langle \rho \rangle^2 \Lambda(\langle T \rangle)$, rather than $\int \rho^2 \Lambda(T) dx$] can overestimate the cooling, no matter how fine the discretization of $\rho(x)$. Rather than attempt to reduce this problem by limiting the local cooling rate (i.e., Blondin, Konigl, & Fryxell 1989; Stone & Norman 1993), we refine the grid until the cooling behavior of the particular simulation, wall shocks in this case, approaches an asymptote. We then determine a resolution that brings the cooling errors down to acceptable levels, where the local gradients are small enough that the zone averages are adequate.

Determining the asymptotic behavior of large simulations (such as shock oscillation frequencies) is expensive in computing terms. Nevertheless, the low-resolution models in any resolution study are generally quick, and time is mostly spent on the high-resolution models. This procedure would need to be implemented for different simulation geometries, i.e., a supernova blast model. There is no universal resolution that applies to all problems.

2.2. Normalization

In computational hydrodynamics, it is common practice for adiabatic simulations to normalize the flow variables, for example, by setting the sound speed equal to unity. However, with the introduction of a nonequilibrium cooling function, which is not scale free, renormalization of the hydrodynamical variables is more complex. In some cases, numericists, utilizing the precision of 64 bit arithmetic, do not normalize the dynamical variables. Nevertheless, using cgs units directly, can be excessively demanding on numerical precision, so we have chosen to normalize the simulations internally as follows.

Assuming $\gamma = \frac{5}{3}$, the ideal atomic gas ratio of specific heats, we adopt fiducial values for the spatial scale x_0 , the velocity v_0 , and the density ρ_0 . Implied scaling factors for the pressure P_0 , time τ_0 , temperature T_0 , and cooling Λ_0 are

$$P_0 = \rho_0 v_0^2, \quad (13)$$

$$\tau_0 = x_0 / v_0, \quad (14)$$

$$T_0 = [(\mu m_{\text{amu}}) / k] v_0^2, \quad (15)$$

$$\Lambda_0 = P_0 / (\tau_0 \rho_0^2), \quad (16)$$

where $\mu = 0.6224$ is the mean molecular weight per particle (electrons included), m_{amu} is the atomic mass unit, and k is Boltzmann's constant. These factors are then used as multiplicative factors to change the internal scaled quantities to physical quantities,

$$a_{\text{phys}} = A_0 a_{\text{int}}, \quad (17)$$

where a_{phys} is the physical quantity, a_{int} is the internal normalized quantity, and A_0 is the appropriate scaling factor given above.

3. ONE-DIMENSIONAL WALL-SHOCK FLOWS

In this paper, we have adopted the well-parameterized and theoretically well-understood wall-shock model (Chevalier & Imamura 1982, hereafter CI82; Strickland & Blondin 1995, hereafter SB95). This provides a simple characteristic shock model that can be studied in either one or more dimensions.

A wall shock is set up on a grid with a \sim Mach 15 inflow (150 km s^{-1} in the NEQ models) on one side and with a dense layer on the other. The density of this layer was fixed at the density that would be produced through a 150 km s^{-1} isothermal shock with initial density $\rho_0, \rho_0 \mathcal{M}^2$, where \mathcal{M} is the inflow Mach number. A fast reverse shock forms at the dense layer interface, while a slow forward shock moves into the dense layer and moves off the grid. An outflow boundary is applied in this region, and this prevents reflected waves from interfering with the simulation. The outflow velocity was prescribed in such a way as to maintain the dense layer with an approximately constant thickness during the course

of the simulations, effectively tracking the buildup of dense gas. This is slightly different from the ideal reflecting wall condition of SB95, where the shock was launched from the wall itself. The present simulation is established in such a way so we can study the effect on the shock of dense cooled layer that is inevitably built up in any physical situation.

There exist analytic solutions for the structure of a wall shock with power-law cooling (CI82) that enable us to determine the required resolution to achieve the desired numerical accuracy. We present comparisons with the analytical, time-independent models and the numerical, time-dependent models for power-law cooling of SB95. We then demonstrate, in one dimension, the differences from the power-law solutions that occur when we utilize an NEQ cooling function. We also estimate the resolution that is required for realistic two-dimensional simulations.

The one-dimensional solutions presented in this paper build on the substantial results reported by SB95 in the following ways:

1. We verify our methodology by comparing the results of our simulations with analytical power-law cooling models. Dynamical models with power-law cooling are compared with time-independent similarity models of CI82, with good agreement. The correct time-dependent stability of cooling of different indices is also confirmed.

2. The time-dependent dynamical models are then extended using the nonanalytical NEQ function derived from MAPPINGS III. These one-dimensional tests are important for determining the grid size required for adequate resolution of astrophysical radiative shocks, where an analytic estimate is unavailable. Adequate resolution is estimated for a uniform grid by examining asymptotic behavior with increasing resolution.

3. We conduct resolution studies of radiative shocks, examining in particular the asymptotic behavior of the bounce amplitudes and oscillation frequencies of one-dimensional NEQ shocks. These tests thus extend the tests for power-law cooling carried out by SB95.

3.1. Power-Law Cooling Solutions

The local cooling function of an interstellar plasma is often described in terms of a local power law with temperature

$$\Lambda(T) \propto \rho^2 T^\alpha, \quad (18)$$

where ρ is the density, T is the temperature, and α is the power-law index. In some circumstances, this representation can be quite accurate. For example, at high temperature, electron bremsstrahlung cooling dominates the cooling giving a power-law cooling function with $\alpha = \frac{1}{2}$. In the theoretical case that the cooling function can be represented as a power law throughout the whole temperature range, CI82 demonstrated that there exist analytical radiative wall-shock similarity solutions that depend only on the particular value of α .

For the cases $\alpha = 2.0$ and 1.0 , we found that steady-flow solutions can be obtained using the *ppmlr* code in one dimension, with 1000 elements of resolution in a uniform grid. The scaled velocity profiles that result are compared with the analytical solutions from CI82 in Figure 1. The $\alpha = 2.0$ model in particular shows a very good fit to the theoretical profile, when a small (0.5%) density undershoot




FIG. 1.—Time-independent one-dimensional power-law cooling tests. 2048 cell one-dimensional time-averaged dynamical models (*solid curves*) are compared with analytical similarity solutions from CI82 (*dashed lines*). (a) $\alpha = 2.0$; (b) $\alpha = 1.0$. See text for details.

in the postshock density is taken into account, which slightly biases the cooling and offsets the velocity curve above the strict similarity solution. This is taken as evidence that the code is integrating the cooling correctly. The dynamical curve for $\alpha = 1.0$ show significant deviations from the theoretical curve at low velocities, resulting in an underestimation of the velocity by approximately 0.1 dex in the region within 1% of the wall. This low velocity corresponds to an overestimate of the density (by mass conservation) and a consequent excess of cooling through the ρ^2 term in the cooling equation. The error arises from interpolation errors and numerical diffusion, as the piecewise parabolic interpolation fails to give a good estimate of the average density in the regions where it rapidly rises. This error can be reduced by simply increasing the resolution of the simulation, but it can never be completely eliminated using a finite grid and a low-order interpolation scheme. In both cases the curves are terminated by the lower temperature cutoff of 10^4 K.

In simulations with strong cooling, in the sense that $\alpha < 1.0$, this density error can lead to erroneously strong cooling. This directly impacts on estimates of cooling time-scales and cooling lengths. If the density error in the innermost cells is large enough, then a single cell in the simulation can dominate the overall cooling. In simulations with a fixed grid the only solution is to increase the resolution and measure the changes in the implied cooling times and lengths to determine when the error is reduced to acceptable levels. Strickland & Blondin (1995) have performed similar simulations to these, but at lower spatial resolution.

In Figure 2 we show the time-dependent behavior of 2048 cell models for $\alpha = 2.0$, 1.0, and 0.0. The expected stable, marginally stable, and unstable behaviors of these shocks is evident. In particular, the $\alpha = 0.0$ model shows a periodic behavior that remains invariant from one cycle to the next,




FIG. 2.—Time-space plots of the density evolution of one-dimensional power-law cooling tests (a) $\alpha = 2.0$, (b) $\alpha = 1.0$, and (c) $\alpha = 0.0$. Here the time axis is horizontal and increasing going to the right, the spatial extent of the shock lies in the vertical direction at any instant in time.

once it has been fully established. This is caused by excessive cooling, which robs the postshock gas of pressure support, causing the shock front to first stall, then to collapse back onto the dense cooled layer before being launched anew. We also see strong secondary shock formation, due to reflections off the dense layer at the onset of strong collapse, that are not visible in the SB95 simulations, due to the different wall boundary conditions we have here.

3.2. Resolution

Numerical diffusion on any finite hydrodynamical grid inevitably causes low-density cells that are adjacent to very dense zones to acquire excessive mass. The interpolation scheme essentially imposes a limit on the accuracy with which high-density contrasts can be modeled. The usual approach is to increase the resolution until the total volume of the simulation with poorly estimated hydrodynamic variables has been reduced to an “acceptable” level.

For radiative losses, the cooling rate is proportional to the density squared. Therefore, numerical mass diffusion of mass causes excessive cooling in low-density cells adjacent to those of high density. Under the approximately isobaric conditions in a radiative shock, the highest density cells have the lowest temperatures in which we have set the cooling rate equal to be zero. The cells adjacent to these dense cells are radiative, and since they have an incorrectly high density, they form a thin layer with anomalously strong cooling (see the two-dimensional models shown in Figs. 3c and 3d). To reduce the impact of excessive cooling in this thin layer on the overall hydrodynamic simulation, the total cooling in a cell in this layer of thickness Δs , at location X , must be much less than the total cooling that has already occurred along the flow up to that point, i.e.,

$$\rho^2 \Lambda(T_X) \Delta s < \int_X^0 \rho^2 \Lambda(T) ds. \quad (19)$$




FIG. 3.—Snapshot of the cooling rate in the region of the formation of a cool filament. The time, $\sim 6.6 \times 10^{10}$ s, corresponds to panel 3 in Fig. 12. Overdense regions at (a) the 30–50 cell scale are strongly cooling, while (b) larger low-density voids are cooling slowly. As the filament approaches 10^4 K (c), it becomes edge brightened and the dense core becomes less radiative. Eventually the center of the knot or filament ceases cooling (d) and is embedded in a strongly cooling envelope. The cooled filaments are well resolved at the 10 cell scale.

The grid resolution that is required to satisfy this condition depends upon both the cooling function and the range of temperatures involved. In order to investigate the appropriate grid size, a series of one-dimensional bouncing wall shock models with increasing resolution were constructed.

3.3. One-dimensional NEQ Cooling

In this test, grids with resolutions of 256, 512, 1024, 2048, 4096, and 8192 cells were essayed, and the extent of the “bounces,” where the shock forms, grows, and collapses after cooling takes over, was measured. Errors resulting from numerical diffusion result in an early collapse and reduced amplitude, and also an increased frequency of the bounce—all the result of unphysical cooling. The aim of this test is to discover at what resolution the shock behavior asymptotes sufficiently closely to its limit and to quantify the errors at this and other resolutions.

Figure 4 shows the density evolution of the wall shock simulation at different resolutions. Both the amplitudes and periods of the oscillations increase with resolution, with the 2048, 4096, and 8192 cell simulations tending toward the limiting maximum amplitude.

The oscillations are analyzed as power spectra in Figure 5. This can be compared to the similar plot in Figure 1 of SB95. Here, for NEQ cooling, the power spectra

show that the principal frequency approaches a limit of $1.148 \times 10^{-11} \text{ s}^{-1}$; this corresponds to a period of 2.94 cooling times for the NEQ function for a hypothetical steady 150 km s^{-1} shock, or $\omega = 0.34$ in the CI82 notation. The noise level at the highest frequencies drops rapidly to $\sim 10^{-6}$ of the peak power of the fundamental at the resolution of 2048 cells and remains at that level at higher spatial resolutions. This characteristic is shown in the right-hand panel of Figure 6. At the same time, the fundamental frequency decreases rapidly with resolution, below a resolution of 2048 cells, and then stabilizes (see Fig. 6, *left-hand panel*). For the 2048 cell simulation, the amplitude of the bounce is within 10% of the limit. For all of these considerations, the 2048 cell resolution was adopted as the standard resolution along the flow direction for all the subsequent one-dimensional and two-dimensional simulations, since it provides an acceptable compromise between accuracy and performance and also provides results in a practical length of computational time in the case of the two-dimensional simulations.

At a given shock velocity, the choice of spatial resolution determines what range of isobarically cooling spatial structures can be resolved. For a shock velocity of 200 km s^{-1} the peak postshock temperature, T_{max} , is approximately 600,000 K, assuming it is propagating into an ionized interstellar medium of solar composition. The subsequent cool-

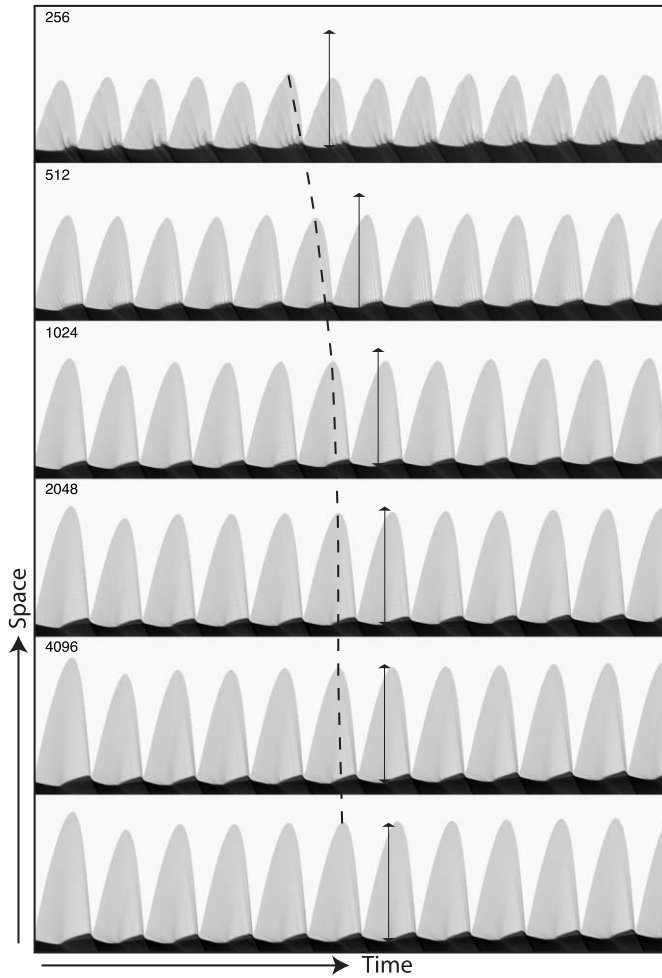


FIG. 4.—Amplitude and period as a function of resolution. The $\log(\text{density})$ (gray) variable from one-dimensional simulations is shown as a function of time in the horizontal direction and spatial coordinate in the vertical direction. Each row represents a simulation with the number of resolution elements indicated on the left. For each simulation the amplitude and period of the pulses rapidly reaches steady values as the initial transients disappear. The temporal location of the sixth bounce is indicated by the dashed line, and the amplitude of the seventh bounce in the highest resolution simulation is marked on the lower resolution simulations for comparison. The lower three panels show how the solution asymptotes as the simulation becomes well resolved.

ing under roughly isobaric conditions down to the minimum temperature (T_{\min} , 10,000 K), therefore provides roughly a 60-fold increase in density. Consider a structure of size x at T_{\max} that we wish to resolve by a minimum of 10 grid elements when it has cooled to near T_{\min} . In the immediate postshock region such a structure must have dimensions of order 600 elements. Such structures are therefore well resolved and contained in a 2048 cell grid. Since the postshock temperature increases as the square of the shock velocity, then simulations of similar structures in faster shocks will require grid resolutions that increase as the square of the velocity.

3.4. Velocity Stability Test

We have also tested the long-term behavior of NEQ shocks in a regime where they are expected to be stable. Figure 7 shows the temporal and spatial evolution of the shock structure for both a 150 km s⁻¹ inflow (corre-

sponding to ~ 200 km s⁻¹ initial shock velocity) and a 75 km s⁻¹ inflow (~ 100 km s⁻¹ initial shock). Both of these simulations are made at the same grid resolution (2048 cells), and the spatial extents are scaled to a similar amplitude in the two cases to produce cooling layers with similar numbers of grid cells. The 100 km s⁻¹ shock shows the development of a steady flow with long-term stability properties, as expected from the form of the cooling function. The 200 km s⁻¹ shock shows the overstable pulsation documented above, with a uniform asymptotic amplitude. This gives us confidence that the code properly reflects the stability properties expected of the cooling plasmas.

4. TWO-DIMENSIONAL SIMULATIONS: THE INITIAL CONFIGURATION

To study the thermal and dynamical behavior of two-dimensional radiative shocks, we have chosen to advect into the shock front an initial density perturbation that contains a spectrum of spatial fluctuations, with a wide range of wavenumbers. The introduction of density perturbations is fundamental to creating the new behavior in these two-dimensional models. In the absence of perturbations, the one-dimensional symmetry of a uniform flow is maintained indefinitely in two-dimensional by the code reproducing the one-dimensional results. All the two-dimensional models presented here include the density fluctuations that are not present in the one-dimensional models.

We study the evolution of this spectrum through the wall-shock models, similar to those described above. The sensitivity of the models to the input perturbation spectrum is tested by introducing density fluctuations with two different, but prescribed, power spectra. We then determine the long-term behavior under both types of initial conditions and compare both to one-dimensional steady-flow models computed with MAPPINGS III.

4.1. The Initial Perturbations

We define two-dimensional 2048×2048 fluctuations in the Fourier domain. We first establish a perturbation spectrum with constant amplitude and random phase throughout Fourier space. (This corresponds to white noise when inverse-transformed to the spatial domain.) We then multiply the Fourier components by a power law or Gaussian (in wavenumber) to produce spatial fluctuations with a well specified power spectrum. The spatial fluctuations, obtained from the inverse Fourier transform, are periodic in space as a consequence of the finite Fourier domain. This has the desirable property that as fluctuations are advected onto the computational grid, the fluctuations “wrap around” in each direction so that there are no sharp discontinuities in either dimension, consistent with our assumed periodic y boundaries and the continual introduction of the fluctuations over time. By specifying the power spectrum of the fluctuations directly in Fourier space, we are able to strictly define and control the characteristics of the input fluctuations.

Let k be the wavenumber in units of the reciprocal of one grid cell and let $P(k)$ be the power spectrum of the density fluctuations. Our first spectrum is a power law, $P(k) \propto k^\beta$ with an index $\beta = -2.0$, and the second is a




FIG. 5.—Power spectra of the oscillations as a function of spatial resolution. The location of the one-dimensional shock front was located and interpolated to fractional cell values with a local parabolic fit to the derivative of the pressure variable. The oscillations, as seen in Fig. 4, were then analyzed with fast Fourier transform techniques to produce a power spectrum similar to those in Fig. 1 of SB95.

Gaussian with $P(k) \propto \exp(-k^2/\beta)$, where $\beta = 8$. These spectra were further modified with a low-wavenumber cutoff leaving $k > 3$; this limits the size of the largest structures to less than the total extent of the grid by a factor of 4. The timescale for the largest structures to

cross the computational grid is $l/4v$, or $\sim 3 \times 10^9$ s, which compares with 8.5×10^{10} s for the fundamental frequency of the bounces, or ~ 25.5 times higher frequency. These advected fluctuations are therefore unlikely to resonate directly with the fundamental frequency.




FIG. 6.—Shock properties as a function of resolution for the one-dimensional NEQ wall shock. (a) The fundamental frequency. (b) The high-frequency noise level (power at the highest frequencies as a fraction of the peak power). At 2048 cell resolution and above, the fundamental frequency is steady and the noise in the simulation is independent of resolution. At a resolution of 2048 cells both the principal frequency and the high-frequency power are close to the asymptotic limits, with $\log(f) = -10.93$ and $\log(\text{relative power}) = -6.18$, respectively.




FIG. 7.—One-dimensional x - t diagrams of the density in radiative shock models that use the MAPPINGS III nonequilibrium cooling function. (a) Unsteady, 150 km s^{-1} shock and (b) steady 75 km s^{-1} shock. The time axis is vertical, increasing going down the page, as in Fig. 2 and in SB95. The stability analysis of the cooling function leads us to expect that only case (a) will give rise to a steady-flow shock, and this is confirmed in the simulation.

Subsamples of the scaled density fluctuation fields are shown in Figure 8, showing the power-law and Gaussian fluctuation spectra models.

Both fluctuation fields were scaled to provide a 5% rms fluctuation about a mean of 1.0 and were used to perturb the initial density variable in the simulations. Throughout the simulation, copies of these fields were used to update the inflow boundary densities, resulting in the structures being advected continuously onto the grid at the inflow velocity. The periodic nature of the fluctuation fields means that discontinuities do not arise in either spatial direction. The gas pressure in the initial conditions and on the boundary was constant, resulting in similar small ($\sim 5\%$ rms) temperature variations, since $T \propto P/\rho$.

In the analysis of the simulations, the evolution of the density structures was analyzed in Fourier space. That analysis reveals that during the advection, the fluctuation structure was maintained down to the $10^{-4.5}$ level. The power-law density spectrum was preserved up to $1000k$, and the Gaussian spectrum extended to 60 wavenumbers before becoming grid noise-limited. (This can be seen in the lower panels of Figs. 9 and 10.)

4.2. Initial Conditions

The initial data for the two-dimensional simulations are established in the two-dimensional equivalent of the one-dimensional wall shock above, as shown in Figure 11. The inflow velocity at the left-hand side of the grid is 150 km s^{-1} with a mean density of 1.0 particles (electrons and ions) per cubic centimeter, or about $0.6 \text{ H atoms cm}^{-3}$. The physical parameters of the simulation are specified in cgs units but were renormalized within the code as for the one-dimensional model (see § 2.2) to improve computational accuracy and to




FIG. 8.—Two density fluctuation models. The left-hand panel shows part of the power-law fluctuation field, and the right-hand panel shows the same portion with a Gaussian fluctuation spectrum. The same random number seed was used to generate the phases in both models, so there is a common overall structure with the differences attributable to the different spatial power.

avoid numerical overflows. The grid cells are nearly square, with a size that varies slightly between the power-law and Gaussian models to optimize the coverage of the grid and to ensure that the results are independent of the exact coordinates used. The grid size for the model with power-law density fluctuations was set at $2.0 \times 10^{17} \text{ cm}$ on a side, and the Gaussian model was computed on a grid of $2.56 \times 10^{17} \text{ cm}$. Each cell thus represents between 1.0 and $1.28 \times 10^{14} \text{ cm}^2$, respectively. The x dimension (2050 cells) was chosen to optimize memory access speed during the y -sweeps of the simulation. The y -dimension was chosen to be a power of 2 (2048 cells) to match the periodic two-dimensional density perturbation structures that are advected onto the grid in the x -direction. Thus, the top and bottom edges of the simulation have periodic boundary conditions, simulating a shock with infinite extent in the y -direction.

The initial inflow is assumed to be in photoionization equilibrium with the upstream photons being generated in the shock, and it is effectively adiabatic. Its temperature is set to 7500 K. This is deliberately chosen to be 25% below the 10,000 K cutoff in the cooling function, with the aim of preventing all regions of the density perturbed inflow from fluctuating in temperature to greater than 10,000 K before entering the shock. Above 10,000 K, cooling switches on, and this could alter the spectrum of the perturbations before they enter the shock. In all but the most extreme points in the initial density fluctuations this was achieved; less than 0.01% of the incoming cells have any preshock cooling, and even then, this is very weak.

The right-hand boundary, where the postshock cooled gas accumulates, is prescribed to be a free outflow boundary, and the last $0.2 \times 10^{17} \text{ cm}$ of the grid was filled with gas at a density $\rho = \rho_0 M^2$, as in the one-dimensional tests.

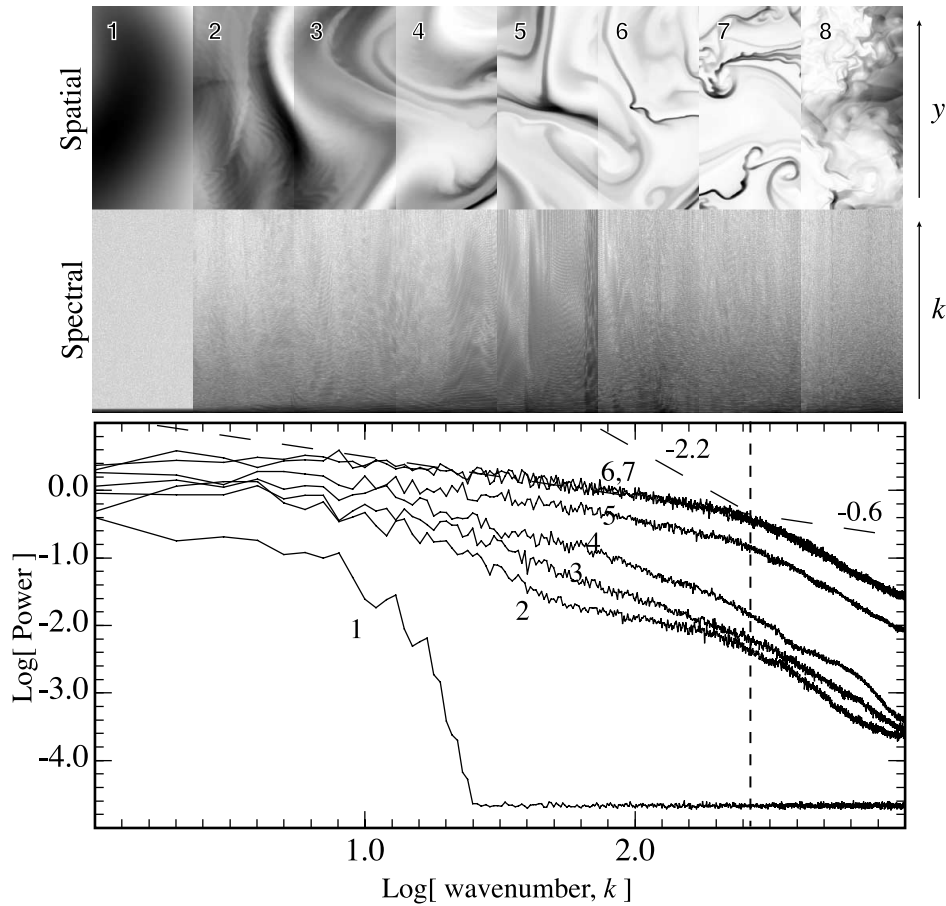


FIG. 9.—Power spectrum of the density fluctuations in the post-shock gas of the Gaussian fluctuation model. The curves in the lower panel correspond to the numbered panes in the upper panel. Two power laws are fitted to the density spectrum from panels 6 and 7. The first is fitted between $\log(k) = 1.0$ and $\log(k) = 2.4$ and has a slope of -0.6 , the second has a slope of -2.2 between $\log(k) = 2.5$ and $\log(k) = 3.0$, and the crossover point is at $\log(k) \approx 2.42$. See text for details.

An outflow velocity of 2.72 km s^{-1} was applied at the right-hand boundary to keep the base of the shock in approximately the same location over long periods of time, effectively preventing the dense layer from building up on the grid by allowing it to drift subsonically out of the right-hand side.

5. TWO-DIMENSIONAL RADIATIVE SHOCK MODELS

In the light of what was learned from the one-dimensional tests, we carried out two-dimensional simulations on a 2050×2048 grid on the ANU Supercomputing Facility's Fujitsu VPP300 computer. The choice of the grid resolution is a compromise between accuracy and performance, and with this, we expect that numerical diffusion in the simulation will affect the results at about the 10%–20% level, in both timescales and lengths, but that all relevant structures will be well resolved.

5.1. Time Evolution of the Shocks

The two-dimensional shocks evolved somewhat similarly to the one-dimensional overstable shock simulations above, but with some important differences. As noted when determining the optimum resolution (see § 3.3), increased cooling resulting from numerical diffusion leads to decreased “bounce” amplitudes and increased fundamental frequen-

cies. Figure 12 shows snapshots of the evolution of the wall shock for the power-law initial fluctuations over time. The average position (averaged along the y -direction) over time is shown in Figure 13.

The first bounce is quite large, and the hot shocked region collapses rapidly with the onset of thermally unstable cooling, allowing the formation of knots and filaments of dense material. This first bounce amplitude is similar to that seen in a well-resolved one-dimensional model. However, subsequent bounces are reduced in amplitude, by a factor of ~ 0.5 . We attribute this, at least in part, to the more complex geometry involving developing filaments and the unevenly dense shock tail; this leads to strongly cooling zones with temperatures $\sim 100,000 \text{ K}$. The filaments and voids are probably initially formed by thermal instabilities, increasing the separation between hotter and slightly cooler gas. Figures 9 and 10 (see § 5.2) show the large-scale fluctuations increasing in contrast as they evolve through the shock to form filaments. The power spectra suggest that the structures down to the cooling cutoff are not very turbulent (or at least not dominated by turbulence) but that once cooling no longer dominates structure formation (below ~ 10 cells), normal dynamical instabilities and turbulence dominate.

In two-dimensions, then, when the geometry becomes more complex, a range of new possibilities arise:

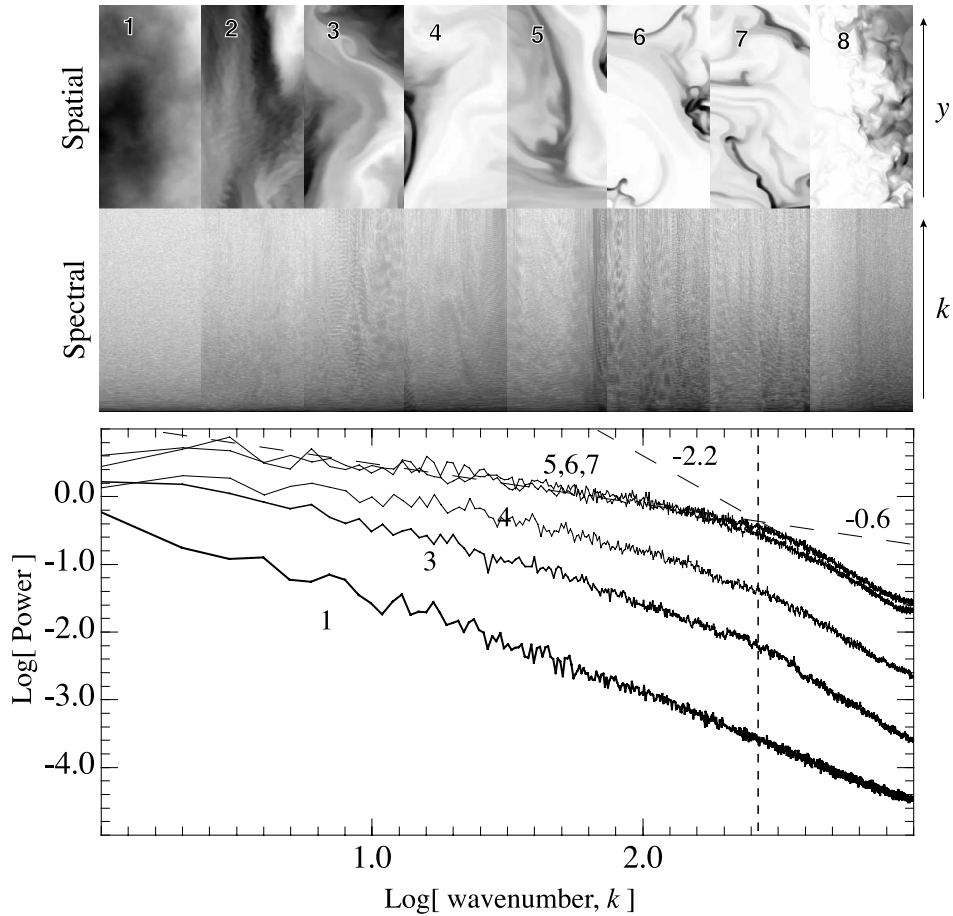


FIG. 10.—Power spectrum of the density fluctuations in the postshock gas of the power-law fluctuation model. The same power-law fits from the Gaussian models are overlaid here without refitting.

1. The more complex (possibly fractal) postshock region can produce a larger fraction of gas in the intermediate ($\sim 100,000$ K) rapidly cooling temperature range by virtue of the increased cold-hot gas boundary. More efficient cooling can reduce the overall cooling timescales and hence the amplitudes of the bounces. If this process dominates, then quantitative estimates of the increased cooling must take the resolution and numerical cooling errors into consideration to separate physical effects from numerical ones. This purely geometrical boundary effect may be less efficient in a real physical system than in the finite resolution numerical ones here.

2. The conversion of kinetic energy into thermal pressure may be less efficient, as for example where an oblique shock thermalizes only a fraction of the flow kinetic energy, resulting in redirection of kinetic energy into the y -dimension. Reduced thermalization may also be reflected in the reduced amplitude of the bounces.

3. Changes in the structure, or geometry, of the dense adiabatic base layer may also increase the ability of this layer to absorb energy, along the lines of item 2 above. As the flow deflects off a prominence of the dense layer, an increasing fraction of energy is transferred to the dense layer. In essence it can become an absorbing “sponge,” reducing the amount of thermalized energy available to drive the relaunch of the main shock, resulting in a lower shock velocity, reduced postshock temperatures, decreased

cooling timescales, and subsequently reduced bounces amplitudes.

Essentially, in the two-dimensional case more dynamical modes are available than in a one-dimensional model. These modes can take the form of motions in filaments, voids, and oblique shocks. All of these possibilities are investigated in the following sections to determine the degree to which they occur, if at all, and whether any are indeed the cause of this dramatic change in behavior in the two-dimensional models.

After a total of approximately 50,000 hydrodynamical cycles in the simulations for both the power-law and the Gaussian fluctuation models, the cycle of shock collapse and relaunching was still continuing. The total length of this simulation covered only three bounces, and the long-term behavior may become less coherent. Nevertheless, for the present we restrict ourselves to analyzing the medium-term properties of the shocks. In terms of their global structure, there appears to be very little difference in the amplitudes or frequencies of the bounces between the two different density fluctuation models. This indicates that the results are not strongly dependent on the choice of the initial fluctuations.

5.2. Analysis of the Shock Structure

Taking the shock structure at a point just past the peak of the first bounce (6.6×10^{10} s, near panel 3 in Fig. 12), before the dense layer becomes too disordered, but after filaments




FIG. 11.—Grid setup: the grid size is 2050×2048 cells, 2.0×10^{17} cm on a side. There is a hypersonic inflow on the left at 150 km s^{-1} and subsonic outflow on the right of about 0.3 km s^{-1} , tuned to keep the thickness of the dense layer approximately constant. The y -axis has periodic boundary conditions, approximating an infinite wall shock.

and knots have formed by the primary thermal instability (i.e., Fig. 14), we analyze the spatial structure of the simulations using Fourier transforms. Dividing up the simulation into 64 cell wide slices, parallel to the shock face, we construct the Fourier transform and thence determine the power spectrum of the density along each y -column in the slice. We then average the 64 power spectra to get a mean power spectrum for the slice. Moving from left to right, slices analyze first preshock gas (testing the quality of the advection of the fluctuations onto the shock), then the hot postshock region, and finally through the cooling layers to the dense cooled base of the shock. In this region we encounter material with the largest postshock time, where any instability, thermal or otherwise, will be most developed.

The results are shown in the multipanel Figures 9 and 10 for the Gaussian and power-law models, respectively. Interestingly, the most evolved slices (panels 5–7 in each case) show that the structures both tend toward a broken power law, for both initial spectra. The low-wavenumber end of the spectrum is driven by the development of postshock thermal instabilities. This leads to dense filaments and low-density voids; the power spectrum is quite flat, with an index $\alpha \approx -0.6$. Ultimately, this region of the spectrum is terminated by the formation of high-wavenumber structures ~ 10 pixels in extent, that are so dense that they have cooled below the 10,000 K limiting temperature of the cooling




FIG. 12.—Two-dimensional shock time-evolution snapshots. Panels 1–12 show the density variable in the power-law density spectrum shock model at evenly spaced time intervals (2×10^{10} s) throughout the simulation. The postshock gas is at times relatively smooth (1, 5, 8, 10), during the approximately adiabatic buildup of the shock before cooling initiated collapse occurs. The visible fluctuations result from shock compressed initial fluctuations. Subsequently, (2, 6, 9), the fluctuation contrast increases and dense filaments form (3, 7) along with low-density voids. The shock then collapses with the loss of internal pressure support (4).




FIG. 13.—Evolution of the averaged shock-front location. The two-dimensional simulations were summed along the y -axis, and the shock location was fitted. (a) The light solid curve shows the power-law model, (b) the heavy solid curve shows the Gaussian model, and (c) the dashed curve shows the 2048 resolution comparison one-dimensional model.

function. More or less conventional two-dimensional turbulence then takes over, with the power law steepening to an index ≈ -2.2 at $k \approx 260$. There appears to be no universal spectral index for two-dimensional turbulence. Nevertheless, this index is in the range of indices derived under a number of physically different conditions (i.e., Nazarenko & Laval 2000).

5.3. Cooling in the Two-dimensional Models

Figure 3 shows a close up of a region of the power-law model shown in Figure 14 at a location and time when a dense filament has just formed. The figure shows the cooling rate: black represents the strongest cooling, white the least. The scale shows that the cooling is indeed strong at the 30 or so pixel scale, and by the time the filament itself has cooled and formed, the cooling has diminished completely in the core of the filament. This gives the filament an edge-brightened appearance. The filament is typically 10 pixels across, and the spectral analysis shows that at this scale, and on down to the grid resolution, the density spectrum is a steeper power law.

The bulk of the cooling in the complex postshock structure emanates from a volume of gas below 100,000 K. In a steady or one-dimensional dynamical model this would correspond to a smooth column in the two-dimensional simulation. Because of the convoluted structure of this layer in these two-dimensional simulations, this zone may have a much larger effective area than for the corresponding one-dimensional or steady-flow model. This zone appears to take on a fractal structure, and to measure the fractal index we use a *coastline measuring method*, by laying down boxes of increasing resolution and determining the number of boxes that contain cells within the temperature range of interest. At the resolution of the grid this gives the effective area of the cooling zone.

The rate at which the number of boxes containing cooling varies with the box size defines a fractal dimension of the structure. If the structure displays self-similarity over the range of resolutions measured then the slope of $\log[N(r)]$ versus $\log(1/r)$ will be a straight line with a slope corresponding to the capacity dimension D_0 , which is defined in terms of numbers of squares $N(r)$ of side r that “cover” the fractal (see Peitgen, Jürgens, & Saupe 1992 and references therein),

$$D_0 = \lim_{r \rightarrow 0} \frac{\log[N(r)]}{\log(1/r)}. \quad (20)$$

For a smooth curve, this D_0 is equal to 1 [$N(r) \propto r^{-1}$], and for a smooth two-dimensional shape it is equal to 2 [$N(r) \propto r^{-2}$]. In each case measured here, the $\log[N(r)]$ versus $\log(1/r)$ curves were indeed straight lines with a slope tending toward 1.3–1.35 on average, over 3 decades of resolution, indicating the presence of true self-similar fractals.

Figure 15 shows the fractal dimension determined for the cooling zone with temperatures between 100,000 and 20,000 K as a function of time. It typically lies between 1.0 and 2.0, between a one-dimensional line and a two-dimensional area, and each measurement was consistent with a power-law structure with self-similarity. There was no significant difference in the fractal structures measured for the Gaussian and the power-law fluctuation models, and D_0 would seem to be independent of the initial fluctuation spectrum. It is important to note that this does not allow us to infer what fractal dimension a three-dimensional model would have; this is still a purely two-dimensional result.

5.4. The Global Energy Budget Analysis

Let us now consider the dynamical energy budget in the two-dimensional region. This is informative since it tells us where the incoming energy flux is converted to radiation and other forms of energy. We can also verify that energy is properly conserved by code.

We consider two control surfaces bounding the computational domain, \mathcal{D} (see Fig. 11): L_1 along the inflow boundary to the left and L_2 on the right, at the subsonic outflow boundary. We take the usual symbols ρ , v_x , v , P , $U = P/(\gamma - 1)$, and $h = \gamma P/\rho(\gamma - 1)$ to represent the density, x component of velocity, speed, pressure, internal energy, and specific enthalpy of the inflow. We use a zero subscript to represent the inflow and a 1 subscript to represent the outflow. The total energy on the grid at any time is given by

$$\begin{aligned} \int_{L_1} \rho_0 \left(\frac{1}{2} v_0^2 + h_0 \right) v_{x0} > dy - \int_{L_2} \rho_1 \left(\frac{1}{2} v_1^2 + h_1 \right) v_{x1} dy \\ + \frac{d}{dt} \int_{\mathcal{D}} \left(\frac{1}{2} \rho v^2 + U \right) dx dy \\ - \int_{\mathcal{D}} \rho^2 \Lambda(T) dx dy = 0. \quad (21) \end{aligned}$$

The first two terms represent the energy flux per unit length (normal to \mathcal{D}) into and out of the grid area, through the two control surfaces at the left- and right-hand edges of the grid, respectively. The third term represents the rate of change of the total energy on the grid (kinetic plus thermal), and the fourth term represents the radiative losses. The periodic boundary condition in y ensures that there is no net flow of energy through the upper and lower surfaces. The area




FIG. 14.—Snapshot of the cooling rate, $\log(\rho^2\Lambda)$, in the postshock gas near the point of collapse. Black represents strong cooling, white represents adiabatic gas. Successive points through the shock region are indicated. (a) The preshock region is adiabatic and hence white. (b) Small fluctuations grow in the moderately cooling (gray) immediate postshock gas. (c) Stronger cooling regions begin to separate. (d) Some compact filaments have formed and stop cooling in their centers as that gas falls below 10^4 K. (e) Many filaments have formed and the gas is turbulent as a result of hydrodynamical instabilities, just before impacting with the cool adiabatic dense layer.

integrals are easily evaluated by summing the zone-averaged values over all zones in \mathcal{D} , consistent with the Godunov formalism being employed.

If we represent the first two terms (giving the net energy flux into \mathcal{D}) by F , the total energy in \mathcal{D} by E_{tot} and the total cooling by C , then Figure 16 shows that when the left- and right-hand sides of $F + dE_{\text{tot}}/dt = C$ are evaluated independently by summing over the entire grid at each time step, they are indeed equal. We are confident, therefore, that the code is behaving consistently and conserving energy correctly.

When the two-dimensional results are scaled to unit width and compared to a one-dimensional model, with the same initial conditions and resolution, it is obvious that the two-dimensional shock “bounces” are smaller (see also Fig. 13 above). Figure 16 shows the total energy on the grid, E_{tot} , the net flux of energy onto the grid, F , and the cooling rate C to facilitate comparison of the one- and two-dimensional models. While the first “bounce” is similar in both cases, and the two-dimensional model initially behaves in an essentially one-dimensional manner, the subsequent evolution is different. The amplitude of the two-

dimensional bounces is much less, and the cooling losses, while maintaining a similar average (as expected), show smaller fluctuations. We attribute this to effects associated with the higher dimensionality.

Figures 17 and 18 show the breakdown of the total energy into its kinetic and internal components for the one-dimensional and two-dimensional models, respectively. The individual panels refer to the entire grid (*top*), the hot ($>20,000$ K) shocked region (*middle*), and the dense cool layer (*bottom*). It is apparent that in the one-dimensional model, the kinetic energy of the inflow is first converted to thermal pressure in the shock. As the shock cools, the kinetic energy of the collapsing gas increases until the bounce occurs and the stored kinetic energy is reconverted into thermal energy. The cycle then repeats, and the distribution of energy between the thermal and kinetic components is similar at the corresponding phases of each cycle.

The two-dimensional model displays important differences. First, the radiative losses never return to the sharp initial minimum. This is a geometric effect, because at each cycle, the newly forming shock is launched at different locations and times along the dense layer, as a result of the




FIG. 15.—Fractal covering dimension of the postshock gas between 30,000 and 100,000 K as a function of time. The lighter curve is the model with power-law input fluctuations, and the heavier curve is the Gaussian fluctuation model. The dimension asymptotes toward a value between 1.3 and 1.35, irrespective of the initial fluctuation spectrum.

uneven distribution of the dense layer. There is no longer an instant in time at which there is no shock, in contrast to the one-dimensional model in which there is an instant of complete collapse.

More importantly, on subsequent bounces, some of the inflowing energy remains as kinetic energy, either in the postshock gas or in the turbulent motions of the dense layer itself, and this accounts for approximately half of the kinetic energy on the grid. The internal energy in postshock gas is lower, resulting in both a lower postshock temperature and in an increased cooling efficiency (C/E_P). This ensures a much shorter cooling timescale for the shocked gas—by a factor of 4 or more compared to the one-dimensional or the initial two-dimensional bounce (see Fig. 19).

With a shorter cooling timescale the shock cannot progress as far against the incoming flow before it recollapses, producing a smaller bounce amplitude and higher bounce frequency. The total cooling still equals to the total net energy flux and the changes in the internal energy between the control surfaces. Moreover, the mean luminosity is similar to or a little less than the average net energy flux since energy is stored in the kinetic energy reservoir of the dense layer. The rate of change of the kinetic energy in the dense layer $d/dt E_{K,\text{dens}}$ and the increase in energy of the gas flowing off the grid is sufficient to appreciably reduce the both the maximum and the mean luminosities of the shock.

This transfer of energy to the dense layer in two dimensions takes place as the fast shock launches from, and dense filaments collide with, the “spongy” dense layer. Transverse motions (i.e., in the y -direction) and the accelerations




FIG. 16.—Time evolution of the energy budgets for a one-dimensional shock (*upper panel*) and two-dimensional shock (*lower panel*). Each panel displays the net energy flux, F , onto the grid and the rate of change of the total energy on the grid, dE/dt . The comparison between the summed total cooling C and the computed difference between the integrated fluxes and rates of change on the grid, $F - dE/dt$, is very close.

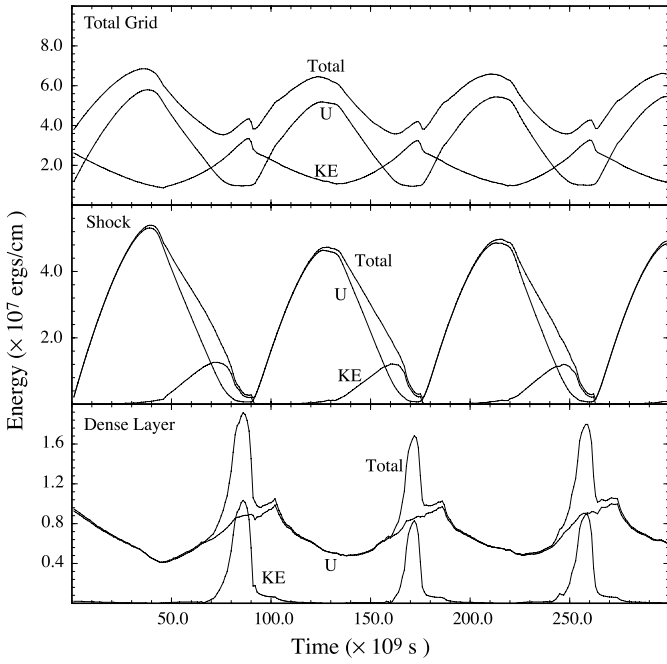


FIG. 17.—Time evolution of the energy components in a one-dimensional shock. *Upper panel:* The total kinetic and internal energy on the entire grid. *Middle panel:* The same for the postshock gas at $T > 20,000$ K. *Lower panel:* The energies contained in the dense cool layer. See text for details.

during the bounce result in the fractal and turbulent dense region seen, for example, on the right in Figure 14. This layer represents the photoionized “tail” of the radiative shock and is, in effect, the “piston head” driving the shock, which is on average stationary in the frame of the grid.

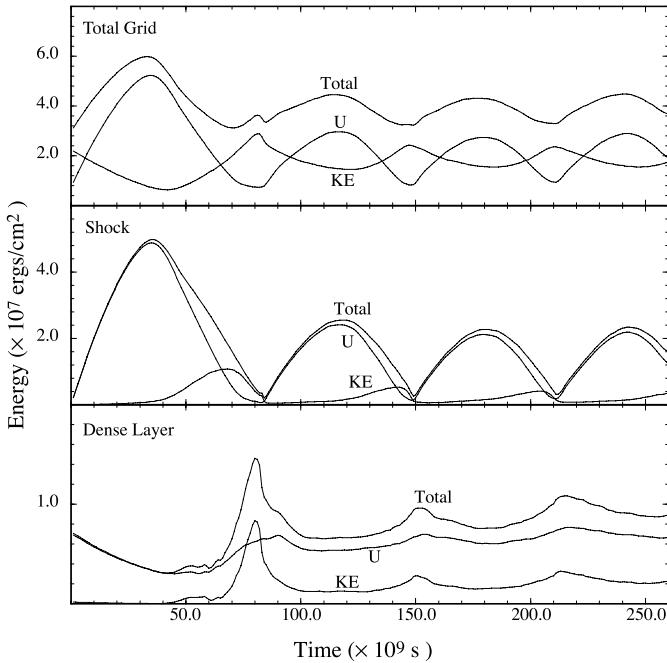


FIG. 18.—Time evolution of the energy components of a two-dimensional shock. *Upper panel:* The total kinetic and internal energy on the entire grid. *Middle panel:* Same, but for just the postshock gas at $T > 20,000$ K. *Lower panel:* The energies contained in the dense cool layer. Compared with Fig. 17, the dense layer and to a lesser extent the postshock gas contain more kinetic and less peak thermal energy. See text for details.




FIG. 19.—Time evolution of the cooling timescale ($U/\rho^2\Lambda$) averaged over the grid, compared to the crossing timescale (grid length divided by the inflow velocity $1/v \sim 1.5 \times 10^{10}$ s) for one-dimensional and two-dimensional dynamical models. In the two-dimensional cases, for the second and subsequent bounces, the cooling timescale is at most twice the crossing time; in the one-dimensional case and the initial two-dimensional bounce it is 4.5–5 times longer. When the cooling time is less than the crossing time of the flow, the shock lose pressure support and collapses. When the cooling time is long compared to the crossing time for an extended period, larger and slower “bounces” result.

When this dense layer remains at about the same thickness as the tail of a real shock, the amount of kinetic energy it contains in the simulation is realistic. However, if allowed to become too great in volume compared to the shock, the simulation would overestimate this energy reservoir.

Steady, one-dimensional MAPPINGS III models suggest that the thickness of the photoionized tail ($10,000 > T > 5000$ K) of a $150\text{--}200 \text{ km s}^{-1}$ shock is of order $(1\text{--}3) \times 10^{17} \text{ cm}$ for the densities in our simulations. Thus, the dense layer in these two-dimensional simulations, which grows to approximately $5 \times 10^{16} \text{ cm}$, has not yet become too thick to be unphysical. The radiation from the hot shocked regions should be sufficient to keep the dense layer from cooling throughout the simulations here.

The dense, cool gas carries away an important fraction of the shock energy in the form of small-scale turbulence. This increased transfer of kinetic energy is a result of the increased degrees of freedom for the two-dimensional flow compared to the steady and one-dimensional models. The difference between two-dimensional and future three-dimensional models may not be so dramatic, as they both have the freedom to transfer energy to turbulent kinetic energy in the postshock and dense layers. Nevertheless, the quantitative measure of the efficiency of this transfer requires high-resolution three-dimensional modeling to correctly estimate the coupling between the energy in the shock across the boundary into the dense gas. Nonetheless, this demonstration that shocks can directly inject small-scale turbulence into the interstellar medium (ISM) is important because it is now generally believed (see Elmegreen 2002; Elmegreen, Kim, & Staveley-Smith 2001) that star

formation is moderated by the turbulent cascade and dissipation on small scales. The turbulence at large scales is believed to come from violent shocks such as supernova remnants and stellar winds. However, the effect of the direct input of small-scale turbulence, which amongst other things, will help the mixing of their chemical products in the ISM, has not hitherto been considered.

6. COMPARISON WITH STEADY-FLOW MODELS

In most astrophysical situations, it is rare to be able to spatially resolve the internal structure in individual shocks, the exception being in local SNRs (i.e., Patnaude et al. 2002) and some young stellar object (YSO) jets (Bally et al. 2002; Hartigan et al. 2000). In most cases, we observe an unresolved ensemble of shocks whose emergent spectrum is an emission-weighted average over the ensemble. Even if they are characterized by an average velocity, then, if they are pulsating like the models here, they will be observed at many different phases in such an average. In the past, steady-flow models, such as those produced by the one-dimensional MAPPINGS III code have been applied to situations such as active galaxy narrow-line and extended emission-line regions both in the nearby and distant universe (Dopita et al. 1997; Allen et al. 1999; Bicknell, Dopita, & O'Dea 1997; Bicknell et al. 2000; Best, Röttgering, & Longair

2000a, 2000b). Thus, in order to compare the dynamical models with the steady models, we construct time-averaged histograms of observable quantities to compare with the steady histograms of the same quantities in a MAPPINGS III model derived for a shock of the same mean velocity.

In these models, the inflow of 150 km s^{-1} results in a shock with a mean velocity of 150 km s^{-1} but with an instantaneous velocity between zero and 200 km s^{-1} . We now compare the time averages of these models with a steady 150 km s^{-1} MAPPINGS III model.

6.1. Time-averaged Distributions

6.1.1. Density

Figure 20 (*upper left-hand panel*) shows the histogram obtained by analyzing ~ 250 two-dimensional data frames for the density variable, averaged over three complete expansion and collapse cycles of the wall shock, for both the power-law and Gaussian fluctuation models.

Histograms of the gas density were constructed for each frame of the simulation. By counting only the regions where the temperature exceeds $10,000 \text{ K}$ and the gas is radiating, we are able to analyze just the postshock region, through the thermally unstable regime down to the point where photoionization is expected to halt cooling. This is the domain where the present simulation are most physically consistent.

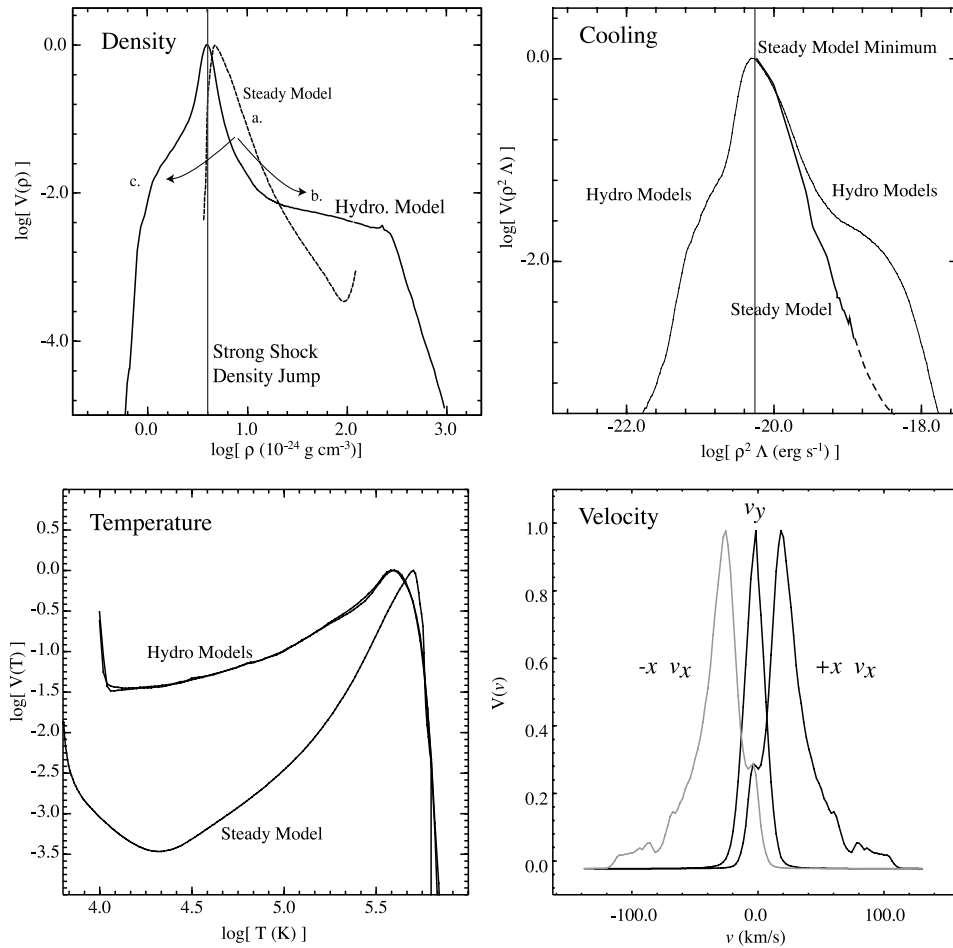


FIG. 20.—Normalized volume distribution of important flow variables derived from a steady-flow MAPPINGS III model and the time-averaged distribution for two-dimensional dynamical models. *Upper left*: density; *upper right*: cooling; *lower left*: temperature; *lower right*: velocity distributions along the x -axis and along the y -axis. See text for details.

The density intervals were chosen to be 0.02 dex logarithmic intervals. Density units are multiples of the preshock density, and the histogram is normalized to a maximum of 1.0 (0.0 in the log). The value of the histogram at any density is then the fraction of the shocked cooling gas by volume present in the time-averaged, or ensemble, sense. Both the Gaussian and power-law models produce essentially identical results.

When these histograms are compared to the steady-flow MAPPINGS III shock model, differences are immediately apparent. In the steady model, the minimum density is determined by the immediate postshock gas; as the temperature decreases monotonically down to 10,000 K, the density can only increase in the shock flow region. In the dynamical models, the existence of the thermal instability allows the increasing separation of both low- and high-density regions in the postshock flow. Not only do dense filaments form, but low-density voids also appear as a result of the physical redistribution of gas as it flows into dense condensations. The voids cool slowly because of the lower the density, but they all collapse as the shock retreats into the dense layer at the end of a cycle.

The two arrows in Figure 20 (*upper left-hand panel*, [a] to [b] and [a] to [c]) show notionally how the fraction of material at intermediate temperatures is redistributed to increase the fraction of high-density gas on the right of the figure and to produce low-density gas (by volume) on the left. There is even an appreciable fraction of postshock gas with a density below that of the preshock gas. Estimates of postshock densities using line diagnostics from steady models (such as [Ar IV] lines) may therefore be misleading compared to the dynamical situation in the hydrodynamical simulations and perhaps in real ISM shocks.

6.1.2. Cooling

The average distribution of cooling in the dynamical shock models is shown in Figure 20 (*upper right-hand panel*). This is constructed in the same way as the density distribution above, using the cooling rate from the same regions. The time-averaged cooling distribution by volume is approximately symmetrically distributed about the mean cooling rate. This is in contrast to the steady model, which shows a one-sided distribution with decreasing fractions of ever increasing cooling rates (*dashed line*). This is partly a consequence of the fact that in the steady model the density distribution down to 10,000 K is a monotonically increasing function in the steady model, and the weakest cooling is at the beginning of the shock and the cooling rate never falls below this rate. Low-density voids in the hydrodynamical models provide opportunities for volumes of gas to cool slowly, through the density-squared dependence of the cooling rate. The dynamical models also include substantial volumes of gas with stronger cooling, above $10^{-19.5}$ ergs cm⁻³ s⁻¹, compared to the steady model. The complex dense filaments, formed via the thermal instability, constitute the more efficiently cooling gas. Hence, the long-term time-averaged luminosity is the same as for a steady-flow model, but the emission originating from individual temperature and density zones can be quite different.

6.1.3. Temperature

The averaged thermal distribution in the two-dimensional dynamical shock is shown in Figure 20 (*lower left-hand panel*).

Compared to the steady flow model, the dynamical shock demonstrates an increased volume of intermediate temperature gas, below 300,000 K, consistent with this region having a complex spatial structure. There is a deficit of very hot gas close to the postshock temperature at (a), partly because the launching phase of the bounces occupies a relatively small fraction of the time. At most other times the shock is slower and the temperatures are consequently, on average, lower. This duty-cycle effect also contributes to the higher average fraction of cooler gas. On average there can be as much as 10–100 times more gas at the critical emitting temperatures of 30,000–100,000 K than a simple steady model would suggest. This has implications for shock diagnostics that derive from species that emit in this region. Dynamical models will be relatively bright in their optical and UV lines, and steady-flow models would overestimate the density or shock area required to achieve the same luminosity for a given shock velocity. Shock velocities would tend to be underestimated and preshock densities would be overestimated in order to achieve the correct luminosity. The systematic differences between a grid of steady and unsteady dynamical shock velocities is important to determine. Since the structure of the critical region is subject to turbulence, quantitative spectra will have to wait for three-dimensional models.

6.1.4. Velocity

Since these two-dimensional simulation fail to capture the complexity that would be expected in a three-dimensional model, it is probably inappropriate to attempt detailed spectral line synthesis from them. Nevertheless, is it useful to look at the distribution of velocities along two lines of sight, along the x -axis, parallel to the flow, and along the y -axis, perpendicular to it, since this demonstrates the importance of the generation of turbulence in the cooling layer. Figure 20 (*lower right-hand panel*) illustrates these projections. A nearly symmetrical distribution is obtained from the y -axis projection wherein the velocity distribution is determined by the internal motions of the turbulent gas. The x -axis velocity distribution is asymmetric as expected, but if an ensemble of shocks were viewed from a number of different angles, the observed profile would be a sum of the symmetrical y distribution and broad $+x$ (looking into the shock; *solid line*) and $-x$ (looking out of the shock; *gray line*) profiles. On convolution over all angles, we can expect a broad, roughly symmetrical line profile from $\pm v_{\max} \sim v_{\text{shock}}$ that falls between the extremes of the $+x$ and $-x$ profiles. This figure suggests that the shock velocity is better estimated from the width of the base of an emission line, such as the full width at 5% or 10% of a line peak, rather than half maximum. A more detailed treatment of emission-line profiles will be performed when three-dimensional dynamical models of this resolution are carried out.

7. SUMMARY

Observations are now beginning to resolve interstellar radiative shocks prompting parallel theoretical developments that can be used to construct credible models of the observations. In this second paper of our series on the computation of radiative shocks we have utilized

one-dimensional tests in order to establish the robustness and reliability of our code for the simulation of multidimensional radiative shocks and for estimating the resolution necessary to accurately model the temporal behavior of such shocks. The results give us confidence that the goal of realistic multidimensional models can be achieved while at the same time cautioning that high resolution is required in order to resolve the high-density contrast that inevitably occur in such situations.

We have shown in § 3 that with adequate resolution the modified VH-1 code correctly reproduces one-dimensional analytical solutions with a power-law cooling function. This is an important result since the inclusion of cooling into a directionally split code is subject to some judgment and could conceivably introduce unexpected consequences. The comparison shown in Figure 1 shows excellent agreement between the velocity profiles of one-dimensional numerical and analytical models; similar agreement is obtained for the other dynamical variables. We have also demonstrated excellent agreement between the oscillation frequencies of numerical and analytical solutions in agreement with similar results obtained by SB95.

Real interstellar shocks do not have a power-law cooling function. Therefore it has been advisable to conduct similar tests for realistic plasma cooling functions. We have chosen an NEQ cooling function corresponding to plasma cooling from a high temperature. Using this cooling function, we have shown that in one-dimensional, the NEQ function promotes both stable and overstable cooling, at different shock velocities and in line with expectations based on the local slope of the cooling function.

Our resolution study of the NEQ shocks has produced some constraints that have important implications for multidimensional studies. The cooling regions of radiative shocks are to a large extent isobaric so that cooling results in large density changes. In one-dimensional shocks these occur at the tail of the shock leading to contact discontinuities in both density and temperature. The density² and temperature dependence of the cooling function implies that, unless these discontinuities are well resolved, the numerically calculated cooling can lead to serious errors. This is evident in the resolution study presented in § 3.3, where we have shown that 2000 grid points are required to resolve cooling in a Mach 15 interstellar shock. In one dimension adaptive mesh methods could be effective. However, in higher dimensions high-density features, such as cooling filaments as well as the tail of the shock, occur throughout the volume and it is problematical whether adaptive mesh refinement techniques would be useful. Another approach may be to incorporate more effective tracking of contact discontinuities. That may be an approach that could be effectively followed in future. Nevertheless, for the present

the most direct approach is to use approximately grid cells for the spatial region that we have used in these tests. Note that, in effect, we have derived this result only for a single velocity (150 km s^{-1}) and one Mach number (15). We expect similar results to hold for other shock velocities. Indeed, it would be straightforward to redo these tests for other velocities of interest.

Having laid the groundwork for computation of multidimensional shocks, we have described in §§ 4 and 5 our work on two-dimensional simulations. The high-resolution, two-dimensional simulations with 4×10^6 computational cells that we have described in this paper have demonstrated important qualitative and quantitative differences between dynamical and steady-flow models:

1. The thermal instability of the postshock gas can generate complex cooling structures, including dense filaments and low-density voids.
2. The cooling zones in the dynamical models are more efficient than one-dimensional or steady-flow models, as a result of the fractal structures that evolve.
3. The structures that do evolve are independent of the initial density fluctuation spectrum, with filaments and possibly turbulence forming rapidly within a single cooling timescale. The resulting fluctuation spectra are not simple power laws but rather broken power laws. The break in the density power laws occurs at an intermediate scale, consistent with the radiative dissipation from the largest initial structures dominating over the numerical dissipation at the smallest grid scale.
4. Multidimensional dynamically unstable radiative shocks will, in general, be characterized by more efficient cooling, both because of a more complex and larger volume of strongly cooling gas and because of the reduced thermalization of kinetic energy into thermal pressure with a resultant lower postshock temperature and higher cooling rate.
5. The long-term average luminosity of multidimensional is approximately the same as that derived from one-dimensional, steady-flow models. However, the formation of a fractal distribution of filaments and voids via thermal instability implies that shock diagnostics from species from different temperature zones are likely to give inconsistent velocity fits when compared to steady models. For example, this may be a source of ambiguity in Danforth et al. (2001).

This research was supported by an ARC Large Grant A69905341 and ARC Discovery project DP0208445. M. D. further acknowledges the support of the ANU and the Australian Research Council through his ARC Australian Federation Fellowship. The authors would like to thank the anonymous referee for the constructive comments which benefited the final form of this paper.

REFERENCES

- Allen, M. G., Dopita, M. A., Tsvetanov, Z. I., & Sutherland, R. S. 1999, *ApJ*, 511, 686
 Anders, E., & Grevesse, N. 1989, *Geochim. Cosmochim. Acta*, 53, 197
 Bally, J., Heathcote, S., Reipurth, B., Morse, J., Hartigan, P., & Schwartz, R. 2002, *AJ*, 123, 2627
 Best, P. N., Röttgering, H. J. A., & Longair, M. S. 2000a, *MNRAS*, 311, 1
 ———. 2000b, *MNRAS*, 311, 23
 Bicknell, G. V., Dopita, M. A., & O'Dea, C. P. O. 1997, *ApJ*, 485, 112
 Bicknell, G. V., Sutherland, R. S., van Breugel, W. J. M., Dopita, M. A., Dey, A., & Miley, G. K. 2000, *ApJ*, 540, 678
 Binette, L., Dopita, M. A., & Tuohy, I. R. 1985, *ApJ*, 297, 476
 Blondin, J. M., Konigl, A., & Fryxell, B. A. 1989, *ApJ*, 337, L37
 Blondin, J. M., & Lukin, E. A. 1993, *ApJS*, 88, 589
 Chevalier, R. A., & Imamura, J. N. 1982, *ApJ*, 261, 543
 Colella, P., & Woodward, P. R. 1984, *J. Comput. Phys.*, 54, 174
 Cox, D. P. 1972, *ApJ*, 178, 143
 Daltabuit, E., & Cox, D. 1972, *ApJ*, 173, L13
 Danforth, C. W., Blair, W., & Raymond, J. 2001, *AJ*, 122, 938
 Dopita, M. A. 1977, *ApJS*, 33, 437
 Dopita, M. A., Binette, L., & Tuohy, I. R. 1984, *ApJ*, 282, 142

- Dopita, M. A., Koratkar, A. P., Allen, M. G., Tsvetanov, Z. I., Ford, H. C., Bicknell, G. V., & Sutherland, R. S. 1997, *ApJ*, 490, 202
- Dopita, M. A., & Sutherland, R. S. 1995, *ApJ*, 455, 468
- . 1996, *ApJS*, 102, 161
- . 2002, *Astrophysics of the Diffuse Universe* (Berlin: Springer)
- Elmegreen, B. G. 2002, *ApJ*, 564, 773
- Elmegreen, B. G., Kim, S., & Staveley-Smith, L. 2001, *ApJ*, 548, 749
- Ferland, G., et al. 1994, in *STScI Symp. Ser. 8, The Analysis of Emission Lines*, ed. R. E. Williams, M. Livio, D. E. Osterbrock, & M. J. Seaton (Baltimore: STScI), 83
- Field, G. B. 1965, *ApJ*, 142, 531
- Ghavamian, P., Raymond, J., Hartigan, P., & Blair, W. 2000, *ApJ*, 535, 266
- Hartigan, P., Morse, J., Palunas, P., Bally, J., & Devine, D. 2000, *AJ*, 119, 1872
- Innes, D. E. 1992, *A&A*, 256, 660
- Innes, D. E., Giddings, J. R., & Falle, S. A. E. G. 1987a, *MNRAS*, 226, 67
- . 1987b, *MNRAS*, 227, 1021
- Nazarenko, S., & Laval, J.-P. 2000, *J. Fluid Mech.*, 408, 301
- Parker, E. N. 1953, *ApJ*, 117, 431
- Patnaude, D. J., Fesen, R. A., Raymond, J. C., Levenson, N. A., Graham, J. R., & Wallace, D. J. 2002, *AJ*, 124, 2118
- Peitgen, H.-O., Jürgens, H., & Saupe, D. 1992, *Chaos and Fractals* (New York: Springer)
- Raymond, J. C. 1979, *ApJS*, 39, 1
- Raymond, J. C., & Smith, B. W. 1977, *ApJS*, 35, 419
- Stone, J. M., & Norman, M. L. 1993, *ApJ*, 413, 198
- Strickland, R., & Blondin, J. M. 1995, *ApJ*, 449, 727
- Sutherland, R. S. 1993, Ph.D. thesis, Australian Natl. Univ.
- Sutherland, R. S., Bisset, D. K., & Bicknell, G. V. 2003, *ApJS*, 147, 187
- Sutherland, R. S., & Dopita, M. A. 1993, *ApJS*, 88, 253