THE USE OF REALTIME ANIMATION GRAPHICS IN THE ANALYSIS OF METEOROLOGICAL MODEL DATA

Kelvin K. Droegemeier
School of Meteorology, University of Oklahoma
Norman, Oklahoma 73019 USA

1. INTRODUCTION

From the first numerical integration of the barotropic vorticity equation (Charney *et al.*, 1950) to the sophisticated ECMWF operational model, meteorologists have pioneered the application of digital computers to the solution of nonlinear, time-dependent and multi-dimensional fluid flows, and have thus been preeminent in continually stressing the need for more powerful machines. Due largely to technological breakthroughs, raw computing power has increased by some five orders of magnitude since the early 1950's, and continues to grow at a phenomenal rate. In only two to three more years, it is estimated that supercomputers will have high-speed central memories in excess of a billion words, and be capable of performing some 10 to 50 *trillion* floating point computations every second. Even prior to the availability of such machines, however, the modeling community at large is facing a quite severe problem: how does one deal with the enormous amounts of data produced by numerical models run on *today's* supercomputers? In the discussion to follow, this question is addressed through examination of techniques for visualizing complex flow calculations. In addition, the adaptation of such techniques to an operational environment is discussed, and suggestions for future work are offered.

2. THE NATURE OF THE PROBLEM

The enormous volume of data from today's numerical models can be illustrated through a simple example in which a convective storm system is simulated in three dimensions. Assume that the model has 10 dependent variables at each gridpoint, and that the computational domain is 500 km on a side and 20 km tall. In order to explicitly resolve cloud-scale processes, a 1 km grid mesh is used in the horizontal with 500 m resolution in the vertical. In addition, assume that the model is core-contained in a large-memory supercomputer, and that a two timelevel numerical scheme is employed. If the experiment were run with a 10 s timestep for a period of 6 hours, and the data saved every 60 s (the temporal resolution needed for an adequate animation sequence), this would result in the need to store approximately 18 trillion numbers, or 144 gigabytes of data using a compression of four words to one. Assuming this data could be printed in character form on standard fanfold computer paper, the resulting output would weigh some 15 tons and be nearly 250 meters tall - all for a *single experiment*! While this may seem to be an unrealistic example, it is exactly the type of computation which can be done, albeit consuming a tremendous amount of CPU and wall clock time, on *today's* supercomputers (e.g., the 512 megaword Cray-2 or the ETA-10).

In nearly every area of meteorological modeling, the flows in question are quite complicated and, most importantly, *highly time-dependent*. For example, even though operational prediction models are run with timesteps on the order of 10 minutes, forecast data are output every 6 to 12 hours. Consequently, forecasters are required to mentally interpolate the intervening sequence of events which are known, but are not made available. The human brain-eye system, which is remarkably adept at integrating enormous amounts of data (with an estimated transfer rate of 1 to 3 gigabaud; Winkler *et al.* 1987), is therefore drastically underutilized when less than 5 percent of the available information is displayed using static and archaic contour plots. In order to overcome this problem, we must begin to effectively utilize the three primary elements of graphical visualization: *color*, *motion*, and *dimensionality*.

The use of color in graphical displays has now been accepted as genuinely valuable by the scientific community. The latter two elements have received increasing attention during the past few years, largely because of the availability of sophisticated computer workstations capable of animating sequences of data and performing realtime rendering of solid objects. By using such devices to animate 2-D color images of data at each timestep of an operational model run, a much larger amount of information could be perused by the forecaster, and in a much shorter period of time, than with conventional contour displays. Most importantly, the forecaster would see a complete *time history* of the flow, and would thus have a better opportunity to understand *how and why* the various weather features arise (see section 4). Unfortunately, this type of data display requires substantial amounts of disk storage and/or very high speed data channels, neither of which are currently available in an operational forecast environment.

It is important to note that the animation of complex flow calculations is typically performed using 2-D images; the use of 3-D images, either in volume visualization or solid-surface rendering form, appears to be a formidable task for a number of reasons. First, although the representation of relatively smooth 3-D objects, such as the arm of a robotics device, poses little problem for today's rendering algorithms, the meteorological counterpart may consist of complex and highly time-evolving 3-D fields of temperature, velocity, or humidity. Consequently, in contrast to the robot example, where motion is created by simply translating and rotating the existing solid objects, animating the meteorological data would require calculating the position and structure of the entire 3-D data surface at each time step in the sequence. This process not only requires an enormous amount of storage, but also supercomputing power to perform the rendering calculations. In addition, as discussed in section 4, the creation of 3-D animation displays in a manner appropriate for human interpretation appears to be much more difficult than with the 2-D counterpart. Exceptions to this include, for example, the 3-D visualization of flow over an aircraft or through an evolving thunderstorm using colored tracer particles.

In the operational community, the ECMWF has taken the lead in applying color, and now animation, to the display of forecast model data. It is important, however, that the fidelity of graphical output keep pace with continued increases in model resolution. Practical considerations such as limitations in communications bandwidth often prohibit the transmission of more than 50% of the actual model gridpoint data for plotting. However, we must recognize that the impact of a three-fold increase in model resolution on the final product (i.e., the forecaster's worded message) may be substantially reduced if the model data continue to be displayed using standard contour drawings at coarse intervals (e.g., 60 m for the 500 mb height field). Alternatives such as the transmission of raster images appear to be rather expensive at this point in time, and remote users will likely be unable to secure the expensive equipment needed to create and display images from raw data on-site. Nevertheless, it is generally believed that enhancements in computer workstation and communications technology will allow affordable, high bandwidth transmission and image processing capabilities in an operational environment within the next 2 to 4 years.

3. A PROTOTYPE ADVANCED METEOROLOGICAL IMAGING SYSTEM

During the past few years, several major graphical display facilities have been created to aid the interpretation of complex numerical model calculations. These include the Ultra-Speed Graphics Project at Los Alamos National Laboratory (Winkler *et al.*, 1987), the Advanced Visualization Facility at the National Center for Supercomputing Applications, University of Illinois (*Cray Channels*, 1987), and major facilities at the Minnesota Supercomputer Institute and NASA Ames Research Center. Each of these sites operates a supercomputer in close proximity to the graphical display system; unfortunately, such an arrangement is the exception rather than the rule, particularly in the operational community where distributed forecast offices obtain data from a central facility which is often more than 2000 km distant.

The College of Geosciences at the University of Oklahoma operates a graphical display facility functionally similar in hardware, though scaled down in performance and without a supercomputer on-site, to those at the Los Alamos National Laboratory and the Minnesota Supercomputer Institute. As shown in Figure 1, the Geosciences Computing Network (GCN) is built around a Digital Equipment Corporation VAX 11/785 computer with 32 megabytes of central memory and 3.6 gigabytes of on-line disk storage, along with several MicroVAX's and color workstations. The devices used to create the animation sequences shown at the Workshop are Gould IP8500 color workstations and associated realtime digital disks (the IRIS and Sun workstations were only recently acquired, but will also be used for advanced image processing). Access to remote supercomputers, primarily at the National Center for Atmospheric Research (NCAR), the University of Illinois, and the University of Minnesota, is provided through a local Ethernet backbone connected to NSFnet via a 56 kilobit leased phone line (MIDnet).

Most of the animation sequences shown at the Workshop were from calculations performed at NCAR using two and three-dimensional cloud models (Droegemeier and Wilhelmson, 1985a,b, 1987; Droegemeier and Davies-Jones, 1987; Droegemeier, 1988; and Carpenter *et al.*, 1988). During each model run, a history dataset is created whereby model results are saved at pre-determined intervals (e.g., 10 to 60 seconds, depending upon the phenomenon being modeled). Using an algorithm developed by Dr. Joseph Klemp at NCAR, the raw data are then packed (4 64-bit words to 1) and either written to tape or transmitted over NSFnet to the GCN. The data are then placed on the VAX as disk files mapped directly to memory, and an interactive software package (SIGMA), also developed by Dr. Klemp and further enhanced at OU, is used to convert each 2-D slice of model data (e.g., a vertical section of the temperature field through the center of a developing cloud) into a raster image of arbitrary size with a gray scale value of 0 to 255 assigned to each pixel. Depending upon the resolution of the data, anywhere from 1 to 3 minutes of VAX CPU time may be required to produce a single quarter-megabyte, 512 by 512 pixel raster image. Note that, in contrast to the centers mentioned above, where raster images are typically created on a resident supercomputer, we perform such processing locally, as might be the case in an operational forecast center, using raw model data.

Once raster processing is completed, the sequence of images is loaded onto the Gould realtime disk, which in its present configuration holds approximately 800 images. A false color map is then created by the user from a palate of 16.7 million colors (256 of which are displayable), typically so as to accentuate certain features in the flow field. The images can then be viewed sequentially at rates ranging from 4 to 30 frames per second, and hardcopy video output created in realtime using an NTSC converter and standard video recorder.

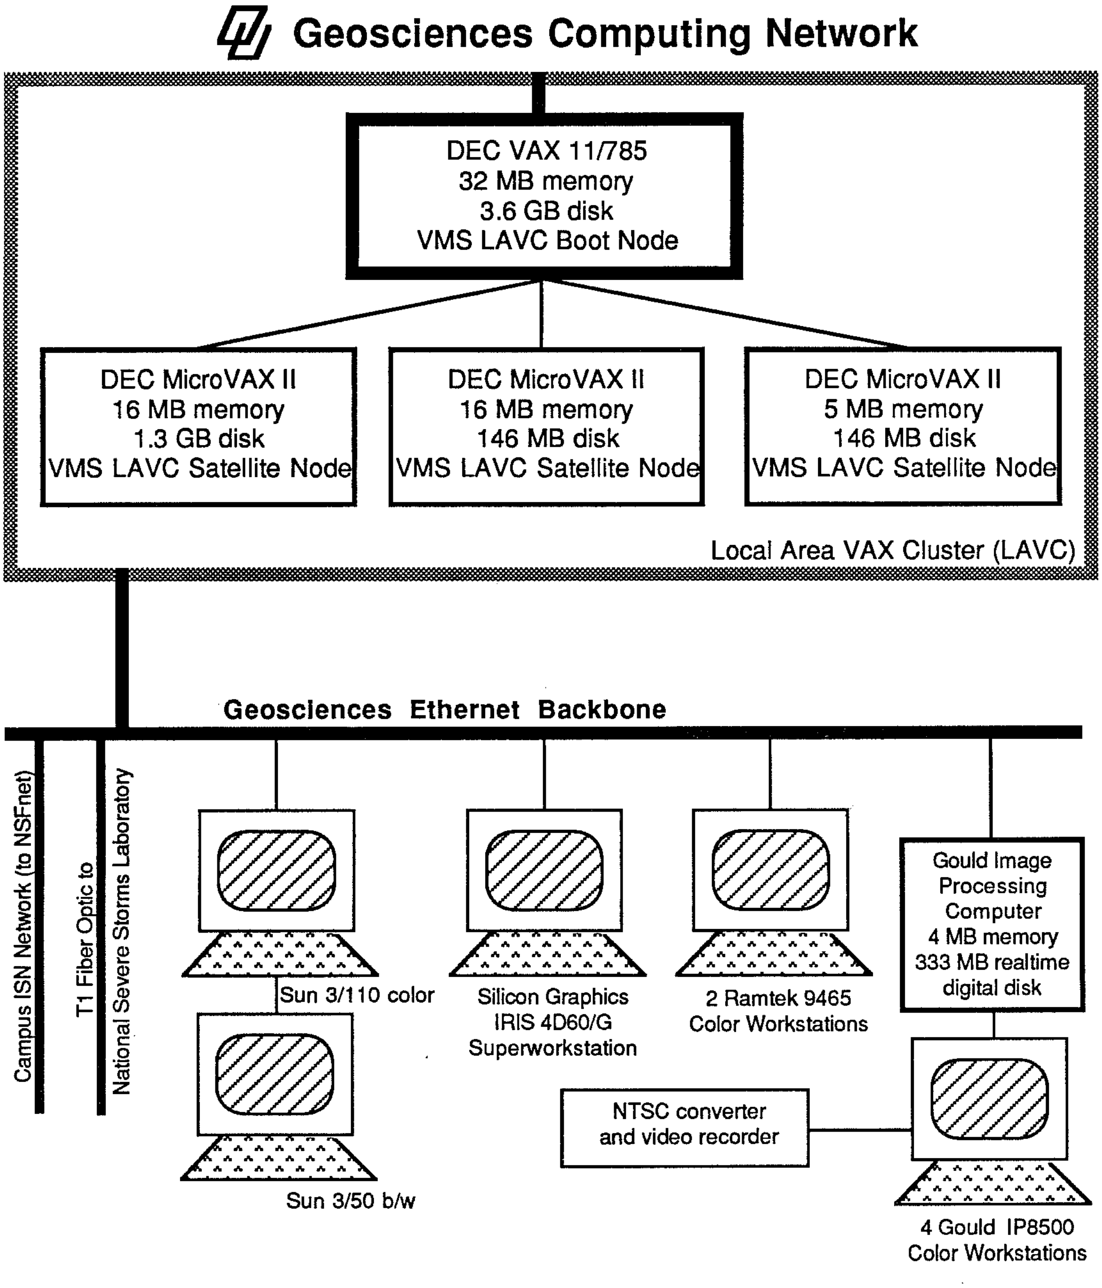

Figure 1. *Schematic of the Geosciences Computing Network (GCN) at the College of Geosciences, University of Oklahoma. See the text for further details.*

Using printed media, it is impossible to illustrate the unbelievable amount of information conveyed through use of color animation displays. For a number of excellent still-image samples of animated data, the reader is refered to pages 589-643 of the *Proceedings of the 3rd International Symposium of Science and Engineering on Cray Supercomputers.* Other discussions of imaging can be found in Smarr (1985, 1987), McCormick *et al.* (1987), Zabusky (1987), Droegemeier (1987), and Winkler *et al.* (1987a,b).

4. THE QUEST FOR ANIMATION IN AN OPERATIONAL SETTING

While color line drawings are obviously superior to their monochrome counterparts for displaying meteorological data, particularly the temperature-coded wind vectors of the ECMWF MAGICS system, it is clear that color raster images and animation displays are needed to fully understand and exploit the capabilities of today's sophisticated computer models. In addition to the display of fundamental variables in the model, it would seem that derived quantities such as low-level moisture convergence, thermal forcing and advection fields, and frontogenesis parameters would provide the forecaster with critical information which, if processed locally from raw data, could be customized to form products appropriate for the local climatology (e.g., orographic forcing parameters for mountainous localities and low-level thermal forcing fields for more uniform arid regions). In addition, without actually using a 3-D image, information characterizing the volume of the atmosphere could be obtained by examining 2-D sequences of forcing terms from the quasi-geostrophic equations, e.g., the differential vorticity advection.

Most of the animation systems currently used by the research community are either too expensive or require powerful computing and/or communications facilities to be feasible in today's operational environemnt. As technology evolves, however, particularly with respect to animation graphics, the distinction between research and operations will become less clear; both communities will be using advanced display devices in an effort to better understand complex phenomena. The one distinction which will remain, however, is the time constraint present in an operational setting. Ironically, it is the research community which, though not under operational time constraints, has fostered display systems which not only allow one to see numerical calculations evolve in realtime, but also provide the capability of interactively changing the calculations as they are performed. Although the latter may be unnecessary for operational forecasting, continued interaction between the research and operational communities is clearly needed to develop display systems appropriate for current and future prediction models.

5. RECOMMENDATONS

In developing a plan of action, we must first determine which techniques are appropriate for visualizing data from operational meteorological models, e.g., 2-D or 3-D displays, volume visualization or solid-surface rendering, animation or still images, transparent or opaque surfaces, etc. To address this issue, data from operational models should immediately be made available for experimentation using both 2-D and 3-D display techniques. It is important that meteorologists work closely with computer scientists and visualization "technologists" in this effort since a particular display may be visually appealing to one community, yet be entirely inappropriate for conveying meaningful information (the necessary constraint when devising visualization methods for scientific data). In addition, one cannot always anticipate the appropriateness of a technique until it has actually been implemented, and thus experimentation is a critical first step toward visualization of operational model data.

Another important aspect to consider is the potential impact of advanced display systems on the functionality of forecast offices. For example, a forecaster may be quickly overwhelmed if he or she has the opportunity to examine model output at any location in space and time over the entire globe! Furthermore, the integration of observed data (e.g., satellite images) with model output in an efficient manner also poses substantial logistical problems. While these and many additional issues are indeed important, each can be addressed with careful thought and experimentation. *Irrespective of the strategy used, we must work to ensure that continued improvements in operational models are accompanied by equivalent enhancements in the display of the data they produce.*

6. ACKNOWLEDGMENTS

The author is grateful for the invitation to visit the ECMWF and present the contents of this paper. The Geosciences Computing Network was established by a grant from the Shell Foundation, and generous hardware donations from Sun Microsystems, Inc., Silicon Graphics Computer Systems, Digital Equipment Corporation, Seismic Services Corporation, and Apple Computer Corporation have contributed to its success. The simulations shown at the conference were performed on the Cray-1A and Cray X-MP/48 supercomputers at the National Center for Atmospheric Research (which is sponsored by the National Science Foundation), and on the Cray-2 at the Minnesota Supercomputer Institute. The thunderstorm outflow simulation movie was produced by Dr. Craig Upson on the Cray X-MP/22 at Digital Productions. I have benefitted from several enlightening discussions on computer graphics with Craig Upson, Mike Norman, and Robert Wilhelmson of the National Center for Supercomputing Applications, Paul Woodward of the Minnesota Supercomputer Institute, Karl-Heinz A. Winkler of the Los Alamos National Laboratory, and J. Daabeck of the ECMWF. This research was supported by Cray Research, Inc., by the National Science Foundation under Grants ATM86-04402 and ATM86-57013, and by the National Oceanic and Atmospheric Administration under Grant NA85RAH05046 to the Central Oklahoma Mesoscale Modeling and Analysis Project.

7. REFERENCES

Carpenter, R.L. Jr., K.K. Droegemeier, P.R. Woodward, and C.E. Hane, 1988: Application of the Piecewise Parabolic Method (PPM) to meteorological modeling. *Proc. 6th Conference on Numerical Weather Prediction*, Amer. Meteor. Soc., Feb. 22-26, Baltimore.

Charney, J.G., R. Fjørtøft, and J. von Neumann, 1950: Numerical integration of the barotropic vorticity equation. *Tellus, 2*, 237-254.

Cray Research, Inc., 1987: University center emphasizes graphics. *Cray Channels, Summer 1987*, 39-40.

Droegemeier, K.K., and R.B. Wilhelmson, 1985a: Three-dimensional modeling of convection produced by interacting thunderstorm outflows, Part I: Control simulation and low-level moisture variations. *J. Atmos. Sci., 42*, 2381-2403.

________, and R.B. Wilhelmson, 1985b: Three-dimensional modeling of convection produced by interacting thunderstorm outflows, Part II: Variations in vertical wind shear. *J. Atmos. Sci., 42*, 2404-2414.

Droegemeier, K.K., 1987: Numerical simulation of thunderstorm outflows and microbursts. *Cray Channels, Summer*, 18-23.

________, and R.B. Wilhelmson, 1987: Numerical simulation of thunderstorm outflow dynamics. *J. Atmos. Sci., 44*, 1180-1210.

________, and R.P. Davies-Jones, 1987: Simulation of thunderstorm microbursts with a supercompressible numerical model. *Proc, 5th International Conference on Numerical Methods in Laminar and Turbulent Flow*, July 6-10, Montreal, Quebec, 1386-1397.

Droegemeier, K.K., 1988: Simulation of microburst vorticity dynamics. *Proc, 15th Conference on Severe Local Storms*, Amer. Meteor. Soc., Feb. 22-26, Baltimore.

McCormick, B.H., T.A. DeFanti, and M.D. Brown, 1987: *Visualization in Scientific Computing*. Report to the National Science Foundation by the ACM SIGGRAPH Panel on Graphics, Image Processing, and Workstations. Also in SIGGRAPH *Computer Graphics Newsletter, 21*, October, 1987.

Smarr, L.L., 1985: An approach to complexity: Numerical computations. *Science, 228*, 403-408.

________, 1987: The computational science revolution: technology, methodology, and sociology. *Proc., Conf. on High Speed Computing: Scientific Applications and Algorithm Design*, University of Illinois, April, 1986, Robert Wilhelmson, ed., 31pp.

Winkler, K.-H.A., J.W. Chalmers, D.E. Tolmie, P.R. Woodward, and N.J. Zabusky, 1987a: Ultra-speed graphics for computational science. *Cray Channels, Summer 1987*, 4-9.

________, ________, S.W. Hodson, P.R. Woodward, and N.J. Zabusky, 1987b: A numerical laboratory. *Physics Today*, October, p. 28.

Zabusky, N.J., 1987: Grappling with complexity. *Physics Today*, October, p. 25.