

# Visualizing Corporate Data

---

Stephen G. Eick

Daniel E. Fyock

Visualization is an emerging technology for understanding large, complex, information-rich data sets. Just as spreadsheets revolutionized our ability to understand small amounts of data, visualization is revolutionizing the way we understand large data sets. AT&T has developed a suite of applications, based on a common software infrastructure, to analyze strategic data sets and solve key business problems. Even as these software tools are being used internally, AT&T is also selling them in the commercial marketplace. The four case studies presented in this paper demonstrate the technology's general applicability and its use within AT&T to address strategic business problems and motivate its guiding research principles.

## Introduction

During the last decade, businesses have experienced an explosive growth in their ability to generate, collect, and store data. This has occurred because of advances in technology, including the availability of larger, cheaper disks; the widespread use and standardization of database management systems; the ubiquity of bar codes; and the computerization of business and government. As businesses, both large and small, become more data driven, data is being recognized as a strategic corporate asset. And as people gain widespread access to data warehouses, another problem is emerging: How can information hidden within the data be extracted? The ability to understand this latent information and to act on it before competitors do lies at the heart of successful competition.

In the past several years, the Visualization Research Group has developed a number of novel techniques and software systems that depict large quantities of data visually. *Visualization* presents information in pictorial form and uses human recognition capabilities to detect patterns within it.<sup>1,2</sup> It links the two most powerful information processing systems—the human mind and the modern computer. Humans have evolved a sophisticated visual system, and they excel at

pattern detection. Researchers at AT&T have exploited this ability by encoding data using characteristics such as color, position, and texture to render the result on a graphics workstation.

AT&T's visualization systems present users with multiple representations of their data embedded in views. Each view is an interactive environment that enables users to explore and analyze data. Within the confines of the data set being analyzed, users can control visual characteristics such as color, shape, and zoom level. They can also link multiple views to observe how the data are interrelated. By visualizing data, users gain a better understanding of their "core business," enabling them to improve their business processes and make better informed business decisions. The target user of visualization is either a knowledge worker struggling with data or a worker who needs to gain insight from a large data set.

## Knowledge Discovery Using Visualization

Traditional tools for analyzing data, such as spreadsheets, ad hoc queries, statistical analysis, and summaries, are no longer sufficient for the volumes of data that businesses want to analyze. Spreadsheets cannot

---

handle massive data sets. Ad hoc queries lack task support and have no facilities for repetitive tasks. Statistical analysis, though powerful, is a reduction technique and, as such, may obscure important information. Furthermore, because most users are not highly trained statisticians, they lack the insight that business experts have for analyzing large amounts of data and using those results to solve business problems. Because of aggregation, summaries, such as statistical analyses, may mask information hidden in individual data elements. Visualization can overcome all these problems.

*Knowledge discovery* is the overall process of extracting insights from data. The process involves:

- *Problem definition*, delineating the overall framework for the discovery,
- *Data access*, retrieving the data from its warehouse,
- *Data cleaning*, ensuring its quality and guaranteeing its integrity,
- *Data mining*, extracting information,
- *Data presentation*, showing the results to the decision makers, and
- *Business impact*, specifying how the discovery has resulted in a business decision, modified process, or confirmed result.

The discovery process uses many technologies, including database management, statistics, machine learning, and artificial intelligence.

As the following case studies illustrate, visualization is used throughout the analysis process, from exploring the initial data, to formulating the intermediate model, and, finally, to presenting the results. Visualization can show what data are available and pinpoint *outliers*—things out of the norm—that need cleaning. It also enables humans to mine information using their pattern perception, can communicate the findings efficiently and effectively, and enhances the other technologies by engaging the user in the analysis process.

#### Case Studies

This section describes four visualization case studies within AT&T—international calling fraud, communities of interest in local calling, customer retention, and program differences.

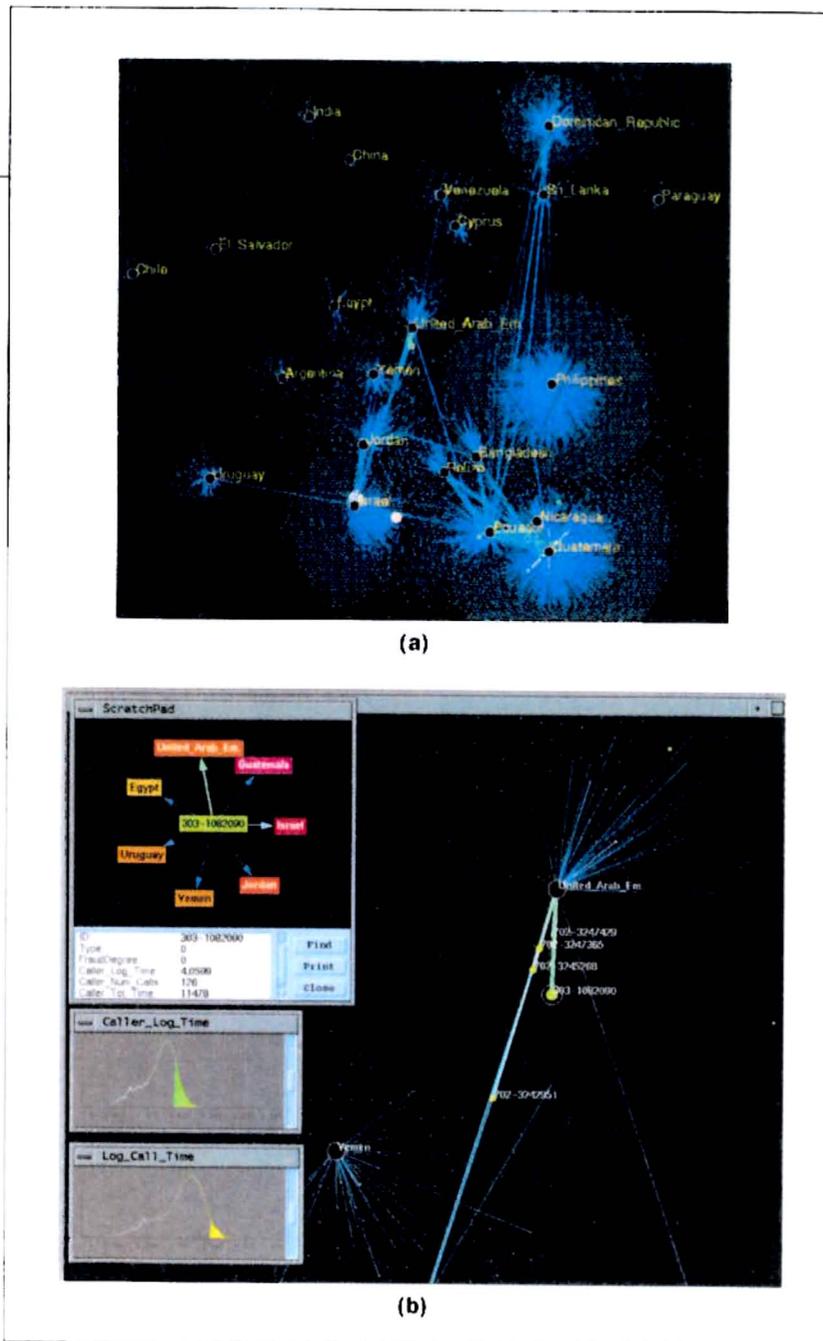
**International Calling Fraud.** Telecommunications fraud is an industry-wide problem. It is also an important issue for AT&T. Each day AT&T handles millions of inter-

national long distance phone calls, the vast majority of which are legitimate. Although the percentage of fraudulent calls is small, this number still represents many thousands of calls, constituting a significant expense and loss of revenue to AT&T. Legitimate customers pay for losses owing to fraud in two ways—through higher tariffs, and through inconvenience when facilities are tied up by fraudulent activity. Detecting and preventing fraud is a critical business challenge for AT&T.

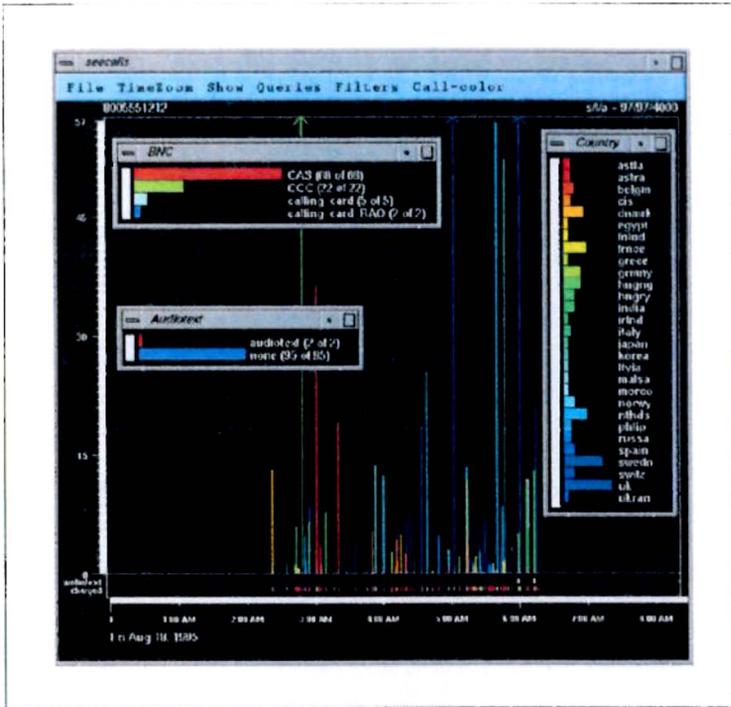
The fraud problem, simply stated, is how to identify the few fraudulent calls within the vast majority of legitimate calls. On a daily basis, network switching and monitoring systems collect gigabytes of data, only a small portion of which represents fraudulent activity. To make matters worse, fraud is dynamic; as soon as fraudsters realize they are in danger of being detected, they begin inventing techniques to circumvent the security measures that have been put in place. Furthermore, because international long distance calls cross national boundaries and encompass multiple network service providers, it is extremely difficult to obtain the complete end-to-end data needed to identify repeat offenders.

The visualization approach for detecting international calling fraud involves displaying customers' calling activity to uncover unusual calling patterns. Figure 1a shows international calling patterns during an eight-hour period to 23 countries frequently targeted by fraudsters. The node size and color both encode the total number of calls made by a customer; unfilled nodes correspond to countries and filled nodes to customers. Links originating from the nodes encode the distribution of the calls made to each particular country. Several unusual patterns that are readily visible became candidates for further investigation. In Figure 1b the user has zoomed in on a suspicious calling pattern, later confirmed as fraud, for a customer who made 126 calls from the U.S. to various Middle Eastern countries.

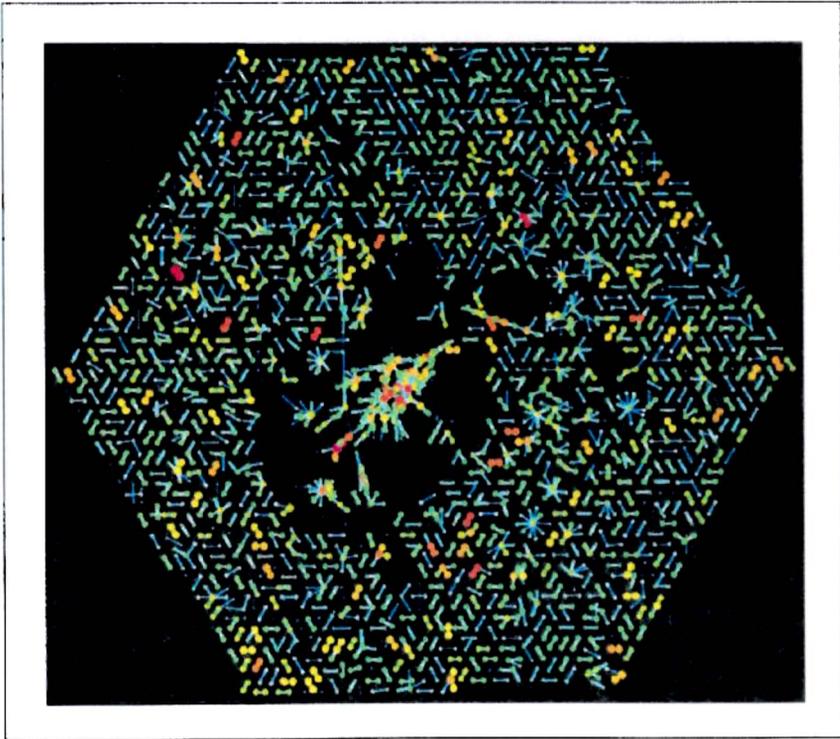
Another example of fraud is third-party billing fraud. Figure 2 shows 97 third-party calls made during a four-hour period. These calls, billed to others, were made to destinations around the world from a single multiline phone system with



**Figure 1. (a) A NicheWorks display showing eight hours of international calling. The nodes represent customers, and lines from the nodes to the countries encode the total number of calls. The large brightly colored nodes are unusual, and further investigation revealed that some calls were fraudulent. (b) A suspicious calling pattern, later confirmed as fraud, for a customer making 126 international calls during a holiday weekend.**



**Figure 2.** Each vertical spike represents an international call that was billed to a third party. During this four-hour period, 97 of the calls made were subsequently determined to be fraudulent.



**Figure 3.** NicheWorks visualization of local exchange calling patterns. Each customer making or receiving a call is represented by a node, with lines between the nodes depicting individual conversations. The nodes are positioned to show calling patterns, with the cluster in the center representing callers who all talk to each other. The time spent talking is coded by a node's color and size. The link colors code the length of individual conversations.

originating number 8001231234 (disguised for purposes of this paper). The horizontal axis, or x-axis, represents time; each vertical spike represents a call and its duration. The three longest calls, represented by the green and dark blue spikes with arrows at the top, lasted 102, 227, and 333 minutes, respectively. The green spikes indicate calls to Hong Kong, dark blue to Sweden, red to Australia, and so on, as shown by the color key on the right. The bar plot shown at the top left represents the distribution of calls by billing type: for example, five calls were billed to a calling card.

The "visual signature" shown in Figure 2 is characteristic of fraud. Many of the calls overlapped, and nearly all the calls were charged to third parties. Of these calls, two were to offshore Audiotex services, which typically are sex lines, a favorite target for hackers.

**Communities of Interest in Local Calling.** Identifying customers who are candidates for purchasing advanced services is a common marketing problem for telephone service providers. Services such as call waiting, conference calling, and call forwarding can be highly profitable to the service provider. For many specialized services, the cost of marketing to all customers is prohibitive, making the service uneconomical. To be successful, the provider must identify customers with calling profiles that will gain high value from these services, and market directly to these customers.

This case study describes a local phone company that was investigating a new telephone features package and wanted to target potential customers for the service. Figure 3 shows the four-hour calling pattern for customers in a calling area. Each customer is represented by a node, with the color and the size of the node coding the total usage and number of calls, respectively. The lines between the nodes indicate calls between customers; a line's thickness and color coding represent the total talk time. The nodes are positioned algorithmically to demonstrate communities of interest.<sup>3</sup>

Figure 3 shows clear patterns. The center cluster represents a group of customers who call each other frequently; this often denotes a virtual telecommunications business section of a calling area. The peripheral patterns represent customers who are connected less frequently, perhaps residential customers. Some customers have star-shaped calling patterns—nodes with several radial links—indicating that they are acting as

calling hubs. Households with teenagers, which often exhibit this type of calling pattern, can be targeted as likely candidates for a second teen line or call waiting.

**Customer Retention.** In the telecommunications industry, analysts estimate that over 1 million customers per month switch long distance carriers. To manage its customer base, AT&T must be able to understand, track, and characterize customer migration patterns (that is, who is leaving, who is joining, what services they use, and so on).

The market managers responsible for AT&T's customer base have access to a wealth of tracking data. This data, often presented in numeric tables, is difficult to understand. Visualization can help by providing a dynamic interface that displays patterns and relationships among these data in a graphical format. Figure 4a shows retention patterns, banded by revenue, for a sample data set of 1,152 customers, starting in October 1993.

At the left of the screen in Figure 4a, a smoothed banding distribution of the average monthly revenue for the sample set is delineated by color-coded revenue thresholds. The rows of bars, color coordinated with their respective revenue segments of the smoothed distribution, portray the retention during the next 11 months. The number of customers retained is displayed at the base of each bar. For added appeal and clarity to the user, the height of each row of bars is scaled according to its respective calling revenue segment in the smoothed distribution. In this sample, customers who spend between \$10 and \$40 per month are leaving AT&T at the fastest rate. The option buttons along the top of the display give the various types of calling revenue and retention time intervals, enabling the user to modify the display interactively.

The categorize view, Figure 4b, breaks down the individual average monthly service revenues for the 1,152 customers shown in Figure 4a by call type (for example, domestic and international) and by AT&T product (that is, direct dial, credit card, or operator assisted). The detailed charts are grouped according to margin totals for each row and column.

**Program Differences.** Corporations are accumulating large databases of text that contain hundreds of thousands of documents. Of particular interest to AT&T are text databases that store computer source code for AT&T products and services. To merge related files, computer users must understand the similarities and differences between them. For example, a maintenance programmer



(a)



(b)

Figure 4. Retention rates by (a) calling revenue thresholds and (b) both call type and product for a sample data set of 1,152 customers who joined AT&T in October 1993.

fixing a recently introduced bug may want to compare the newer version of a program with its older working predecessors to help locate a problem. On UNIX<sup>®</sup> operating systems, *diff*, the standard tool for comparing files, is one of the most frequently executed commands. Diff is used not only for comparing source code and other ASCII files, but also as the engine for many other system commands.

Although extensive algorithmic work has been conducted to calculate program differences efficiently, little attention has been paid to helping users understand the differences. Two types of comparison are important: many differences between two files, or a few differences among many files. Figure 5a is a visual display of file differences with two interactive views, one for side-by-side differences between two files, and another for comparing many files within directories. Differences are color-coded to highlight text that has been added (green), deleted (red), changed (yellow), or remained the same (grey). Users may navigate and browse the differences using the compact graphical representations shown in Figure 5a.

The text files in this figure come from a study looking for code that is common and unique within two versions of a software product. Figure 5a compares one subsystem comprising over 2,000 files across 80 modules. The larger background window is a file browser showing the side-by-side file differences. The two text areas are synchronized to align the text identified as common to both files. Text is scrolled by manipulating the controller in the middle. The center of the text browser displays a reduced textual representation of the two files, showing where the differences occur.<sup>4</sup> This is useful for identifying code that has been moved or copied within a file. The smaller window displays differences by directory level across the top, all files within one level of the directory in the middle, and side-by-side comparisons of all files within one level at the bottom right. A selector on the left allows the user to focus solely on deleted, added, or changed code. Files may be selected from this view and displayed in the file browser window.

In production software systems, the source code is kept in version control databases (such as *rcs*<sup>5</sup> and *sccs*<sup>6</sup>), which contain a complete history of the source code changes. These databases are a rich resource for identifying frequently changed, bug-prone, overly complex, and problematic code. Figure 5b shows *change hot*

*spots* in a particular directory—areas of code that have been modified repeatedly over time. Six versions of the file are depicted, taken every two months from July 1994 through May 1995. Each file is represented by a rectangle whose height encodes the number of lines in the file. Through time this file monotonically increased in size, as is common with software. The color indicates how frequently the code in the files was changed; bright spots indicate areas of code that have been modified repeatedly, and darker areas represent more stable code. The scale on the left maps the number of times a piece of code was modified to a color scale. The visualization clearly shows that one area of code was replaced several times throughout the process. This is problematic; it indicates that the code might benefit from reengineering or might require stringent testing.

#### Software and Technology

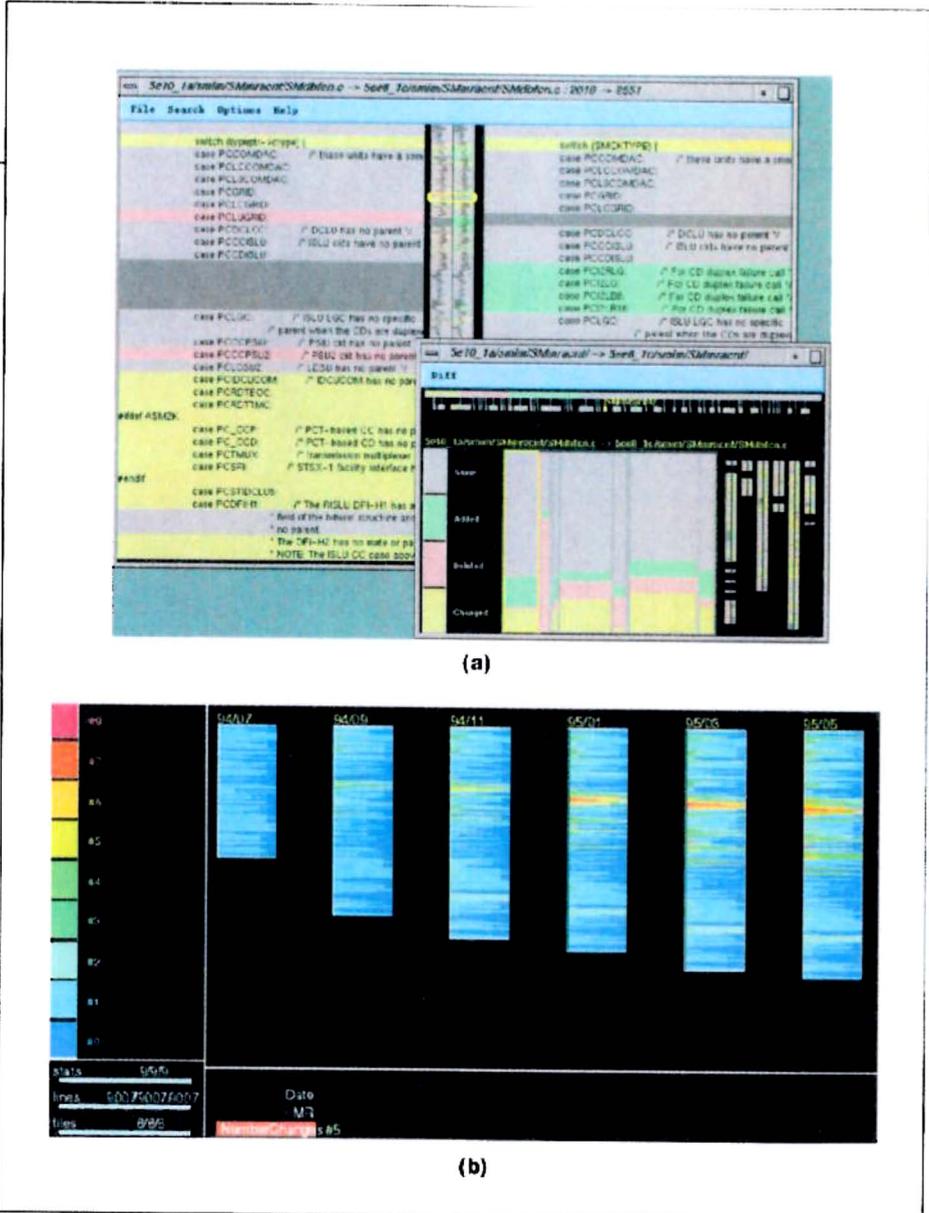
AT&T's visualization technology consists of about a dozen interactive views, each tuned to display one class of data. There are views for showing:

- Tabular data,
- Relationships using node and links display (see Figure 1a),
- Geographic networks,
- Hierarchical structures,
- Market segments,
- Customer retention (see Figure 4a),
- Text (see Figure 5a), and
- Univariate statistical views.

Each view has a common look and feel. A *linking* mechanism enables interactive operations in one view to propagate to the others. Each view functions both as a display for conveying information and as a control panel.

Applications are created by combining one or more views and populating them with data. Users then apply some of the more interesting applications within the suite to solve business-critical problems. These include:

- SeeSoft<sup>®</sup>, text files and source code;
- SeeSlice, dynamic program slices and code coverage;
- SeeData, relational database integrity constraints;
- SeeLog, computer log files and audit trails;
- NicheWorks, communities of interest in networks;
- SeeNet, geographic networks;
- SeeCalls, customer calling signatures;
- Segmentation Spreadsheet, market segments;



**Figure 5. (a) Differences between two versions of a software product. The larger window shows the synchronized display of two compared files. The smaller window shows the directory structure and amount of change in each file. (b) The monthly change history of a source code file. Each column represents a version of a file in time. Yellow and red spots indicate areas in the code that have been modified repeatedly.**

- SeeRetention, customer retention;
- SeeSales, inventory movement and process/personnel productivity; and
- SeeDiff, text file differences.

Underlying all these views is a common software substrate embodied in an object-oriented, cross-platform (Microsoft Windows,\* OpenGL,\* and X11), C++ library that handles the interaction and graphics. The V2 C++ library hides platform and operating system differences, handles display rendering in a portable manner, contains placement and identification algorithms, and provides a standard "look and feel." It also facilitates view linking, enables printing and session management, and includes many utility classes for data management, statistics, and mathematics. The library currently contains about 15,000 lines of code, and as the foundation for data visualization, it provides the core and common functions in its system and tools.

#### Guiding Principles

During the past several years, in the process of handcrafting many views for particular data sets, AT&T researchers have formulated some design guidelines for building novel, effective, information-rich visualizations of business data. These heuristic guidelines are firmly rooted in practical experience, are linked to perception, and have enabled AT&T researchers to create many innovative, effective displays of data. They have been designed to produce systems that allow a user to discover relationships in the data that enable them to take action based on discovery.

**Task Focus.** Because the needs of each user are unique, the best visualizations are task-oriented; they help frame, as well as answer, interesting questions. Building a successful visualization involves understanding the tasks and analysis goals of users, as well as incorporating their domain knowledge.

**Domain-Specific Representation.** The representation determines how the items in the data set are rendered on the computer display. The best representations are domain specific, for example, scatter plots for statistical data, maps for spatial data, and node and link diagrams for network data. Inventing a representation for a new domain is a difficult, creative, and iterative process. Representations should take full advantage of perceptual cues (for example, size, position, color, and depth),

motion, and perhaps even sound. By understanding user tasks, the representations try to provide local detail in a global context.

**Color.** Color processing in the human vision system is an independent perceptual process, making color a natural choice for encoding information. Using color to show details allows additional information to be layered onto the display. Colorful displays also engage the user and are visually pleasing, increasing the appeal of the system.

**High Information Density.** Representations using compact, color-coded glyphs pack a lot of information into an image and thereby display a large data set. A high-resolution 1280- × 1024-pixel workstation monitor contains more than 1,310,720 pixels. Using every pixel to display data maximizes the information content of the visualization.

In some cases, it is possible to display an entire data set on a single screen, thereby eliminating the difficult navigation problems associated with panning and zooming interfaces that focus on small portions of a data set. However, information-dense displays often become cluttered with too much detail.

**Interactive Filters.** One approach to solving the problem of display clutter is to use interactive filters that reduce the amount of information shown on the display. Humans have sophisticated pattern recognition capabilities, owing, perhaps, to evolution, and are very efficient at manipulating interactive controls to reduce visual clutter. Because humans can effortlessly solve complex computational problems involved with determining when a display is too busy for easy interpretation, interactive filters are a natural interface, a good match for the innate ability of people.

**Multiple Linked Views.** The power of a visualization system's representations is magnified through the use of interaction and linked views. Whether a view is custom or standard (for example, color keys, bar charts, box plots, histograms, or scatter plots), it functions as both a display and a control panel. Selecting and filtering data in one view instantly propagates to the other views, thereby providing additional insights. Linking multiple views interactively provides an integrated visualization far more powerful than the sum of the individual views.

**Drill Down.** When users discover interesting patterns, they need to access the actual data values. As a user touches any item in the display with the mouse, the

data values associated with that item are displayed; mouse clicks are not necessary.

**Animation and Motion.** Many commonly encountered business data sets have a temporal aspect. Animation, in which each frame represents a single time period, is an ideal tool for analyzing large, time-oriented data sets. Human perception is fine-tuned for motion detection, making animation a natural technique for scanning large data sets. However, for animation to be effective, the frames must change smoothly and continuously, as they do in a motion picture. Big or unexpected changes are jarring and stand out perceptually, resulting in a loss of continuity.

#### **Commercialization by QUEST**

To take advantage of the full potential of the data visualization prototypes, both within AT&T and in the commercial marketplace, a partnership was formed between two AT&T Bell Laboratories organizations—Research and Quality, Engineering, Software, and Technology (QUEST)—and several AT&T business units, including Capital Corporation, Network Systems, and Global Information Systems. QUEST has assumed key responsibility for technology transfer, new enhancements, distribution, and marketing of a high-quality product.

#### **Future Directions**

The role of the Visualization Research Group in its partnership with QUEST is to ensure a full pipeline of new, state-of-the-art views and innovative techniques that can be commercialized by QUEST. Research is currently investigating five areas of great potential: three-dimensional (3D) representations, sound, geography, time-oriented views, and built-in pattern recognition.

**3D Representations.** For business data visualization, 3D has been underexploited, but it is an active research area. The depth dimension is a natural way to pack more information onto the screen without overloading it. Eventually, after 3D representations of abstract data are fully explored, it may be possible to increase the information content—the ratio of information to pixels<sup>7</sup>—by a factor of 10. It is difficult, however, to find natural embeddings of nongeometric 3D data that are easier to interpret than their two-dimensional (2D) counterparts.

Perceptually, it is possible for 3D displays to increase the information density beyond that possible

with 2D displays by enabling our minds to create *virtual pixels*. For example, to a person watching a pedestrian walking behind a tall picket fence, the exact position of the pedestrian is clear, even though his or her position can only be seen in the gaps between the pickets. From the fragmentary input and motion, our minds create the position. When this idea is applied to network visualization, however, users become visually confused by line crossings depicting too many complex 2D node and link displays. If, instead, a 3D representation is used to visualize the same network, it becomes easier to distinguish the links passing behind or in front of one another.

Figure 6 shows one frame from an animation of a 3D network visualization of worldwide Internet traffic.<sup>8</sup> The height, color, and curvature of the arcs between nodes encode the packet counts between routers.

**Sound.** Sound is another underexploited medium for encoding data.<sup>9</sup> There are many aspects to sound—such as pitch, timbre, and loudness—and many ways that sound can encode data. Sound is fundamentally different from a visual display in four ways:

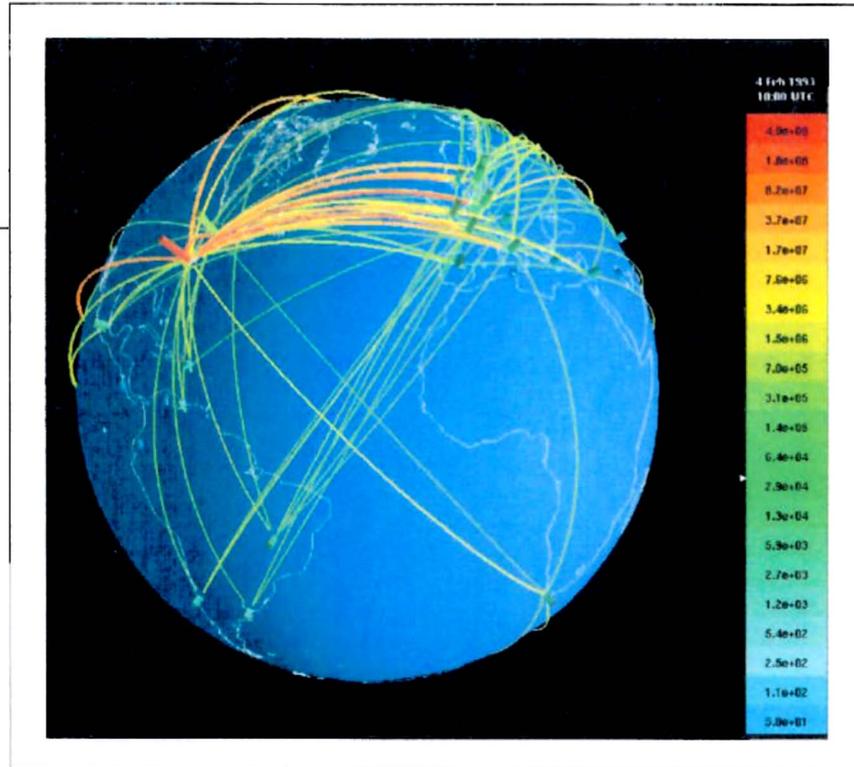
- It arrives through an independent channel;
- Its bandwidth is lower than that of vision;
- It is immediate instead of persistent, as with visual representations;
- It is serial in time.

Sound works well for alerting, particularly for monitoring tasks, with its ability to cut through visual clutter. Because much communication is conducted through speech, using voice to convey detail is natural. In animations, researchers have found that using voice to announce the passage of time is effective because it enables users to concentrate on the data display without having to monitor the current time.

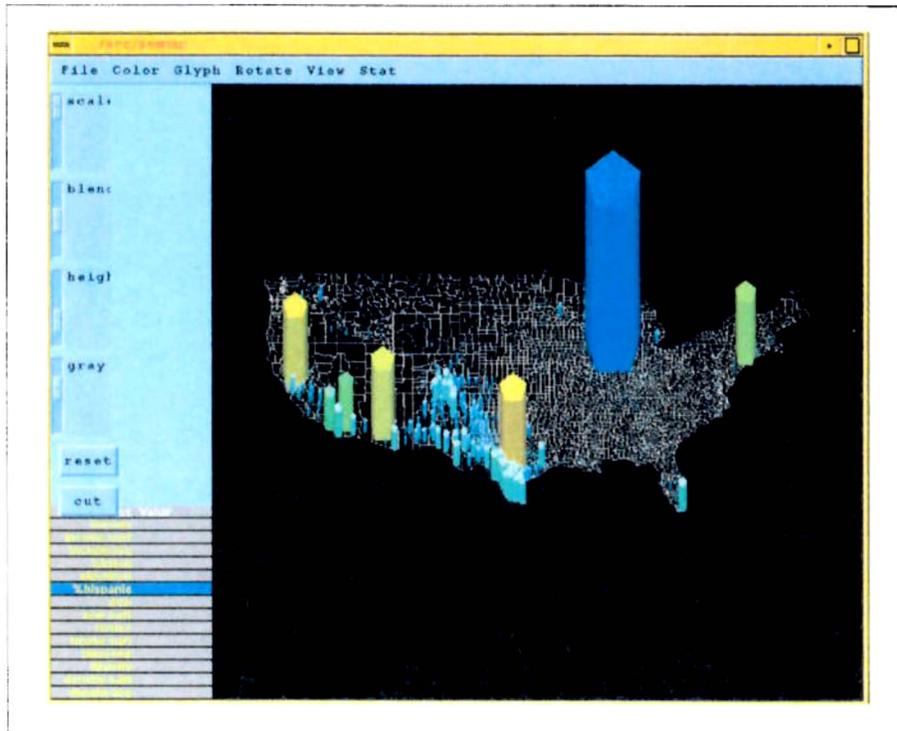
**Geographic Representations.** Many corporate data sets have a geographic component, such as demographic information. Developing spatial views with geographic information seems to be a promising direction for future research. Figure 7 shows an experimental 3D view of geographic data.

**Time-Oriented Views.** Many business data sets evolve through time, suggesting that it is advantageous to tailor the views in relation to time. Although AT&T researchers have developed several time-oriented 2D and 3D views, it is possible to use other, more sophisti-

**Figure 6. One frame from an animation showing worldwide Internet traffic over a two-hour period.**



**Figure 7. A prototype view showing the percentage of the population that is Hispanic, by county, using 1994 census data. The large glyph represents Cook County in Illinois, which includes the city of Chicago and many of its immediate surrounding suburbs.**



cated methods to take advantage of this naturally occurring phenomenon.

**Built-In Pattern Recognition.** AT&T researchers are trying to build automated tools to allow a user who identifies an interesting pattern to detect all other instances of the pattern automatically and to feed this information to other programs for action.

### Conclusion

Corporate databases have been recognized as strategic assets, and a successful corporation will make full use of its data resources to gain competitive advantage to better manage its business. Visualization is a key technology for extracting information from data; therefore, it is becoming increasingly important in our information-rich society. It complements other analytic, model-based approaches and exploits human pattern perception. Visualization can help users to navigate and explore the fast-growing number of data warehouses far more easily, and to rapidly discover the information hidden within volumes of data.

During the past several years, AT&T researchers have developed an innovative technology that permits interactive analysis of large corporate data sets. They have built a suite of applications based on innovative, novel views using a common software infrastructure. To demonstrate the usefulness of data visualization, they have employed data visualization within AT&T, focusing on key business unit problems. In partnership, AT&T plans to continue exploiting this technology internally, to address strategic business problems, and externally, to leverage this technology in the commercial marketplace.

### Acknowledgments

The authors gratefully acknowledge the many people who contributed to this paper and who have worked on data visualization: the case studies were provided by Tom Ball, Ken Cox, Debbie Fisher, Samir Gulve, and Graham Wills; and the careful proofreading was done by Anselm Spoerri and Diane Compton.

### \*Trademarks

UNIX is a registered trademark of Novell in the United States and other countries, licensed exclusively through X/Open Company Limited.

Windows is a trademark of Microsoft Corporation.

OpenGL is a trademark of Silicon Graphics.

### References

1. Carrie R. Smith, "Visualization: An unconventional look at data," *Inside Gartner Group This Week*, August 16, 1995, pp. 9-13.
2. Joel J. Hadary, Stephen G. Eick, Daniel E. Fyock, and Mary Zajac, "Database visualization brings hidden patterns to light," *Systems Reengineering Economics Letter*, Vol. 2, No. 4, April 1995, pp. 4-5.
3. Stephen G. Eick and Graham J. Wills, "Navigating large networks with hierarchies," *Visualization '93 Conference Proceedings*, San Jose, California, October 25-29, 1993, pp. 204-210.
4. Stephen G. Eick, "Graphically displaying text," *Journal of Computational and Graphical Statistics*, Vol. 3, No. 2, June 1994, pp. 127-142.
5. Walter F. Tichy, "RCS—a system for version control," *Software—Practice and Experience*, Vol. 15, No. 7, 1985, pp. 637-654.
6. Marc J. Rochkind, "The source code control system," *IEEE Transactions on Software Engineering*, Vol. 1, No. 4, 1975, pp. 364-370.
7. Edward R. Tufte, *The Visual Display of Quantitative Information*, Graphics Press, Cheshire, Connecticut, 1983, p. 162.
8. Kenneth C. Cox and Stephen G. Eick, "3D displays of Internet traffic," *Information Visualization Symposium*, Atlanta, Georgia, October 30, 1995, pp. 129-139.
9. Tara M. Madhyastha and Daniel A. Reed, "Data Sonification: Do you see what I hear?," *IEEE Software*, Vol. 12, No. 2, March 1995, pp. 45-56.

(Manuscript approved December 1995)

**Stephen G. Eick** is technical manager of the Data Visualization Research Group at AT&T Bell Laboratories in Naperville, Illinois. His research focuses on using novel interactive visualizations to extract the information latent in large databases. This involves inventing techniques, developing software tools, and building an infrastructure to mine knowledge from corporate databases and use it to competitive and commercial advantage. His research group has developed a suite of visualizations, including tools for visualizing geographic and abstract networks, software source code, bodies of text, log files, program slices, and relational databases, among others. Mr. Eick holds several software patents. He joined AT&T in 1985, after receiving a B.A. in mathematics from Kalamazoo College, Michigan, an M.A. in mathematics from the University of Wisconsin at Madison, and a Ph.D. in statistics from the University of Minnesota, Minneapolis.



---

**Daniel E. Fyock** is technical manager of the Specialized Software Solutions Group in the Enterprise Software Solutions Department at AT&T Bell Laboratories in Naperville, Illinois.



He has contributed to and/or led a number of innovative projects, including operating systems and computer networking design, development, and support; project management; organization reengineering; and establishment of new business opportunities. Currently, Mr. Fyock has responsibilities related to the emerging data visualization field. These include the commercialization of research data visualization prototypes into high-quality, commercial-grade software products; market management, design and development of features requested by customers; and ongoing customer care activities. Mr. Fyock joined AT&T in 1970, after receiving B.S. and M.S. degrees in mathematics from DePaul University, Chicago, Illinois.