

A numerical laboratory

One can simulate and visualize the complex evolution of fluid flow by doing 'computer experiments,' which have emerged as a third method for investigating nature, complementing traditional experimental and theoretical work.

Karl-Heinz A. Winkler, Jay W. Chalmers, Stephen W. Hodson, Paul R. Woodward and Norman J. Zabusky

In a series of talks in 1946, John von Neumann envisioned the use of high-speed computers to generate solutions to nonlinear problems, particularly in fluid dynamics. He pointed out that scientists were conducting expensive and difficult experiments to observe physical behavior even when the underlying principles and governing equations were known. "The purpose of the experiment is not to verify a proposed theory but to replace a computation from an unquestioned theory by direct measurements," he wrote.¹ "Thus wind tunnels are used at present, at least in large part, as computing devices of the so-called analogy type to integrate the nonlinear partial differential equations of fluid dynamics."

In von Neumann's day computing equipment was not equal to the task of replacing these experiments with numerical computations, but today, four decades later, the situation has changed dramatically. The array of computing equipment now available, ranging from personal computers to supercomputer mainframes, is making

von Neumann's optimistic vision a reality. During the last three years we have been working to establish a means of simulating physical systems on large computers in a manner that allows something resembling an experimental approach. We think of the computing and graphics environment that we have set up as a "numerical laboratory."

We begin this article with a discussion of the similarities and differences between computer experiments and traditional experimental and theoretical work. Then we describe the facilities of a numerical laboratory and give an example of an experiment carried out in one. We conclude with our view of the direction in which this approach may lead us over the next few years.

Computations as experiments

Numerical computations share characteristics with both analytical theories and laboratory experiments. Like analytical theories, numerical computations are based on theoretical concepts and attempt to predict the behavior of physical systems using abstract mathematical equations. Numerical work is not unlike theoretical work in the many areas of physics where we think we know the governing laws of nature but not the behavior that these laws imply. Representing a fluid flow

field on a computer by 200 cells in which the velocity changes only linearly is not so different from representing it on a piece of paper by a truncated Bessel series expansion with only two or three terms. However, numerical work is rarely used to uncover new fundamental laws of nature, and in this respect computational physics is fundamentally different from theoretical physics.


Although numerical computations are not aimed at the discovery of new laws of nature, they can be used to discover previously unknown phenomena.² In this respect computations are like physical experiments. Given appropriate numerical tools, a researcher can explore the behavior of physical systems, as predicted by a set of governing equations, and look for interesting new effects. In spirit each new computation is very much like an experiment. The researcher simply wants to find out what the simulated system will do under some new or unusual set of circumstances. Like the laboratory experimenter, the computational physicist must perform a large number of numerical experiments to get a feel for the general behavior of a physical system. He does not obtain an algebraic expression quantifying this behavior in general; instead, he must study many cases separately.

Another feature of numerical compu-

Karl-Heinz Winkler, Jay Chalmers and Stephen Hodson are staff members at Los Alamos National Laboratory, in Los Alamos, New Mexico. Paul Woodward is a professor of astronomy and a member of the Supercomputer Institute at the University of Minnesota in Minneapolis. Norman Zabusky is a professor of mathematics at the University of Pittsburgh.

Numerical simulation of a gas jet penetrating a dense medium. A cylindrical jet of gas emanates from an orifice at the right. The colors indicate velocities: Blue is the fastest, yellow is intermediate and red is stationary. The jet enters at 10 000 times the speed of sound, and the high pressure set up in the surrounding cocoon of very hot gas generates a shock within the jet. This shock heats the jet gas and reduces its Mach number from 10 000 to about 7 or 8. A pinching oscillation of the jet is set up, and the jet gas is ultimately forced back toward the orifice due to the great resistance of the ambient medium, which is a gas initially 100 000 times denser than the entering jet gas. Many very rapidly spinning vortices form within the hot cocoon.

Figure 1



tations that is shared with laboratory experiments is the inescapable presence of experimental uncertainty. Laboratory measurements are never exact, and experimenters expend much effort keeping uncertainties as small as possible. An experimenter must know the possible range of his measurements, or his data will not be of much value. For the computational physicist too, uncertainties are a source of constant concern, and there is a continual search for techniques to reduce uncertainties at a reasonable cost in programming or computer time. Unless uncertainties are kept under control, the computational approach cannot uncover new physical phenomena, its most demanding task. Theoretical physicists, by contrast, are blessed with the possibility of generating exact solutions. However, all too often this blessing is mixed. The solutions may be generated for an extreme idealization of a real system and may have only limited applicability.

The numerical laboratory


Let us take the similarity between numerical computations and laboratory experiments seriously and find what would be required to set up a laboratory for numerical experiments. We will restrict our attention to fluid dynamics, a field in which numerical computations have played a major role due to

the nonlinearity and hence analytic intractability of the governing equations. The most common method for observing the behavior of laboratory flows is to make photographs using a variety of techniques that bring out specific features.³ Shadowgraphs are sensitive to density changes in the fluid, and one can sometimes observe streamlines by introducing smoke into the flow at specific points. Movies of such flow visualizations often allow one to appreciate the time evolution of complex structures in the flow.

Computational fluid dynamicists naturally want to use similar techniques to display their results. Displays not only make comparisons with laboratory data much easier, but also are useful in getting at the fundamental physics of wave interactions, surface instabilities, vortex generation and other phenomena that may be

involved in the flow. Displays allow the enormous amount of raw data that a numerical experiment produces within the central processing unit of the computer to be communicated to the researcher in the form that the human visual system and brain are best adapted to appreciate.

The flow in figure 1 is an example.^{4,5} The figure shows graphically the Mach numbers obtained from a computer simulation of the propagation of a Mach-10 000 jet through an ambient gas whose density is 100 000 times as great. In the hot, diffuse cocoon of gas surrounding the jet are many vortices. Near the boundary of the jet and the ambient gas these vortices cause some of the ambient gas to be entrained into the flow in the cocoon. It would be extremely difficult to get a useful impression of this very complicated flow without generating images like



Numerical laboratory scheme. This is the present configuration of the system at Los Alamos National Laboratory. Figure 2

figure 1. Even such images are not sufficient, because the flow is evolving in time. From a few snapshots like this, one can glean only the roughest impression of that time evolution. To really see the flow, one must generate a movie from such images. In addition to displays of the Mach number, displays of the density, pressure, vorticity, entropy and several other flow variables are very enlightening.⁶

Data communication requirements. The desire to generate animated flow visualizations of several variables in the manner of laboratory experiments sets up several rather difficult, but achievable, requirements for an effective numerical laboratory. These requirements are mainly a consequence of the enormous amount of data that must be processed. Numerical computations proceed by breaking up the simulated fluid flow into a large number of individual pieces, or computational zones. One assumes that each zone has a simple structure (in our computations a biparabolic structure) so that the local dynamics can be computed. To produce equally accurate representations of the flow, different numerical methods require different numbers of

computational zones, the number being smaller when the complexity of the flow allowed within an individual zone is greater. In our work we use the piecewise parabolic method, which allows relatively great complexity in the computational zones.⁷⁻⁹ Work with most other methods would therefore tend to generate more voluminous raw data if the computations were carried out to the same level of accuracy.


As an example we consider a typical two-dimensional piecewise parabolic simulation with 250 000 computational zones and ask, "What is the sustained data generation rate of the most robust version of the code executing alone on a Cray X-MP computer with four processors?" Let us assume that we make a dump for a movie frame every four time steps and that we reduce the word sizes of the five computational variables as follows: We compress the fluid density ρ , the fluid pressure P and the fractional volume f of one fluid relative to another from 64-bit Cray words to 16 bits each, and we compress the velocity components u_z and u_r to 32 bits each. These new word lengths still allow us to compute secondary flow variables with sufficient accuracy for movies. We find

for the five compressed primary flow variables alone a sustained data generation rate of 3.47 megabits/sec, where a megabit is defined as 1024×1024 bits.

While the program is running on the Cray, we compute on a second machine up to 15 additional flow variables such as entropy, vorticity and kinetic luminosity, and transform them, as well as the original five variables, into visual images of typically $512 \times 512 \times 8$ bits each. This generates an additional data stream of up to 5.2 megabits/sec, depending on the symmetry and aspect ratio of the images. Of course, we exploit obvious redundancy in the data to increase the efficiency of communication and storage. We also make additional movies by zooming into specific features of interest, which again increases the data generation rate. In summary, the resulting sustained data generation rate of a typical numerical experiment is on the order of 5–10 megabits/sec.

To put this number into proper perspective, let us compare it with the sustained data collection rate of the Very Large Array radiotelescope in New Mexico. This, the largest instrument of its kind, features a sustained data rate of 0.02 megabit/sec, which is 5–10% of its burst data rate.¹⁰ Even if the VLA were to operate continuously at its burst data rate, it still would generate less than a tenth the amount of data used in the numerical laboratory we envision here.

The data rate estimated above for a typical numerical experiment places rather severe constraints on the channel of communication from the mainframe computer on which the experiment is done to the researcher's numerical laboratory, where the data must ultimately be turned into scientific understanding. Running the simulation for 3 hours of machine time, or 12 hours of processor time, produces approximately 1400 dumps of data in the numerical experiment of our example. The storage requirement for a single such run is therefore 5–10 gigabytes. The data flow in an uninterrupted numerical experiment is analogous to an astrophysical accretion flow: What goes in one end has to come out the other or things begin to pile up and disaster strikes. Very few computer centers will allow this much data to



Shock hits bubble. These calculated images show the interaction of a Mach-2 shock in air with a spherical bubble of a gas 2.86 times denser than air. At each of eight times the flow is represented by displays of (from left to right) the logarithm of the density, the arcsinh of the divergence of the velocity, the arcsinh of the vorticity and the logarithm of the entropy. The distortion of the bubble and the instability of its surface are best seen in the density sequence; the many shock waves, in the velocity divergence sequence; the many vortices that develop, in the vorticity sequence; and the flow within the supersonic ring vortex generated behind the bubble, in the entropy sequence. Time increases from bottom to top in each of the four sequences. The left ends of the color bars at the bottom correspond to low values and the right ends to high values of the displayed variables.

Figure 3

pile up within their domains in such a short time, so disaster takes the form of outlawing the activity. Consequently, the communication channel from the mainframe computer center to the numerical laboratory must sustain 10-megabit/sec data transfers.

Data storage requirements. We saw above that a typical numerical experiment generates 5–10 gigabytes of data. If the researcher is to work with all the data from his run in an interactive way, his numerical laboratory must contain at least 10 gigabytes of fast random-access storage—that is, magnetic disks. If he wishes to archive his data for later reference, some sort of long-term storage is required as well. Ten gigabytes of data occupies approximately 60 high-density 6250-bit/inch tapes, making tape handling a major inconvenience. A better solution to the storage problem has to be found. The new 8-mm digital videotape cartridges hold 2 gigabytes of data, so just five small cartridges will hold the 10 gigabytes produced by a numerical experiment. These cartridges increase storage density by a factor of 150 over 6250-bit/inch tapes, indeed a dramatic improvement. The numerical experiment given above as an example fills up about one such cartridge per hour of machine time, while the VLA could store the data from approximately nine days of observation time on one cartridge.

The speed with which one can transfer data to and from archival storage is also an important consideration. One would like to be able to shift attention quickly from one simulation to another by dumping one run to archival storage and bringing another out of storage. At present transfer rates, this process can take the better part of a day and is a real impediment to research. Transfer rates in the range of a few megabytes per second are very desirable and seem achievable with current technology. They would allow the exchange of two 10-gigabyte data sets within a couple of hours.

Data visualization requirements. The key remaining question is how to interact with such an enormous database and how to extract the relevant physical information. Looking at the printed numbers is out of the question. Our visual systems are designed to process large amounts of data most efficiently

through pattern recognition; in fact, many individuals receive insight and retain concepts through seeing. We also have the ability to abstract and generalize, enabling us to compress information significantly by separating what is important from what is unimportant. In the past, the communication of scientific results was based largely on the verbal and analytical abilities of the left side of the brain. In the future, with the enormous amount of information thrown at us by numerical experiments, we will have to rely to a greater extent on the nonverbal and synthetic abilities of the right side of the brain.

In our numerical laboratory we expect to interact with the computed data principally through visual images. We can place an upper limit on the number and resolution of the images that we will require by recognizing that human visual perception is an intelligent system with limited bandwidth.¹¹ The upper limit for the data rate of the human visual system is estimated to be a few gigabits per second. This estimate is based on limitations of the field of view, angular resolution, color resolution and viewing frequency of the human brain-eye system.¹² We can think of this limit as 8–12 frames per second of 3600×3600-pixel images, where each pixel contains 18 bits of color information. At 8–12 Hz humans begin to lose the ability to distinguish individual frames and get a movie-like impression, so we prefer higher display rates, and we must adjust the required data rates for display devices accordingly. Because of the limitations of mainframe computer speed, researchers rarely use 3600×3600 computational grids. Images of 2048×2048 pixels would be sufficient to display all the information contained in a snapshot for a single flow variable in an affordable numerical experiment, but present video technology limits us to images of about 1280×1024 pixels for movies.


Scientists disagree as to the kind of images that give the best description of computed data. We feel that at present two-dimensional raster displays with 8 bits of color information are a wise compromise. Such images are excellent at depicting numerical simulations in two space dimensions, as the

figures in this article attest. The ability to display 24 bits of color would allow more complex displays that would show more than one flow variable at a time. (For an example, see the figures in reference 5.) However, animating 1280×1024×24-bit images at movie rates is currently very expensive, although feasible. The images we use are constructed from two-dimensional arrays of values of a variable on the computational grid; these values are mapped directly onto the pixels of a video screen. The images are almost trivial to construct once the numerical simulation code is written. Two-dimensional arrays are the natural form in which the flow variable data reside in the memory of the mainframe computer. We have found that the computer time for constructing the images and performing the input-output operations to get them to the video display is less than 5% of the computational time for the simulation itself. This is true in part because constructing the images is easy in the parallel and vector modes preferred by today's fastest computers. Another advantage of the two-dimensional raster display is that the researcher has no need for complicated graphics packages. The image construction work involves little computation but quite a lot of input and output, and can be done interactively on a relatively inexpensive machine.

Through our numerical laboratory, we hope to exploit fully the human visual system's unique abilities by achieving an "impedance match" between its input capacity and the display rate of a graphical output device. We believe the result will be an unprecedented, revolutionary gain in productivity. We should try to satisfy human needs rather than adapt humans to a particular piece of hardware. The goal is to create a computational environment in which the researcher is limited not by the idiosyncrasies of his surroundings, but only by his own imagination. Extraordinary convenience and simplicity in performing computational tasks will not necessarily encourage laziness but are prerequisites for creativity and productivity.

Our present numerical laboratory

Figure 2 shows the present configuration of our numerical laboratory,



Supersonic ring vortex like the one generated in the flow in figure 3 but displayed in a close-up view. This vortex was generated by a Mach-3 shock, and it is stronger and more persistent than the one in figure 3. The same density, divergence of velocity, vorticity and entropy functions as in figure 3 are displayed counterclockwise beginning at the lower left. In the vorticity display, slip surfaces with opposite signs of vorticity roll up to form small vortices that tend to come together in pairs. The supersonic spin velocity of the large vortex is made obvious by the presence of shocks within it oriented like spokes on a wheel. The supersonic propagation of this ring vortex through the ambient air is made clear by the bow shock that the vortex drives in front of it. The core of the ring vortex represents a density minimum because the strong centrifugal forces from its supersonic spinning tend to evacuate this region.

Figure 4

built up over the last year at Los Alamos National Laboratory. The Supercomputer Institute of the University of Minnesota has meanwhile established a similar system using to a large part identical hardware and software. The present configuration is intended as a "proof of principle" of the validity of the basic ideas underlying a fully functioning numerical laboratory. It has been assembled with off-the-shelf equipment. We took a conservative approach in selecting the equipment

but allowed for future growth and enhancement. With this approach, we have already obtained an immediate improvement in performance of three to four orders of magnitude over the data rate of 9.6 kilobits/sec that is widely used between supercomputers and devices that display raster color images. We can display in a smooth way 800 digital raster images of $1024 \times 1024 \times 8$ bits at up to 8 frames per second, or 3200 images of $512 \times 512 \times 8$ bits at up to 30 frames per

second. This is a sustained data rate of 60 megabits/sec.

We achieved this data rate by separating the data generation, which we do on centrally located supercomputers, from the truly interactive visualization process, which we do on a locally available and controllable image support processor. This separation is very important because it gives the user complete control over the interactive part of the work, detached from the normally heavy workload of the supercomputers. It also enables the user to view the same movies again and again without tying up precious supercomputer resources.

Our image support processor consists of a Gould Power Node 9080 dual processor, a Gould IP8500 image array processor, a set of 4 real-time disks and 16 normal system disks, providing a total of about 12 gigabytes of storage capacity. This system permits us to store up to 100 movies of varying lengths. A 300-ft cable in full duplex mode connects the image support processor to a Cray X-MP worker machine through two sets of high-speed port interfaces operating at a sustained data rate of 27 megabits/sec. This data rate is high enough to allow continuous monitoring of the numerical fluid dynamics experiment that we will describe below. The image support processor is used primarily to receive large sets of compressed solution numbers from several Cray machines, turn them into individual images at a rate of up to 100 000 images per day and display the images as movies on digital display devices. Storing the images in digital form, completely decoupled from a particular color representation, allows tremendous flexibility in digging out the flow structures that are hidden in the numbers. Also, the displays produced by digital display devices typically exceed by a factor of 15-40 the video quality standard given in the National Television Standard Code.

We develop most of our computer programs using a variety of scientific workstations, which are connected to the image support processor through an ETHERNET link. The image support processor, with its substantial disk capacity, also acts as a file server for these workstations. UNIX is our oper-

ating system of choice on all machines. The network is complemented with the PRONET-80 token ring network, which operates at a burst data rate of 80 megabits/sec. We hope the network will eventually allow us to display a $1024 \times 1024 \times 8$ -bit image in 1 second at the workstations, compared with the approximately 10 seconds required via the ETHERNET network. The limiting factor is the interface connecting the workstations to the PRONET-80.

Important to the functioning of the whole system has been the acquisition of four Exabyte 5.25-inch tape drives that use the new 8-mm digital storage technology. These tape drives have been integrated into the scientific workstations. Unfortunately, the sustained data rates of the new drives are currently limited to the data rate of normal 6250-bit/inch tape: about 2 megabits/sec. Therefore, an important addition needed to increase user productivity further will be the acquisition of a high-capacity, high-speed data storage and retrieval system matching the existing disk controller speeds. This addition would remove the remaining bottleneck in the system by matching on-line and off-line storage, allowing us to deal with almost limitless databases and bringing the entire facility to bear on each computational task. Currently, digital laser disks do not fulfill our storage needs because of their low data rates and high storage costs. In this context, staying in touch with the rapidly developing technology around us is essential. The new, small Video Home System technology, for example, may offer even higher storage density than the 8-mm digital storage technology.

One of the major pieces of hardware, specially developed for our project, is the 42-megabit/sec sustained link between the Gould Power Node 9080 and the Gould IP8500 image array processor. Together with the enhancements we have made to the Unix operating system, this link allows us to access images through four system disk controllers in parallel at a rate of twenty-one $512 \times 512 \times 8$ frames per second and to deposit them on the real-time disks. From there, we can show them repeatedly as movies without interfering with the work-

load on the Gould Power Node 9080 itself. Only this enhancement gives us the ability to interact with our entire database on a human time scale.

A numerical lab experiment

One can best appreciate the arguments we have made for the usefulness of a numerical laboratory by considering a specific example of its use. We will describe recent work on the generation of supersonic vortices in shock-bubble interactions, inspired by a laboratory experiment. Jean-Francois Luc Haas and Bradford Sturtevant at Caltech have observed the interaction of shock waves in air with cylindrical and spherical bubbles of various gases of various densities.¹³ These idealized experiments were conceived to study the more general process of vortex generation when a shock passes through an inhomogeneous medium. The shadowgraph pictures from the experiments reveal a wealth of interesting shock interactions. The surface of the bubble is less distinct than the shocks in these pictures because it is subject to instabilities that cause the bubble gas to mix somewhat with the surrounding air in a layer near the bubble surface. The overall distortion of the bubble is easy to see in the shadowgraphs, and the generation of vortex pairs for the cylindrical bubbles, or ring vortices for the spherical bubbles, shows up clearly.

The experiments of Haas and Sturtevant were first simulated numerically by J. Michael Picone and Jay P. Boris at the Naval Research Laboratory.¹⁴ They used the FAST2D computer code^{15,16} to solve Euler's equations governing inviscid fluid flow in two spatial dimensions. Their results capture the large-scale features of the experimental flows rather well, but the computations lack sufficient resolution to permit studying the flows in detail. References 17 and 18 discuss similar flows in a rather different context. With the high resolving power of the piecewise parabolic code, the great speed of the Cray X-MP/416 and the enormous capacity for data analysis and computed flow visualization provided by our high-speed graphics system, we were able to study these shock-bubble interactions in more de-

tail than even the laboratory at Caltech permits.

The Euler equations describe a compressible gas and can be written in the conservative integral form

$$\begin{aligned} \frac{\partial}{\partial t} \int_V \rho \, dV + \int_{\partial V} \rho (\mathbf{u} \cdot d\mathbf{S}) &= 0 \\ \frac{\partial}{\partial t} \int_V \rho \mathbf{u} \, dV + \int_{\partial V} \rho \mathbf{u} (\mathbf{u} \cdot d\mathbf{S}) \\ &+ \int_{\partial V} P \, d\mathbf{S} = 0 \\ \frac{\partial}{\partial t} \int_V E \, dV + \int_{\partial V} E (\mathbf{u} \cdot d\mathbf{S}) \\ &+ \int_{\partial V} P (\mathbf{u} \cdot d\mathbf{S}) = 0 \end{aligned}$$

Here ρ represents the density of the gas, E the sum of kinetic and internal energy, P the gas pressure, $\mathbf{u}$ the velocity and ∂V the boundary of the volume V . The equations allow for the description of discontinuities—shock fronts, contact discontinuities and slip lines—in the flow variables.

We solve these equations numerically with the piecewise parabolic method, in which we discretize the equations in a form that conserves mass, total energy and linear momentum; calculate zone averages of all the variables; interpolate with parabolas, and enforce monotonicity to guarantee numerical stability.⁷⁻⁹ In every zone we update the conservation laws explicitly, computing fluxes from an approximate solution of the Riemann problem at zone interfaces. The piecewise parabolic method is accurate to second order in the smooth part of the flow and allows for an accurate description of all flow discontinuities without spreading them over more than one or two zones, as von Neumann-Richtmyer-type artificial viscosity methods would. The method is complex and therefore computationally intensive, requiring about 3000–4000 floating point operations to update the flow variables in a single zone. The code for the piecewise parabolic method is highly vectorized and can update approximately 130 000 zones per second on a Cray X-MP/416 employing all four CPUs.

Figure 3 shows the results of one of our numerical experiments. We display four variables and use logarithmic and arcsinh scalings to bring out subtle low-level features in the flow that would otherwise be missed. The variables are shown at eight different times carefully selected from our movie representation of the simulation. The last frame shown (at the top of each sequence) is image number 290 in a movie

sequence of over a thousand images; this illustrates our ability to focus on any time interval in which an interesting phenomenon may appear.

The experiment shown in figure 3 goes beyond the present capabilities of the Caltech experimental laboratory and studies the particularly interesting behavior at higher shock Mach numbers. In our experiment the incident shock has a Mach number of 2, while the Caltech work is now limited to Mach numbers under 1.3. It is in this regime of stronger incident shocks that we observe the formation behind the bubble of a strong, growing, supersonic ring vortex. This vortex stands out dramatically in the displays of vorticity and entropy.

We have matched the initial density difference in our experiment to that in Haas and Sturtevant's work at Caltech with a fluorocarbon that is 2.86 times denser than the surrounding air. The fluorocarbon's greater density causes bubbles of it to elongate under gravity in the physical laboratory. In our numerical laboratory we can of course set up any desired initial bubble geometry, and we have set the parameters of our numerical experiment to resemble those in the Caltech experiments.

At Caltech the lengths in the experimental setup are measured in centimeters. That these lengths are in centimeters rather than meters or light years matters only if viscous effects are important. In the Caltech experiment viscous effects may cause boundary layers to form along the shock tube walls, affect the thickness of the shock or affect the manner in which surface instabilities mix bubble gas with the surrounding air. In our numerical experiment viscous effects enter through the use of a computational grid and through the stabilizing numerical error terms that are related to that grid and that make the computation feasible. We have used a mesh of 1200×400 square computational cells, or zones, to describe our shock tube; we obtain the results for half of these zones by symmetry. The initial bubble radius is resolved with 112 zones. Our computational grid allows for very thin shock waves and slip lines in the gas, as is evident in figure 3.


One of the most striking features of

Shock hits vortex. This sequence of eight snapshots shows the computed interaction of a cylindrical shock wave and the supersonic ring vortex generated in the flow in figure 3. This interaction, which evolves from bottom to top in the figure, takes place at a later time than is shown in figure 3. The ring vortex, seen in cross section, stands out in this display of the arcsinh of the velocity divergence because of the strong rarefaction (blue) followed immediately by a shock (red) that joins the two vortex cores. This structure is hardly modified at all by the shock interaction shown here. The supersonic vortex is indeed a rather robust structure, at least when axial symmetry is enforced, as it is in this simulation.

Figure 5

the flow shown in figure 3 is the growth of ripples along the front surface of the bubble. This surface is initially a smooth transition from bubble to ambient gas over a distance of three zones. The incident shock immediately compresses this surface layer to a thickness of about one zone. As the shock runs along the bubble surface it introduces small perturbations because the calculation is performed on a grid. These perturbations would not grow if the bubble surface were physically stable. However, at the front surface we have a Richtmyer-Meshkov instability, while more to the side a Kelvin-Helmholtz instability develops. The Kelvin-Helmholtz instability is the familiar instability of a slip surface; it leads the wind to produce water waves and the waving of flags. The less familiar Richtmyer-Meshkov instability occurs when a fluid is accelerated by a shock from a more diffuse fluid. At the boundary between these fluids the shock acceleration produces a transient effective gravity pointing from the denser to the lighter fluid. The instability of the surface under this acceleration is then analogous to the Rayleigh-Taylor instability, which causes a denser fluid to fall when it is superposed on a lighter fluid in a gravitational field. We should note that the back surface of the bubble in figure 3 is accelerated in such a way that the effective gravity points from the lighter to the denser fluid. This part of the bubble surface is physically more stable, and is beautifully smooth, as it should be.

Discovering new phenomena. The computation in figure 3 is a good example of the use of a numerical laboratory in another respect. It shows that numerical experiments, like real physical experiments, can lead to the discovery of unexpected phenomena. In this case the generation of the strong supersonic vortex ring that travels along the axis of symmetry just behind the incident



shock wave was quite a surprise. The phenomenon is even more dramatic in figure 4. Here a Mach-3 shock impinging upon a bubble 2.86 times denser than air generates a very strong and persistent ring vortex. The figure shows the structure of this vortex in detail.

The generation of the strong supersonic ring vortex is apparently related to the special configuration of shocks at the back of the bubble, visible in the third and fourth frames from the bottom of each of the four sequences in figure 3. Here the shocks racing around the bubble collide at an oblique angle, and a very complicated structure involving a number of shocks and slip surfaces develops. By the sixth frame this structure is well resolved. A core of hot, diffuse (high entropy) gas spins supersonically in a ring vortex. The velocity divergence plot brings out a shock wave that passes through the entire vortical structure. In the center of the vortex the gas accelerates strongly along the axis of symmetry until it is suddenly compressed in a shock disk. This shock has been set up by the supersonic collision of the gas in the vortex with the ambient air. The collision drives a shock into the air as well, and this shock protrudes ahead of the original shock that set all the flow in motion. The curved shock in the air ahead of the vortex joins the original shock front in a short shock segment that has kinks where it joins the vortex bow shock and also where it joins the original shock. Each kink represents a three-shock intersection from which a slip line must emerge. Such slip lines have opposite signs of vorticity, and in figure 4 they can be seen rolling up to form small vortices that tend to gather into counter-rotating pairs.

The structure of the supersonic vortex rings in figures 3 and 4 is reminiscent of similar phenomena observed in Cartesian geometry when strong shocks hit reflecting surfaces at oblique angles. In some cases such shocks experience complicated reflections that are termed complex, Mach or double Mach reflections. Examples of double Mach reflection for shocks in air can be found in references 7 and 16. In these examples a jet along the plane of symmetry, a reflecting wall, is generated along with an associated vortex. In subsequent work at Livermore¹⁹ treating a number of cases

using different equations of state, jets are obtained that are strong enough to drive the shock forward at the reflecting wall, producing a curved bow shock like those in figures 3 and 4.

Figure 5 shows the interaction of a shock wave and the supersonic ring vortex of figure 3. Here we display the divergence of velocity to focus attention on the shock fronts. In the first frame, at the bottom of the figure, we see the vortex mainly through the strong rarefaction (blue) at its center, followed immediately by a shock compression (red). In the final frame of the figure the vortex presents almost exactly this same appearance. Considering all the shock interactions that have occurred during the intervening frames this is rather amazing. Clearly the supersonic vortex ring is quite a robust structure, at least so long as we maintain strict axial symmetry, as our two-dimensional calculation requires. A thorough investigation of the stability of such a vortex ring would require three-dimensional simulations of the response of the vortex ring to nonaxisymmetric perturbations.

Future numerical laboratories

The example presented above demonstrates our recent progress in numerical experimentation. However, as we have already pointed out, many issues remain unresolved with our present numerical laboratory. We still have to increase data communication rates by at least another order of magnitude for the displays to approach the physiological limit of human visualization. Three-dimensional computations are necessary to overcome the simplifying assumptions of two-dimensional symmetry. Only in three spatial dimensions do we have a chance to capture all the essential physics of fluid flow. The transition to turbulence, for example, can only be described accurately in three spatial dimensions because two-dimensional simulations have the tendency to propagate energy from smaller modes to larger ones, contrary to laboratory experiments. However, requirements for the simulation and visualization of three-dimensional time-dependent continuum physics problems are considerably greater than those for two-dimensional problems.

As part of our project at Los Alamos,


we have looked into how to prepare for the significantly faster machines that we can expect in the near future from the commercial sector. Harnessing the computational power of machines that are one or two orders of magnitude faster will be impossible with existing computer networks. These networks are barely able to cope with the workload of our current numerical experiments. Consequently, we are putting tremendous emphasis on upgrading the existing network. A futuristic view of a possible attempt is shown in figure 6. This network would connect many supercomputers with thousands of users through high-speed fiberoptic lines. These lines would also provide the data stream from supercomputers to our numerical laboratory. The basic ideas for building this high-speed equipment are discussed in reference 10.

Visualization of faster data streams obviously requires more computational power. This power could be provided, for example, by the massively parallel Connection Machine from Thinking Machines Corporation, which also features a gigabit/sec frame buffer for visualization purposes. As this machine is based on a single-instruction, multiple-data architecture, it also would fit nicely with our computational algorithms for fluid flows. Using the Connection Machine optimally requires a fast host computer. Coupling it to a Cray X-MP would combine the advantages of a fast scalar and pipeline machine with those of a massively parallel machine. This would further integrate the simulation and visualization of fluid dynamics phenomena.

We believe that recent technological advances make it possible for computational physicists to work in a manner more closely related to laboratory experiments. A researcher can now see the many results, computed in the mainframe CPU, that in the past had to be ignored. These results can now be animated as color movies to achieve a match between man and machine. The equipment for exploiting this new mode of working forms a numerical laboratory, and such a laboratory opens up exciting new avenues for discovery in computational science.

* * *

Getting a project like ours off the ground in a matter of months required the help of many dedicated people. It is therefore a



Future numerical laboratory. Most of today's numerical laboratory, represented in figure 2, would fit into the rectangle labeled "Image support processor." Figure 6

pleasure to acknowledge the help and support above and beyond the call of duty of the following people: Robert W. Selden initiated the funding; E. Staley Hadden, Lawrence B. Warner, Suzanne W. Peterson and Carla R. Walsh provided operational support; Eldon J. Linnebur provided staff support; William Spack took care of matters involving the Department of Energy; Norman R. Morse provided financial and technical support from the computing and communications division at Los Alamos; and Hoyt E. Hart made an extraordinary effort in providing customer support. In particular, we thank Donald E. Tolmie and Michael McGowen for developing and implementing high-speed communications hardware. We thankfully acknowledge the help of Michael Gehmeyr in preparing the manuscript. This work was supported by the Department of Energy through the Institutional Research and Development Program at Los Alamos National Laboratory as well as through a grant to the University of Minnesota from the Office of Energy Research under contract DE-FG02-87ER25035. Our work at Minnesota was also supported by grants from the University of Minnesota, by grant AST-8611404 from the National Science Foundation, by equipment grant AFOSR-86-0239 from the Air Force Office of Scientific Research and by generous hardware donations from Sun Microsystems and Gould Inc. Our work at Pittsburgh was also supported by the Army Research Office, by the Office of Naval Research and by a hardware donation from Sun Microsystems.

References

1. H. H. Goldstine, J. von Neumann, in

- John von Neumann, *Collected Works*, vol. V, Pergamon, New York (1963), p. 1.
2. For an example, see N. J. Zabusky, *PHYSICS TODAY*, July 1984, p. 36.
3. For an impressive compendium of such photographs, see M. Van Dyke, *An Album of Fluid Motion*, Parabolic, Stanford, Calif. (1982).
4. K.-H. A. Winkler, S. W. Hodson, J. W. Chalmers, M. McGowen, D. E. Tolmie, P. R. Woodward, N. J. Zabusky, *Cray Channels*, Summer 1987, p. 4.
5. P. R. Woodward, D. H. Porter, M. Ondrechen, J. Pedelty, K.-H. A. Winkler, J. W. Chalmers, S. W. Hodson, N. Zabusky, in *Proc. Third Int. Symp. on Science and Engineering on Cray Supercomputers*, W. Porter, ed., Cray Research Inc, Minneapolis, Minn., in press.
6. For a detailed discussion of the usefulness of these auxiliary displays, see K.-H. A. Winkler, M. L. Norman, in *Astrophysical Radiation Hydrodynamics*, K.-H. A. Winkler, M. L. Norman, eds., Reidel, Dordrecht, The Netherlands (1986), p. 223.
7. P. R. Woodward, P. Colella, *J. Comput. Phys.* **54**, 115 (1984).
8. P. Colella, P. R. Woodward, *J. Comput. Phys.* **54**, 174 (1984).
9. P. R. Woodward, in *Astrophysical Radiation Hydrodynamics*, K.-H. A. Winkler, M. L. Norman, eds., Reidel, Dordrecht, The Netherlands (1986), p. 245.
10. *Cray Channels*, Summer 1987, p. 37.
11. P. Thompson, in *Fundamentals of Human-Computer Interaction*, A. Monk, ed., Academic, New York (1984), p. 5.
12. K.-H. A. Winkler, M. L. Norman, J. L. Norton, in *Supercomputers: Algorithms, Architectures, and Scientific Computation*, F. A. Matsen, T. Tajima, eds., U. of Texas P., Austin (1986), p. 415.
13. J.-F. Haas, B. Sturtevant, *Interaction of Weak Shock Waves with Cylindrical and Spherical Gas Inhomogeneities*, Caltech Graduate Aeronautical Laboratories preprint (March 1987); *J. Fluid Mech.*, in press.
14. J. M. Picone, J. P. Boris, *Vorticity Generation by Shock Propagation Through Bubbles in a Gas*, Naval Research Laboratory preprint (April 1987); *J. Fluid Mech.*, in press.
15. J. M. Picone, J. P. Boris, *Phys. Fluids* **26**, 365 (1983).
16. D. L. Book, J. P. Boris, A. L. Kuhl, E. S. Oran, J. M. Picone, S. T. Zalesak, in *Proc. Seventh Int. Conf. on Numerical Methods in Fluid Dynamics (Lecture Notes in Physics, vol. 141)*, W. C. Reynolds, R. W. MacCormack, eds., Springer-Verlag, Berlin (1981), p. 84.
17. P. R. Woodward, *Astrophys. J.* **207**, 484 (1976).
18. P. R. Woodward, in *Early Solar System Processes, Proc. Int. School of Phys., "Enrico Fermi" course no. 73*, D. Lal, ed., Italian Physics Society, Bologna (1980), p. 1.
19. P. Colella, H. M. Glaz, in *Proc. Ninth Int. Conf. on Numerical Methods in Fluid Dynamics (Lecture Notes in Physics, vol. 218)*, Soubbaramayer, J. P. Boujot, eds., Springer-Verlag, Berlin (1985), p. 154. □