Visualizing Spatial and Temporal Variability in Coastal Observatories

Walter H. Jiménez* Wagner T. Corrêa† Cláudio T. Silva‡ António M. Baptista§

OGI School of Science & Engineering, Oregon Health & Science University

Abstract

In this paper, we describe a set of 3D and 4D visualization tools and techniques for CORIE, a complex environmental observation and forecasting system (EOFS) for the Columbia River. The Columbia River, a complex and highly variable estuary, is the target of numerous cross-disciplinary ecosystem research projects and is at the heart of multiple sustainable development issues with long reaching implications for the Pacific Northwest. However, there has been until recently no comprehensive and objective system available for modeling this environment, and as a consequence, researchers and agencies have had inadequate tools for evaluating the effects of natural resource management decisions. CORIE was designed to address this gap and is a major step towards the vision of a scalable, multi-use, real-time EOFS.

Although CORIE already had a rich set of visualization tools, most of them produced 2D visualizations and did not allow for interactive visualization. Our work adds advanced interactive 3D tools to CORIE, which can be used for further inspection of the simulated and measured data.

CR Categories: I.3.2 [Graphics Systems]: Computer Graphics—Computing Methodologies; I.3.8 [Applications]

Keywords: coastal observatories, environmental observation and forecasting systems, coasts, estuaries, Columbia River

1 INTRODUCTION

Paradigms for modern modeling and visualization of complex ecosystems are changing quickly, creating enormous opportunities for scientists and society. For instance, powerful and integrative modeling and visualization systems are at the core of environmental observation and forecasting systems (EOFS), which seek to generate and deliver quantifiably reliable information about the environment at the right time and in the right form to the right users. As they mature, EOFS are revolutionizing the way scientists share information and represent an unprecedented opportunity to break traditional information barriers between scientists and society at large.

EOFS have implicit assumptions of robustness and computational efficiency in complex simulations of natural processes. Coastal observatories, focused on specific estuaries or other coastal regions, are examples where modeling systems have advanced dramatically towards the EOFS vision. Various coastal observatories are, in particular, developing unprecedented "4D virtual realities" of physical processes, some of them in quasi real-time. However, the success of coastal observatories poses novel challenges for both quality control and interpretation of simulations.

Some of the challenges result from the breaking of traditional modeling cycles. Tight production schedules, dictated by real-time forecasts and multi-decade simulation databases, lead to a very large number of complex runs being produced on a daily basis, which are beyond the capacity of analysis of many modeling teams.

In this paper, however, we rather focus on challenges that result from the level of detail contained within each simulation. Using as a reference a specific coastal observatory [Baptista 2002; Baptista et al. 1999], we describe the visualization core of an emerging strategy to generate multi-dimensional representations of complex coastal circulation processes. The strategy is driven by the needs of scientific research and regional management, for the Columbia River estuary and plume.

2 THE COASTAL OBSERVATORY

The Columbia River is the target of numerous cross-disciplinary ecosystem studies and is at the heart of multiple sustainable development issues with long reaching implications for the Pacific Northwest of the United States. Urgent research efforts are, in particular, focusing on the role of the estuary and plume on the life cycles of fish species protected by the Endangered Species Act. These efforts are being conducted in an often-conflicting context of navigation improvements, evolving hydropower management strategies, and ecosystem restoration efforts. A key challenge within those efforts is to separate natural from man-made effects in a system characterized by extreme variability in both space and time.

CORIE, a coastal observatory developed and maintained since 1996 [Baptista 2002; Baptista et al. 1999], is integral to the region's strategy, because of its ability to characterize key physical variables in 4D detail, and to understand their response to natural and anthropogenic factors. Designed as multi-purpose infrastructure, CORIE includes an observation network, an advanced modeling system, and an information management system.

The observation network covers the estuary extensively and the plume sparingly, through fixed stations typically with real-time telemetry. At each station, variable combinations of in-situ sensors measure water temperature, salinity, pressure, velocity, acoustic

*Department of Computer Science & Engineering, OGI School of Science and Technology, 20000 NW Walker Road, Beaverton, OR, 97006; wherrer@cse.ogi.edu.

†Work performed while at OHSU. Currently at Department of Computer Science, Princeton University, 35 Olden St., Princeton, NJ 08540; wtcorrea@cs.princeton.edu.

‡Work performed while at OHSU. Currently at Scientific Computing and Imaging Institute, School of Computing, University of Utah, 50 S. Central Campus Dr., Salt Lake City, UT 84112; csilva@cs.utah.edu.

§Department of Environmental and Biomolecular Systems, OGI School of Science and Technology, 20000 NW Walker Road, Beaverton, OR, 97006; baptista@ccalmr.ogi.edu.

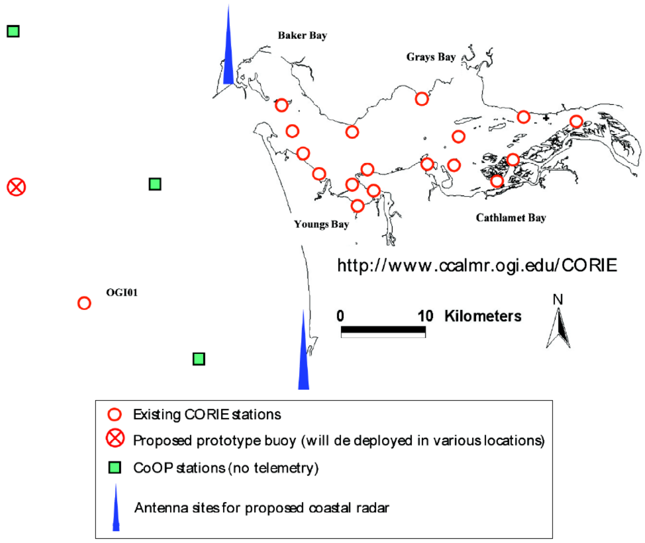

Figure 1: The CORIE observation network is an adaptive system, inherently collaborative with other programs. For instance, recently proposed CORIE stations and stations of the CoOP program are shown in the figure side by side with existing CORIE stations.

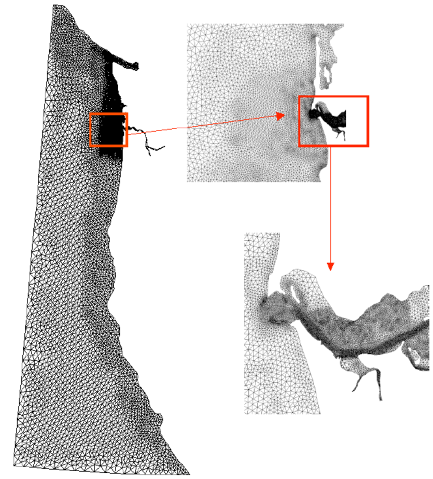

Figure 2: The CORIE computational domain extends along the continental shelf from British Columbia to California. An horizontally unstructured grid enables enhanced refinement in areas of primary interest, including the near-field plume, the estuary entrance and the navigation channel.

backscatter, wind speed and direction. A vessel used regularly by a local community college, for seamanship training, is also equipped with hull-installed sensors for temperature, salinity and velocity, designed to function unattended. In addition, we occasionally conduct releases of Lagrangian drifters, which passively follow the water flow at prescribed depths.

The modeling system is designed to characterize through computer simulation the complex circulation in the Columbia River estuary and plume, through quasi real-time forecasts, retrospective simulations of past conditions ("hindcasts") or exploratory simulations of future conditions ("scenarios"). In addition to numerical codes, the modeling infrastructure of CORIE includes modules that specifically deal with characterization of forcings (bathymetry; and river, atmospheric and ocean forcings), quality control of simulations, and generation of post-processing products (see Section 3). The modeling domain is shown in Figure 2, and extends from the downstream-most dam in the river to the vast continental shelf (British Columbia to California) that is affected by the powerful freshwater plume.

The information management system organizes both primary data (observations and simulations) and post-processing products (including an extensive array of visualizations). It also coordinates the generation and automated web publishing of most post-processing products.

3 VISUALIZATION SYSTEM

An integral part of CORIE is a set of visualization capabilities (see `http://www.ccalmr.ogi.edu/CORIE`). In fact, the only way to really understand the data products generated by the system is by using some form of visual representation.

The thousands of visualization products that are generated daily by CORIE on an automated or semi-automated basis already offer an extensive insight on the overall circulation of the Columbia River, and on modeling error and uncertainty.

However, at this time, there is a mismatch between the simulation capabilities of the system, which is based on high-resolution time-varying 3D unstructured grids, and the visualization component that for the large part only generates 1D or 2D plots (some 3D information can be inferred from depth plots, which slice the 3D data).

Part of the problem is that the visualization of time-varying 3D unstructured grids is not a well-developed area of research. Most state-of-the-art commercial visualization tools (e.g., VTK, IBM OpenDX) have extremely limited capabilities for handling large 3D unstructured data, and the support for time-varying data is also limited. (Most works that target the visualization of time-varying data have so far far focused on regular grids).

Our ongoing work involves extending current systems to support time-varying datasets with particular emphasis on the visualization of CORIE data products. For most of our work, we are using VTK augmented with custom code, such as the volume renderer described in Corrêa et al [Corrêa et al. 2003].

3.1 Pre-processing of model output

In order to visualize the CORIE output data, it is necessary to convert its format into one that is suitable for our rendering algorithms. The volume rendering of the salinity scalar values requires as input an unstructured grid of tetrahedron where each vertex is associated to one salinity scalar value, while the rendering of the bathymetry needs a grid of triangles representing the ground surface. Finally, for visualizing the velocity field it is necessary to have an unstructured grid of points with vector attributes associated to each one.

The CORIE output files have the geometry and one attribute value associated with each point of the data set. The geometry consists of one horizontal grid of triangles where the vertices and triangles are called nodes and elements respectively, as we can see in Figure 3. Each node has associated a vertical list of points, called column of points, in which each point is located at a specific level of depth. The number of points in each column is variable, depending on the depth of water under the node. The attribute value can

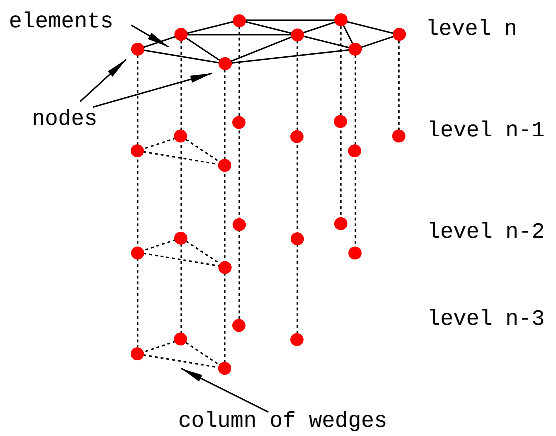

Figure 3: The CORIE output data.

be a scalar as salinity or temperature, or a vector as velocity. A set of 96 attribute data samples collected during one day is stored in a single file, which can reach 500MB in size each.

In Figure 3, we can see that under each element there is a column of wedges. Every wedge can be subdivided into three tetrahedra, thus we can create the unstructured grid of tetrahedra needed for the volume-rendering algorithm. The wedge subdivision is done by using the face diagonals of its rectangular faces. As each rectangular face has two diagonals, the wedge can be subdivided into two triangles in two different ways. In order to get two adjacent wedges consistently subdivided (i.e., their shared face subdivided by the same diagonal) we use the diagonal starting in the node with smallest index (see Figure 4) as described in Max et al [Max et al. 2000].

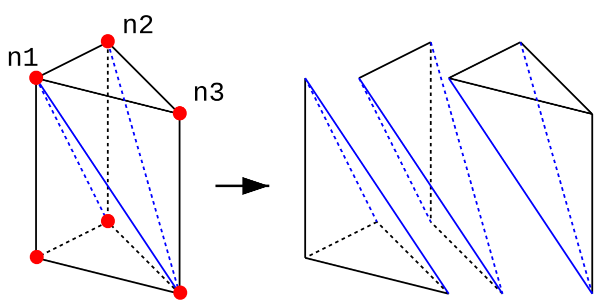

Figure 4: One wedge can be subdivided into three tetrahedra.

The grid of triangles representing the ground surface, called the bathymetry, is constructed using as vertices the points of greater depth in each column of points, thus each node has its corresponding vertex in the ground surface. The triangles are defined by the three vertices corresponding to the nodes that define the elements (see Figure 5). (When rendering bathymetry, we often applied different colors to each depth level in order to enhance bathymetry visualization.)

Finally, for visualizing the velocity field, the unstructured grid of points needed as input is directly created from the nodes and their associated columns of points.

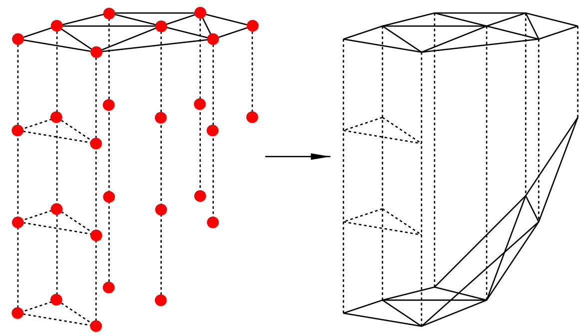

Figure 5: The ground surface (bathymetry).

3.2 Visualization strategies

3.2.1 Data handling issues

CORIE simulations are both large and very long. Typical grids have on the order of six million cells, and simulations are commonly conducted in increments of one week for multiple months or years, with dumps of one time step for each 15 minutes of simulation. Output files are typically structured as daily files, each with 96 time steps.

In order to be able to handle the large runs, we added two preprocessing capabilities to the existing repertoire:

- The ability to extract a subset of the data, so the user can visualize specific 3D regions of interest;
- And the ability to merge data sets of many consecutive runs (days) into one file, in a way that we can visualize the behavior of scalar and vector fields for a larger period of time.

3.2.2 Unstructured-grid volume rendering

There are several scalar and vector fields in CORIE that can be effectively visualized with direct volume rendering. In volume rendering, the 3D scalar file to be visualized is modeled as a cloud-like material, which both attenuates light along the viewing ray and add light into it [Max 1995]. To create an image, the effects of the material must be integrated along the viewing ray through each pixel. This requires a separate integral for the contribution along the ray segments inside each cell. A key property of volume rendering is that it generates holistic views of datasets since the images actually have contribution from potentially all the cells in the dataset. By varying the optical properties assigned to particular cells, it is also possible to generate isosurface renderings using direct volume rendering techniques. Thus volume rendering can be seen as a truly general rendering approach. Unfortunately, this generality comes at a cost. Volume rendering can be quite time consuming, and substantial efforts have been devoted to optimization techniques [Levoy 1990; LaMar et al. 1999; Williams 1992; Silva et al. 1996]. The fastest techniques for volume rendering at this time are the hardware-assisted cell projection algorithms [Shirley and Tuchman 1990; Max et al. 1990], and currently there have been proposals that show how to use a new class of programmable hardware to achieve even faster rendering rates [Wylie et al. 2002; Weiler et al. 2002]. Currently, we are using a custom built volume rendering engine [Corrêa et al. 2003], which works out-of-core and has been augmented with some functionality for streaming the time-varying data when needed. Our work follows along the lines of Lum et al [Lum et al. 2002], but we primarily target unstructured data.

On a 2.53Ghz Pentium4 machine with an Nvidia GeForce 4200, our code renders about 700-800 thousand tetrahedral per second, when not bottlenecked by disk I/O. We are currently exploiting several options for faster rendering, including the use of faster disks, and the use of programmable-hardware techniques [Wylie et al. 2002; Weiler et al. 2002].

3.2.3 Representations for specific data types

Visualization strategies were developed for bathymetry, salinity fields, velocity fields, and drifters. In all cases, the user can interact with the 3D model using the mouse and with some keyboard commands. With mouse movements the model can be rotated, translated and scaled in order to get appropriated images that allow the scientists to observe and better analyze the model's behavior.

On the other hand, keyboard commands present a set of tools like the capacity of playing the entire animation or going forward and backward through it step-by-step. Besides, it is possible to modify the color and transparency maps for the salinity and velocity fields. Those maps are lists of tuples that associate an RGB color and transparency to a salinity scalar value or to the norm of a velocity vector. The maps are stored as ASCII files and can be modified and reloaded at any time.

Another keyboard command allows the user to toggle between parallel and perspective projections and while in perspective it is possible to do virtual flight through the model.

Bathymetry

The bathymetry is the foundation layer of the body of water and is represented as a grid of triangles as described in section 3.1. We applied different colors to each depth level in order to enhance bathymetry visualization.

Salinity fields

For visualizing the scalar values of the salinity field, the CORIE output data is previously converted into an unstructured grid of tetrahedra and is visualized with direct volume rendering. The use of volume rendering allows for the study of the fine detail between high and low salinity regions (shown in blue and yellow respectively in Figure 8), and the interface region shown in red. Surface-based techniques, such as isosurfaces, would make it difficult to get an overall picture of the behavior of the river.

Velocity fields

The flow vectors for the velocity field are visualized as a set of oriented lines. All oriented lines have the same length, their colors represent vector magnitude and its orientation illustrates flow direction. The map of colors can be modified interactively as described before.

By choosing an specific 3D region of the velocity field as required by the user, it is possible to view all or just a small set of horizontal layers of the field. It is well-known to be very difficult to analyze the complete 3D vector field as a set of oriented lines, so we find it more convenient to visualize only a small subset of horizontal or vertical layers.

In the future, we plan to explore using more advanced vector visualization techniques, e.g., [van Wijk 2002; Telea and van Wijk 1999; van Wijk 1991; Interrante and Grosch 1998; Shen and Kao 1998].

Drifters

Each drifter released in the river has its corresponding computer simulated one and both are represented in the visualization as small spheres, each one with a different color. The path they follow along the time is drawn as line segments.

4 SCIENTIFIC IMPLICATIONS

By creating the ability to flexibly visualize 3D features, we dramatically increase our ability to model and understand estuarine and plume processes in the Columbia River.

The advantage starts with an improved visual understanding of the topology and bathymetry of the estuary and its insertion in the continental shelf (Figure 6). The macro-dynamics of the estuary is controlled by two deep channels, but a myriad of smaller channels plays an important role in the function of the ecosystem.

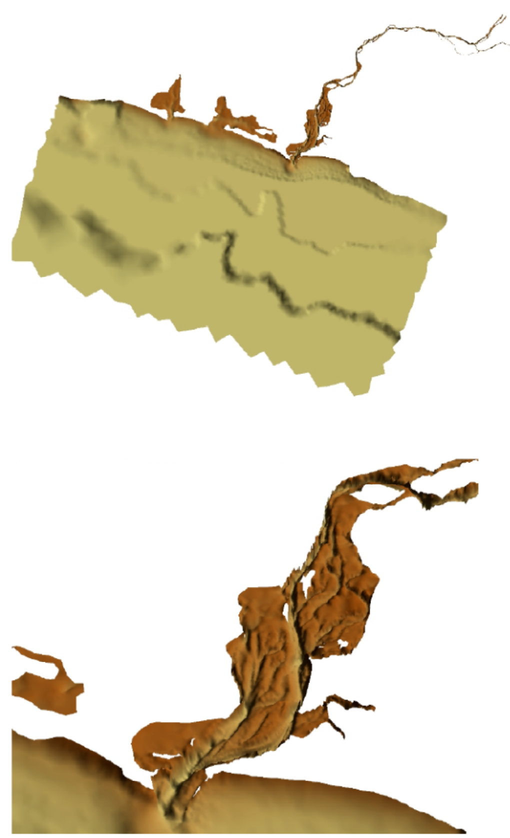

Figure 6: Columbia River topology and bathymetry (bottom) and insertion of the estuary in the continental shelf (top).

The split of the estuary into two main channels has substantial implications on the intrusion of ocean water, a determining factor of the ecological function of the system. As shown in Figure 7 (and corresponding animation[1]), the spatial structure of the intrusion is complex in both space and time. The ability to follow that structure

[1]A set of animations corresponding to the different visualization dis-

in visual detail is invaluable to evaluate the quality of model simulations, to gain insight on physical processes such as the formation of estuarine turbidity maxima, and in conveying this information to biologists and ecologists.

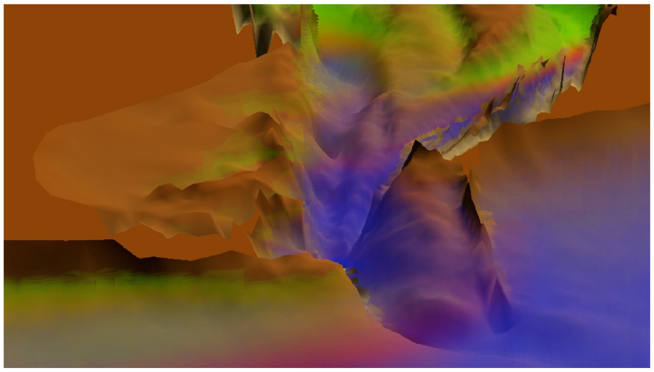

Figure 7: 3D detail of the salinity intrusion during flood tide, showing the split of ocean water entering the estuary between the Navigation Channel (at right) and the North Channel (at left). Ocean water is represented in blue, fresh water in green.

Salinity intrusion is only one of the “faces” of the water exchange across the mouth of the estuary. A complementary “face” is the creation of a dynamic plume of freshwater in the continental shelf, which near-shore manifestation is captured in Figure 8 (and corresponding animation[1]).

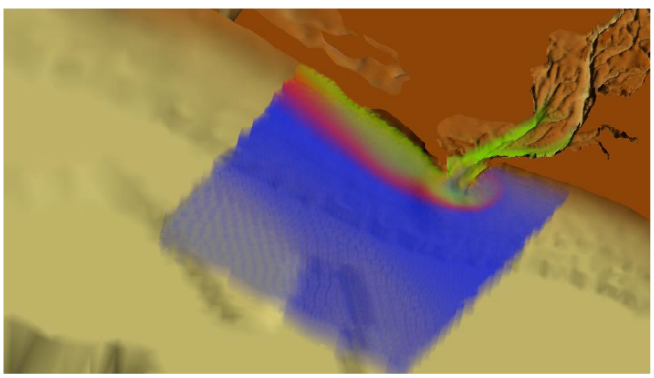

Figure 8: 3D detail of the freshwater plume during the ebb tide, showing a sharp density front (right side of the image).

Of particular interest in Figure 8 is a clearly visible density front, which identification is enhanced by the use of the color red in a preselected salinity range. Density fronts trap nutrients and plankton, becoming natural attractors for fish. Their dynamic nature poses difficulties to fisheries researchers, interested in sampling the distinct environments in each side of the front. Gradients of primary physical variables, such as salinity, are often useful in enhancing front identification (e.g., Figure 9).

The dynamic environment of the Columbia River provides a number of different trapping mechanisms, in addition to fronts. Of particular importance are eddies (Figure 10), which form at various locations and times during the tidal cycle, and are also evident in residual circulation fields (obtained by averaging instantaneous fields).

cussed in the main text are available in the proceedings DVD and from the project webpage.

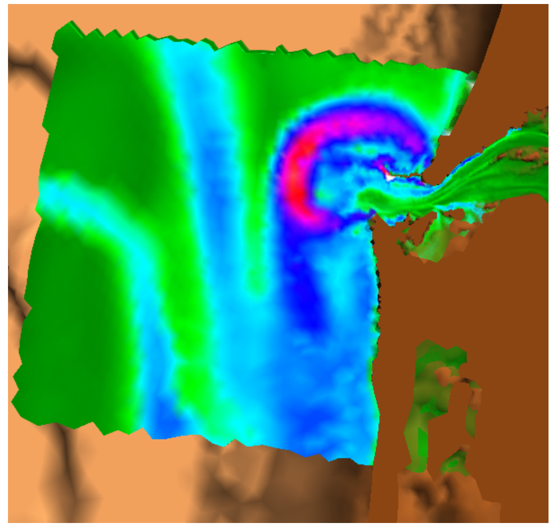

Figure 9: Maximum gradients of salinity reveal potential locations of ecologically significant fronts in the Columbia River plume.

Although circulation models add enormous insight to our understanding of the complex dynamics of the Columbia River, models are just a representation of the reality. Characterizing errors and uncertainties in this representation requires multiple approaches, including 3D visualization. An example is shown in Figure 11 (and corresponding animation[1]), where observed and simulated trajectories of a drifter released in the estuary are compared. The 3D animation makes it intuitively clear that observations and simulations remain close until the real and virtual drifters follow different channels in the upstream end of their progression into the estuary.

5 CONCLUSIONS

In this paper, we describe our initial efforts in closing the gap between the simulation capabilities of CORIE and its visualization tools. Although CORIE has sophisticated 3D modeling and simulation components, its visualization tools are mostly based on 2D tools that generate canned animations, and do not allow for direct “interactive visualization”. Here, we describe several new tools for looking at CORIE data that generate 4D (i.e., time-varying 3D) visualizations and allow for interactive exploration of the data.

There are important directions we intend to pursue in future work to both address limitations as well as extend the existing tools:

Deliver real-time frame rates: even though our visualization tools are efficient, they are not able to render the full resolution data at real-time frame rates.

Ubiquitous visualization platform: our tools are not machine-scalable, i.e., it is not possible to adjust their performance to the platform being used by a given user. Adaptability is especially important in the context of EOFS, since users might be out on the field, without access to high-end visualization machines.

Specification of visualization products: currently it is hard to assemble and manage complex visualization pipelines. We need to go through many steps, run several programs on a variety of machines to generate data products to be visualized. This process can be very

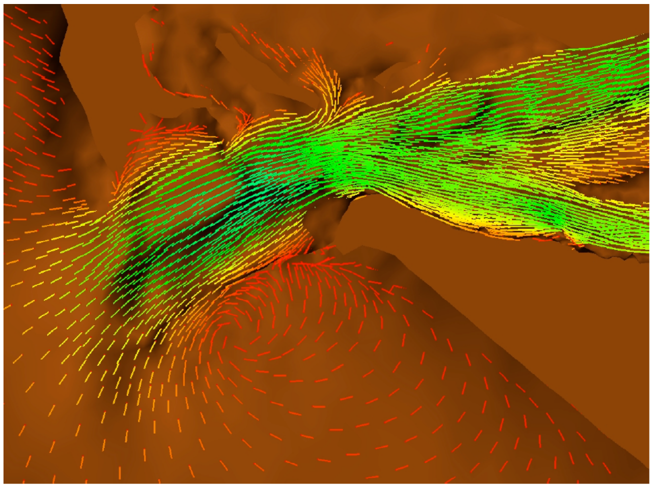

Figure 10: Instantaneous view of the velocity field during ebb tide. An eddy can be observed south of the entrance, in sharp contrast with the linear flow patterns in the channels.

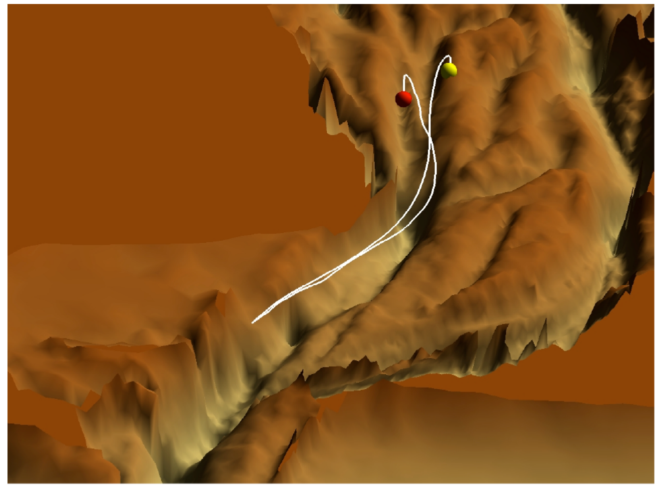

Figure 11: Observed and simulated trajectories of a Lagrangian drifter released near the mouth of the estuary at the beginning of a flood tide. The red sphere represents the real drifter, and the yellow the simulated drifter.

complex, time consuming, hard to manage manually, and it can be rather tedious especially if it has to be performed over and over again. This greatly hinders the ability of the scientist to perform complex and extensive analyses.

Acknowledgments

This work was partially funded by the National Science Foundation under grant ITR-ACI 0121475 and the Department of Energy and Sandia National Laboratory.

References

BAPTISTA, A. M., WILKIN, M., PEARSON, P., TURNER, P., MCCANDLISH, C., AND BARRETT, P. 1999. Coastal and estuarine forecast systems a multi-purpose infrastructure for the columbia river. *Earth System Monitor 9*, 3.

BAPTISTA, A. M. 2002. Modern paradigms for modeling and visualization of environmental systems. In *Encyclopedia of Physical Science and Technology*, Academic Press, R. A. Meyers, Ed., vol. 5, 565–581.

CORRÊA, W. T., JIMÉNEZ, W. H., AND SILVA, C. T. 2003. An out-of-core cell-projection technique for rendering large unstructured grids. Manuscript in preparation.

INTERRANTE, V., AND GROSCH, C. 1998. Visualizing 3d flow. *IEEE Computer Graphics & Applications 18*, 4, 49–53.

LAMAR, E. C., HAMANN, B., AND JOY, K. I. 1999. Multiresolution techniques for interactive texture-based volume visualization. In *IEEE Visualization '99*, IEEE, 355–362.

LEVOY, M. 1990. Efficient ray tracing of volume data. *ACM Transactions on Graphics 9*, 3, 245–261.

LUM, E. B., MA, K.-L., AND CLYNE, J. 2002. A hardware-assisted scalable solution for interactive volume rendering of time-varying data. *IEEE Transactions on Visualization and Computer Graphics 8*, 3, 286–301.

MAX, N., HANRAHAN, P., AND CRAWFIS, R. 1990. Area and volume coherence for efficient visualization of 3D scalar functions. *Computer Graphics 24*, 5 (nov), 27–33.

MAX, N., WILLIAMS, P., AND SILVA, C. T. 2000. Approximate volume rendering for curvilinear and unstructured grids by hardware-assisted polyhedron projection. *International Journal of Imaging Systems and Technology 11*, 1, 53–61.

MAX, N. 1995. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics 1*, 2 (jun), 99–108.

SHEN, H.-W., AND KAO, D. L. 1998. A new line integral convolution algorithm for visualizing time-varying flow fields. *IEEE Transactions on Visualization and Computer Graphics 4*, 2, 98–108.

SHIRLEY, P., AND TUCHMAN, A. 1990. A polygonal approximation for direct scalar volume rendering. In *Proc. San Diego Workshop on Volume Visualization*, 63–70.

SILVA, C. T., MITCHELL, J. S. B., AND KAUFMAN, A. E. 1996. Fast rendering of irregular grids. In *1996 Volume Visualization Symposium*, IEEE, 15–22.

TELEA, A., AND VAN WIJK, J. J. 1999. Simplified representation of vector fields. In *IEEE Visualization '99*, 35–42.

VAN WIJK, J. J. 1991. Spot noise-texture synthesis for data visualization. In *Computer Graphics (Proceedings of SIGGRAPH 91)*, vol. 25, 309–318.

VAN WIJK, J. J. 2002. Image based flow visualization. *ACM Transactions on Graphics 21*, 3 (July), 745–754.

WEILER, M., KRAUS, M., AND ERTL, T. 2002. Hardware-based view-independent cell projection. In *Proceedings of the 2002 IEEE symposium on Volume visualization and graphics*, IEEE Press, 13–22.

WILLIAMS, P. L. 1992. Visibility-ordering meshed polyhedra. *ACM Transactions on Graphics 11*, 2, 103–126.

WYLIE, B., MORELAND, K., FISK, L. A., AND CROSSNO, P. 2002. Tetrahedral projection using vertex shaders. In *Proceedings of the 2002 IEEE symposium on Volume visualization and graphics*, IEEE Press, 7–12.