Visualizations of Earth Process for the American Museum of Natural History

Abstract

The American Museum of Natural History in New York City has built a new exhibit space—The Hall of Planet Earth. This hall highlights earth processes using various exhibits including actual rocks and core samples, demonstration models, and video display stations.

One specific scientific area that the museum wants to highlight is that of modeling and simulation. Los Alamos has a long history in this area through our involvement in programs such as the DOE Grand Challenges and the Institute for Geophysics and Planetary Physics. Because of this, we were asked to participate in the design of, and provide content for, five exhibits designed to showcase modeling and simulation of individual earth processes.

This paper briefly describes the scientific visualizations developed and used to model atmosphere, ocean, and mantle processes for the American Museum of Natural History's exhibit.

1 Introduction

The American Museum of Natural History in New York City is currently building a new exhibit space—The Hall of Planet Earth. This hall will highlight earth processes using various exhibits including actual rocks and core samples, demonstration models, and video display stations.

One specific scientific area that the museum wants to highlight is that of modeling and simulation. Los Alamos has a long history in this area through our involvement in programs such as the DOE Grand Challenges and the Institute for Geophysics and Planetary Physics. Because of this, we were asked to participate in the design of, and provide content for, five exhibits designed to showcase modeling and simulation of individual earth processes.

The modeling and simulation exhibits will consist of five video display stations distributed throughout the hall. Each video station will play 4-5 minutes of pre-recorded video when triggered by a museum visitor. The first few minutes of each video will explain the details of modeling a specific earth process through graphic animations, textual overlays, and interviews with the simulation scientists. The final 1-2 minutes of each video use actual scientific visualizations of the simulation data to explain specific features under study. The museum specified the use of scientific visualization, as opposed to artists renditions, to convey the process that scientists use to understand simulation results.

In the following sections we'll describe three of the five visualizations that Los Alamos delivered to the museum—an atmospheric simulation of a severe winter storm, a global ocean model, and the process of mantle convection. In each section we'll briefly describe the model and the visualization tools and techniques used to produce these animations.

2 Atmospheric Model

The atmospheric model was used to simulate the development of one of the strongest storms to hit the eastern United States this century and tracks its development from Brownsville to Newfoundland. This storm, through a combination of heavy snow, high winds, severe storms, and coastal flooding, claimed dozens of lives and caused over 2 billion dollars in damages. The storm also produced one of the largest areal coverage of deep snow ever, paralyzing the eastern seaboard, and its effects were felt deep into the tropics including in Cuba and the Yucatan. The Regional Atmospheric Modeling System (RAMS) [1], originally developed at Colorado State University, uses measurements from weather stations all over the country and numerical calculations to predict evolving weather patterns.

Model output includes temperature, pressure, wind vectors, and species of condensate such as ice crystals, high-elevation snow, snow and rain. Two animations were created to visualize the dynamics of the simulated storm system. A overhead view animation details the life cycle of the storm. A side view highlights the storm's intense development phase. To create these animation sequences for the scientists and museum we used IBM's Visualization Data Explorer (DX) product [2]. Data Explorer provides a full collection of visualization operators and allows for fast program creation via a data-flow program graph editor.

Both animations depict three and half days of simulated time from 12 PM, March 11, 1993 to 12 AM, March 15, 1993. The overhead view animation details the winds near the jet stream level using stream ribbons. The extent of the clouds associated with the storm are shown using volume rendering. Contours of surface pressure are used to show how the storm intensified over the eastern seaboard, producing hurricane force winds in some locations. The areal extent of the rain/snow is depicted using scalar color mappings as the storm propagates from Texas to Maine. Local temperatures are also reported using numerical values. Figure 1 shows a frame from this animation.
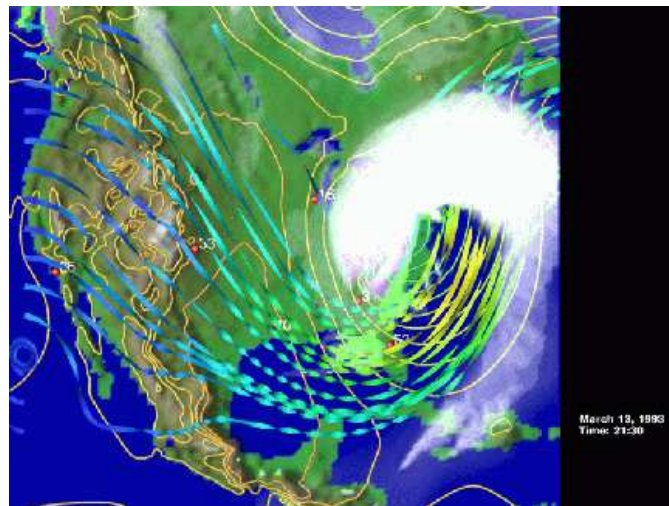

Figure 1: Overhead view of storm.

A side view highlights the storm's intense development phase. The view shows the strong vertical lifting associated with the low pressure at the center of the storm by using stream ribbons which originate at the surface. This lifting produces the heavy clouds and rain/snow shown using volume rendering. Figure 2 shows a frame from this animation.
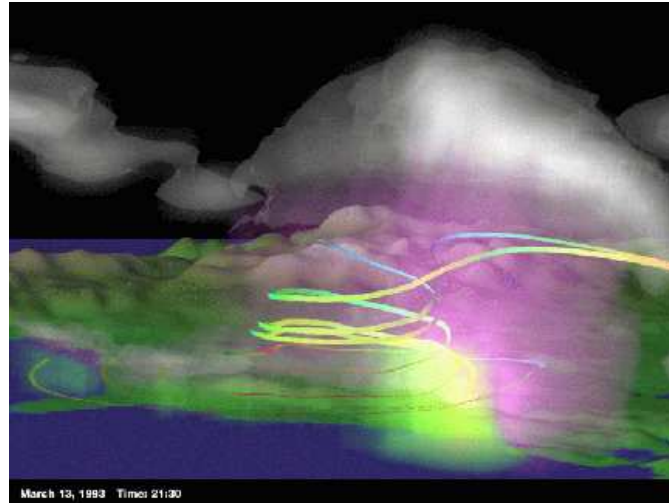

March 13, 1993 Time: 21:30

Figure 2: Side view of storm's vertical development.

The importance of these animations is in being able to see how all of the different variables that define atmospheric structure interact to produce such an extraordinary event. From this type of visualization scientists can better understand how subtle aspects of atmospheric dynamics can come together at the right time to produce a killer storm.

3   Ocean Model

The Earth's climate is determined by a complicated interaction between the ocean, sea ice, atmosphere, and biosphere. Computer models that simulate numerically the behavior of this system are one of the best means we have for projecting future climate and the impact of humanity's activities on it. Present-day general circulation models (GCMs) are able to simulate satisfactorily many aspects of the current climate, though a new generation of models is needed that have finer spatial resolution and that more realistically treat the physical processes that control our climate. To meet these objectives, we need GCMs that run on massively parallel computers. As part of the DOE's Grand Challenge program, scientists at Los Alamos have developed one such model: a global ocean circulation model named the Parallel Ocean Program (POP).

The POP ocean simulation, running on the Laboratory's SGI Origin 2000 parallel computers, employs a global grid containing 1280 uniformly spaced points in longitude and 896 variably spaced points from 78° N to 78° S latitude, yielding a spatial resolution ranging from 31 km at the Equator to 7 km at 78° latitude.

The model uses 20 non-uniformly spaced depth levels and realistic bottom topography (bathymetry). Observed surface winds from the period 1985-1995 and realistic monthly mean heat and salt fluxes are used to force the model. Additionally, we run the model using roughly the same grid spacing over only the North Atlantic, resulting in much greater spatial resolution for that area.

As POP runs it periodically writes data files representing the progress of the simulation. Although the simulation computes on a 30 minute time-step, these files are written every three days of simulated time. At each three day time-step, one file is written for each variable being computed: salinity, temperature, sea-surface height, and flow vectors).

Historically, we've visualized this sequential collection of data files using video technology. These video visualizations, while useful for viewing the progress of the simulation, have a serious drawback—they are static and can't be modified without creating a new video. Because large simulations are run infrequently, these video animations have long lifespans—their shortcomings become increasingly apparent as time goes on.

To address these and other limitations we developed an interactive ocean model rendering tool called POPTEX [3]. This tool duplicates the benefit of video visualizations (putting the results of the simulation into motion) while adding capabilities that enable dynamic, flexible, and interactive exploration of their data. To do this we used the powerful combination of hardware features available on the Laboratory's SGI Origin 2000 and its Infinite Reality (iR) graphics pipes.

Specifically, we exploit the iR's fast texture mapping capabilities [4] to provide the desired interactivity. The results of the simulation are converted to 8-bit texture images which are then mapped through an editable texture lookup table (TLUT) onto the globe. The TLUT itself is implemented in the iR's hardware and can be loaded almost instantaneously. The main advantage of POPTEX though, is its animation capability. The collection of textures in main memory can be continuously streamed into texture memory at an observed maximum rate of 72 million texels per second. This results in a maximum frame rate of 60Hz. At this rate, ten years of simulated time pass in just 21 seconds. More useful than end-to-end animation though, is the ability to choose a period of time and selectively animate over only that range—at any speed, forward of backward, pausing or changing the rate as desired.

Although most of the variables we visualize (sea-surface height, temperature, and salinity) are mapped to colors, we have experimented with some alternative mappings. For example, we've used the hillshading technique [5] to display sea-surface height in shaded relief. Figure 3 shows sea-surface height in the North Atlantic using this technique. Visualizations of sea-surface height are of interest in many areas of the world. For example, the strong eddies seen in the Caribbean and Gulf of Mexico can affect the operations of oil drilling platforms.

For the museum project, we added code to visualize the surface currents of the ocean. Our first attempt was to advect particles (or drifters) through the vector field, leaving a dissipating trail behind them as they progress through the flow field. Figure 4 shows an example of this technique applied to the Agulhus current that flows around southern Africa. Both the visualization researchers and the simulation scientists we were encouraged by the results of this technique—the drifters effectively tracked the eddies in the flow field. We previewed an animation of this technique to the museum staff fully expecting an equally positive response.
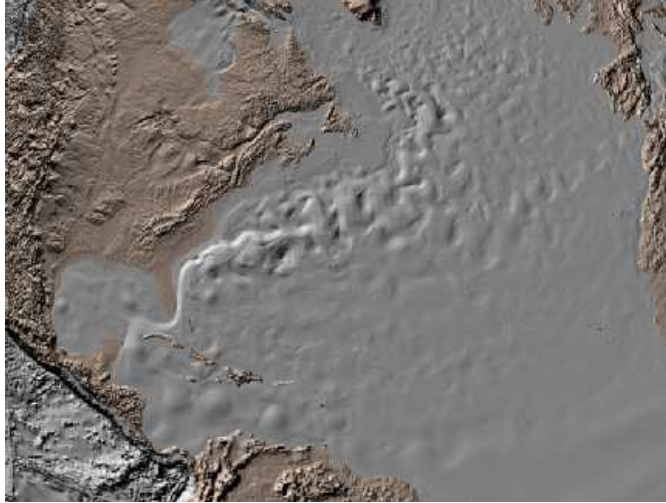
Figure 3: Hillshaded sea-surface height in North Atlantic.

Unfortunately, that's not what we heard. Their initial response to the drifters was that they looked "creepy" or "like bugs". We spent quite a bit of time trying alternate color schemes and line styles—none of which appeared a great deal more appealing. In the end, we stayed with the original depiction since that's what the ocean scientists will be using on a day-to-day basis.

4   Mantle Model

Solid state convection within the Earth's mantle determines one of the longest time scales of our planet. The Earth's mantle, the 2900 km thick silicate shell that extends from the iron core to the Earth's surface, though solid, is deforming slowly by viscous creep over long time periods. While gradual in human terms, the vigor of this subsolidus convection is impressive, producing flow velocities of 1-10 cm/year. Plate tectonics, the piecewise continuous movement of the Earth's surface, is the prime manifestation of this internal deformation, but ultimately all large scale geological activity of our planet, such as mountain building and continental drift, must be explained dynamically by mass displacements within the mantle.

A major problem for researchers in computational mantle dynamics is to resolve the Earth's outer 100 km deep skin, or lithosphere. This lithosphere is an integral part of the mantle and thus a 100 km wide spatial resolution has to be achieved throughout the volume. The resulting computational problem requires numerical discretizations with approximately 10-100 million grid points to resolve the mantle volume on scales of 50 km or less. Mantle convection researchers at Los Alamos use the 3D spherical mantle dynamics code TERRA, which solves the Navier-Stokes equations in the infinite Prandtl number limit using a multigrid approach [6]. A message passing version of TERRA runs on a wide variety of parallel platforms, from clusters of Linux PCs through large parallel machine such as the SGI/Cray Origin 2000 and the SGI/Cray T3E [7]. The large memories of these machines has allowed scientists to investigate convection using a numerical grid of more than 10 million finite elements and thus allows them to resolve a
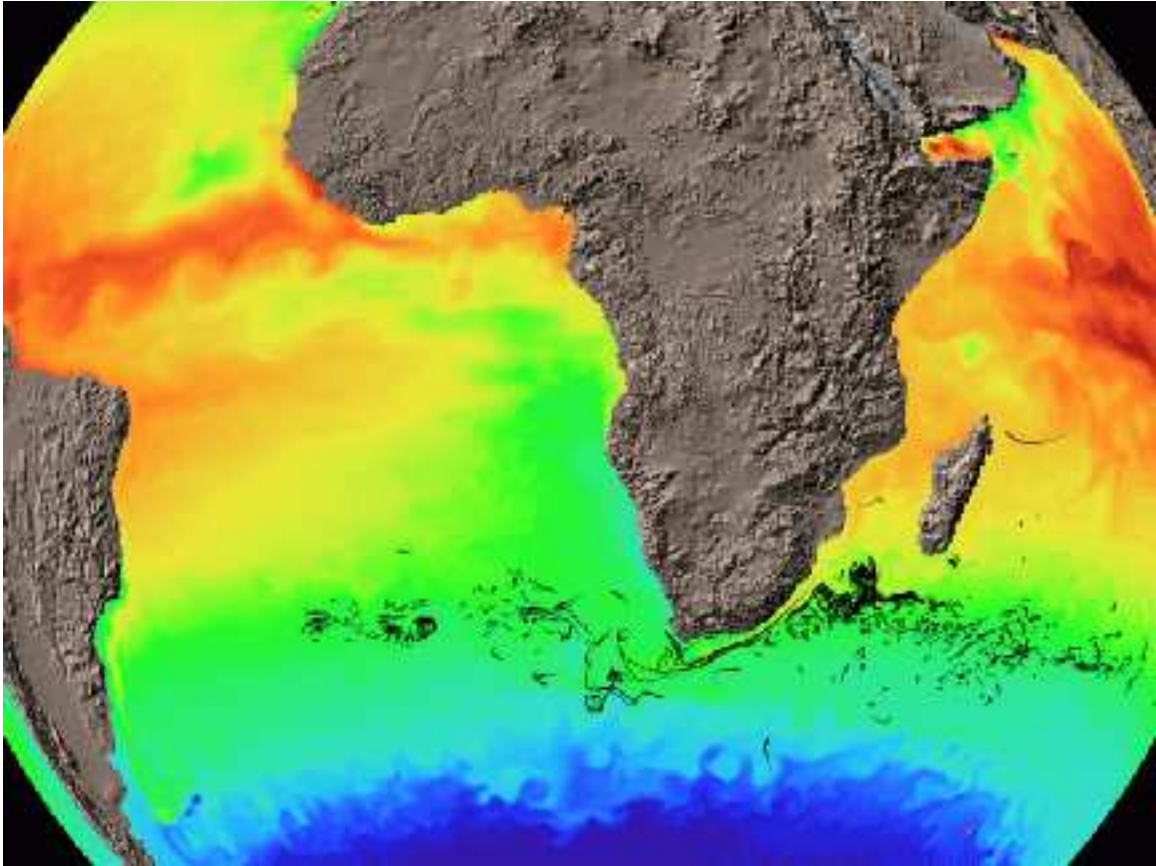
5

Figure 4: Drifters following Ahulhus current.

large range of dynamical length scales within the mantle.

When the TERRA visualization effort was first begun, simulations were primarily run on a 256 processor Cray T3D system and being able to run visualization codes on the same platform as the simulation was a big advantage. We therefore developed purely software visualization tools that ran on the parallel computer where the data was generated. It allowed for both a rapid and high resolution display of simulation results too large for visualization on even the high-end graphics workstations of the time and avoided time consuming data transfers between the simulation host and the visualization computers. In today's environment at the Advanced Computing Laboratory, where our main platform for computation is a 2048 processor cluster of Origin 2000 systems that can include hardware graphics accelerators, an OpenGL solution might have better performance. Still, the parallel software tools are portable and scalable and can run efficiently on any platform where the simulation code runs.

The parallel visualization tools consist of an isosurface extractor, a parallel software polygon renderer, and a parallel slicer that can interpolate arbitrary planar slices through field data. These tools use a message passing and active message programming model [8]. The tools operate directly on the TERRA grid structure. While the TERRA grid is not a structured grid, the recursive subdivision basis of the grid allows the grid geometry to be implicitly represented rather than explicitly stored, saving memory and allowing for efficient geometric queries of the grid.

Our software parallel renderer uses a sort-middle based rendering algorithm. Both the data domain and the image are partitioned evenly among the processors. Each processor first handles the geometric processing for the portion of the data it holds: isosurface extraction, arbitrary slicing and geometric transformation. The resulting geometric primitives are partitioned into scanline segments according to the portion of screen space they cover and sent to the processor responsible for that portion of the image using an active message communications model. When the active message arrives at its destination processor, a handler function is invoked that completes the rasterization of the primitives it contains. Opaque scanline segments are directly z-buffered. Transparent scanline segments are buffered and sorted and composited after all processors complete geometric processing. Arbitrary slicing is handled through software based texture mapping which maps pixels in the slice plane back into the field grid for color lookup. Isosurfaces are extracted using a parallel version of the NOISE algorithm [9]. More details about the TERRA visualization tools can be found in [10].

Figure 5 shows a frame from the mantle convection animation that will be used at the American Museum of Natural History. This frame shows the temperature field from the simulation colormaped red (hot) to blue (cold). The outer blue transparent isosurface is a relative low temperature and indicates where cold material moves back toward the interior of the mantle. The inner orange isosurface is a relative high temperature and indicates hot material moving outward. The visualization tools run at interactive rates (3-5 FPS). While the video produced for the museum was rendered in batch mode, the rendering rate was still over 3 frames per second, excluding image write time.
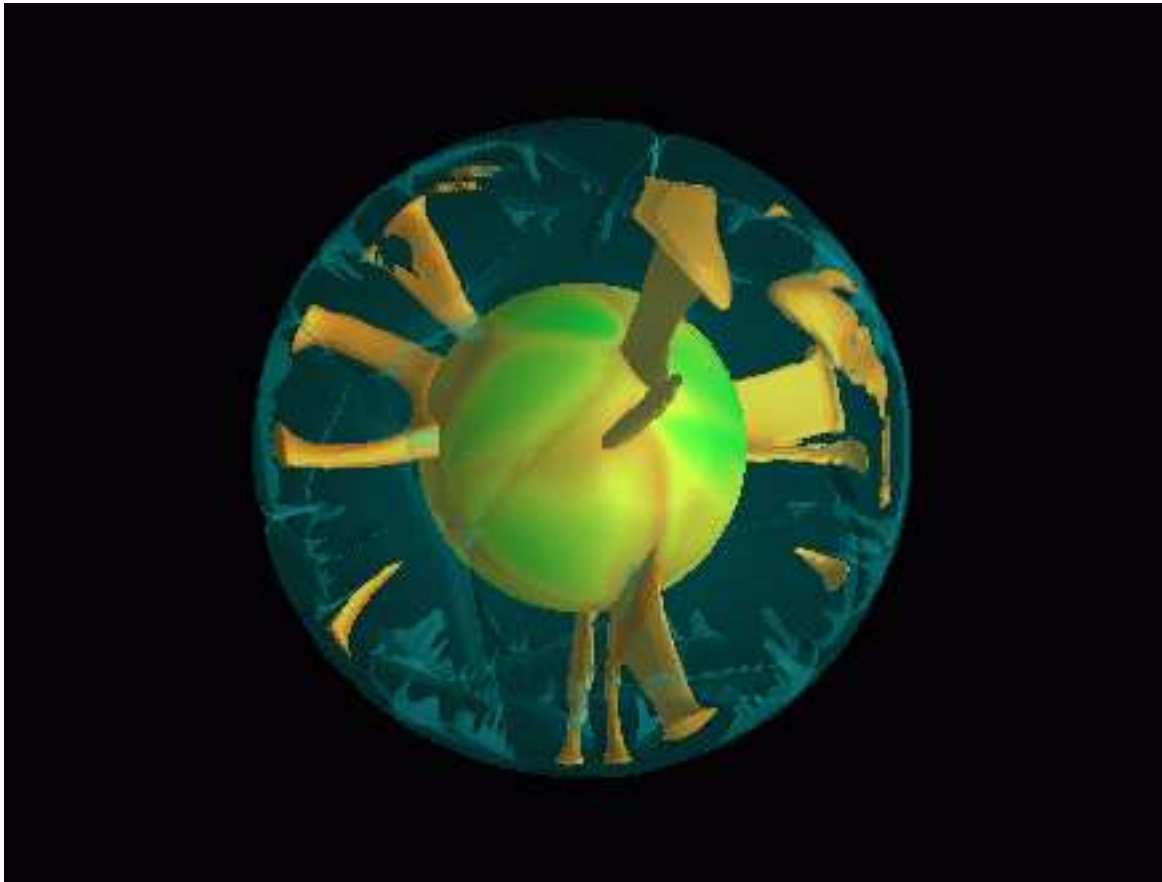
Figure 5: Mantle convection animation frame.

5 Acknowledgements

6 References

[1] R. A. Pielke et al., "A Comprehensive Meteorological Modeling System - RAMS," Meteorology of Atmospheric Physics, Vol. 49, 1992, pp. 69-91.

[2] G. Abram and L. Treinish, "An Extended Data-Flow Architecture for Data Analysis and Visualization," IEEE Visualization 95 Conf. Proc., IEEE Computer Society Press, Los Alamitos, Calif., Oct. 1995, pp. 263-270.

[3] John S. Montrym, Daniel R. Baum, David L. Dgnam, and Christopher J. Migdal, "InfiniteReality: A Real-Time Graphics System," SIGGRAPH 97 Conf. Proc., ACM SIGGRAPH, Addison-Wesley, August 1997, pp. 293-302.

[4] Allen McPherson and Mathew Maltrud, "POPTEX: Interactive Ocean Model Visualization Using Texture Mapping Hardware," IEEE Visualization 98 Conf. Proc., IEEE Computer Society Press, Los Alamitos, Calif., Oct. 1998, pp. 471-474.

[5] B. K. P. Horn, "Hillshading and the Reflectance Map," Proceedings of the IEEE, 1981, 169(1):14-47.

[6] John R. Baumgardner. "Three dimensional treatment of convective flow in the Earth's mantle," J. Stat. Phys, 39(5-6):501-511, 1985.

[7] Hans-Peter Bunge and John R. Baumgardner. "Mantle convection modeling on parallel virtual machines," Computers in Physics, 9(2):207-215, 1995.

[8] James S. Painter, Patrick McCormick, Michael Krogh, Charles Hansen, and Guillame Colin de Verdire. "The ACL message passing library," EPFL Supercomputing Review, 7, November 1995.

[9] Yarden Livnat, Han-Wei Shen, and Christopher R. Johnson. "A near optimal isosurface extraction algorithm using the span space," IEEE Transactions on Visualization and Computer Graphics, 2(1):73-84, 1996.

[10] James S. Painter, Hans-Peter Bunge and Yarden Livnat, "Case Study: Mantle Convection Visualization on the Cray T3D,", In Proceedings of IEEE Visualization '96, pages 409-412, San Francisco, CA, October 1996.

mailto:mcpherson@lanl.gov Allen McPherson