Case Study

3D Displays of Internet Traffic

Kenneth C. Cox and Stephen G. Eick*
AT&T Bell Laboratories

Abstract

The explosive growth in world-wide communications, especially the Internet, has highlighted the need for techniques to visualize network traffic. The traditional node and link network displays work well for small datasets but become visually cluttered and uninterpretable for large datasets. A natural 3D metaphor for displaying world-wide network data is to position the nodes on a globe and draw arcs between them coding the traffic. This technique has several advantages of over the traditional 2D displays: it naturally reduces line crossing clutter, provides an intuitive model for navigation and indication of time, and retains the geographic context. Coupling these strengths with some novel interaction techniques involving the globe surface translucency and arc heights illustrates the usefulness for this class of displays.

1 Introduction

Recently, stimulated by the explosive growth of the Internet, a large variety of world-wide networking data has become available for analysis. Besides the Internet, other sources of world-wide network data include international telecommunications traffic, financial flows, trading patterns, and national migration patterns [Tob87]. The key questions involving this data often relate to the structure of the information and how it varies through time.

At its most basic level the Internet or any network consists of nodes and links. The nodes and links may represent physical objects such as machines, or nonphysical objects such as hypertext pages. Statistics, possibly time-varying, may be associated with the nodes and links. These statistics may be raw measurements, such as the number of accesses of a particular home page in a particular time-period, the number of times a hypertext link is accessed, or computed aggregates, such as an average link utilization.

The most common way of visualizing network information is as a diagram with lines drawn between glyphs representing the nodes; see, for example, Figure 1. The color, thickness, line type or other visual characteristic may encode the link statistic and the glyph size or shape may encode the node statistic. Node and link displays work well for small, sparse networks with few nodes or links, but become cluttered and jumbled by line crossings for any reasonably-sized network.[1] The clutter comes from the long lines connecting distant pairs of nodes that cause overplotting. (For examples of particularly complex graph displays see [SW93].)

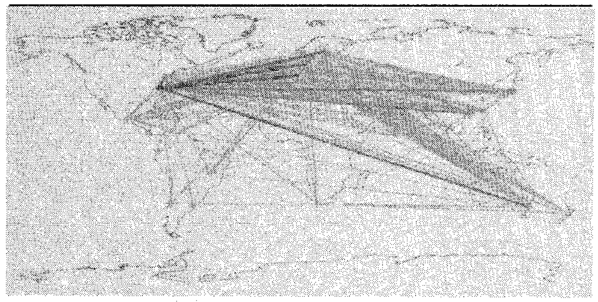

Figure 1: World-wide internet traffic over a two-hour period, with the color and thickness of the lines encoding the traffic.

There are several possible solutions to the display clutter problem:

- using curves, perhaps spline-based, instead of straight lines to connect distant nodes [GNV88],
- thresholding to show only the most important lines,
- shortening the lines from the middle, and
- using interactive techniques [BEW95].

Other approaches involve the node positioning:

- placing the nodes to minimize the overplotting [EW93],
- using fisheye-motivated, focus-and-context based techniques to distort less important distant regions [SB94],
- depending on the task [Cas91], and
- as a graph in 3D [FPF88].

Although none of these techniques is entirely satisfactory, the 3D displays for networks are appealing because in 3D the extra dimension may eliminate many

*Correspondence Contact: AT&T Bell Laboratories - Rm 1G-351, 1000 East Warrenville Road, Naperville, IL 60566, email: eick@research.att.com

[1] A small sparse network may contain a tens to hundreds of nodes and about the same number of links.

See Color Plates, page 152.

of the line crossings, and thereby reduce the clutter. General 3D network displays, however, are often confusing to navigate.

For world-wide networks there is a natural 3D embedding: position the nodes geographically on a globe and draw lines or arcs between them (see Figure 3). This results in a pleasing display, looking somewhat like international airline routes, that retains the spatial information associated with the nodes and also eliminates the line crossings associated with 2D displays.

The remainder of this paper explores globe-based network displays in detail. Section 2 describes how we encode information using glyphs and arcs. Section 3 discusses some techniques for interacting with the display. Finally, Section 4 discusses our results.

2 Internet Display

Figure 3 shows one frame from an animation of Internet traffic between fifty countries over the NFSNET/ANSnet[2] backbone for one two-hour period during the week of February 1-7, 1993. The dataset contains the packet counts, by two-hour period, between each pair of countries.[3]

Each country is represented by a box-shaped glyph that is both scaled and colored to encode the total packet count for all links emanating from the country. The glyphs are positioned at the locations of the countries' capitals and extend perpendicular to the surface of the globe. The color-coded arcs between the countries show the inter-country traffic, with the higher and redder arcs indicating the larger traffic flows. The globe is illuminated by a light which is positioned to indicate the angle of the sun for the frame of the time-series data that is displayed.

2.1 Surface

Drawing a world map on the surface of the sphere in Figure 3 converts it into a globe, thereby providing spatial context for the location of each of the nodes. Our map contains only the continental outlines, avoiding possibly excessive detail that would obscure network information.

The surface of the globe in Figure 3 is an opaque blue and thus obscures those portions of arcs and nodes which lie *behind* and *within* the sphere. By interactively varying the translucency of the surface the user may control how much of the display is obscured.

2.2 Nodes

The boxes on the surface of the globe in Figure 3 encode the node statistic by scaling in the radial direction. The visual effect is pleasing; the boxes appear to be small towers standing on the sphere, with the tower height and color of the tower tied to the statistic. The largest boxes correspond to the nodes having the largest statistics, thereby focusing attention on the important nodes. Other glyphs and data encodings are possible. We have experimented with cylinders and pyramids, negative scalings where the glyph descends toward the center of the sphere, and encoding of two values using the glyph height and position above or below the surface. Our preliminary results show that the latter technique of information encoding can be readily understood by viewers.

2.3 Arcs

The arcs are the analog of the lines in the traditional 2D node and link displays. They touch the spherical surface at each end and reach a maximal radial height in the center. Tying the height of the arc to the link statistic ensures that the most important arcs corresponding to the largest values of the statistic are always visible and never obscured by arcs corresponding to lesser values.

In extending the straight lines of the 2D network displays to the arcs of the 3D display, we have many choices for the paths that the arcs may take. A "straight line" on a sphere is a great circle, as is used in Figure 3. More generally, for any projection of the sphere to a plane, the straight line between two points in the plane corresponds to a path for an arc between the points on the sphere. The *polar* projection generates the great circle paths. Figure 2 shows arcs drawn as straight lines using a rectangular coordinate system (i.e., one in which latitude and longitude are treated as X and Y coordinates in a Cartesian system) and as great circle paths, and Figure 4 shows the Internet data drawn using rectangular paths. We have found the great circle paths to be both more attractive and easier for the viewer to interpret, since the arcs are usually spaced over a larger portion of the globe.

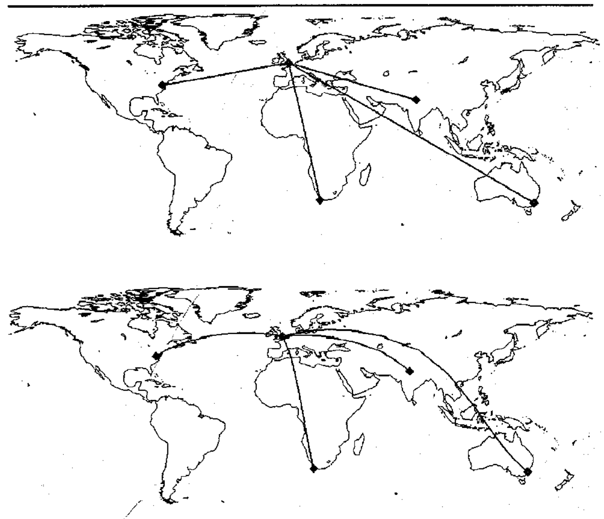

Figure 2: Rectangular (top) and great-circle (bottom) paths.

[2]Thanks to Hans-Werner Braun at the San Diego Supercomputer Center and the NSFNET partnership for making the data available. It is in /pub/scsc/anr/data, available by anonymous ftp from ftp.sdsc.edu.

[3]50*50*84 = 210,000 total counts.

Proceedings of the Proceedings on Information Visualization (INFOVIS '95)
0-8186-7201-3/95 $10.00 © 1995 IEEE

The arcs are scaled to reach a maximum height above the surface according to the link statistic. In our implementation we scale according to the function:

$$R(1 + H \sin(x\pi))$$

where H is the maximum height, R the globe's radius, and x varies between 0 and 1 along the arc's path. Thus the arc's elevation above the surfaces reaches a maximum H half way between the two endpoints. If the height H is negative, the arcs will lie within the sphere, and thus the arc will be obscured unless the surface is translucent.

3 Interactions

By dynamically manipulating the network display it is possible to obtain greater insights into the data. Rotating the sphere (our implementation uses spaceball like interface) enables a user to see where arcs terminate that are obscured by the sphere. Animation shows the evolution of time-varying network data, and allows the user to move back and forth through the time series. Each figure in this paper is one frame from a time series.

Besides the usual 3D geometric interactions, the sphere lends itself to some novel controls. For example, the globe may be used as a clipping surface to obscure glyphs or arcs within it. By tieing the calculation of radial heights to controls such as sliders, the user can cause some of the glyphs and arcs to be drawn inside the sphere and thus obscured as shown in Figure 5.

In conjunction with such thresholding, it is convenient to allow the user to interactively vary the opacity of the sphere, thereby partially or wholly reveal the glyphs and arcs inside the sphere. This can also be used to show the nodes and arcs on the reverse side of the sphere that would otherwise be occluded, thus aiding in navigation (see Figure 6).

4 Discussion

Displaying networks in 3D is a useful mechanism for solving the display clutter problem with 2D displays. The difficulty with general 3D network displays is that they are often confusing, difficult to navigate around, and cause the user lose a sense of global context. Restricting the display to a sphere captures many of the advantages of 3D network displays while simultaneously helping the user maintain context.

Users are accustomed to globes, so navigation is simplified and there is little chance of the user becoming disoriented. The number of line crossings, and hence the amount of visual confusion, is also reduced by the three-dimensional embedding and by the presence of the surface, which acts as a background. Incorporating user interface controls such as thresholding and translucency can further reduce the visual complexity of the display, and thereby lead to greater insights.

References

[BEW95] Richard A. Becker, Stephen G. Eick, and Allan R. Wilks. Visualizing network data. *IEEE Transactions on Visualization and Graphics*, 1(1):16–28, 1995.

[Cas91] Stephen M. Casner. A task-analytic approach to the automated design of graphics presentations. *ACM Transactions on Graphics*, 10(2):111–151, 1991.

[EW93] Stephen G. Eick and Graham J. Wills. Navigating large networks with hierarchies. In *Visualization '93 Conference Proceedings*, pages 204–210, San Jose, California, 25-29 October 1993.

[FPF88] Kim M. Fairchild, Steven E. Poltrock, and George W. Furnas. Three-dimensional graphic representations of large knowledge bases. *Cognitive Science and Its Applications for Human Computer Interaction*, pages 201–233, 1988.

[GNV88] E. R. Gansner, S. C. North, and K.-P. Vo. Dag–A program that draws directed graphs. *Software–Practice and Experience*, 17(1):1047–1062, 1988.

[SB94] Manojit Sarkar and Marc H. Brown. Graphical fisheye views. *Communications of the ACM*, 37(12):73–84, December 1994.

[SW93] Michael F. Schwartz and David C. M. Wood. Discovering shared interests using graph analysis. *Communications of the ACM*, 36(8):78–89, 1993.

[Tob87] Waldo R. Tobler. Experiments in migration mapping by computer. *The American Cartographer*, 14(2):155–163, 1987.

Proceedings of the Proceedings on Information Visualization (INFOVIS '95)
0-8186-7201-3/95 $10.00