The September 1, 2001, version of the 1-D PPM Gas Dynamics Interface GAS1D

A Guide for Users

Paul R. Woodward, LCSE, University of Minnesota, 9/9/01

Introduction

The LCSE 1-D Eulerian Gas Dynamics Experiment, GAS1D, is a Microsoft-Windows-based utility program that allows simple experiments in gas dynamics to be performed with immediate visual feedback that is both qualitative and fully quantitative. This program was originally designed for a course, "Methods in Computational Physics," at the University of Minnesota. It was then adapted for use in modifying and testing new numerical methods for gas dynamical simulations, and finally, it was redesigned for non-technical student use. The success of this last redesign is yet to be evaluated. (Send your suggestions to paul@lcse.umn.edu.) The numerical method, PPM (the Piecewise-Parabolic Method), imbedded in this program gives excellent fidelity to the gas dynamics simulations, so that the non-technical user can fairly safely suppose the results of the simulations that are presented to be an accurate representation of what a real gas, in this case air, would actually do.

Visualizing 1-D fluid flow for liquids or gases:

Since we live in 3 dimensions, it is perhaps a bit difficult to visualize a one-dimensional fluid flow. But we all have experience of rivers. Even though rivers are three dimensional, it is fairly clear that the variations in flow speed, water temperature, and, say, silt content, measured along the length of a river like the Mississippi, the Colorado, or the Nile are more important

(usually) than the variations we might find if we were to take detailed measurements across the river's width at any point. In this sense, we could make a one-dimensional model of a river, and for many purposes it could give us a completely adequate representation. Rivers are good examples of approximately 1-D fluid systems, but it is harder to find examples from everyday experience of 1-D gas dynamical systems. Nevertheless, they are not entirely unfamiliar. A flute or an organ pipe is perhaps the most common example. These instruments involve air confined to move within a tube with hard walls. The air motion is therefore channeled to take place primarily in the single direction along the length of the tube. Just as in the case of the river, and actually even more appropriately, we can characterize the motions of the air in the tube by measuring the variations in the air velocity, pressure, and density along the length of the tube. We all know that signals propagate along these tubes, and those signals are sound waves, which make musical notes we can hear at a distance.

Shock tubes to produce 1-D gas flows:

Laboratory experimenters, in an effort to explore gas dynamical phenomena, have developed techniques to construct 1-D gas flows. They have done this because these 1-D gas dynamical systems are the simplest to describe theoretically. A typical laboratory setup, which is actually simulated in the LCSE gas dynamics experiment program, is a shock tube. This is a long thin tube that contains gas inside. The portion of the tube that is observed is transparent, so that the behavior of the gas inside can be photographed. Because the tube is much longer than its diameter, and because the viscosity of the gas inside is small enough that the disturbances set up by friction along the wall of the tube are not important compared to the disturbances that propagate along the length of the tube, we can consider this shock tube as a good approximation of a 1-D system. This 1-D character of the gas flow inside the tube means that, essentially, the state of the gas varies only along the length of the tube, and not across its diameter.

The shock tube is thus like a big flute or organ pipe made of transparent material, like plexiglass. This pipe is made very strong, because, unlike with the musical instruments, we will

not blow softly into this pipe near one end. Instead, we may choose to blow strongly enough into this tube, one way or another, so that the pressure pulse that travels along the tube is as strong as the shock from an explosion. We might also blow strongly into a shock tube so that the velocity of the gas inside exceeds the speed of sound. This might allow us to investigate the behavior of objects in supersonic flight. What objects fly this fast? Some aircraft, space ships, and, of course, speeding bullets. But these uses of shock tubes exploit the ability of the tube to create a uniform gas flow in one direction, but they then break the 1-D nature of that flow by inserting an obstacle into the flow. The ultimate measurements in such applications of shock tubes are 2-D or 3-D in character, so we will not discuss them here.

The "state" of a gas in a 1-D flow:

What is the "state" of the gas in our shock tube? This is determined fundamentally by 3 variables: the density, denoted by the Greek symbol rho (ρ), the pressure, p, and the velocity, u, along the length of the tube. These are all familiar quantities from everyday life, although we are more accustomed to stating the temperature of the air around us than its pressure. Even though we know what the density of air is, conceptually, the nearly incompressible behavior of air in everyday experience does not present many occasions in which we care to know the precise density of the air around us. However, if we were to scale a high mountain or to pilot a jet aircraft, we would begin to care about the density of the air. In our shock tube, the density of the air (and it is indeed air that is inside our shock tube) will vary. These variations can even be large. The reason for this is that we are simulating a shock tube, a device in which large changes in air pressure can be supported, so that shock waves can propagate. Shock waves are sudden compressions of a gas that propagate like sound signals. When fireworks explode on the 4th of July or at the State Fair, the sudden explosions send out shock waves in all directions, which we hear as the characteristic bangs. These are sound pulses, rather than continuous sound vibrations, like the continuous notes sounded by a flute. Shocks involve much larger pressure changes than singing or talking, and they are also strong enough to compress the air so that its density changes

significantly. This is why the value of the density in our shock tube is reported by the LCSE gas dynamics experiment program, GAS1D, on its form and is displayed in one of its 4 plotting windows.

Specifying a 1-D gas dynamical experiment in GAS1D:

Within the context of GAS1D, a single gas dynamics experiment is specified by means of the state of the gas in the shock tube at the initial time (always defined as time 0) and by the behavior that is desired at the boundaries of the section of the tube that will be displayed by GAS1D. This tube section always begins at the left at location $x = 0$ and ends at the right at location $x = 1$. Thus the length of the displayed section of our shock tube is always unity. Don't worry that this shock tube length measures one unit but that unit (centimeters, inches, feet, meters, miles, light years) is not specified. It turns out that the behavior of the gas in our tube will be the same regardless of which units we use for its length, so long as some specific conditions apply. One such condition, relevant at the moment, is that the viscosity of the air in the tube should not affect its motion. This means that friction with the walls must be unimportant. Understandably, this requires that the diameter of the tube should be substantial (i.e. not 1 micron or something like that). This in turn means that the length of our tube should not be 1 inch, but 1 meter would be fine, and 1 mile, if we could build such a shock tube, would be dandy.

Special units of measurement used in GAS1D:

The GAS1D program performs its computations in a special internal set of units. Because of the nature of gas dynamical fluid flow, it is possible for us to rescale the results presented by GAS1D to arrive at valid predictions for real world experiments, should we desire to do so. However, for exploring 1-D gas dynamical phenomena with GAS1D, it is rarely necessary to perform such conversions of the internal units used by the program; we need only to understand what they mean. First, as we have already stated, the length of the section of the tube in which results are computed for us by the program is always 1. A natural unit of time is the

interval required for a sound signal to traverse our section of the shock tube. To determine this time interval, of course, we need to know the sound speed in the tube. Then the time interval is just the length of the tube, 1, divided by the average speed of sound in the tube. At present, GAS1D does not use the average speed of sound in the tube to compute this time interval. Instead, it uses the sound speed at the left-most end of the tube, which is a little easier to obtain.

What determines the sound speed? In your 1-D gas flow, GAS1D allows you to prescribe values for the gas density, pressure, and velocity at selected locations along the shock tube. The values you pick for the density and pressure will determine the sound speed, c, via the formula $c^2 = 1.4\, p/\rho$. You need not be concerned with this formula. You only need to know that if the ratio of your pressure and density values is near 1, then you will have sound speeds around 1. GAS1D does this arithmetic for you. Nevertheless, it is convenient to know that a pressure of 1 and a density of 1.4 prescribed at the left-hand end of your tube will result in a sound speed of 1 there, and hence a sound traversal time of 1. The reason that a sound traversal time of your shock tube is important is that this is the time that it will take for signals to propagate across the system, and therefore it tends to be the time it takes for something significant to happen in your experiment.

Specifying how long to run your experiment:

When you specify a length of time over which to observe your gas flow experiment, you express it in terms of the sound traversal time intervals just discussed. You enter that number in the text box to the left of the label "Time Units to Run" on the upper right portion of the GAS1D form. GAS1D computes how long this is and stops your experiment at the appropriate time. As your experiment runs, plots of the density, pressure, and velocity distributions will be displayed at intervals that you determine by entering the number of such plots you want in each sound crossing time interval in the text box to the left of the label "Plots / Time Unit" in the upper right portion of the GAS1D form. We will discuss this graphical output more in a moment. First, let's discuss how you set up the initial state of the gas in your shock tube.

The Operation of the Windows Graphical User Interface to the Program

You start up the GAS1D program by double clicking on its executable file, which is called GAS1D.exe and can be downloaded from the LCSE Web site by first going to the URL www.lcse.umn.edu/seminar, locating this file in the file list, right clicking on it, and selecting the option "save link as" from the pop-up menu that appears (at least with the Netscape Web browsers). Upon start-up, you will first be prompted for your name. If you simply press "Enter," then you will be assigned the default name: Alfred E. Newman. It is important, however, to supply your name at this prompt, since it will be used in generating a default title for your gas dynamics experiment, which will help to identify it if you choose to save your results (as has been done for the examples here) in a WORD document for printing or for sending via E-mail to your collaborator or to your instructor.

Once you supply your name, you are presented with a graphical user interface (or "GUI") which allows you to compose a particular gas dynamics experiment to perform. By default, you are presented with a standard test problem, all set and ready to go. This test problem is well-known in the small community of applied mathematicians who design numerical methods for gas dynamical simulations on computers. It is a good illustration of the built-in features of the GUI. The window which you see on your screen is duplicated on the next page. In the following text, we will explain the various subwindows, labels, and buttons that appear on this Windows form.

The plotting windows:

Most noticeable on the form are the 4 subwindows that are used for displaying the results of your gas dynamics experiment. Beneath each window is a label identifying the variable that will be plotted in the window. For example, the upper right window shows a plot of the pressure. To the left and right of the label "Pressure" are text boxes in which the minimum and maximum values of the pressure, respectively, are given that correspond to the bottom and top of the pressure scale displayed. To see the actual value of any pressure value along the horizontal extent of this

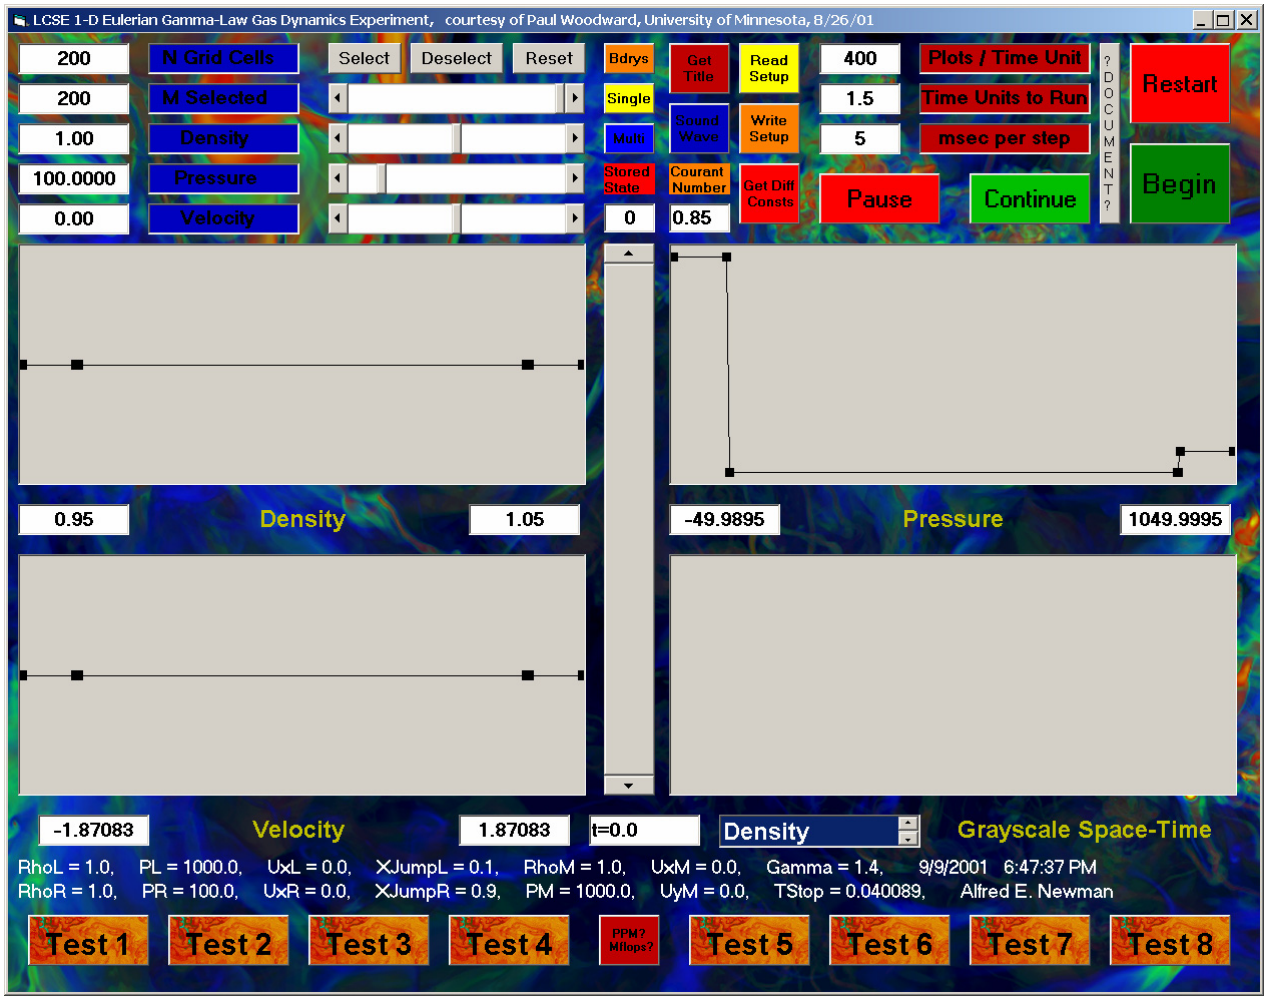

window, just place the cursor at that horizontal position within the window and click the left mouse button. The result of doing this is shown directly below the picture of the initial screen (on the following page). An unfilled square is drawn at the location on the pressure curve where you clicked, and at the upper left on the form the numerical values are given for the density, pressure, and velocity of the gas at that location. Thus you can see that the pressure value in the low-pressure area in the middle of this default problem setup is actually 0.01, a fact that would have been nearly impossible to determine from just the minimum and maximum pressures listed in the text boxes below the pressure plot without an accurate measuring device and some careful arithmetic.

Each of the 3 plots, labeled "Density," "Pressure," and "Velocity," on this form are intended to give you a quick and qualitative idea of the distributions of those variables with

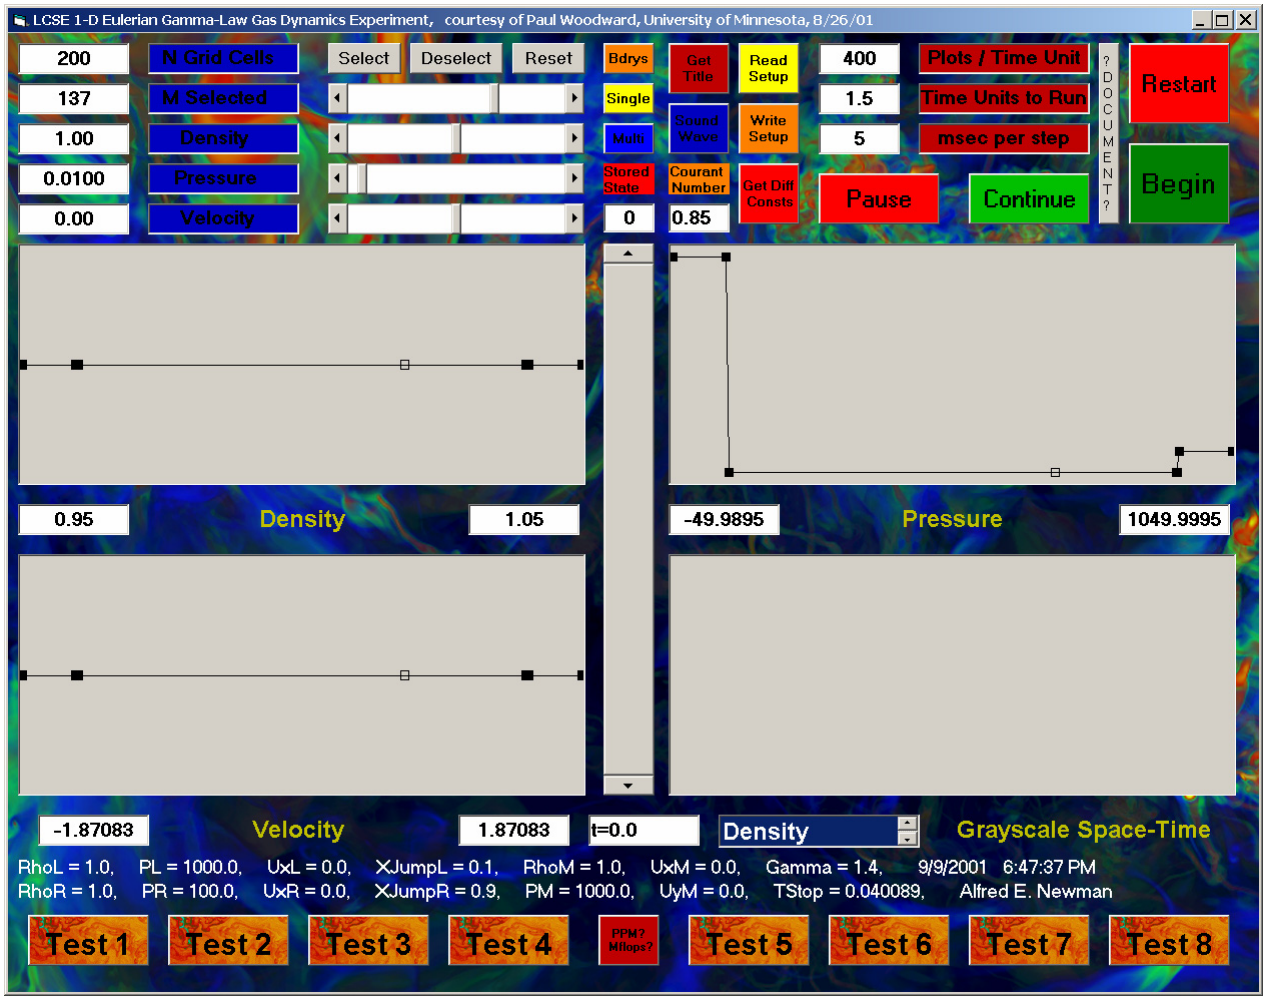

distance, the horizontal axis in each plot, in your gas dynamics experiment. These plots are not intended to be highly quantitative. However, the GAS1D GUI can present you with highly quantitative information easily upon demand. By clicking the left mouse button with the cursor positioned at any horizontal location within any of the 4 plotting windows on this form, you will get a precise read-out of the appropriate numerical values.

The "selected" grid cell:

Above the text boxes in the upper left portion of the form, where these precise values are given, is another text box which is labeled "M Selected." In this text box appears the grid cell number of your selected cell. In the example shown above, your selected cell is the one you chose by clicking in the plotting window. It is grid cell number 137. Its horizontal location along

your shock tube is not given, but you can figure it out from this grid cell number if you like. Here it is important to know a bit about how the simulation program actually works. It artificially divides the section of the shock tube into a number of length intervals, called grid cells (or just cells), which each have equal length along the tube. The number of such cells used in the computation is shown in the text box labeled "N Grid Cells" at the very top left on the GAS1D form. In the example shown, this number takes its default value of 200.

The program performs a large number of computations on each grid cell, using data from that cell and its neighboring cells at an initial time level to update the state (i.e. the density, pressure, and velocity) of the gas in the cell to a new time level. All you need to know about this process is that the more grid cells you have, the more accurately the program will simulate real gas behavior. This should be pretty obvious. What is not so obvious is that the amount of time it takes the computer to run your experiment will tend to increase with the square of the number of grid cells you choose to use. With the default value of 200 cells, your experiment should run very fast. With 2000 cells, you will get beautiful, clear results, but your experiment will take a little time. So you could do exploratory experiments with, say, 200 grid cells, and when you find an experiment that exhibits some interesting behavior, you could run it carefully with 800 or 2000 grid cells.

Setting up your initial state:

Now let's talk a bit about how you can set up your own gas dynamics experiment using GAS1D. Of course, you could specify the gas state in every one of your grid cells, but this would take an enormous amount of time. The default problem for which the initial state is displayed in the screen snap shots given above, is defined by the values of the gas state in just 6 grid cells. These are the special cells that are indicated by the solid boxes drawn on the density, pressure, and velocity plots. That there are 6 of these special cells is most easily seen on the pressure plot. The special cell numbers are 1, 20, 21, 180, 181, and 200. You could set up this same problem by entering these numbers successively in the "M Selected" text box, which has

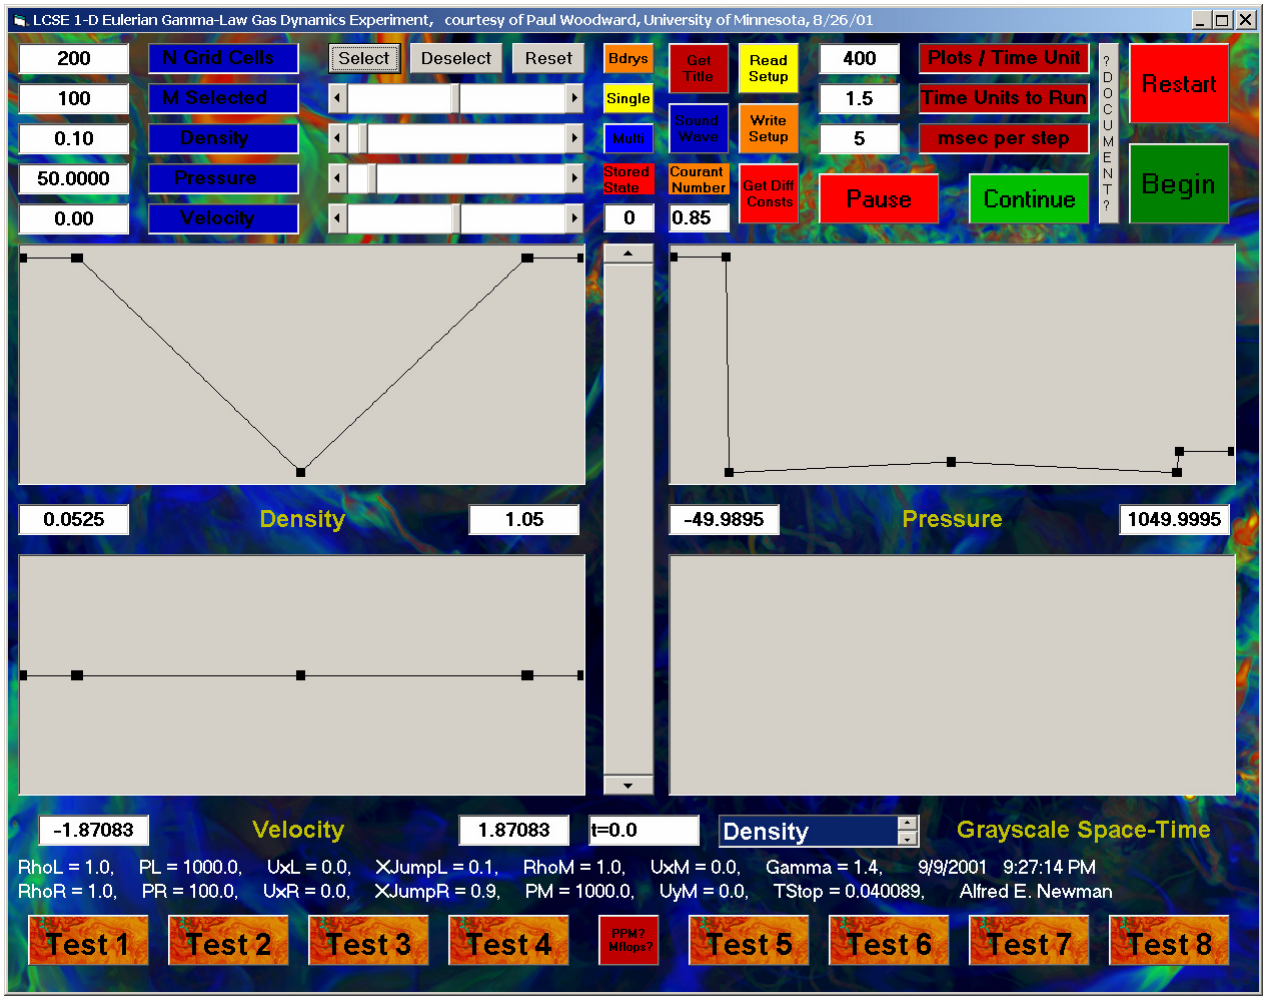

the effect of setting the selected cell number to that value (although this cell is not then immediately "selected," since you must click on the "Select" button at the top of the form to do that). Once the selected cell number is specified in this way, you may enter any desired values into the "Density," "Pressure," and "Velocity" text boxes to designate the state you want the gas to have in this selected cell. When you are satisfied with these state values, you can click on the button labeled "Select" near the top of the form. This will reset the state values in the cell you chose, and it will also reset them between this cell and the nearest special, previously selected cell on either side of your current selection.

Let's show how this works by choosing cell 100 and setting its pressure and density to the values 50 and 0.1, respectively, and by then clicking on the select button. After this action, the

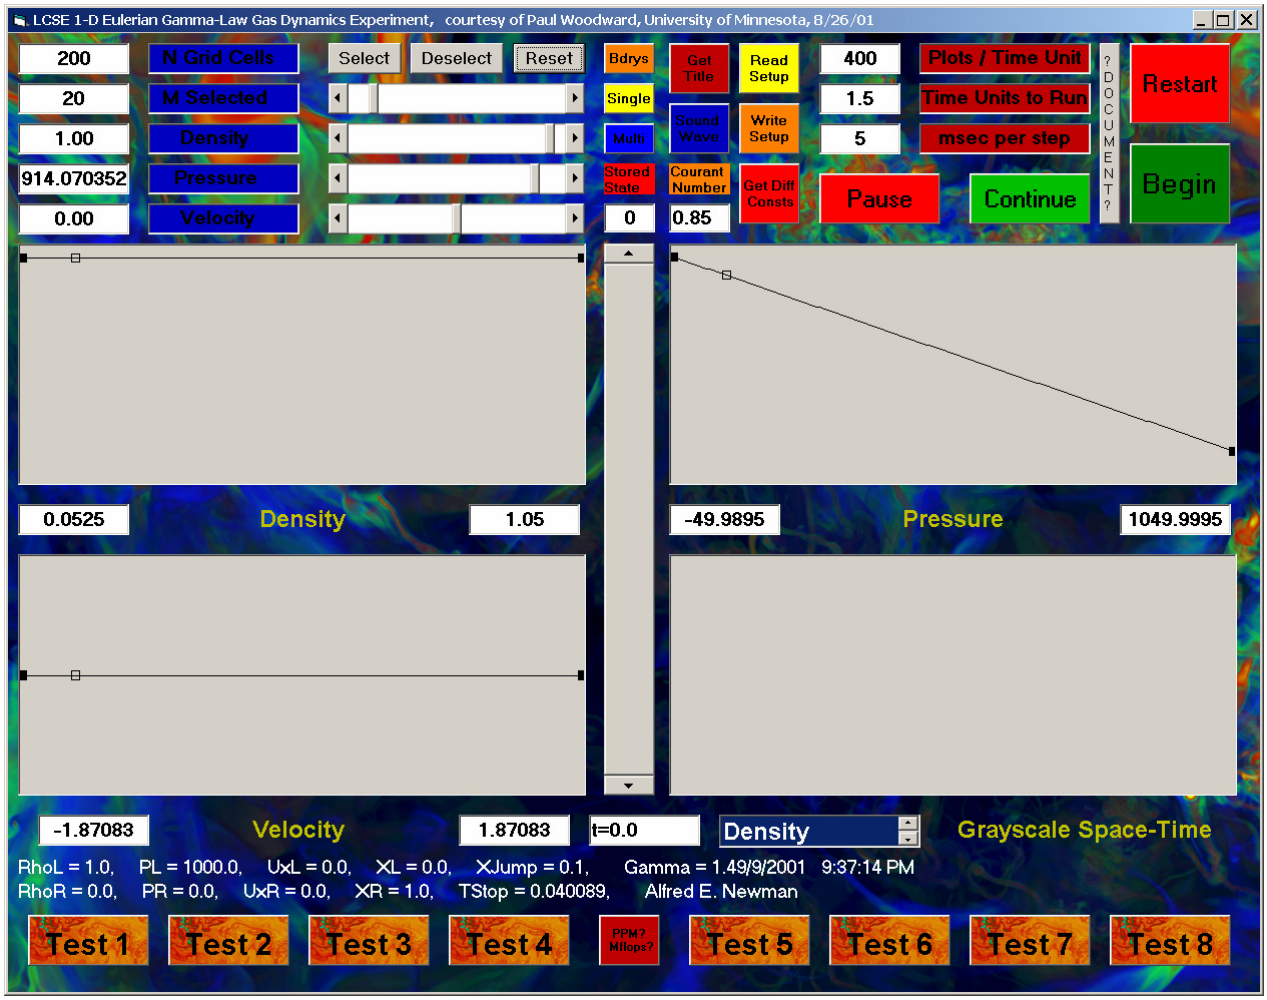

state of the gas is as shown in the screen snap shot on the previous page. In between each special cell that has been officially "selected" GAS1D has computed intermediate states by linear interpolation. That is, each state variable – density, pressure, and velocity – varies along a straight line between the values at the special points. Now let's use this feature of GAS1D to set up a particularly simple experiment. We first want to clear the previously selected grid cells. This is easily done by clicking on the "Reset" button at the top of the form. (If we wanted to clear only specific selected cells, we could go to each such cell in turn and then click on the "Deselect" button.) After clicking Reset, our form looks as shown on this page.

Our clicking on the Reset button has eliminated all the previously selected grid cells, except for the ones at each end of the shock tube grid. What remains is a state in which the density is

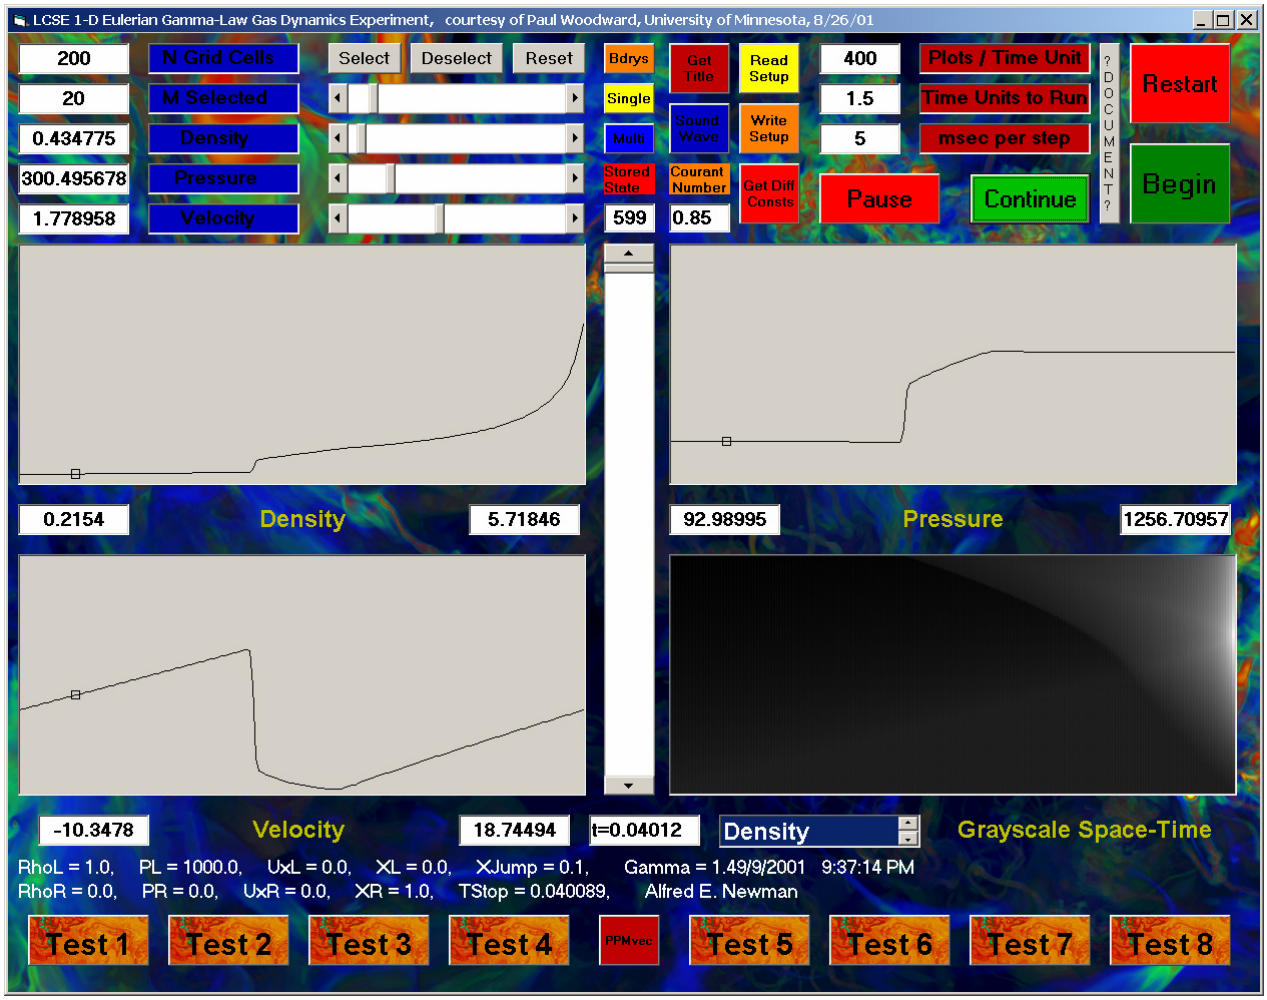

everywhere 1, the velocity is everywhere 0, and the pressure varies along a straight line ramp from a value of 1000 at the left-hand end of the shock tube to a value of 100 at the right-hand end. How the gas will respond to this initial pressure ramp depends upon what happens at the boundaries of our shock tube. By default, each end of our shock tube is a reflecting wall. This means that the end of the tube is sealed, and no gas can enter or escape. This is called a reflecting wall, because, if the wall is hard, sound waves will reflect when they strike it. We can see this behavior by simply running our experiment. To do this, we simply click on the big, green button labeled "Begin."

A few moments later, our screen looks like the snap shot on this page. We chose (in this case by default) to run our experiment for 1.5 sound traversal time intervals. Because the

pressure was so high, a value of 1000, at the left-hand end of our shock tube, this request resulted in our experiment being halted at a time of 0.04, which is displayed near the lower center of the GAS1D form. In the 3 plotting windows we have been discussing, we can see the density, pressure, and velocity distributions with length along the tube at the end of our experiment. The density and pressure are much higher on the right than on the left, which is the natural result of our initial pressure ramp. The high pressures in the left-hand end of our shock tube have pressed the gas up against the right-hand reflecting wall. Although the density was initially uniform in our tube, at the end of our experiment it is much larger at the right-hand end of the tube. This makes sense. The pressure is also higher at the right, which indicates that perhaps all is not over in this flow, and we should expect the gas to slosh back and forth within the shock tube for some time.

We chose to run our experiment for 1.5 sound traversal intervals, and it is pretty clear that this is about the time that has elapsed. We can see this by inspecting the grayscale plot in the bottom right plotting window on the GAS1D form. Beneath this plotting window is the label "Grayscale Space-Time." The "Grayscale" part of this indicates that the display in the window is produced in various shades of gray. White represents the highest value in this plot, and black represents the lowest one. It is a space-time plot, because the vertical axis represents time. The horizontal axis, as in the other plotting windows on the form, represents length along our shock tube. The vertical axis represents time, starting at the bottom at the initial time, time 0, and ending at the top at the final time, time 0.04.

The grayscale values plotted in this window represent the density of the gas in the shock tube. We can see this from the label below the window on the left that says "Density." This label appears in a list box. By clicking on the arrows at the right in this list box, we can change the selection of the variable to plot in this grayscale space-time window. If for example, we want to plot the pressure, we need only to go down one item in the list and then to click once with the left mouse button inside the list box. This causes the plot to be repainted (with pressure values). If

we now choose any particular point within the grayscale plot window and click once on it, the time level we have chosen by the vertical position of the cursor in the grayscale plot is selected, and the density, pressure, and velocity plots are redrawn with the data appropriate to this time in the course of our experiment. By choosing a horizontal position in the grayscale plot to click on, we have selected a particular grid cell, which is then indicated in each plot by an unfilled box, and for which detailed values of the grid cell number, density, pressure, and velocity are given in the textboxes at the top left on the GAS1D form. Just try this a few times. I think you will see that it is a very powerful technique to acquaint yourself with details of the evolution of your experiment.

In fact, with a bit of experience, you should be able to glean most of this information about the time history of your experiment from the grayscale space-time display. As long as you are concentrating on this window, you might as well make GAS1D draw it bigger. To do this, simply double click anywhere inside this window. After doing this, your GAS1D form should look like the snap shot on the previous page. The bright spot at the lower left shows the high pressure on this end of the tube at the initial instant. There is a very clear, bright diagonal line in the plot that joins this spot with a bright spot at the middle right. This marks the time when the gas has sloshed maximally to the right, producing a high pressure on that end of the tube. There is another diagonal bright line that leads from this spot at the middle right toward the middle top of the space-time plot. This is a pressure wave moving back toward the left-hand end of the tube. You can also see a sharp edge of this bright region on the right in the plot that appears to have formed about a third to a half of the way through this experiment. This is a shock wave. It corresponds to a sudden increase in the gas pressure (and density). It is clearly moving from the right to the left through the gas (remember that time increases as we go up in this plot). We asked to stop this experiment after 1.5 sound signal traversals of the length of the tube, and it is clear from the bright diagonal ridges in the plot, that trace the progress of the main pressure wave, that this is just about the amount of time that has indeed elapsed.

If you click once inside the grayscale plot window, your display will return to the original 4 separate plots. In this display mode, you have another method available for reviewing the time history of your experiment. You can play a movie of your experiment by using the vertical scroll bar in the center of the GAS1D form. To get this scroll bar activated, just click on the indicator bar showing the time level that is presently displayed in the 3 plotting windows. Alternatively, you may activate the scroll bar by clicking on either of the arrows at the ends of the bar. Once activated, you may play back the states of the gas in the shock tube during the course of your experiment. These states of the gas have been stored in your computer's memory. The particular stored state number that is presently plotted is displayed above the top of the vertical

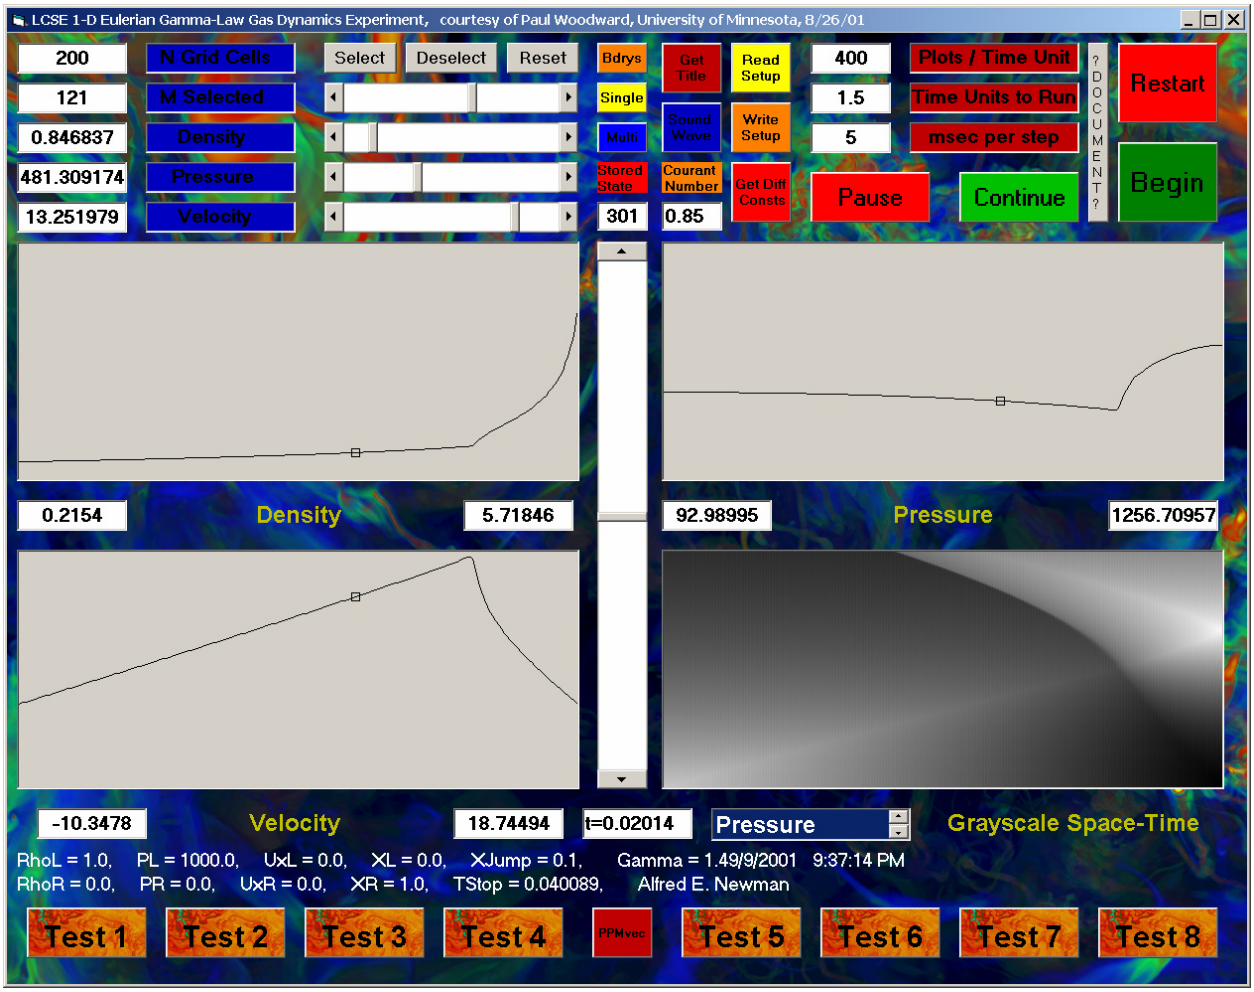

scroll bar in a small text box there. Because (by default) you requested that your experiment be run for 1.5 sound traversal times of the shock tube, and because (again by default) you requested that 400 states be stored for each traversal time interval, you have stored a series of 600 states equally spaced over the time history of your experiment. This is enough states to play back as quite a smooth movie. To play the movie, you can just use the arrow keys on your computer. Holding either of them down causes the movie to play, forward or backward, as fast as your machine can do the required plotting. As your movie plays, the number of each stored state of the gas and the time of that state are displayed in the text boxes above and below the vertical scroll bar.

If you are bored with this experiment and want to set up and run another, just click on the big, red "Restart" button at the top right of the GAS1D form. The display will immediately jump back to the initial state of the experiment just performed. Now that you have pressed the

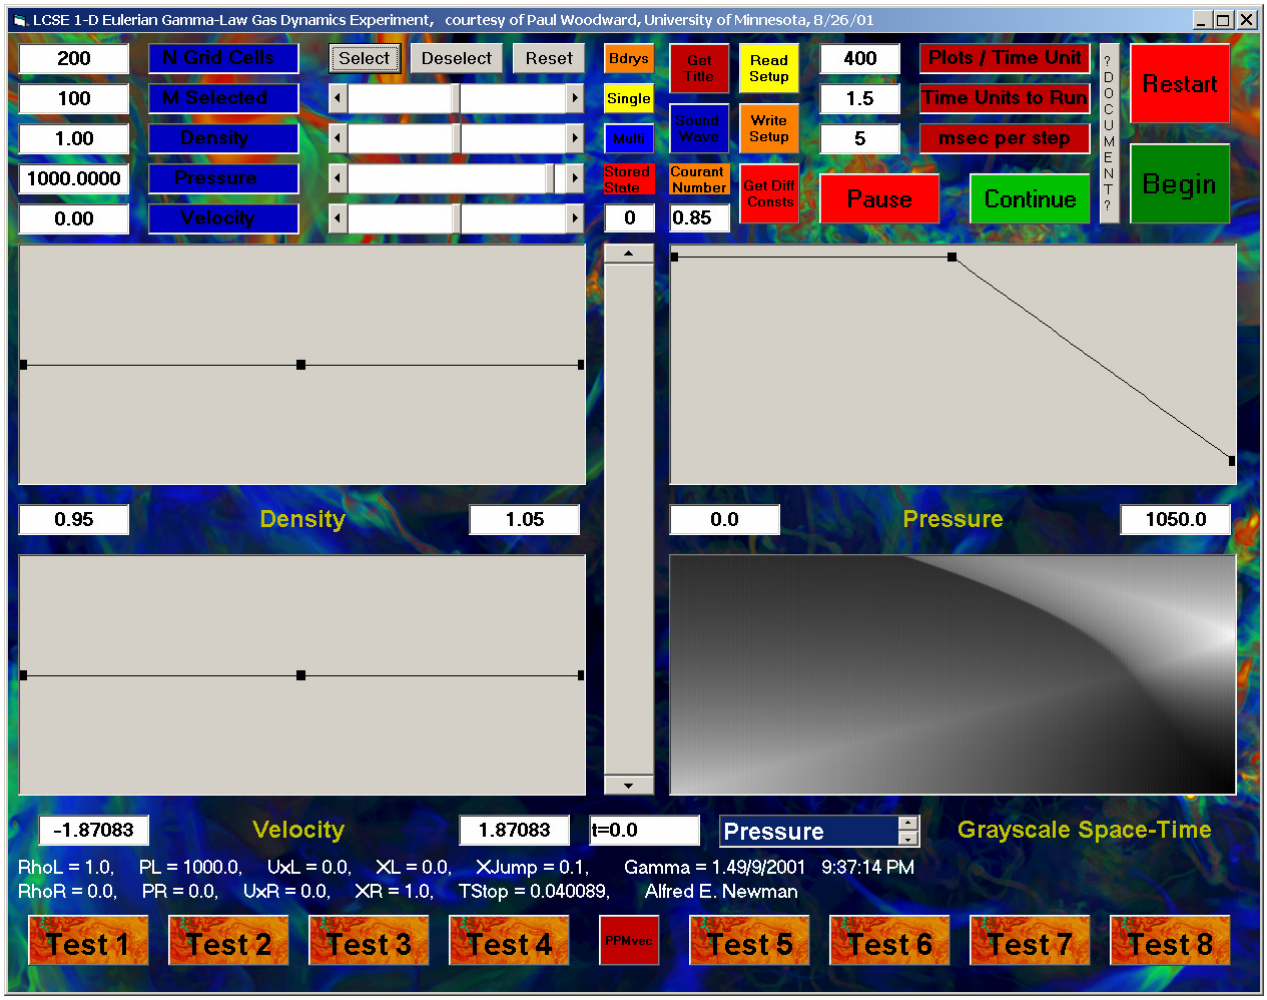

Restart button, several of the other buttons on the form can be activated to help you set up a new experiment to do. You can do a very simple one in which you put the bulk of the available gas in the shock tube on the left-hand side, with a near vacuum on the right. You can do this by introducing just two new selected grid cells. You can choose a cell to select by clicking at the desired horizontal location in any of the plotting windows on the form. Alternatively, you can click on a desired location in the horizontal scroll bar to the right of the "M Selected" label, and then position the scroll bar finely using the arrow keys until the right value appears in the "M Selected" text box. The most precise way to choose a point to select, for those who like typing, is simply to enter the grid cell number into the "M Selected" text box.

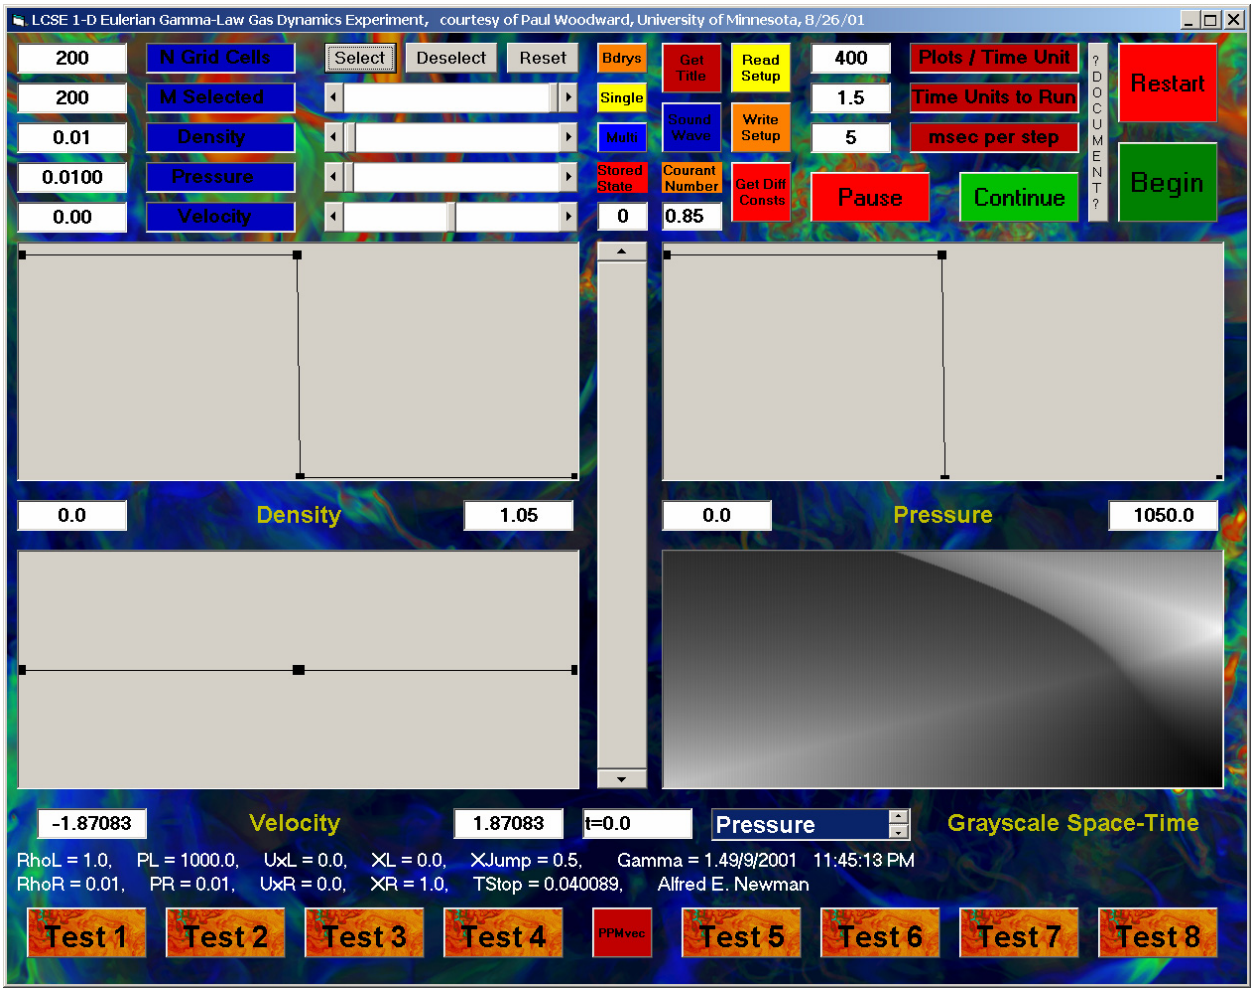

If you enter the value 100 for the grid cell to select, and if you choose to set the density of this cell to 1, the pressure to 1000, and the velocity to 0 and click on the "Select" button, then your form should look like the snap shot on the previous page. Your newly selected special cell is indicated by the filled box on each of your 3 plots. Now you can just use the arrow button on the M Selected horizontal scroll bar, or just modify your entry in the M Selected text box, to choose the adjacent cell on the right, cell number 101, to select. For this cell, you can set the density to 0.01, the pressure to 0.01, and the velocity to 0, then click on the Select button. Finally, you can set these same values for the special, already selected cell at the right-hand end of the shock tube, cell number 200. At this point you are ready to run your new experiment, and your GAS1D form should look like the snap shot on this page.

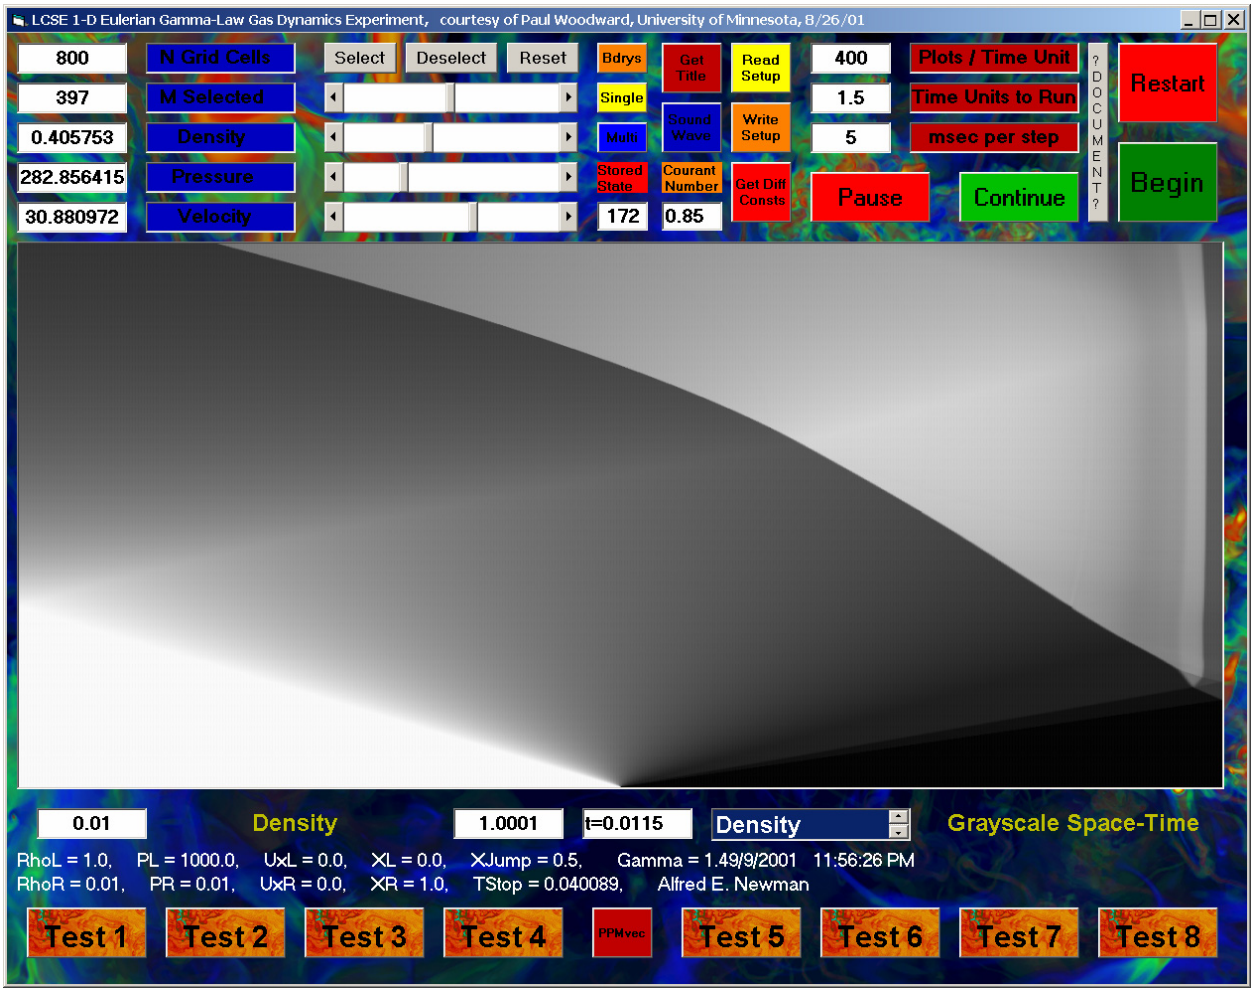

For this experiment, it might be fun to run with a larger number of grid cells, so that we get really crisp, well-resolved results. You can change the grid resolution by entering a new number of grid cells into the text box labeled "N Grid Cells" at the top left on the GAS1D form. Why fool around? Why not enter the number 800, so the grid will be really fine? This will not change the appearance of your form much now, but its appearance will be different later, since all the plots will be better. Before we launch this new experiment, it might be a good idea to save its state, so that we can come back to it instantly in the future without having to go through all the steps documented here. This is easy to do. You only need to click on the "Write Setup" button near the top center of your form. You will be prompted for a file name to save this information in on your computer's hard disk. This is a classic Microsoft Windows "save file" pop-up window, and it has all the features that these windows do in Microsoft WORD. Pick a file name, and

navigate to the desired directory, and the information will be saved. In the future, you can read this information back by clicking on the "Read Setup" button, then navigating to this directory and selecting this file. When the file is read, all the setup information will immediately be restored, and you will again be all set to run this particular experiment.

So why don't you run the experiment now? If you do, the result will look like the snap shot on the previous page, where the density has been selected for the grayscale plot, and that plot has been drawn large. There are some intricate wave reflections and interactions that happen at the right-hand end of the shock tube after the first shock wave strikes, but these are just details. Mostly, the gas just expands into the nearly empty portion of the shock tube, reducing the density at the left, and then of course the pressure builds up at the right end of the tube, because of the resistance of the reflecting wall there, and the gas sloshes back toward the left, with a shock front marking the leading edge of the advancing gas. At the far right in the shock tube, there is a thin layer of squashed gas. This is the gas that was originally occupying the nearly evacuated portion of the shock tube. It has the same pressure as the gas further to the left, but it is less dense. This is because it is very hot, having been through such traumatic compression in multiple shock waves. But of course this is just a detail.

Changing the behavior at the boundaries (the ends) of the shock tube:

The behavior you just observed will change dramatically if you change the behavior at the ends of the shock tube. This behavior at the ends of the tube is referred to by professionals as the "boundary conditions" for this gas dynamics problem. You cannot change the boundary conditions for a running problem in the midst of its evolution. Therefore you must first click on the Restart button, in order to return to the initial time and to set up a new experiment. Now click on the small "Bdrys" button at the top center of the GAS1D form. In a pop-up window, you are prompted for the boundary condition type of the left-hand boundary of the shock tube. You are "reminded" that a value of 4 will give periodic boundary conditions (which we will discuss later), a value of 1 will give a reflecting wall at the left-hand end of the tube, and a value of 2 will give a

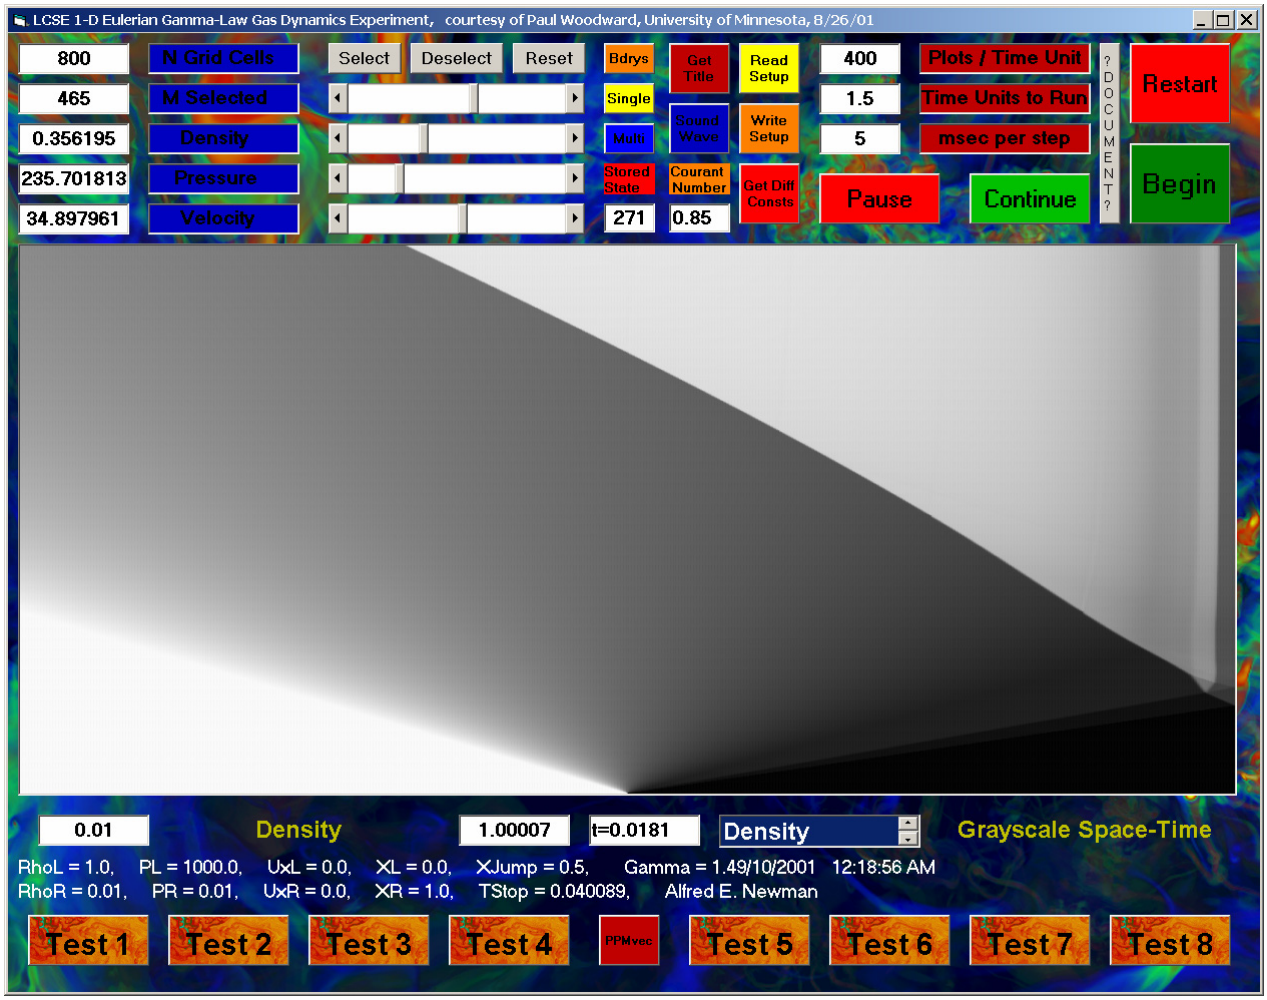

"continuation boundary" there. Just for fun, type 2 in this input box, for a continuation boundary on the left, and then type 1 for a reflecting boundary at the right.

You can see what difference this continuation boundary on the left will make by running this new experiment. The result is shown in the snap shot on this page. It looks like the previous result, except that there is no clear signal reflected from the left-hand boundary. If you play back the movie of this experiment, you will see that once the rarefaction wave reaches the left-hand end of the shock tube, it simply passes right on. No signal is reflected. The rarefaction simply passes on up the shock tube into a region whose behavior we do not simulate. The way that this behavior is achieved, technically, is that the GAS1D program always sets the state of the gas to the left of the left-hand boundary of our section of the shock tube to be exactly the same as the gas state

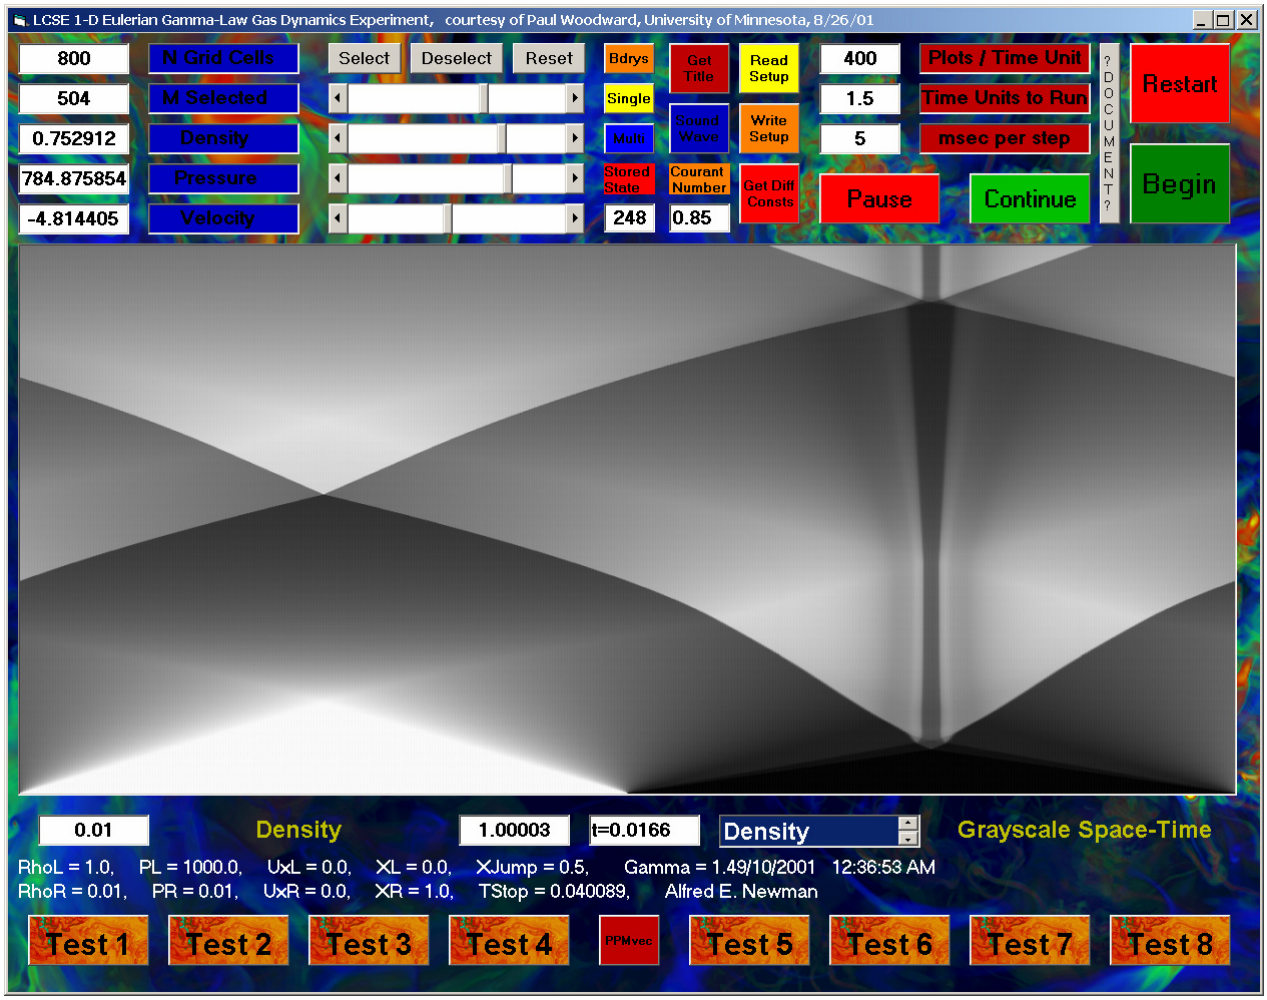

right at the left-hand end of our section. Thus, no signals can enter our section through this boundary that contain information different from that which is already there. This is actually a rather complicated thing. A technical detail. You do not need to understand it (and there are quite a few professionals who don't understand it, so you should not feel bad if you are puzzled).

So much for the continuation boundary. What about the periodic boundary conditions, which you can select by clicking again on the Restart button and then clicking on the Bdrys button? Why not just do this? When the input box pops up, just type in 4 for your boundary type selection. Then click on Begin to run the experiment. Your result should look like the snap shot on this page. It is quite different from the other two. This time signals emanate immediately from both the left and the right boundaries. The gas flow that we get is the same as

we would have if our section of the shock tube were replicated an infinite number of times on each side. That is, the gas state to the left of the left-hand end of the tube is identical to the gas state to the left of the right-hand end of the tube. Thus there is initially a huge jump in pressure and density at each end of our shock tube as well as in its middle. It is evident from the space-time plot of the density that each time a signal leaves our section of the shock tube through the right-hand boundary, the identical signal enters our section of the tube through its left-hand boundary.

Studying sinusoidal sound wave propagation with GAS1D:

Periodic boundary conditions are most commonly used to study the behavior of periodic waves. Examples of periodic waves are water waves on a lake, light waves, or sound waves propagating in air. The GAS1D program is instrumented so that you can study periodic sound wave propagation in air. You can do this by clicking on the Restart button (or by restarting the program from the beginning by double clicking on its executable file) and then by clicking on the "Sound Wave" button near the upper center of the GAS1D form. The last experiment that we performed showed the time development of a standing wave in our shock tube. This was a special kind of standing wave. It began with a "square" waveform and its initial amplitude was huge, resulting in the formation of multiple shock fronts.

When you click on the "Sound Wave" button, the GAS1D program will initialize for you sinusoidal sound wave disturbances for you to study. These would be very difficult to set up by manually specifying gas states in individual grid cells. Instead, GAS1D prompts you for the amplitudes (which can be positive or negative) of the sinusoidal disturbances you want to initialize for sound waves traveling to the right (you are prompted for the amplitude of the "R+ Riemann invariant" disturbance), for sound waves traveling to the left (you are prompted for the amplitude of the "R- Riemann invariant" disturbance), and for sinusoidal disturbances that move along with the gas (you are prompted for the amplitude of the entropy or the "R0 Riemann invariant" disturbance). For all these experiments, initialized in this fashion using the "Sound Wave" button, the sound speed at the left-hand end of our section of the shock tube is

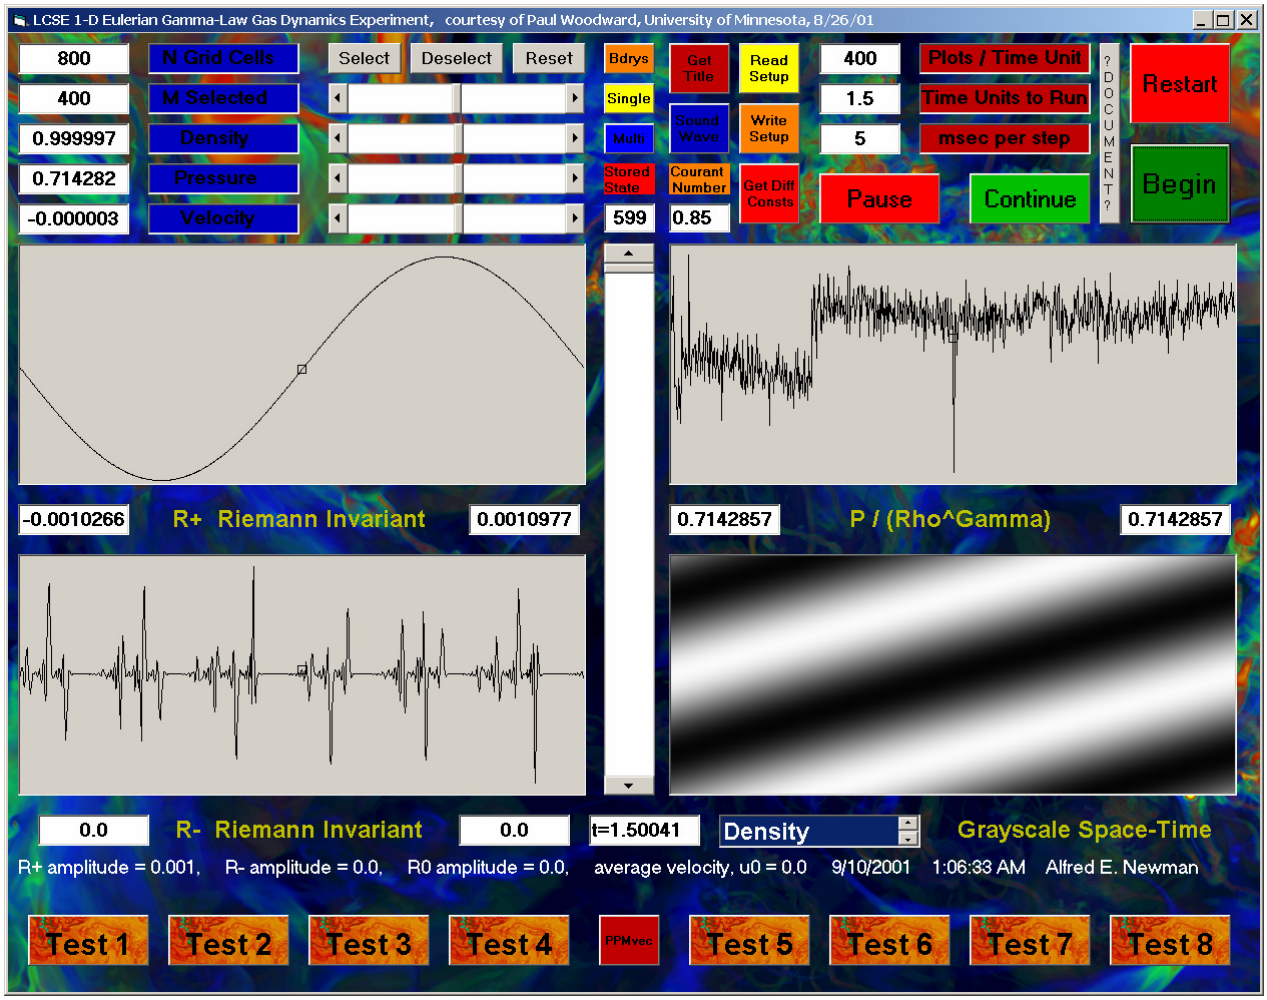

automatically set to a value of 1. This means that a small-amplitude sound wave will traverse our portion of the shock tube, that is, it will propagate for one wavelength down the shock tube, in one unit of time. This is a convenient convention. For this convention, a strong sound wave has an amplitude of 1, and a weak sound wave has an amplitude small compared to 1. This is also convenient for our study of sound wave propagation.

As a first test, why not set up an isolated, rightward moving sound wave of small amplitude, say, with amplitude 0.001? Just click on the Sound Wave button, and enter 0.001 for the R+ Riemann invariant signal amplitude, enter 0 for the amplitudes of the R- and R0 Riemann invariant signals, and enter 0 for the average gas velocity at the final prompt. Thus we expect to see our sine-wave signal simply move to the right down the tube at an average velocity

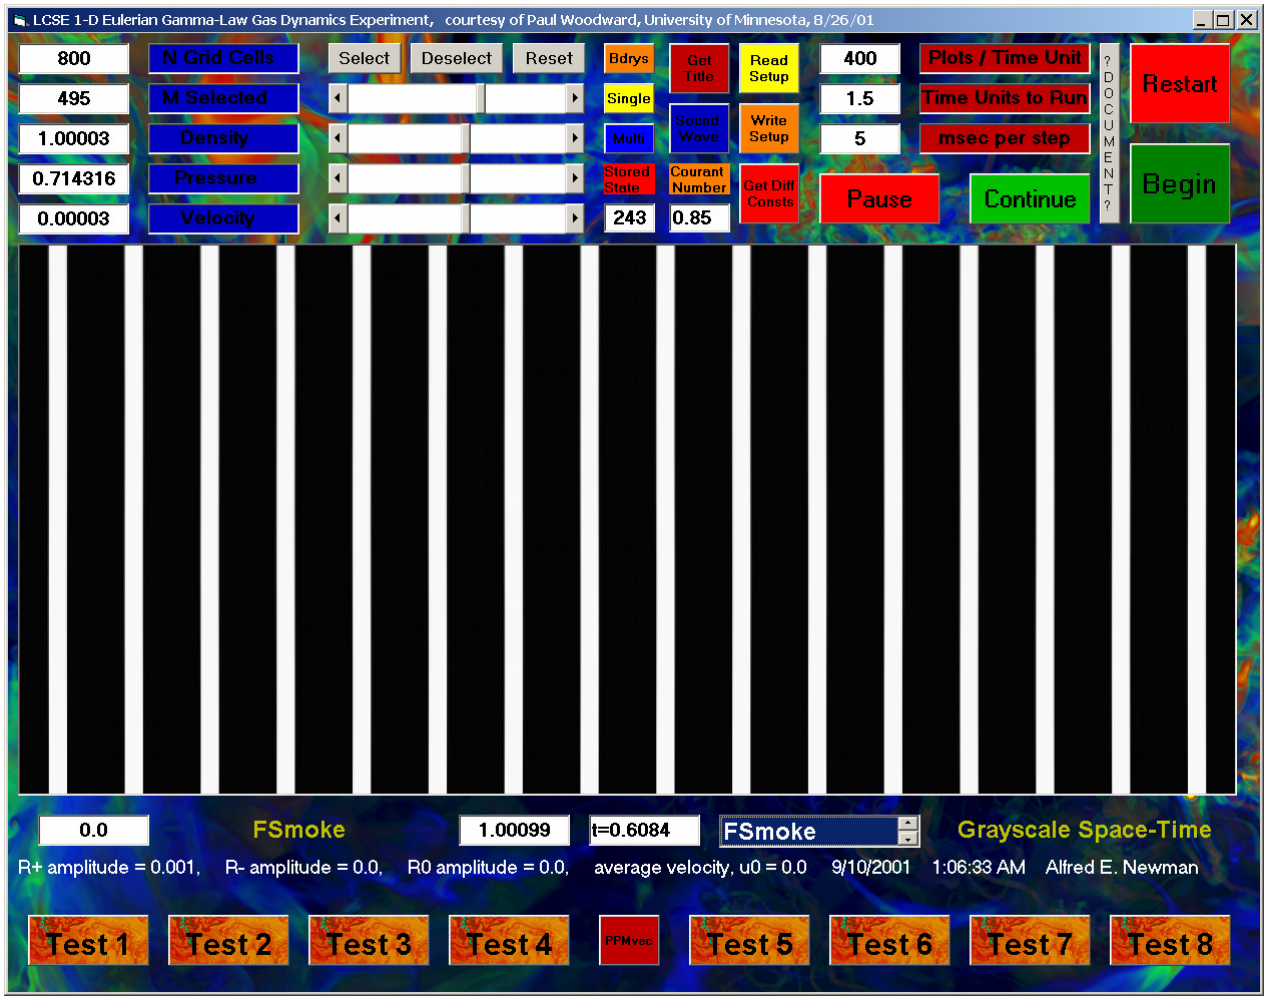

of 1 through a gas which is, on average, at rest in our shock tube. Click on the Begin button to run this experiment. At the end of the experiment, your screen should look like the snap shot on the previous page.

Your screen looks a little different than before, because for these experiments, GAS1D automatically plots the individual sound wave signals instead of the state variables density, pressure, and velocity. This makes it easier to see what happens to your waves as they propagate. In the upper left-hand plotting window, the rightward-moving sound wave signal, the so-called Riemann invariant, R+, is shown. Note that it is still a nearly perfect sine wave, after moving rightward along the shock tube for 1.5 wavelengths. Its amplitude is also essentially unchanged. In the plotting window below this, we can see the sound wave signal that moves to the

left along the shock tube. You requested no signal in this mode, and you get a pretty good numerical approximation to no signal (note that the minimum and maximum values for R- in this plot are both, according to the labels shown, 0.0). At the upper right is plotted the entropy signal. The entropy of the gas is here defined as the quantity $p/\rho^{1.4}$. This entropy is the R0 Riemann invariant, for which you requested no signal. Once again, you see that you get a good approximation to no signal.

This simple sound wave propagation experiment is not so interesting. Nothing exciting happens. However, we can use it to illustrate an interesting capability of the GAS1D program. This is a gas flow visualization technique like that used in many lab experiments (see for example the images in the book **An Album of Fluid Motion**, by van Dyke). GAS1D introduces smoke into your gas flow at the initial instant in 16 equally spaced, equally wide regions of your shock tube. It then follows this smoke as it flows along with the gas in the tube. You can see a plot of these smoke streams in the grayscale space-time plotting window by selecting the display variable "FSmoke" (fractional volume of smoke) in the list box below the window. It is often nice to draw this plot large. For this sound wave experiment, the result is shown on the previous page. The smoke streams just go vertically upward in the diagram. This is because the amplitude of the sound wave is so small that there is hardly any perceptible gas motion as the wave passes by.

Let's do just one more experiment before we quit. Let's try to see how the gas moves as a strong sound wave passes down our shock tube to the right. To do this, click on the Restart button. Then click on the Sound Wave button. At the prompts, set the R+ wave amplitude to 1, the R- and R0 amplitudes to 0, and the average gas velocity to 0. Select FSmoke in the grayscale space-time plot list box and double click in this plotting window. Now click on Begin. The smoke streams should show you what the gas does as the wave passes by. What's your interpretation of the result?

One final button:

There is just one final button you should know about. It is located at the bottom of the GAS1D form in the center. It is labeled "PPM? Mflops?" When you click on this button while setting up an experiment, you are asked to select a numerical method from among a series of choices. Of the choices listed at the prompt, only 2 function properly at this writing – PPMvec (the default) and PPMvec32. These are versions of the PPM difference scheme that are optimized for execution on Intel processors. PPMvec does its arithmetic in 64-bit precision, which makes it run about 2 times slower than PPMvec32, which uses 32-bit precision. For almost any experiment you are likely to set up, PPMvec32 will get you the results in perfectly adequate detail and precision in about half the time required by PPMvec. However, the 32-bit arithmetic of PPMvec32 will cause sinusoidal propagating waves represented by only a few cells per wavelength to be artificially damped in amplitude. This amplitude damping of poorly resolved waves is a desirable feature of a numerical method, but PPMvec32 overdoes it a bit in comparison to the behavior you get with PPMvec. This is a detail, of course, and you need not concern yourself with it.

What about the second prompt that you get when you click on the "PPM? Mflops?" button? This prompt asks you to type in a 1 if you want Mflop/s metering turned on or to type in a 0 if you don't want this additional overhead. What does Mflop/s stand for? One Mflop is one million floating point operations (things like adding, multiplying, or dividing two non-integer numbers). One Mflop/s is a computational speed of one million floating point operations per second. The PPM numerical methods implemented in GAS1D automatically count the number of flops (floating point operations) they perform as they go along. This costs next to nothing. What introduces overhead, and slows your experiment down, is asking GAS1D to use this feature of PPM to estimate how fast your processor is working within the PPM gas dynamics module. That requires making sure that the time meter on your machine, which is not as accurate as it might be, has sufficient amounts of PPM work to do in between each time

measurement. In this way, GAS1D is able to return to you accurate measurements of the processing speed of your computer on this numerical method. To do this, GAS1D will perform redundant PPM work upon each invocation of the numerical scheme, so that enough time elapses to make good performance measurements. This wastes time. So if you aren't interested in how fast your PC computes, why turn on Mflop/s metering?

For the professionals who might, by some odd chance, be reading this document, I should point out that GAS1D calls PPM modules that perform fully 3-D computations. The 1-D experiments that GAS1D sets up do not need 3-D gas dynamics computations, of course, but in order to get meaningful performance numbers from this program, full 3-D PPM modules have been inserted. They perform much more work (about 50% more, in fact) than is actually needed for your 1-D gas dynamics experiments. But they run so fast (420 Mflop/s on my laptop machine), who cares about a little extra work? This is another one of those little details that you are best advised to simply ignore.

Oh, and one more thing. Along the bottom of the GAS1D form you may notice a row of buttons labeled "Test 1," "Test 2," and so on. When you click on these buttons, GAS1D automatically sets up the specified standard test experiment. You should ignore these buttons. These tests were taken from a study carried out recently by investigators at Los Alamos National Laboratory that compared a number of different numerical methods for gas dynamics, of which PPM was one such method. I don't know why these particular tests were chosen, although it is a good bet that either (1) each test caused one of the codes in the study to crash or to exhibit otherwise unacceptable behavior, or (2) each test caused the numerical method favored by the designers of the study to stand out in some significant positive way. Having performed such a comparison study myself in my youth, I am sympathetic to the Los Alamos scientists who worked on this, even though they do not appear to have chose my own scheme, PPM, as their favored choice (we could argue about this, but who wants arguments?).

Enough talk. Go set up some 1-D gas dynamics experiments now and explore gas dynamical phenomena for yourself!

Here's the way it looks at the end of the experiment:

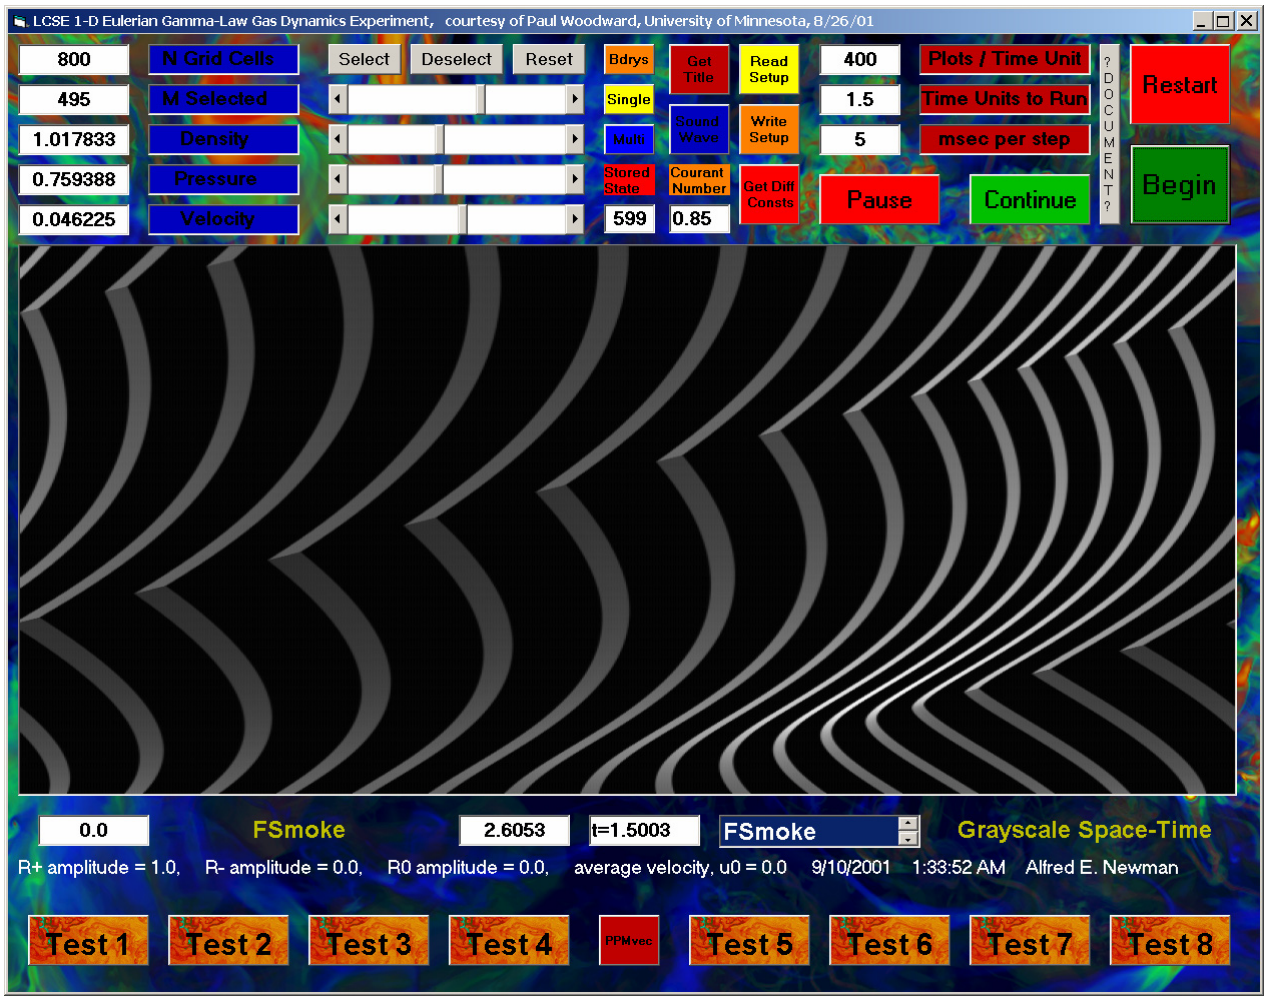

In this case, the space-time diagram covers the entire display area of the form. The user has selected the variable FAir, actually a kind of smoke tracer that moves with the gas, in the list box below the picture. To make this plot small, the user clicks once on a point in the space-time frame of the picture. Then the space-time diagram is repainted in the lower right, and the distributions of the density, pressure and x-velocity are plotted around it. These plots display the variable distributions in space at the time corresponding to the designated point in the space-time diagram window. Also, each such plot shows a selected spatial point, which also corresponds to the point designated in the space-time window. Finally, the values of the variables at the selected point in space at this time are displayed in the text boxes in the upper left portion of the form.

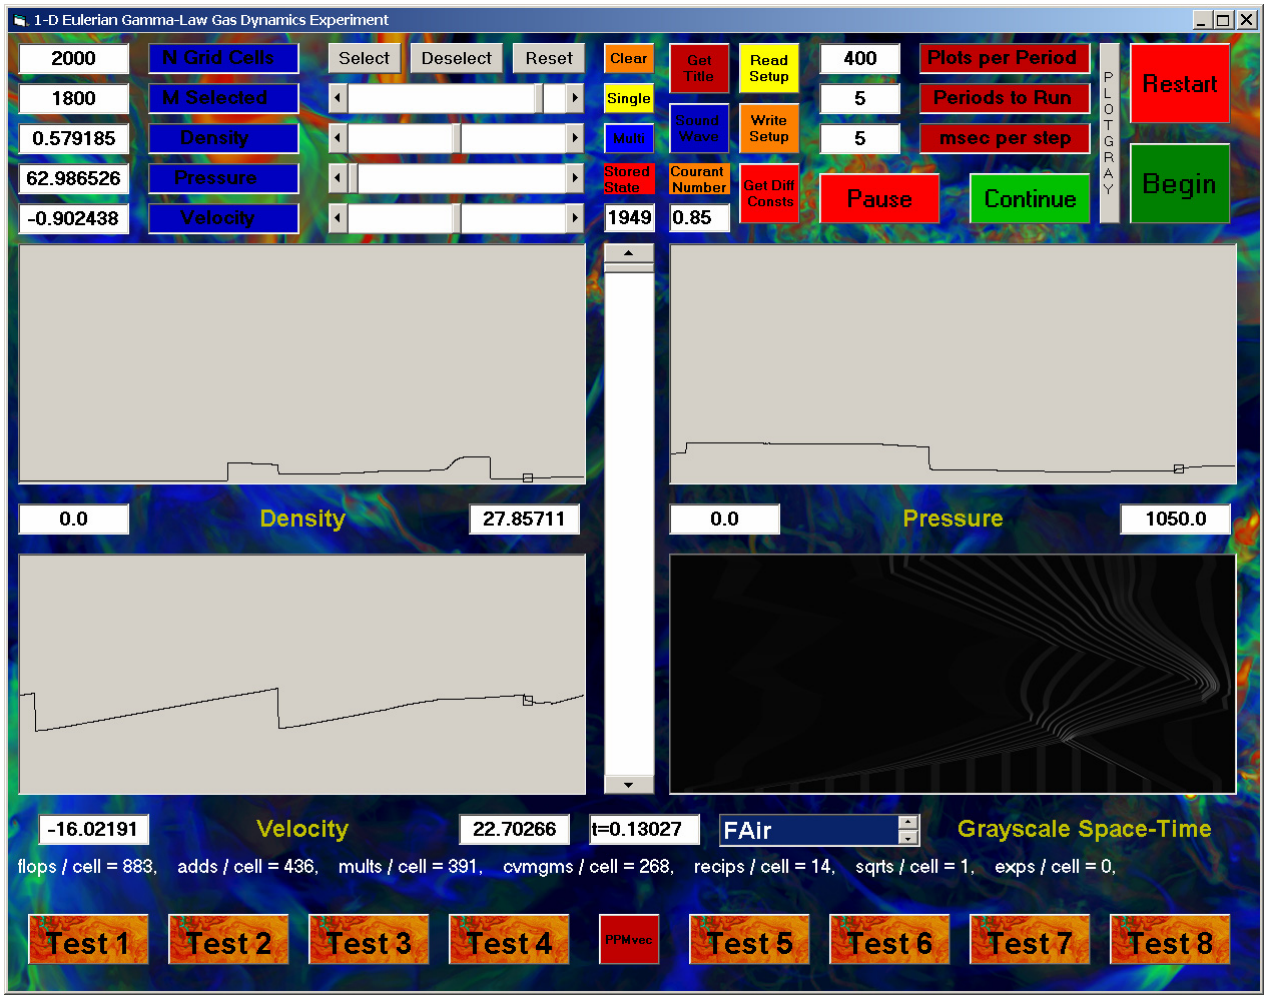

The user can also select a time level to display by positioning the vertical scroll bar at the center of the form. The time level is then displayed below the scroll bar, and the number of the stored state of the gas is displayed just above the vertical scroll bar. In order to operate the scroll bar, however, the user must first click on the "Pause" button, to pause the ongoing computation. Then by holding the mouse button down on the scroll bar arrow buttons (or holding down the keyboard arrow keys once the scroll bar is selected as active), the user can play back the stored states as a sort of movie.

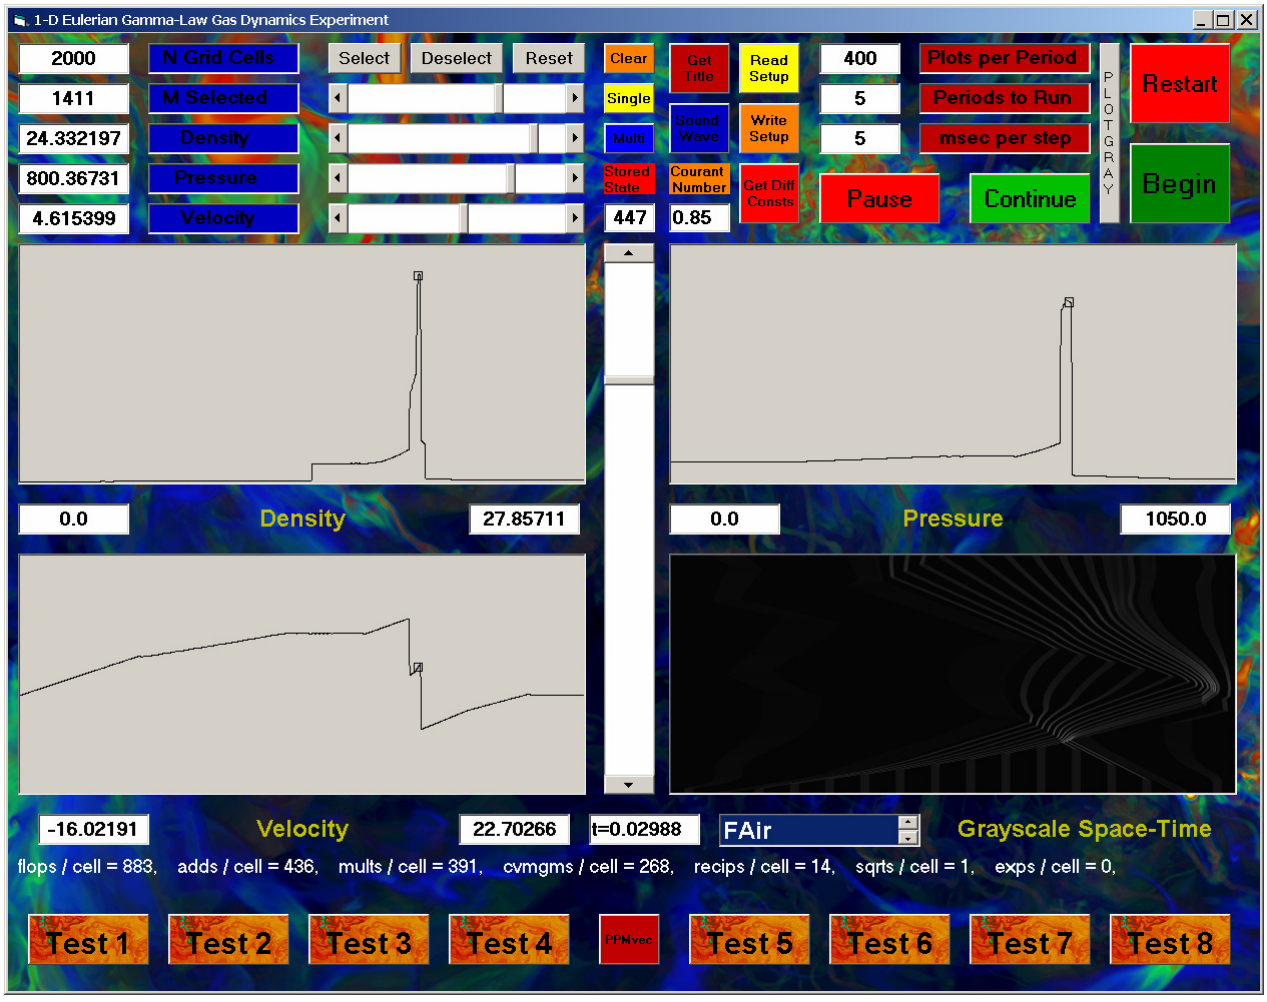

The user can also create any shock tube problem he or she desires by using the horizontal scroll bars and the text boxes for the variable values to set the physical state at selected points, between which a linear interpolation will be performed. Preset test problems can also be invoked via the row of "Test" buttons along the bottom of the form. The red button in the center of this bottom

row of buttons, when clicked, prompts the user for the designated difference scheme and any desire for Mflop/s performance metering (if offered for that difference scheme).

One may restart a problem from the beginning by pausing it and then clicking on the "Restart" button, followed by the "Begin" button.

One can create beautiful patterns of smoke ribbons by running sound wave steepening problems. Here is the result for a sound wave amplitude of unity in the R+ Riemann invariant alone:

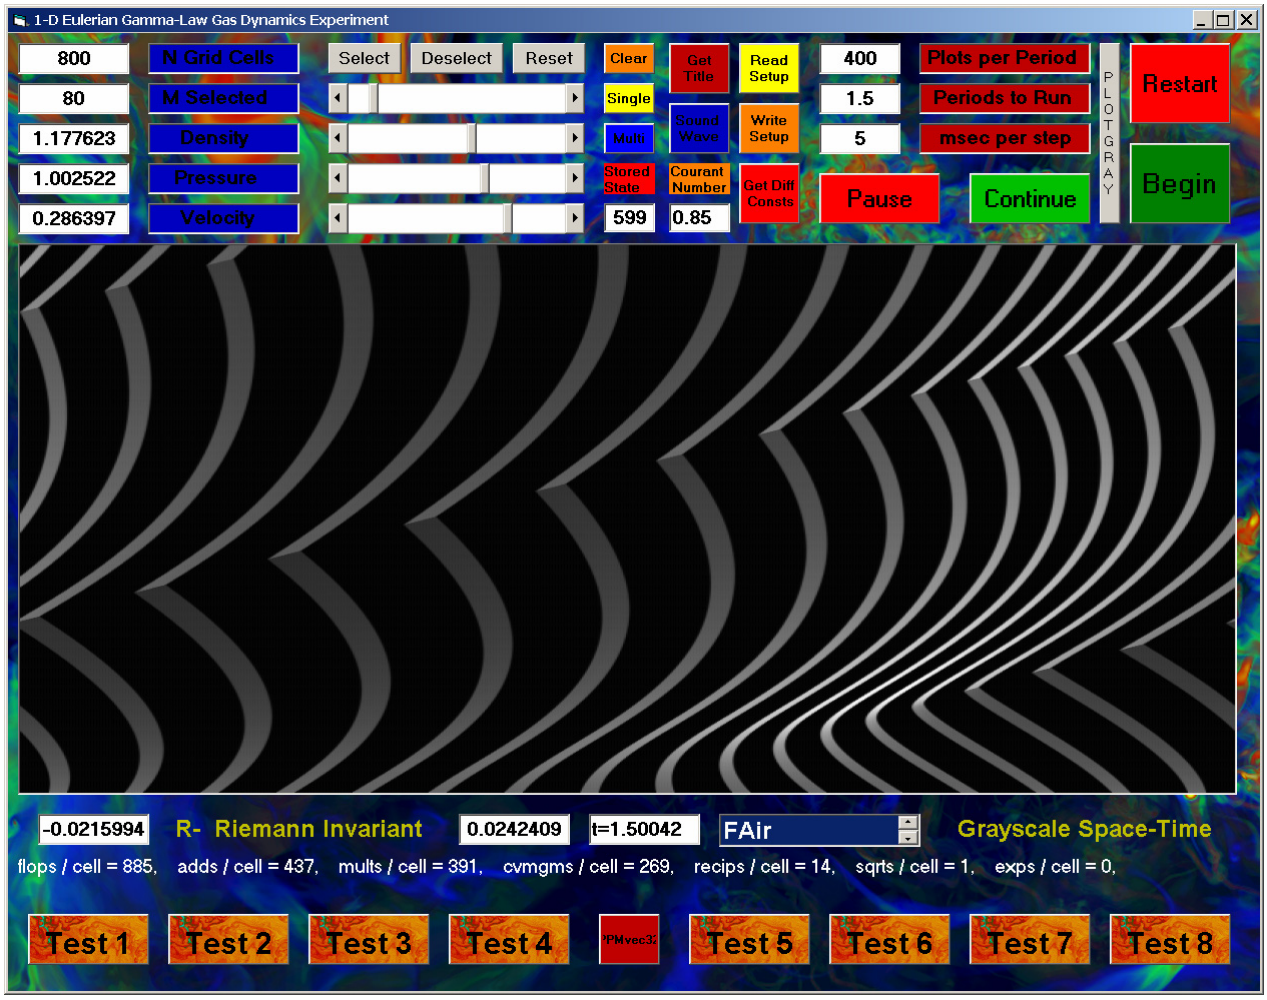

Pretty cool, hunh?

To run this type of problem, just click on the "Sound Wave" button (after first clicking "Restart"), then click "Begin."

Here's a run for a tenth of this wave amplitude:

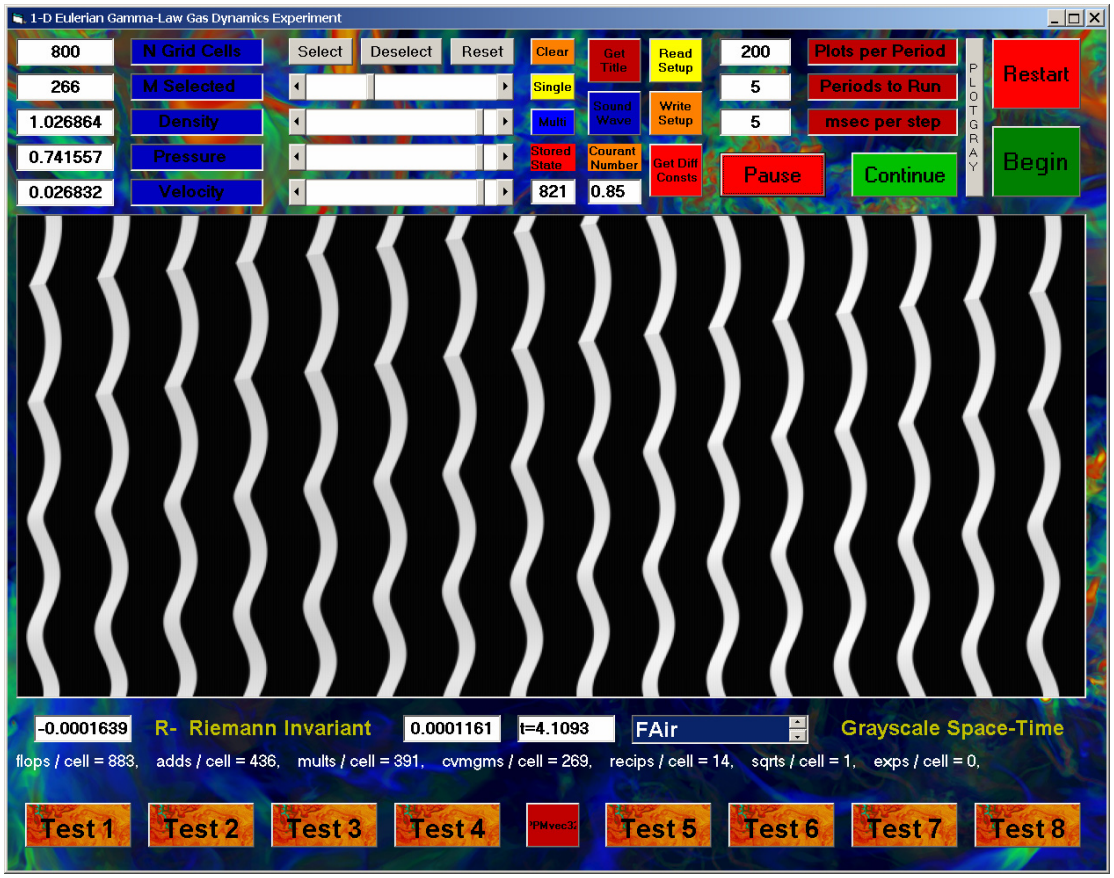
1-D Eulerian Gamma-Law Gas Dynamics Experiment
800
N Grid Cells
266
M Selected
1.026864
Density
0.741557
Pressure
0.026832
Velocity
Select
Deselect
Reset
Clear
Single
Multi
Stored State
821
Get Title
Sound Wave
Courant Number
0.85
Read Setup
Write Setup
Get Diff Consts
200
Plots per Period
5
Periods to Run
5
msec per step
Pause
Continue
PLOTGRAY
Restart
Begin
-0.0001639
R- Riemann Invariant
0.0001161
t=4.1093
FAir
Grayscale Space-Time
flops / cell = 883, adds / cell = 436, mults / cell = 391, cvmgms / cell = 269, recips / cell = 14, sqrts / cell = 1, exps / cell = 0,
Test 1
Test 2
Test 3
Test 4
PMvec32
Test 5
Test 6
Test 7
Test 8

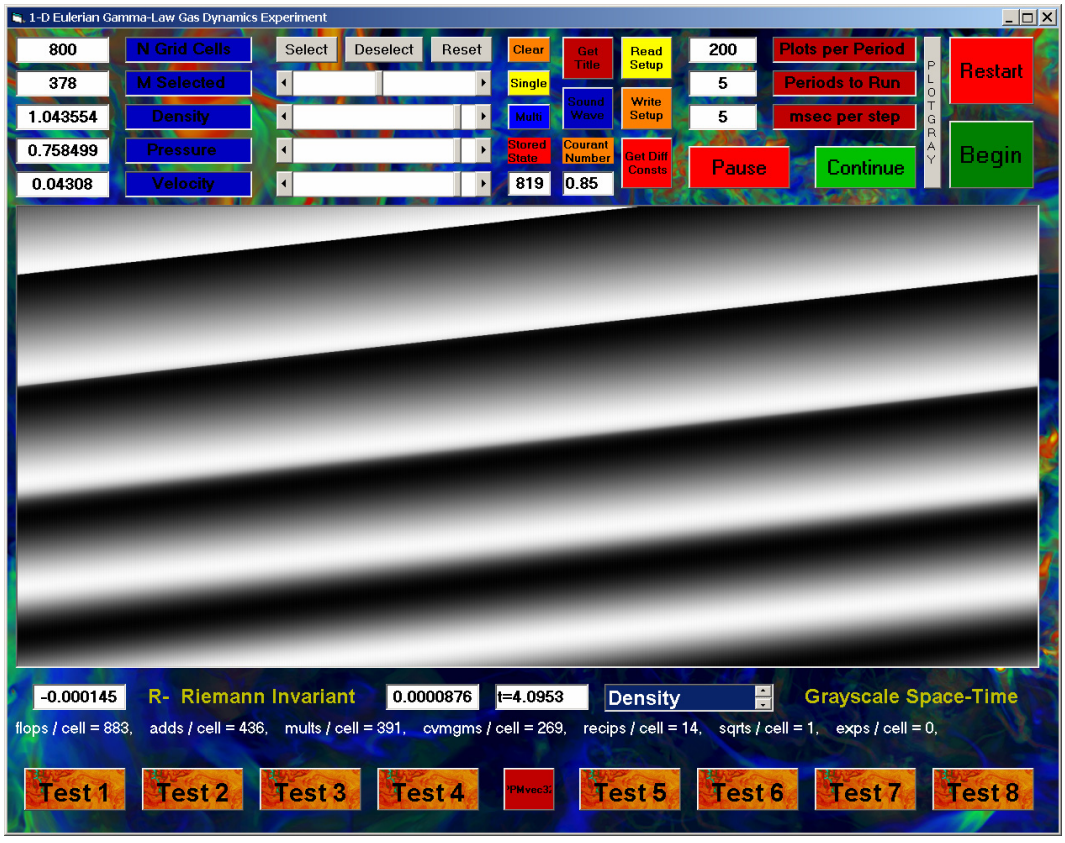
1-D Eulerian Gamma-Law Gas Dynamics Experiment
800
N Grid Cells
378
M Selected
1.043554
Density
0.758499
Pressure
0.04308
Velocity
Select
Deselect
Reset
Clear
Single
Multi
Stored State
819
Get Title
Sound Wave
Courant Number
0.85
Read Setup
Write Setup
Get Diff Consts
200
Plots per Period
5
Periods to Run
5
msec per step
Pause
Continue
PLOTGRAY
Restart
Begin
-0.000145
R- Riemann Invariant
0.0000876
t=4.0953
Density
Grayscale Space-Time
flops / cell = 883, adds / cell = 436, mults / cell = 391, cvmgms / cell = 269, recips / cell = 14, sqrts / cell = 1, exps / cell = 0,
Test 1
Test 2
Test 3
Test 4
PMvec32
Test 5
Test 6
Test 7
Test 8

In the example below, we have set 3 separate initially sinusoidal disturbances going: (1) a sound wave traveling to the right with initial amplitude 0.5, (2) a sound wave traveling to the left with initial amplitude 0.15, and (3) a sinusoidal entropy variation with amplitude 0.01. In this case we perform the simulation in a frame in which the gas is initially, on the average, moving to the left at the average ambient speed of sound. Here's what the smoke trails look like after 3 sound wave traversal times of our periodic system:

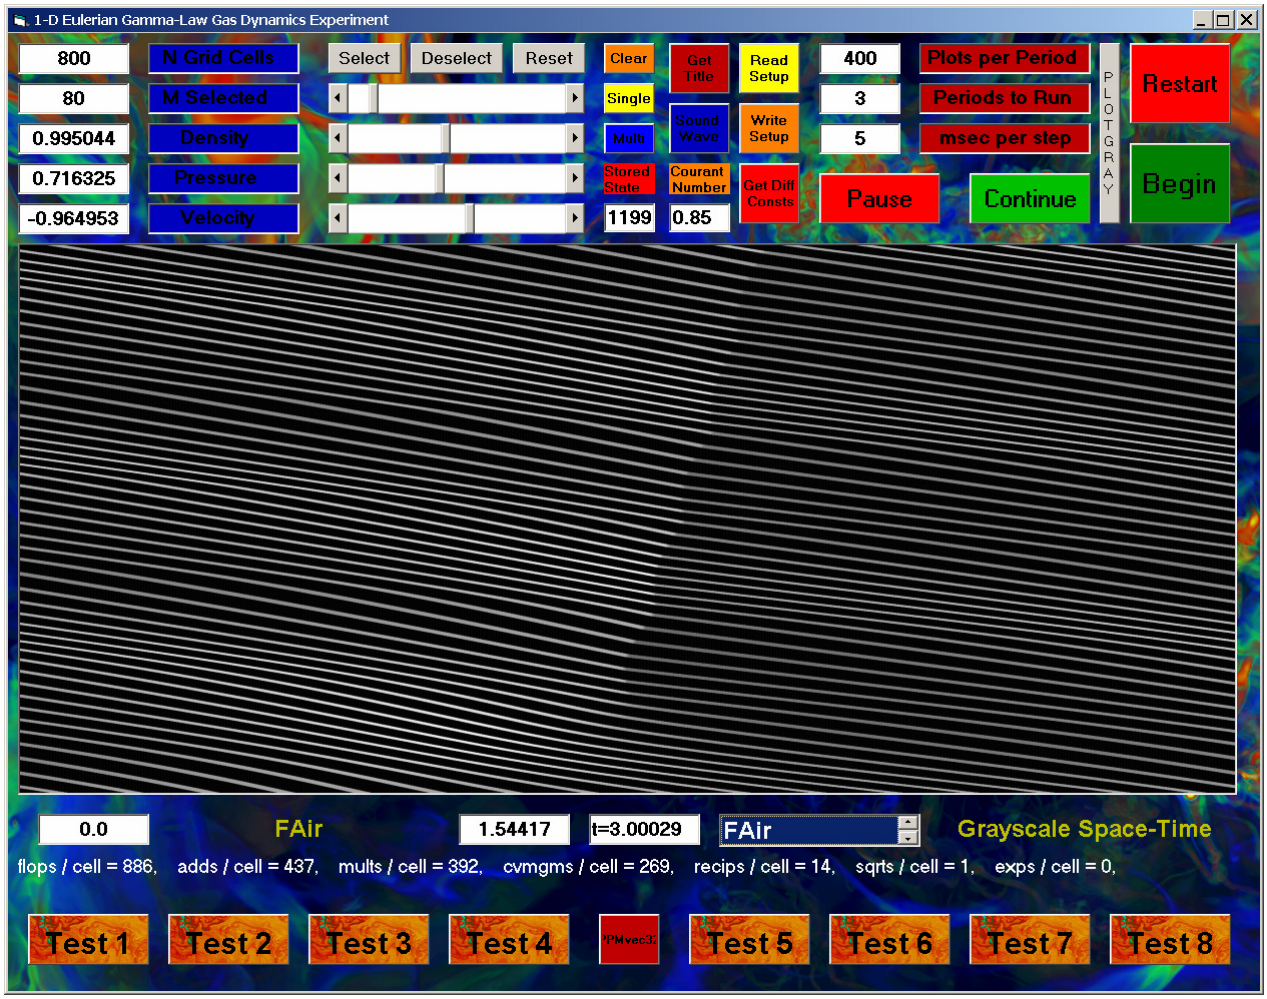

From this pattern, it's easy to pick out the wavefronts for the sound wave moving to the left. You can easily see the periodic compression and expansion of the smoke trails. Regions of nearly equal compression or expansion travel leftward as we go up in the diagram (that is, as time advances). The average slope of the smoke trails shows the motion of the gas to the left at about the speed of sound (a speed of unity). There are just about 3 cycles of these compressions and expansions along

any vertical line in the diagram, and the time advances by 3 units from the top to the bottom in this case. The crease-like feature, where the smoke is suddenly compressed in density marks the slowly advancing position of a shock front that has formed from the steepening of the right-moving sound wave. This steepening is very much like the similar phenomenon of water wave steepening as a wave approaches the beach. The formation of the shock front is akin to the formation of a nearly vertical wall of water. Immediately after this, the water wave breaks, but the shock front in the gas does not do that. Now let's look at the picture of the density from this same simulation:

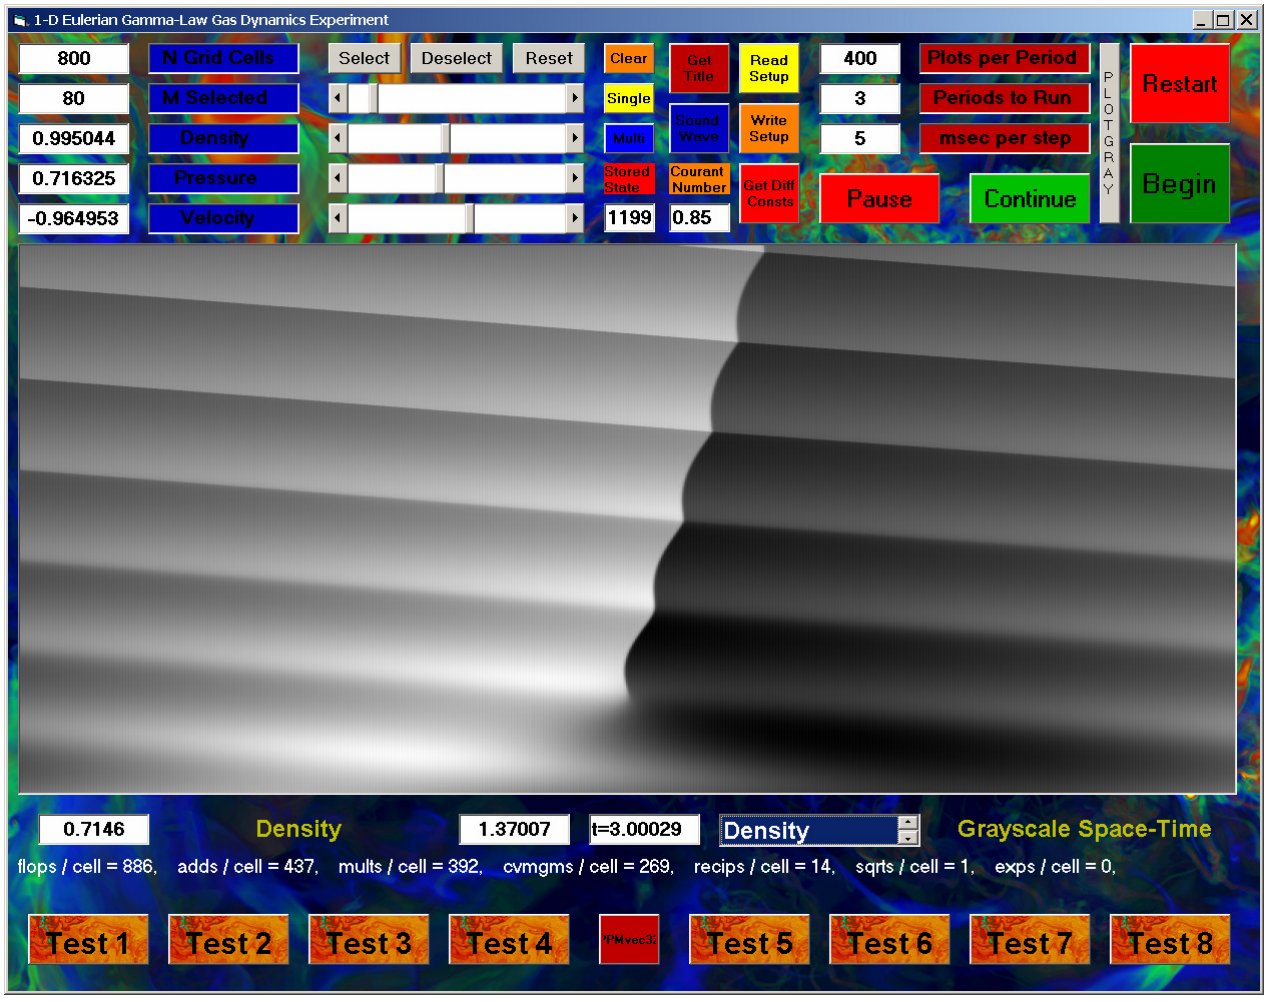

Perhaps this will clarify matters somewhat. We can see the sound wave slowly moving to the right. Its wavelength is as big as our periodic system. The brightest region, where the gas density is largest, is always on the left, and it slowly moves to the right. The crease-like feature is here revealed to be a sudden enhancement in the gas density. Remember that the gas flows

through this shock front from the right to the left (as the smoke trails clearly show). When the gas hits this shock, it is almost instantly compressed substantially. In this density image, we see streaks of high and low density going diagonally from the lower right to the upper left. Remember that time increases upward in this plot. These streaks show the presence of a wave moving to the left. You can count just about 6 cycles of this wave along any vertical line in the plot, which shows that this wave is moving to the left at about twice the speed of sound (remember that time increases by 3 units from the bottom to the top of the plot). So this is a sound wave. It is the one we set up with an amplitude of 0.15, a little under a third of the amplitude of the wave moving slowly to the right.

It turns out that we can approximately separate out the combinations of density, pressure, and velocity that constitute the signals which are carried by the three disturbances we set up at the outset of this problem. These signal quantities, that are carried along with each wave motion, are called Riemann invariants, after the French mathematician Riemann. We can plot them separately to better elucidate the wave motions in the problem. First we plot the signal moving to the right:

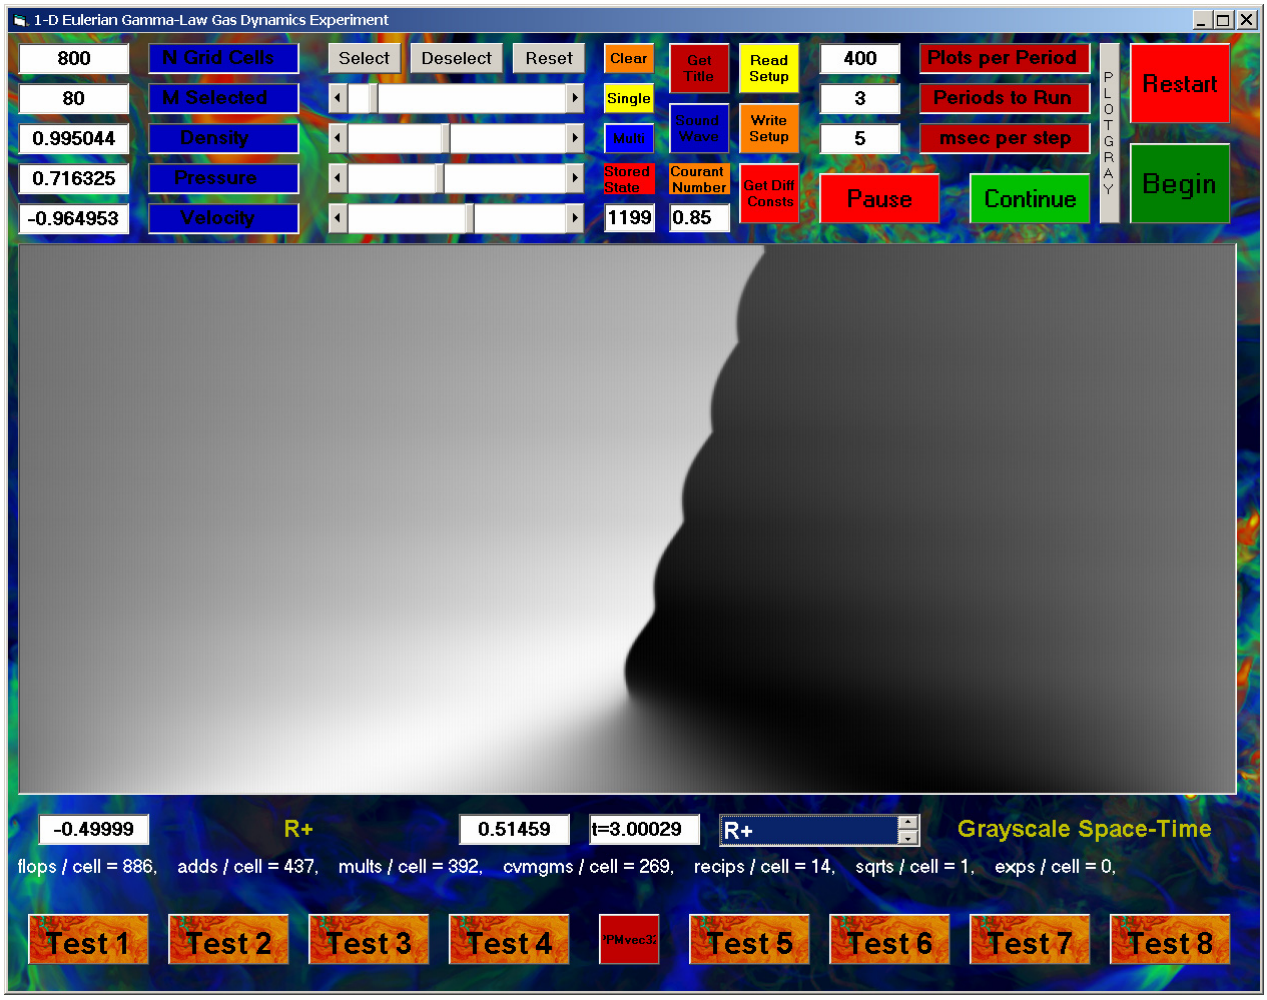

We can see that the peak of this signal is initially on the left and the trough on the right, but that it steepens, like a water wave, so that the trough and the crest of the wave are right together near the middle of the diagram. This is the shock front. It travels just a bit faster than a small-amplitude sound wave, so it moves slowly to the right as we go up in the picture.

Now let's look at the sound wave signal that moves to the left:

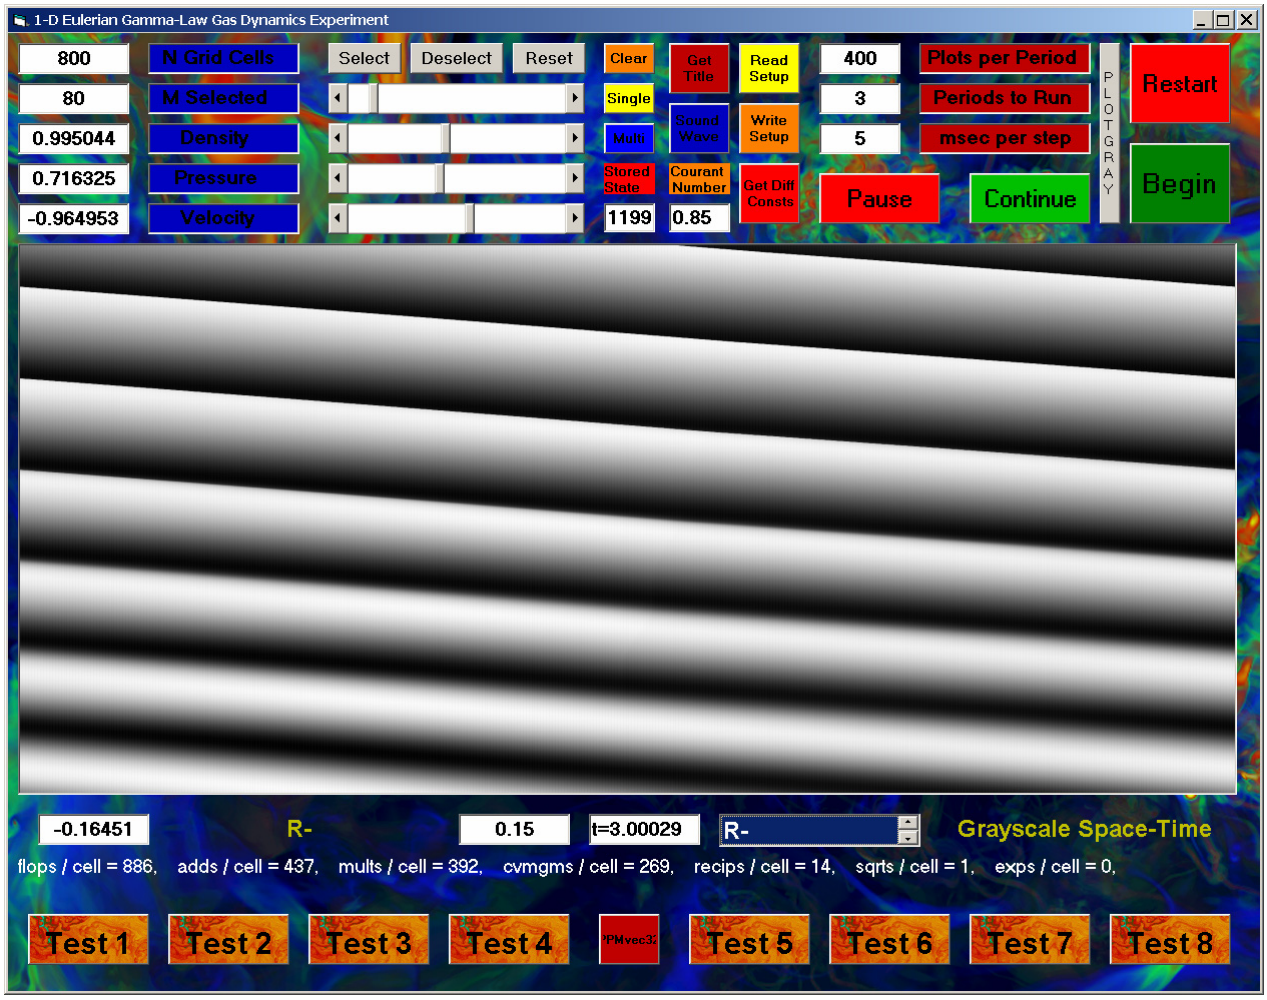

This signal clearly is moving leftward, as we have seen already. About half-way up the diagram, this signal also has steepened so that the crest of this wave is located right behind the trough. This wave too has steepened into a shock, but it has taken a bit longer, since the initial amplitude of this wave is smaller than that of the wave moving to the right.

Finally, let's look at the signal that is carried along with the gas motion. This signal is called the gas entropy. Entropy is that thing that is supposed always to increase. It's related to disorder (surely you have noticed that the entropy of your room always increases). Disorder in a gas is motion of the gas molecules every which way as contrasted to all the molecules tending on the average to move in one direction (the phenomenon that we call wind, which arises from the velocity of the gas). So let's look now at the entropy:

At the bottom of this diagram, we can see the small amplitude sinusoidal disturbance that we introduced initially. That disturbance moves to the left at a velocity of about unity (the ambient sound speed). We can be assured of this by counting the cycles we get along any vertical line in the diagram. We count just about 3 complete cycles, and the time increases by 3 units from the bottom of the diagram to the top.

Regions of higher entropy in a gas are places where the individual molecules move faster. They are thus hotter regions. If two regions of a gas have the same density, then the one with higher entropy will produce a higher pressure, because the molecules will collide at greater velocity with any wall that we might insert in the region. In this entropy space-time diagram, you can clearly see the path of the shock front that formed from the steepening of the sound wave moving slowly to the right. As the gas passes through this shock, the entropy suddenly increases. This represents

a conversion in the shock itself of kinetic energy of ordered gas motion (that is, all the molecules tending to move in the same direction) into disordered kinetic energy of individual molecules moving in all directions. This disordered energy of molecular motion is what we call heat. The diagram shows that the gas retains the same entropy as it moves along, except that its entropy increases each time it passes through a shock front. You can see this phenomenon, although it is not prominent, for the shock wave that formed when the sound wave moving to the left steepened. If you look carefully, you can see the tell-tale sharp edge angled upward and to the left that traces the propagation of this shock through the gas. This sharp edge is oriented more horizontally than are the lines of nearly constant entropy (which follow the gas motion, with the same orientation in the diagram as the smoke trails). This more horizontal orientation results from the motion of this shock to the left relative to the gas. In this entropy diagram, the leftward-moving shock is a subtle feature, visible only in the very upper portion of the diagram. You can find it, if you are having difficulty seeing it, by noticing the location of this feature in the plot of R-, the leftward-moving sound wave signal.

Acknowledgements:

The work that led to the development of the numerical methods in this Gas1D utility was supported by the U. S. Department of Energy's Office of Science. It was also supported in part through the University of Minnesota's Minnesota Supercomputer Institute's support to the Laboratory for Computational Science & Engineering (LCSE).