

DataViewer: A Scene Graph Based Visualization Library

Randy Paffenroth
Applied and Computational Mathematics
California Institute of Technology,
California, USA.
email: redrod@acm.org

Thomas Stone
Veridian MRJ Technology Solutions,
Fairfax, Virginia, USA.
email: tstone@mrj.com

Dana Vrajitoru
Computer and Information Sciences,
Indiana University at South Bend,
Indiana, USA.
email: danav@cs.iusb.edu

John Maddocks
LCVM², EPFL, Switzerland.
email: John.Maddocks@epfl.ch

ABSTRACT

This article outlines the capabilities of a scientific visualization toolkit called DataViewer, and compares it to analogous software. DataViewer was originally designed for the construction of the visualization part of certain computational steering packages, and consequently it is particularly straightforward to closely couple DataViewer with numerical calculations. Rendering is performed through a high-level scene graph which facilitates the easy construction of complex visualizations. DataViewer differs from other such libraries by allowing complex geometrical objects, which efficiently encapsulate large amounts of data, to be used as nodes in the scene graph. Graphics hardware access is through the OpenGL API.

KEY WORDS

Scene Graphs, Scientific Visualization, Steered Computations

1 Introduction

DataViewer (<http://lcvmwww.epfl.ch/DV/>) is a scientific visualization library created to allow the easy development of efficient visualization programs for those who are not, and do not wish to become, graphics programming experts. DataViewer comprises a high-level set of routines for rendering geometric objects. It utilizes the OpenGL graphics API [1] for low level rendering. DataViewer has previously been used in several different visualization projects, e.g. most of the visualization modules of the parameter continuation software package VBM (Visualization of Bifurcation Manifolds, [2, 3], <http://lcvmwww.epfl.ch/VBM/>) have been developed using it.

Most of the goals that have driven the design of DataViewer are comparatively standard. It had to be computationally efficient, but with a high-level easy-to-use interface, portable (at least over UNIX platforms), and open source, freely available to the mathematical and scientific

research communities. A less standard design goal for DataViewer was to facilitate a close coupling between visualization and numerical computations, so as to be able to visualize dynamic data sets easily. By dynamic data sets we here mean data sets that are continually being modified as, for example, in interactively steered codes. Such dynamic data sets require a tighter coupling to the actual computation of the numerical data than is necessary in the more standard mode of visualizing a static, pre-computed data set *a posteriori*.

Rendering in DataViewer is performed through a high-level scene graph based structure which allows the user to construct complex visualizations simply. DataViewer is object oriented and provides a selection of leaf nodes that contain data for visualization, and container nodes that group together other nodes and which share properties. In addition to using only existing node types, more advanced users may easily add new node types to extend DataViewer's capabilities. DataViewer differs from other such libraries in that it allows both simple and complex geometrical objects to be used as single nodes in the scene graph. The complex nodes can be used to visualize large amounts of data more efficiently than is possible by using conglomerations of simple nodes, while the simple nodes can be used when flexibility is required. Section 2 will detail DataViewer's implementation of a scene graph and give examples of geometrical objects that DataViewer provides.

Many packages and algorithms have been developed for rendering three dimensional graphics on modern computers. Because our work requires highly interactive 3D graphics tightly coupled to computation, only certain approaches were of potential interest to us. For example, ray-tracing, e.g. POV-Ray (<http://www.povray.org>) performs well under many circumstances and can produce very realistic images, but is generally too slow and often inappropriately detailed for our purposes of interactive scientific visualization. Similarly our experiences with [3], and evaluations of, general purpose high-level visualization packages such as AVS (<http://www.avsc.com/>),

DataExplorer (<http://www.qbsssoftware.com/>), and Explorer (http://www.nag.com/Welcome_JEC.html) lead us to the conclusion that it was awkward to couple such programs tightly with computations. For the problems of interest to us it seems to be more efficient to visualize by embedding 3D graphics capabilities into a computational code, rather than to embed computational code into a general graphics package. In addition such large software packages as AVS, DataExplorer, and Explorer, typically do not meet our desiderata of low cost, and access to source code.

From all of the available packages for scientific visualization, VTK (<http://www.kitware.com>) represents the software package most closely realizing the DataViewer objectives. In spite of its many advantages, it was not designed for close coupling to numerics, and is in many ways much more sophisticated than our needs. Other software packages like AVS (<http://www.avs.com/>), Open-Inventor (<http://oss.sgi.com/projects/inventor/>), and GeomView (<http://www.geomview.org>), have been considered before starting the development of DataViewer, but none of them met our needs closely enough.

Hence in late 1996, we were left with the prospect of starting to create our own visualization software library. DataViewer 2.x was developed by R. Paffenroth, T. Stone, D. Vrajitoru, and A. Ahearn; it is our attempt to create a library which encompasses all of the needs detailed in the Introduction. The package is still being developed along with new applications in research and education.

This paper is structured as follows. Section 2 describes the design, implementation and specific features of DataViewer. Section 3 introduces some examples of DataViewer applications.

2 DataViewer 2.x Design

The first release, DataViewer1.0, was implemented based on C++ with Motif for the GUI. DataViewer 2.x differs from DataViewer1.0 in two main respects. First, DataViewer 2.x uses Python to create GUIs, which we feel makes implementations much simpler. Second, DataViewer 2.x encapsulates all low-level graphics calls. The user may know nothing about OpenGL and still be able to exploit DataViewer 2.x efficiently. Further detail on the DataViewer 2.x software library may be found at <http://lcvwww.epfl.ch/DV>.

2.1 Implementation

DataViewer 2.x is written entirely in C++, and uses many of the language's object oriented features, such as templates and virtual functions [4]. Each graphical object class in DataViewer 2.x inherits from a single base class called DVobject. Each DVobject defines a virtual *draw* routine that canonically contains a set of OpenGL commands that create the graphical representation of the object. In addition, each DVobject contains a set of properties, such as

color, translation, and rotation. These properties are discussed in greater detail below.

Python gives access to the Tk [5] widget library, which DataViewer utilizes for the design of GUIs. Tk is a library of widgets which can be accessed from a wide variety of languages, including Tcl, Python, C/C++, and Perl. It is a robust and mature code which has implementations on all Unix systems as well as Microsoft Windows and Macintosh. In our opinion, Tk is much easier to use than Motif and makes the GUI more flexible. For users with only a little familiarity with Tk and Python, it is quite straightforward to create customized DataViewer GUIs for their particular problems.

Finally, DataViewer 2.x provides several features that are accessible in any application which uses the library. For example, DataViewer provides stereo viewing capabilities, using hardware such as Crystal Eyes from StereoGraphics Inc, along with access to six degree-of-freedom controls, such as the Magellan 6D Mouse from Logitech Inc.

2.2 The Scene Graph and its Properties

The most important feature of DataViewer 2.x is its representation of the geometrical scene as a graph. Scene trees are a well-known and longstanding notion in computer graphics, either as a structural hierarchy of the view scene [6], or a spatial development of it [7, 8]. Although some recent research has expressed reservations concerning scene graphs [9], especially for photo-realistic rendering, they are still an efficient paradigm for scientific visualization, and we chose to implement them in DataViewer 2.x.

In a scene graph, a geometrical object is represented in a hierarchical fashion using a tree. A given object is divided into pieces, each of which is either a *leaf node* (i.e. a geometrical object which can be represented as an atomic data type) or a *group node* which itself contains other nodes. DataViewer 2.x implements a scene graph by using two main types of nodes.

The first type of node in DataViewer 2.x is the *geometry node*. Geometry nodes are the leaf nodes of DataViewer 2.x. They represent basic geometrical objects. DataViewer 2.x differs from most libraries that use scene graphs in that DataViewer 2.x provides both simple leaf nodes, and in addition leaf nodes that are quite complex geometrical objects. Such objects usually group any number of data to be translated into the same category of geometrical shapes and/or OpenGL primitives. Geometry nodes in DataViewer 2.x include sets of lines, cylinders, polygons, triangle strips, spheres and others. For example, Figure 1 is represented as two leaf nodes of DataViewer 2.x, a simple one containing the planar grid, and a more complex one containing a set of line segments rendered as cylinders with an associated ribbon.

While complex leaf nodes reduce the flexibility of a library, DataViewer 2.x includes complex objects for two reasons. First, in our work certain types of geometry nodes

are commonly used as metaphors for scientific visualization, and we wish these complex objects to be easy to use. For example, Figure 2 represents a surface computed as a surface of rotation generated by a periodic curve satisfying a certain system of ordinary differential equations. Second, each node in DataViewer 2.x has a certain amount of computational overhead imposed at each processing of the scene graph. For complex objects grouped into a single node, this overhead is low, while conglomerations of simpler objects increases the overhead. Concretely, many operations can be performed once for a whole set of basic geometrical shapes. Such operations include the definition of rendering properties such as color or data components to be displayed. In addition, properties such as rotation, translation, and scaling, are composed as the drawing advances in the depth of the tree. These operations take a non-trivial amount of computational effort, and the modification of some of them can even cause the OpenGL rendering pipeline to stall, causing a significant loss of efficiency. Thus, grouping basic objects into single leaf nodes as much as is possible both significantly improves the speed of translating a scene tree into OpenGL commands, and reduces the number of such commands.




Figure 1. A complex geometric object which can be represented in DataViewer 2.x by two leaf nodes. The tube formed from cylinders and ribbon are both included in the definition of the first leaf node. The planar grid forms the second leaf node.

The second type of node in DataViewer 2.x is the *container* node. Container nodes group together other nodes into a set sharing a number of common graphical properties, which may then be treated as a single node. The container node and the nodes that are attached to it are linked by a *parent - children* relationship. DataViewer divides container nodes into two classes. The first class of container nodes is represented by *normal containers*, which have the following functionality: their draw function calls the draw functions of all of the children by passing to them a composition of the geometrical properties received from higher in the tree and the container's own properties. The

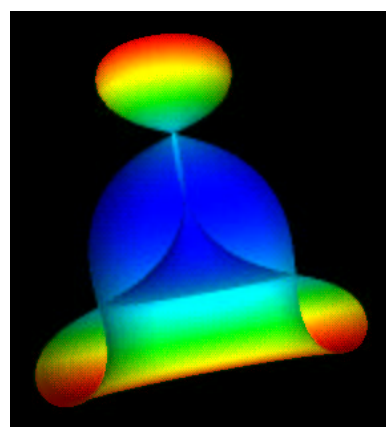


Figure 2. A complex geometric object which can be represented in DataViewer 2.x by a set of triangle strips in a single leaf node.

second class of container nodes includes several kinds of *selection containers*. Their draw function calls the draw functions of only some of the children. Generally, the selection containers are used for special effects, such as animation and level-of-detail.

A scene graph also requires some way to handle *properties*, such as color. For example, the developers of OpenInventor chose to implement properties as leaf nodes themselves. In other words, leaf nodes are not just geometrical objects, they may also be properties. A color node may appear in the scene graph, and each geometrical object is drawn in the color of the color node that is closest to it (depending on how the scene graph tree is traversed). While this approach is quite flexible, it does introduce an order dependence in the scene graph. Depending on the order in which the scene graph is rendered different results may occur. DataViewer therefore took the different approach of including the properties into the geometry and container nodes themselves. Each node may either set the value of a given property itself, in which case objects in the node are rendered using its own value of the property, or the node may inherit the property from its parent. An example of a group of spheres and a set of property definitions is shown in Figure 3. Using this type of property inheritance the scene graph is not order dependent, and there is only a small loss of flexibility. We remark that, subsequent to our development of DataViewer 2.x, Java3D adopted the same property structure.

2.3 Numerical issues

DataViewer was designed to be easily linked to numerical programs. This link can be achieved in several ways.

First of all, any numerical program written in c++ or FORTRAN can be directly integrated by *compilation* in a DataViewer application. This requires programming skills in one of these languages, but the number of lines of code




Figure 3. This figure shows a scene graph on the left and the graphical result of the scene graph on the right. The top container sets the color property to be green. The sphere attached to the top container sets the color property to red, and the first container attached to the top container sets the color property to blue. No other node sets the color property. The sphere marked with 1 is drawn in red since it sets that property itself. The sphere marked with 2 inherits its blue color from the container directly above it. The sphere marked with 3 inherits its green color from the top container.

that are necessary for the integration is minimal. About 5 lines of c++ are required to define a new DataViewer 2.x application. The rest of the needed work concentrates on communicating the data between the numerical code and the graphical objects. This part has been simplified by the implementation of the data structures in DataViewer. For example, an entire array of data can be transformed into a set of lines by one line of code. Here is an example of discretizing a 3 dimensional curve $C(x(s), y(s), z(s), a \leq s \leq b$:

```
LCVMarray_2D<float> data(number_of_points, 3)
step = (b-a) / (number_of_points-1);
for (int i=0; i<number_of_points; i++) {
    s = a+ i * step;
    data[i][0] = x(s);
    data[i][1] = y(s);
    data[i][2] = z(s);
}
DVlines curve(data);
```

The DVlines object translates the data into OpenGL commands creating a line strip, where a line segment links every point $data[i]$ to the next one $data[i+1]$. A parameterized surface can be represented as a discrete set of triangles just as easily.

An example of this kind of integration can be found as part of the VBM program.

Second, more complex stand-alone programs like AUTO [10] can communicate with DataViewer by means of data files. Using Python's system modules, one can run the numerical program from the DataViewer interface, even choose the parameters. This method has the advantage to separate the compilation of the numerical program from the compilation of the application, and is recommended in the case where the numerical program is too complex to be easily integrated at a compilation level into DataViewer, or if

the output of the program is already in the form of a file. It is also a good solution for anybody that is more comfortable with the programming in a scripted language like Python, rather than c++. It requires a file parser to be written either in c++ or Python for the specific output of the program. VBM also contains some examples of this kind of communication.

Third, the users that are not familiar with programming, or that prefer not to do this kind of effort, can use the scene file format to visualize data produced by other programs. This file format is specific to DataViewer and is described in Section 3.2. This is the easiest way of using the library, but it requires a minimal knowledge of the file format, and eventually a script to produce the data in the required form.

2.4 Animation

DataViewer 2.x provides 2 types of animation: *flip-book* and *key-framed*. Both types of animation are implemented using selection containers. Flip-book animation is achieved by simple selection containers that draw only one of their children for each frame. They occupy a relatively large amount of memory space, but are very flexible. In contrast key-framed animation is implemented using rotation and translation interpolation containers. They contain a list of rotation or translation properties corresponding to some particular frame numbers. Their draw function calls the draw functions of all of their children for each frame, but hands them a rotation or translation property that is interpolated according to the frame number. This kind of animation is less flexible, but occupies significantly less memory and requires less user effort to create.

Another feature of DataViewer is the possibility to combine any kind of normal and selection containers to create quite complex scenes. DataViewer 2.x is quite flexible and allows one to balance memory space restrictions and speed requirements when combining large static scenes with small animated objects.

2.5 Runtime Features

Every application written with DataViewer is automatically provided with a number of features that make the program interactive and easy to use.

The first category of runtime features concerns navigation through the 3D scene. With simple mouse and keyboard actions, the user can perform rotations, translations and zooms of the 3D scene. Other more complex actions, such as selecting particular objects from the scene for further manipulation, can be linked to the mouse and keyboard by each application.

The second category of runtime features concerns view and display manipulation via a menu bar associated with DataViewer's main window (see Figure 4). The menu options include control of 6D and stereo de-

VICES, macros for setting the viewpoint, exporting snapshots of the displayed image, etc. An important widget in DataViewer is the animation control that provides navigation through an animated scene and frame by frame view. It is possible to export any particular DataViewer scene to various 2D and 3D formats such as bitmap, POV-Ray (<http://www.povray.org>), etc.

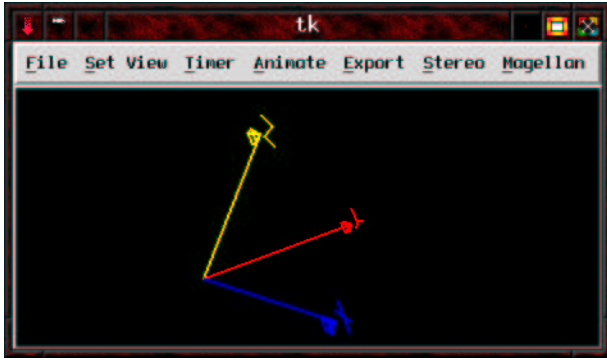


Figure 4. An example of DataViewer’s main window.

3 Applications

Several applications of DataViewer have already been implemented. We briefly mention three.

3.1 VBM

Visualization of Bifurcation Manifolds (VBM, <http://lcvmwww.epfl.ch/VBM/>) [2] is a software package whose goal is to provide tools for computing, manipulating, and visualizing bifurcation manifolds obtained by parameter continuation. As described in [3] VBM is a confluence of the packages MC² <http://lcvmwww.epfl.ch/visual.html> and PCR [11, 12], which represent much of our earlier, and more painful, experiences with graphics packages. VBM allows easy access to large data sets using special projections of the bifurcation diagram. It also provides direct selection of any particular solution, followed by its more detailed visualization in a separate window using a *data probe*. Most of the data probes in VBM have been written using DataViewer, Figure 5 shows an example of a VBM data probe. VBM has been developed by R. Paffenroth, J. Maddocks, R. Manning, D. Vrajitoru, and K. Hoffman.

3.2 Scene File Parser

The Scene File Parser by D. Vrajitoru (http://lcvmwww.epfl.ch/DV/Scene_file/) is an application of DataViewer allowing rapid development, modification and visualization of 3D scenes by means of a file parser (see Figure 6) even with no knowledge of programming.

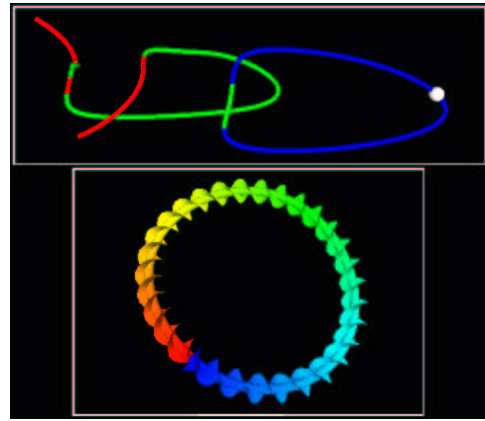


Figure 5. A runtime example of the software package VBM. The top window contains a bifurcation diagram in which each point on the curve represents a solution to a system of differential equations for some parameter values, with color indicating stability properties. The bottom window visualizes the solution corresponding to the point in the bifurcation diagram marked by a white ball.

As DataViewer can also export scenes to this file format, the scene file parser can be used as an input and output mechanism, to take 3D snapshots from DataViewer applications, and for debugging purposes.




Figure 6. An example of a geometrical object generated with the scene file parser.

This application defines a file format that provides special syntax structures for each graphical class in DataViewer and for most of the properties attached to them. This graphical file format uses especially minimal syntax requirements while keeping the structure of the scene clear. Thus, with minimal effort, the user can transform a raw data file produced by any other program, into a coherent scene file that can then be visualized with DataViewer.

For example, a file of raw data composed of real numbers on 3 columns and any number of lines, separated only

by spaces, representing successive 3D points on a curve, can be turned into a scene file by adding the keyword *lines* at the beginning of the file, followed by the character { , and the closing character } at the end of the file:

```
lines {  
  0 0 0  
  1 1 1  
}
```

The previous syntax defines a line segment from the point (0, 0, 0) to the point (1, 1, 1). The *lines* keyword is associated in general with the line strips concept from OpenGL. More details in the scene file documentation.

Following the ideas from the design of DataViewer, the data are grouped into large objects according to the graphical primitives they represent. This procedure decreases memory needs, streamlines the rendering, and minimizes the number of tokens in the scene file. The syntax also provides a tree structure similar to the one described in Section 2.2.

3.3 Slinky

DataViewer2.0 has been used in the latest implementation of a package called Slinky that was built to interactively explore solutions of an initial value problem for a system of ordinary differential equations describing an elastic ribbon. The ODE system governs the shape of an elastic rod that is used as a model of DNA (see Figure 7). The GUI allows the effect of changes in various coefficients and model parameters to be viewed interactively, with no external coding required.




Figure 7. An example of an elastic rod generated with Slinky.

Acknowledgments

It is a pleasure for the authors to be able to thank the many useful suggestions and criticisms from the past and current

members of the Laboratory for Computation and Visualization in Mathematics and Mechanics (or LCVM²) of the EPFL, and the former Laboratory for Computation and Visualization in Mechanics (or LCVM) at the University of Maryland.

References

- [1] J. Neider, T. Davis, and M. Woo, *OpenGL Programming Guide*, Addison-Wesley, 1993.
- [2] R. Paffenroth, *Mathematical Visualization, Parameter Continuation, and Steered Computations*, Ph.D. thesis, University of Maryland, 1998, <http://www.acm.caltech.edu/~redrod>.
- [3] J. H. Maddocks, R. Manning, R. Paffenroth, K. Rogers, and J. Warner, “Interactive computation, parameter continuation, and visualization,” *International Journal of Bifurcation and Chaos*, vol. 7, no. 8, pp. 1699–1715, 1997.
- [4] B. Stroustrup, *The C++ Programming Language*, Addison-Wesley, third edition, 1997.
- [5] J. K. Ousterhout, *Tcl and the Tk Toolkit*, Addison-Wesley, 1994.
- [6] S.R. Clay, “Put: Language-based interactive manipulation of objects,” *IEEE Computer Graphics and Applications*, vol. 16, no. 2, pp. 31–39, March 1996.
- [7] K.R. Subramanian and B.F. Naylor, “Converting discrete images to partitioning trees,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 3, no. 3, pp. 273–288, July–September 1997.
- [8] O. Sudarsky and C. Gotsman, “Dynamic scene occlusion culling,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 5, no. 1, pp. 13–29, January–March 1999.
- [9] W. Bethel, “Scene graph APIs: Wired or tired?,” in *SIGGRAPH’99 Conference Abstracts and Applications*. 1999, Computer Graphics Annual Conference, pp. 136–138, ACM SIGGRAPH.
- [10] R. Paffenroth, “The auto2000 command line user interface,” in *Proceedings of the 9th International Python Conference*, 2001., pp. 233–241.
- [11] G. Domokos and R. C. Paffenroth, “Case study: Visualization for boundary value problems,” in *IEEE Visualization ’94 Conference Proceedings*, R. D. Bergeron and A. E. Kaufman, Eds. 1994, pp. 345–348, IEEE Computer Society Press.
- [12] G. Domokos and R.C. Paffenroth, “PCR a visualization tool for boundary value problems,” WWW page, <http://lcvmwww.epfl.ch/PCR/>.