

Visualizing High-Resolution Climate Data

Sheri A. Voelz¹ and John Taylor^{1,2}

¹Mathematics & Computer Science, Argonne National Laboratory, Argonne, Illinois 60439

²Environmental Research Division, Argonne National Laboratory, Argonne, Illinois 60439
{voelz, jtaylor}@mcs.anl.gov
<http://www-climate.mcs.anl.gov/>

Abstract. The complexity of the physics of the atmosphere makes it hard to evaluate the temporal evolution of weather patterns. We are also limited by the available computing power, disk, and memory space. As the technology in hardware and software advances, new tools are being developed to simulate weather conditions to make predictions more accurate. We also need to be able to visualize the data we obtain from climate model runs, to better understand the relationship between the variables driving the evolution of weather systems. Two tools that have been developed to visualize climate data are Vis5D and Cave5D. This paper discusses the process of taking data in the MM5 format, converting it to a format recognized by Vis5D and Cave5D, and then visualizing the data. It also discusses some of the changes we have made in these programs, including making Cave5D more interactive and rewriting Cave5D and Vis5D to use larger data files. Finally, we discuss future research concerning the use of these programs.

1 Introduction

Weather simulations are important in helping us understand the key components that determine our weather, particularly extreme events. Manipulating the variables within the simulations gives additional insight into how a particular weather pattern is produced. Unfortunately, many variables are involved in these simulations, making the output files large. For example, one day's data in v5d file format can be around one gigabyte. This number can vary depending on the number of variables the file contains, the frequency with which they are written, and the size of the grid.

Two programs that have been developed for visualizing meteorological data are Vis5D and Cave5D. Vis5D takes atmospheric data values and visualizes them. One can easily view multiple variables at the same time to see how they relate and interact with one another. Vis5D can be run on many different platforms, including Linux, Sun, SGI, and HP. Cave5D is a version of Vis5D for virtual reality environments such as the CAVE and the ImmersaDesk. By visualizing in a virtual environment, a researcher can much more rapidly develop an understanding of the model results. A virtual environment offers much greater freedom to explore and manipulate the data. For example, one can enlarge and rotate the data at much larger scales, in order to study the dynamics of a particular area. Thus, Cave5D makes scientific analysis far easier.

2 MM5 Data Format

The MM5 modeling system consists of six programs: TERRAIN, REGRID, RAWINS, INTERP, GRAPH, and MM5. Each of the programs is executed in the above order interdependently. They are composed of a series of scripts that traditionally have been time consuming to modify and execute [1]; recently, however, this process has been simplified by the use of the Climate Workbench, developed at Argonne by Mathematics and Computer Science Division staff, including Veronika Nefedova, John Taylor, and Kevin Reitz [2].

The TERRAIN program defines the domain and map projection, generates the terrain, calculates vegetation and soil categories, and calculates the map-scale factors. The REGRID program calculates first-guess pressure level fields and calculates map-scale factors and the Coriolis parameter. RAWINS is usually used in forecasting and is not used in simulations of past events over long durations. It uses the first-guess fields combined with radiosonde and surface observations to improve the initial state of the simulation. INTERP uses the interpolated pressure-level data calculated in the REGRID program and transforms these data to the model's sigma coordinate. The resulting output from REGRID provides the initial and boundary conditions for an MM5 run. The GRAPH program generates simple line plots from the output of the other programs. Finally, MM5 integrates the mesoscale model for the time period selected producing a series of output files recording the state of the model at fixed time intervals [1].

3 Vis5D

The Vis5D program was developed at the University of Wisconsin-Madison Space Science and Engineering Center, with main contributions from Bill Hibbard, Johan Kellum, and Brian Paul [3]. The Vis5D program provides an interactive environment where users can view five-dimensional gridded data files. The dimensions consist of latitude, longitude, height, time, and an index into a set of physical fields (for example, wind, pressure, temperature, moisture). The most common types of data sets that are viewed using Vis5d are weather and ocean models [3].

3.1 How Vis5D Works

The Vis5D program accepts the v5d and comp5d data formats. MM5 model output data is not in this format. It can, however, be converted from the MM5 format by using Vis5D's import option, NCAR's TOVis5D tool [1], or a similar tool created by the user. The converted data is in a five-dimensional rectangle. The conversion process also creates additional variables needed to run Vis5D. These variables keep track of the number of time steps, variables for the number of rows and columns, information about the grid, the variable names, time stamps, and date stamps [3].

Vis5D starts by reading in the data file. In versions lower than 4.0, the entire file must be loaded into memory. Later versions discarded this restriction and now allow users to run files larger than the allocated size of physical memory. Using virtual memory, Vis5D loads only those data that are needed and discards least-

recently-used data. While the program is displacing data, it performs computations necessary to display isosurfaces, horizontal and vertical slices with contours or in color, volume renderings, and trajectories [3].

3.2 How Vis5D Is Used

Vis5D allows the user to customize the image displayed, while viewing the data. At the command line prompt, many different options can be set. The program allows the user to change window sizes, the way the date is displayed, the amount of memory the program uses, the frame rate, the way the program is loaded, how the display looks, and so on. Most of the arguments can also be changed while the program is running. Selecting the display button of the control panel allows the user to change the default values.

When Vis5D is loaded, the control panel and the display window pop up. Within the control panel the user can change views, chose what data to display and how to display it, animate or step through the data, add a map or topography information, import new data into the current session, save pictures, save configurations, and restore data. An example of the Vis5D interface showing results from our 40 km run of the "Perfect Storm" is shown in Figure 1. The run was computed with MM5v3.4 for October 26, 1991, through November 4, 1991.

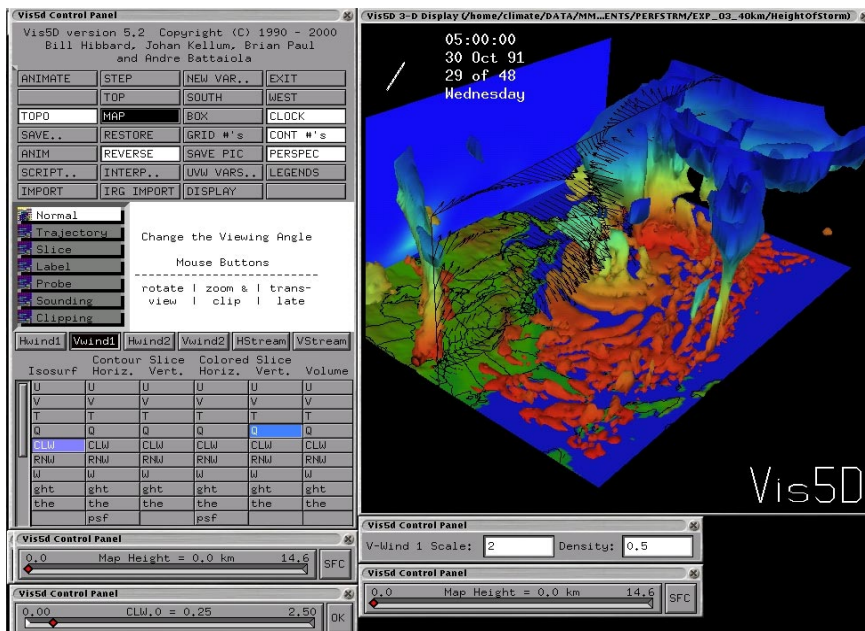


Fig. 1. Vis5D interface displaying data from the "Perfect Storm." The screen shot was taken from our simulation of 5:00 A.M., October 30, 1991.

The Vis5D package also comes with other utility programs to assist in managing Vis5D data files. V5dappend appends v5d data files together. Up to 400 time steps can be appended together without code modification. The v5dinfo tool outputs information about the v5d file. It generates information on the file format, compression size, header size, number of variables, information about the variables and time steps, and the coordinate location of where the data is taken from. Vis5D includes the program v5dstats. This utility computes the minimum, maximum, mean, standard deviation, and number of missing values for each variable in each time step. The package also includes v5dedit, which allows the user to edit a v5d file easily. The user can change a variable's name or unit, the file's times, dates, projection, vertical coordinate system, and low levels. Additionally, the Vis5D package includes v5dimport, which can be used to change the data format of a file to the v5d format.

3.3 Benefits of Using Vis5D

Vis5D allows users to visualize data in a straightforward GUI environment. The program gives users freedom to customize the visualization so every aspect of the simulation can be explored. Graphical object settings can be changed and duplicated to add to its flexibility. For example, the user can change the location of a slice and its colors. Isosurfaces can be displayed at different locations within the data. For example, by changing the value of the variable "rain" (measured in grams), the user can direct Vis5D to draw an isosurface where that particular amount of rainwater is found; higher values will be enclosed within this isosurface. This flexibility is excellent for testing hypotheses relating to a model simulation.

The visualization is presented in a clear, straightforward manner similar to that of a standard weather map. One can add land topography and a polygon map of the region to help locate areas of special interest. Moreover, the isosurfaces, slices, and volumes visualize the data in a simple manner that is easy to comprehend.

3.4 Modifications to Vis5D

Modifications were needed in Vis5D in order to allow us to work with large data sets. In particular, we needed to work with a 1.7-gigabyte v5d file, considered significantly larger than the default settings within the Vis5D code. This file held data from a simulation of the Great Midwest Flood of 1993. The main activity of the flood was during the months of June and July. The model results output by MM5 contained two months of hourly data. In this case, straightforward modifications to Vis5D involved changing the maximum number of time steps VIS5D could handle and increasing the amount of memory it used.

4 Cave5D

Cave5D uses Vis5D code to read in the v5d data file and to compute the desired objects. The basic difference between the two is that Vis5D was written to run on a desktop PC, whereas Cave5D was written for virtual reality environments. With

Cave5D one has the same freedom to change the configuration of the data (although it is more difficult to do so). The virtual environment has the advantage, however, that one can view the data within a larger area than a desktop computer and with higher resolution, thus revealing much greater detail in the model simulation. The results also can be rotated in space, and the size of the image can be changed along the x -, y -, and z -axes. All of these features help scientists to better analyze weather patterns.

The main contributors to the development of Cave5D are Glen Wheless and Cathy Lascara (Center for Coastal Physical Oceanography, Old Dominion University) and Bill Hibbard and Brian Paul (Space Science and Engineering Center, University of Wisconsin) [4].

4.1 How Cave5D Is Used

Even though Cave5D is based on Vis5D, the programs are run differently. Cave5D requires a configuration file, which it reads as ASCII characters so it can be modified using any text editor. This file must contain the path of the v5d input file. Other settings are based on the user's preferences. For example, graphical objects can be added. To generate the graphical objects in a desired form, the Cave5D developers suggest that the user first view the data in Vis5D, then select the settings that look best and save the file as a *.tcl file that can be read by a text editor. This file contains the values that were displayed in Vis5D. These values can then be used in the configuration file [4]. Other options that can be set include memory size, rate, and the display's size, position and color.

Once the Cave5D configuration file is prepared, the program can be loaded and run. (We have written sample instructions for running Cave5D at Argonne.) To bring up the starting panel, the user simply presses the right button on the CAVE wand. The panel allows the user to change the dimensions, animate the data, to speed or slow the animation, and return to the first time step. This panel also contains a button that will switch the panel to one where graphical objects are turned on and off and moved. To rotate the view, the user must hold down the middle button while moving the wand.

4.2 How Cave5D Works

Cave5D begins by performing a CAVE configuration. After configuring the CAVE environment, it processes the configuration file. It then picks out the variables and converts all necessary data from ASCII format. It also identifies which data file to use and which topography and map resolution to use, if applicable.

Cave5D then starts the Vis5D function. First it analyzes the machine(s) that it is running on. Vis5D will count the number of CPU's and will fork off an appropriate number of processes to maximize throughput. Then it starts to analyze and compute the necessary functions to display the data. This process includes initializing the topography, map, and data file and extracting key information from the data file on how the data is to be displayed. Some of this information is used in the next call to draw the clock.

After the clock is drawn in memory, Cave5D creates an "object" view. The object view is a rotation strategy used in manipulating the view of the image. This

function also establishes a navigation matrix that is used to transform the image when maneuvering and scaling in the image in the CAVE.

Next Vis5D functions are called again to compute the graphical objects the user defined in the configuration file. Each object's type is identified and handled accordingly. These computations take up most of the loading time. Objects are placed in a queue and computed in order of their appearance in the configuration file. This strategy was adopted in order to improve the run-time performance of Cave5D.

After this step is completed, the interactive menus are drawn into memory as texture maps. In the Cave5D version 1.4, there are two menus the user can interact with. Plotting pixels in memory draws each of these menus. As the user selects options on these panels, vector coordinates are gathered from the wand position to determine which button was pressed.

Finally, Cave5D does some additional CAVE configurations, and then the program goes into its event loop. Within this loop, it continuously calls CAVE functions to locate the wand and head position in order to correctly redraw the data. It also checks for events such as button pressing. Cave5D will continue in this loop until the user selects the "Close" button on the first menu.

4.3 Benefits of Using Cave5D

Cave5D allows users to easily analyze weather simulations in a virtual environment. Such an environment provides many benefits. In the CAVE one can enlarge the data image and rotate the view to best interpret the simulation. By changing the perspective, one can view the data from a distance or stand in it to see greater detail of the process involved.

Cave5D also allows flexibility. By editing the configuration file, the user can add many options to the simulation run. The benefit is that the users can define the objects that they wish to view when the program is loaded. The weakness is that the users cannot change these values while the program is running. We note that this feature has been changed in the modified version of Cave5D developed at Argonne and scheduled for release as version 2.0.

4.4 Modifications to Cave5D

The first change that was made to Cave5D involved rewriting the code to allow Cave5D to handle large data files with many time steps. Specifically, we converted the 32-bit version into a 64-bit version, modifying variables and changing the compilation flags. This allowed us to run data files in the multigigabyte range with the Cave5D application. We also increased the maximum number of time steps. We were then able to run larger data files containing a large number of time steps successfully. The test data set is discussed in the next section.

One other major modification to Cave5D involved making it more interactive. This was achieved by displaying extra panels that can be used to change the configuration values while in the CAVE. Previously the user had to stop the program, edit the configuration file, and restart Cave5D each time a modification was required. The new approach is far more efficient. Figure 2 displays the extra panels

that were added to Cave5D to allow interactive modification to the configuration values.



Fig. 2. These panel menus were added to Cave5D to increase its functionality. In this example we are going to modify the value of the cloud water isosurface.

The new panels are drawn into memory as follows. The first panel displays the list of the graphical objects that were included in the configuration file. Once the user selects the object to be changed, a new panel is displayed where the user can select which variable to change. Once the variable is changed, the final panel is displayed. This panel shows the object and variable names that were selected, along with the current value. After the user enters in a value and hits the "Enter" button, the entire object is sent back to Vis5D to be recomputed. Each time step of the object is placed inside a queue and is recalculated.

The last major modification to Cave5D involved updating the version of Vis5d. Previously Vis5D version 4.3 was used. Our current Cave5D software uses Vis5D version 5.2 to handle the necessary computations of graphical objects and a majority of the memory management.

5 Midwest Flood of 1993

The Midwest Flood caused fifty fatalities and created damages nearing \$15 billion. The flood resulted from heavy rains June through July 1993 on already saturated ground from the past winter and spring. During the month of July nine of the states that were affected by the flood saw twenty days of rain, when on average they would receive about eight to nine days [5].

By visualizing weather patterns such as this flood, we hope to gain a better understanding of why these events occur. This particular data set initially was difficult to visualize because of the size of the file (1.7 gigabytes). With the modifications to Vis5D and Cave5D, however, we were able to visualize the flood in its entirety. Figure 3 represents the data that was ran for July 6, 1993, of the Midwest Flood. The visualization was created through Vis5D and displays rainwater isosurfaces colored to moisture and a horizontal wind slice. The visualization shows moisture flowing from the Gulf of Mexico. The moisture is fed into the upper atmosphere by convection, adding more rainwater into the storm over the Midwest.




Fig. 3. The Midwest Flood of 1993 on July 6, 1993, at 5:00A.M. These results were part of a 52 km grid resolution MM5v3.4 run.

6 Conclusion and Future Research

Today's technology allows us to run more detailed weather simulation programs, producing even larger data files that are not susceptible to traditional modes of analysis. Visualization, on a desktop or in a virtual reality environment, provides a new way of analyzing these enormous data sets. Our efforts have involved using two new tools, Vis5D and Cave5D, to visualize data from MM5. In addition to converting the MM5 data to a format recognized by these tools, we modified both Cave5D and

Vis5D to handle over a gigabyte of meteorological data. Cave5D was further modified to make it more interactive and user friendly.

Future research will include further modification of both Vis5D and Cave5D. A short-term goal is to be able to work with an entire year of meteorological data. Our long-term goal is to be able to work with terabytes of data. Achieving this latter goal will require rewriting both programs to handle the large amounts of data involved. We also will need to redevelop the way Cave5D computes graphical objects. At present, all graphical objects are computed in their entirety before the program is fully loaded. We need to modify this approach because there is not sufficient memory to hold terabytes of data. Such steps would involve the following:

- Modifying memory management within the Cave5D and Vis5D code to handle the larger terabyte data sets.
- Allowing Cave5D to create graphics within a memory buffer, close to the maximum available memory, to achieve continuous graphics display.
- Modifying Vis5d and Cave5D to enable parallel I/O in order to work efficiently with large data sets and complex visualizations.
- Modifying the sequencing of the graphics calculations in Cave5D.

Acknowledgments

We thank the staff of the Futures Laboratory at the Mathematics and Computer Science Division at Argonne National Laboratory in Argonne, Illinois. This work was supported in part by the Laboratory Director Research and Development funding subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38.

References

1. Pennsylvania State University / National Center for Atmospheric Research, MM5 Home Page <http://www.mmm.ucar.edu/mm5/mm5-home.html>, 1999.
2. Taylor, J. Argonne National Laboratory, Climate Workbench http://wwwclimate.mcs.anl.gov/proj/climate/public_html/climate-workbench.html, 2000.
3. Space Science and Engineering Center University of Wisconsin - Madison, Vis5D Version 5.1 <ftp://www.ssec.wisc.edu/pub/Vis5D/README>, 1999.
4. Old Dominion University, Cave5D Release 1.4 <http://www.ccpo.odu.edu/~Cave5D/Cave5DGuide.html>, 1998.
5. Larson, L.W. Hydrological Research Laboratory, The Great USA Flood of 1993 http://www.nwrhc.noaa.gov/floods/papers/oh_2/great.htm, 1996.