3D Geographic Network Displays

Kenneth C. Cox, Stephen G. Eick*, Taosong He
Bell Laboratories
e-mail: {kcc,eick,taosong}@bell-labs.com

Abstract

Many types of information may be represented as graphs or networks with the nodes corresponding to entities and the links to relationships between entities. Often there is geographical information associated with the network. The traditional way to visualize geographic networks employs node and link displays on a two-dimensional map. These displays are easily overwhelmed, and for large networks become visually cluttered and confusing. To overcome these problems we have invented five novel network views that generalize the traditional displays. Two of the views show the complete network, while the other three concentrate on a portion of a larger network defined by connectivity to a given node. Our new visual metaphors retain many of the well-known advantages of the traditional network maps, while exploiting three-dimensional graphics to address some of the fundamental problems limiting the scalability of two-dimensional displays.

1 Introduction

With the explosive growth of networks and widespread availability of the World Wide Web, huge volumes of information are now accessible in networked databases. Now that people are gaining widespread access to data warehouses, another problem is emerging: how to extract the information latent in the data. Unfortunately, our ability to extract knowledge from large databases has not kept pace with the increasing data volumes.

One approach for extracting useful information from the data stores involves *information visualization*. Information Visualization exploits this capability of human visual system by encoding data using color, shape, position, texture, motion, etc., and rendering the result on a graphics workstation. The fundamental problem in Information Visualization involves inventing the visual metaphors for representing the non-physical data. Unlike the visualization of data about physical phenomena, where a natural representation of the data is available in almost all cases, abstract information usually has no natural form or representation, and the literature and techniques are correspondingly less well-developed. The visual *metaphors*, that is, the representations for making abstract relationships visible, are only now emerging [15].

Graphs are recognized as an important visual metaphor for many different types of information [12] [13] [10]. Graphs consists of nodes and links; there are often statistics associated with the nodes and links. Many types of information may be represented by graphs with the nodes corresponding to entities and the links corresponding to relationships among the entities. The properties of the graphical objects used to draw the graph, such as the color and size of glyphs or the thickness of lines, can be used to encode statistics about the nodes and links. The interesting aspects of a graph often involve its topology, structure, and connectivity, and the positioning or layout of the nodes and links can often be exploited to emphasize these properties.

Our interest and motivation for visualizing graphs comes from analyzing many networks associated with Bell Laboratories. In our usage, a network is a graph where the associated statistics represent *traffic*. The traffic may represent a time-varying statistic, as with the call flow in a telecommunications system throughout the day; a constant, as with network capacity; or a stochastic statistic, as with the number of IP packets sent between routers on a backbone data network.

Networks often have a geographic component, and some of the properties of the networks cannot be understood without reference to this component. The most common technique for visualizing geographic networks involves node and link diagrams [4]. Glyphs (graphical objects) representing the nodes are positioned on a geographic map, with lines drawn between them showing the connections between nodes. The color and thickness of the lines may be used to represent the traffic, with the thicker and brighter lines showing the links carrying the most traffic, with the greatest capacity, and so forth. The glyphs may be colored, shaped, and sized to encode statistics associated with the nodes, for example, the router capacity, utilization, and packet losses.

Node and link diagrams are effective for visualizing small sparse networks containing ten to a hundred nodes. For larger and richer networks, however,

*Correspondence contact: Room 1G-351, Bell Laboratories, 1000 East Warrenville Road, Naperville, IL 60566,

the displays are easily overwhelmed and become visually cluttered with too many line crossings. One approach to solving this problem involves graph layout algorithms. However, they are often not suitable for geographic networks, where the position of the nodes cannot be varied without losing the geographic context. Our focus, then, is on layouts that preserve geographic context.

2 3D Network Displays

One way to solve the clutter problem inherent in 2D network displays is to draw the network in 3D. The idea is that by positioning the nodes in 3D and drawing arcs instead of links, we can eliminate the line crossings that confuse 2D displays [14]. Of course when viewed from any particular angle certain links may appear to cross on a 2D computer screen. The advantage, however, of the 3D representation is that through our preattentive depth perception we will automatically perceive the display correctly.

There has been previous research into 3D network displays; for some of the early papers see [8] [6]. We build on this work and differentiate our results in three important ways. For **geographic context** we concentrate on 3D networks layouts that maintain geographic context. As we indicated above, for many information spaces involving spatial locations, maintaining the geography in the visualization aids in the understanding. For **drill-down network views** we present a suite of novel views that show details on demand for particular user-designated nodes. Finally, we consider the **restricted navigation**. In general navigating a 3D network display is difficult and a user can easily become lost and disoriented. In our views we have carefully restricted the user's ability to navigate and thereby prevent disorientation, while simultaneously striving to maintain the perceptual advantages of 3D layouts.

2.1 Global Networks

The first technique, initially described in [5], positions the nodes geographically on a globe and draws lines or arcs among them. This results in a pleasing and informative 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 top left of Figure 1 shows one frame from an animation of Internet traffic between fifty countries over the NFSNET/ANSnet 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. Each country is represented by a box-shaped glyph that is scaled and colored to encode the total packet count for all links emanating from the country. The glyphs are positioned at the locations of

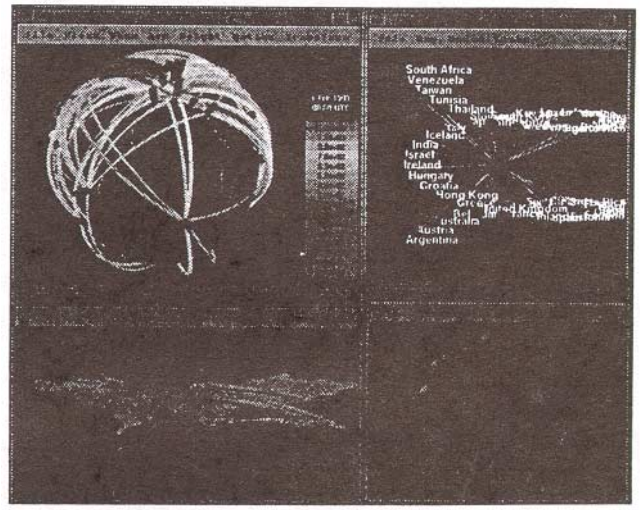

Figure 1: SeeNet3D network visualization system.

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, via the angle of the sun, the time for the frame of the time-series data that is displayed.

The difficulty with general 3D network displays is that they are often confusing and difficult to navigate around, and cause the user to lose a sense of overall context. Restricting the display to a sphere captures many of the advantages of a general 3D network layout while simultaneously helping the user to maintain context. Users are accustomed to globes, so navigation is simplified since the user may rotate the globe interactively 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 globe surface, which acts as a background.

One weakness of drawing arcs around the globe to show traffic is that only the front side of the globe is visible, making it impossible to see where some arcs terminate. We attempt to overcome this occlusion in three ways. The first uses an interactive translucency control that makes the globe partially transparent. Translucency works well for some examples and some viewing orientations, but may still be visually confusing if there are a large number of arcs curving around the globe terminating at many different locations. The second allows the user to select from several different ways of routing the arc paths, including routes that run through the middle of the globe. Appropriate selection of the path can reduce the display clutter, and the routing of paths through the globe is effective when used in combination with the translucency control. Finally, the user may perform filtering of the arcs to select only those with certain attributes. Incorporating user interface controls such as filtering and translucency can further reduce the visual complexity of the display, and thereby lead to greater insights.

2.2 Arc Maps

Another related technique draws arcs among nodes positioned on a flat 2D map embedded in the 3D space (see bottom left of Figure 1). Each node is positioned geographically, as on the globe, and then arcs are drawn between the nodes with the height, color, and thickness of arc encoding the statistic. As with the global display, drawing the arcs in 3D eliminates the line crossings that curse the 2D displays.

Arc maps have several desirable characteristics. Firstly, unlike the global network display discussed before, the geographic context of the network is not restricted to world-wide networks. Arc maps can display a network on any map at any desired resolution. The retention of geographic context makes the displays interpretable. Secondly, although the geographic map is drawn on a 2D plane, it can be arbitrarily positioned and oriented in the space. Therefore the user can interactively navigate the display by translating, rotating, scaling, and visualizing the network from different perspectives under different rendering conditions. Thirdly, as with the global display, drawing arcs greatly eliminates the line crossings associated with planar 2D displays. Finally, the most important links, such as those having the largest traffic flows, are represented by the highest arcs and therefore are visually predominant from different angles.

By parameterizing the height using a slider, as we do in our implementation, a user may interactively and smoothly transit between a traditional 2D node and link map and a 3D arc map. To further reduce the visual cluttering of the display, we have also incorporated another two options into the arc map. Firstly, the glyphs representing the country encode the total packet count for all links emanating from the country using both color and size, drawing visual attention to the most important countries. Secondly, the arcs can be rendered with translucency, thereby avoiding any occlusion of otherwise hidden arc segments. Figure 2 presents the arc map display of internet traffic with these two options.

2.3 Drill-Down Network Views

The visual network display techniques discussed above attempt to show a complete network, with all links and nodes visible simultaneously. For many analysis tasks, however, users need to *drill down* and obtain more details about a particular node or subnetwork. For example, in Figure 2 the largest glyph, colored red and corresponding to the United States, is clearly interesting. Where do the traffic flows from the United States go? Which country has the highest community of interest with the United States? The lowest? From Figure 2 it is difficult to answer fine-grained detail-oriented questions about a particular node, even after rotating and adjusting the viewing parameters.

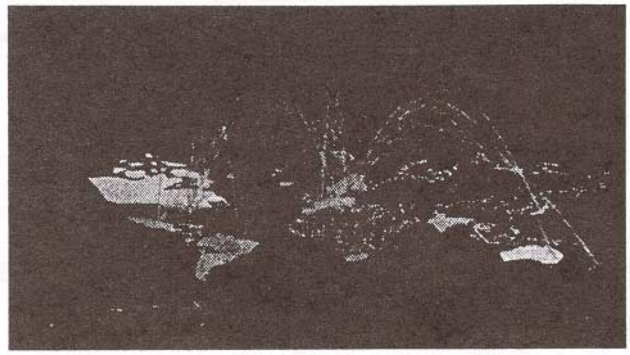

Figure 2: A Partially translucent arcs map showing the world-wide Internet traffic..

Linked views showing "details on demand" are common in the data mining community and well established in the statistical graphics community. This subsection describes three types of network-oriented drill-down display for displaying all links emanating from a designated *focal node*.

The first view, shown in bottom right of Figure 1, looks somewhat like the spokes on a wheel. This spoke view shows all traffic between the USA (the central node) and each other country. The spokes are color-coded and the nodes size-coded to represent the link traffic. In this particular example the nodes are positioned alphabetically in a circle, but the nodes could be ordered according to geography or even placed in geographic position, for example on a polar projection map centered on the focal node. The latter view could also be considered as a version of the arc map in which the graph has been filtered to show only the focal node and the links and nodes directly connected to it.

Spoke displays work well for a small number of nodes (less than 50 or 100) but eventually become overwhelmed as the number of nodes increases and the glyphs begin to overplot. The reason is that spoke displays do not make effective use of screen real-estate because all lines are the same length. We can partially overcome this problem using a 3D layout. The helix view generalizes the spoke view by positioning the nodes on a 3D helix (see top right of Figure 1). When viewed from above, the helix view is a spoke view. By interactively rotating the helix view all nodes come into view sequentially.

Another 3D layout, motivated by the helix display, positions the nodes approximately uniformly around a sphere as with a pincushion (see Figure 3). In this view the nodes are actually arranged along circles of "latitude", with the number of circles and number of nodes per circle chosen so the angle between circles and between nodes on a circle are approximately the same. An alternate placement algorithm tessellates the surface of the sphere and selects points from the tessellation. To be effective, the pincushion view (like the helix view) needs to be viewed interactively with motion.

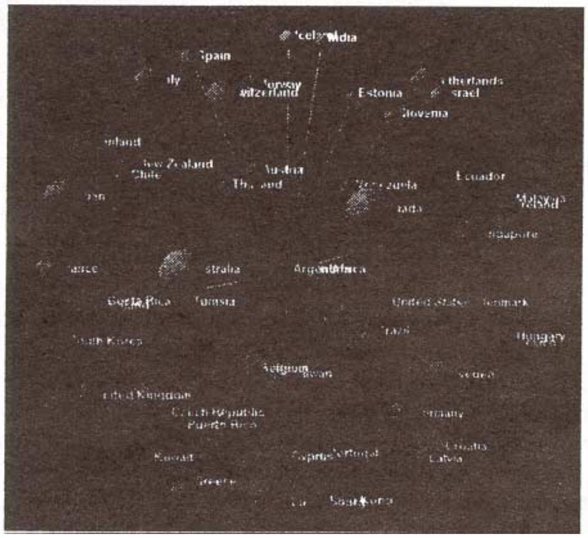

Figure 3: The pincushion view positions the nodes approximately uniformly on the surface of a sphere.

3 Discussion

The figures in this paper are screen images created from the SeeNet3D network visualization system. This system, developed over the last few years, is a comprehensive network visualization designed for visual exploration of large, time-varying communication networks. SeeNet3D is the most recent member of a series of network visualization and exploration environments that we have developed in an on-going research program aimed at visual network data analysis. Earlier members include SeeNet [3] and NicheWorks [7].

SeeNet displays time-oriented geographic network data using a suite of 2D displays and overcomes the display clutter problem using *dynamic parameter adjustment*. The analyst manipulates the display parameters interactively while watching the display change; good parameter focusing is achieved when the display shows meaningful information about the data. The contribution of SeeNet involves the careful choice of orthogonal dynamic controls and the appreciation that interactivity and dynamic manipulation may usefully address display clutter.

The NicheWorks system is most useful for abstract networks where there is no natural geographic layout. It uses a two-pass "spring-based" algorithm with several heuristics that the user may fine tune interactively. The essential idea in the positioning algorithm is that all nodes repel with unit force and the relationships among them form a counter-balancing attraction. Numerically solving the motion equations results in an incremental positioning that places related nodes close to each other [9]. For sparse networks, placing related nodes together has the very desirable effect of reducing the display clutter caused by long connecting links. NicheWorks embeds the placement algorithms in an interactive data manipulation, control, and exploration environment that exploits many techniques from dynamic statistical graphics [2].

The new contribution of SeeNet3D involves the use of 3D geographic node positioning. Ongoing technology trends in processor speeds and display hardware are making the creation of interesting and useful 3D network displays on desk-top equipment feasible. Such 3D network displays, in our opinion, are more engaging and visually interesting than the 2D displays that were previously standard. There is also accumulating evidence that 3D network displays are more effective than 2D displays [1] [16]. SeeNet3D attempts to exploit the engaging aspects of 3D displays while simultaneously using 3D layouts to overcome the fundamental problems inherent in displaying large networks.

The global network view captures the inherent geography in world-wide networks and displays network traffic using a natural and understandable metaphor. It scores well on the engaging criteria and has been used by other researchers for such purposes as visualizing MBone Internet traffic [11]. The arc map view captures many of the benefits of the globe while retaining the advantages of the traditional 2D network displays. Users are able to find the most pleasing viewing angles by interactively manipulating the display. This view may also be easily incorporated into other 3D network display systems, and is not restricted to world-wide networks. The linked drill-down views show all traffic from a designated node. In our implementation the user mouses on an interesting node in either the global or arc view to set the focal node in the drill-down views. The spoke view is conventional and works well for nodes with degree up to perhaps 50 or 100. For nodes with more connections the spokes become too close. The helix and spherical views place the nodes in 3D and use rotation and human perception to address the limited screen space problem and effectively support much larger fanouts than the spoke view. In our experience both the helix and sphere layouts are more effective than the spoke layout, although we prefer the helix layout perhaps because it is more ordered.

4 Implementation

The SeeNet3D system is currently a 5,000 line C++ program built on top of the *Vz* framework. Vz is a visualization platform embodied in an object-oriented, cross-platform (MS Windows, OpenGL, and X11) C++ library. The Vz library provides a foundation for building highly-interactive, linked-view graphical displays. It hides platform and operating system differences; handles display rendering in a portable manner; provides a standard "look and feel"; facilitates the view linking;, and includes many utility classes for data management, statistics, and mathematics. As the foundation for data visualization, the library provides the core and common functions in our system and tools.

We have version of SeeNet3D running on SGI workstations and PCs with MS Windows 95 and Windows NT. On a top-end personal computer (150MZ processor) with graphics accelerator we can render the display in less than a second. This is not yet quite fast enough for interactive performance, which we find requires at least four to five image refreshes per second.

5 Summary

This paper describes five views for showing complex geographic networks embodied in the SeeNet3D system. SeeNet3D is a network visualization environment for the exploration of large, time-varying communication networks. The research focus is on how to overcome some fundamental problems in understanding large and complex networks using 3D graphics technologies, while simultaneously maintaining the benefits of the useful and well-established 2D node and link maps. Two of the views display the complete network. The other three show subnetworks and may be used to "drill down" for details of all links emanating from a designated node.

By applying 3D technologies, we have successfully provided more visual information in a more understandable and engaging manner. Visually, our 3D views avoid some major problems of 2D views and present the graph in a clearer and cleaner way. Equally important, the users have more capability to understand the network through interactive manipulation of the display.

Information Visualization is an important field with great potential. In this paper we have only addressed a few of the issues associated with the 3D display of geographical networks. Many fundamental problems remain to be solved before the whole picture of 3D Information Visualization can be drawn.

Acknowledgments

We gratefully acknowledge and appreciate Audris Mockus's help in providing some of the maps used in our views.

References

[1] K. W. Arthur, K. S. Booth, and C. Ware. Evaluating 3D task performance for fish tank virtual worlds. *ACM Transactions on Information Systems*, 11(3):239–265, 1993.

[2] R. A. Becker, W. S. Cleveland, and A. R. Wilks. Dynamic graphics for data analysis. *Statistical Science*, 2:355–395, 1987.

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

[4] J. Bertin. *Graphics and Graphic Information Processing*. Walter de Gruter & Co., Berlin, 1981.

[5] K. C. Cox and S. G. Eick. 3D displays of internet traffic. In *Proceedings Information Visualization Symposium*, pages 129–131. Atlanta, Georgia, October 1995.

[6] I. F. Cruz and J. P. Twarog. 3D graph drawing with simulated annealing. In *GD'95 Proceedings Lecture Notes in Computer Science*, volume 1027, pages 162–165. Springer-Verlag, 1996.

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

[8] Kim M. Fairchild, Steven E. Poltrock, and George W. Furnas. Semnet: Three-dimensional graphic representations of large knowledge bases. In R. Guindon, editor, *Cognitive Science and Its Applications for Human Computer Interaction*, pages 201–233. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1988.

[9] T. Fruchterman and E. Reingold. Graph drawing by force-directed placement. *Software-Practice and Experience*, 21(11):1129–1164, 1991.

[10] John Lamping and Ramana Rao. Laying out and visualizing large trees using a hpyerbolic space. In *UIST: Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 13–14, 1994.

[11] T. Munzner, E. Hoffman, K. Claffy, and B. Fenner. Visualizing the global topology of the mbone. In *Information Visualization '96*, 1996. To appear.

[12] G. G. Robertson, S. K. Card, and J. D. Mackinlay. Information visualization using 3D interactive animation. *Journal of the ACM*, 36(4):56–71, 1993.

[13] M. Sakar and M. H. Brown. Graphical fisheye views. *Communications of the ACM*, 37(12):73–84, December 1994.

[14] L. Spratt and A. Ambler. Using 3D tubes to solve the intersection line representation problem. In *IEEE Symposium on Visual Languages*, pages 254–261, 1994.

[15] E. R. Tufte. *Envisioning Information*. Graphics Press, Cheshire, Connecticut, 1990.

[16] C. Ware and G. Frank. Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics*, 15(2):121–140, April 1996.