

Visual Behavior Characterization for Intrusion and Misuse Detection

Robert F. Erbacher
Department of Computer Science, LI 67A
University at Albany-SUNY
1400 Washington Avenue
Albany, NY 12222, USA
erbacher@cs.albany.edu

Deborah Frincke
Department of Computer Science
University of Idaho
Moscow, ID 83844-1010
frincke@cs.uidaho.edu

Abstract

As computer and network intrusions become more and more of a concern, the need for better capabilities to assist in the detection and analysis of intrusions also increases. System administrators typically rely on log files to analyze usage and detect misuse. However, as a consequence of the amount of data collected by each machine, multiplied by the tens or hundreds of machines under the system administrator's auspices, the entirety of the data available is neither collected nor analyzed. This is compounded by the need to analyze network traffic data as well. We propose a methodology for analyzing network and computer log information visually based on the analysis of the behavior of the users. Each user's behavior is the key to determining their intent and overriding activity, whether they attempt to hide their actions or not. Proficient hackers will attempt to hide their ultimate activities, which hinders the reliability of log file analysis. Visually analyzing the user's behavior, however, is much more adaptable and difficult to counteract.

Keywords: Information Visualization, Security, Intrusion Detection, Computer Networks

1. Introduction

Intrusions and misuse of computer systems are becoming a major concern of our time [1, 2]. Our nation's infrastructure is heavily network based in all industries. The nation's infrastructure is not capable of dealing with attacks on a local or global scale, leaving network and computer security up to an organization's individual efforts. "... crackers, criminals, and foreign powers are building sophisticated cyber attack capabilities and doing reconnaissance on our networks today." [2] A Chinese newspaper referred to information warfare "as a means of achieving strategic victory over a militarily superior enemy." [3] A 1996 CSI-FBI survey found that \$4.5 billion was lost to business due to compromises in information security, and that the majority of businesses experienced some form of intrusion during the year. Systems of all kinds are vulnerable, as evinced by the 1998 attacks on Pentagon computers and the 2000 attacks on e-commerce. The FBI recently warned that hundreds of machines in the U.S. alone have likely been subverted for future Denial of Service attacks. With growing interest in Internet voting there is a correlated growth in concern over Denial of Service attacks, fraud, viruses, and Trojan horses being used, by both domestic and foreign agents, to disrupt the voting process or changing an individual's vote without their awareness [4, 5].

Break-in successes at Microsoft have raised questions as to the future security of future Microsoft products. Have modification been made that will allow hackers to gain future entry? The concern is that with the millions of lines of code in Microsoft's flag ship products that they may never be able to fully verify the integrity of their applications. This shows the necessity of protecting such information from attack and the need for future protections as hackers may have a new route for entry.

It is clear that even if systems can be made more secure, attack and internal misuse of technologies will evolve, making some form of intrusion/misuse management a necessity for all systems. New tools are needed to aid in the detection and eradication of attacks. This need extends from the fact that the Internet was initially designed under the guise of open communication and security has merely been retrofitted on top of the existing infrastructure and not integrated tightly with the design [5].

2. Current Techniques

Currently, system administrators are required to analyze log files in order to identify an attack [6]. These log files can incorporate hundreds if not thousands of messages per day. The amount of data available results in system administrators not fully collecting or monitoring all available information for all systems under the administrators control, rather focussing on primary systems and servers and examining minimal information from the remaining systems. In fact, system administrators in general do not collect or analyze any data related to Microsoft or Apple based operating systems at all; even though it is becoming possible for these systems to be the targets of break-ins, subversion, and misuse. Log file analysis is becoming the greatest time consumer for system administrators. Identifying actual intrusions and misuse requires that the intentions of the user be known during examination of the user's activity. This is currently unfeasible and results in missed attacks and many false alarms. This situation is only likely to get worse and with the globalization of e-commerce and interest in Internet voting the potential for serious damage increases as well. Ultimately, the goal must be to identify an attempted break-in or attack before the attack is successful. Current log file analysis only reveals that an attack has occurred after the fact. It is also imperative to reduce the number of false alarms and increase the number of actual attacks detected.

Attacks on a system can range from an attempt to gain entry to a system or merely to access a system remotely for inappropriate purposes. With large environments it is difficult for a system administrator to keep track of the configuration of all systems in the environment. This is complicated by the growing population of somewhat savvy users who often install software or change a systems configuration without understanding the full impacts of their modifications on security and reliability. It is rarely through the primary systems, over which the system administrator keeps a careful eye, that an attack first occurs.

By visually representing the log information we can reduce the time that is required to examine and analyze the data enormously. Graphical representations can include enormous amounts of information onto a single page or display compared to textual representations, greatly reducing the analysis task [7]. By incorporating characterizations of the activity of individual users we can improve the accuracy of the resulting analysis. These mechanisms when incorporated into a system administrator's war chest allow the system administrator to collect more information from the environment since it can be made feasible for the extra log information to be fully analyzed, instead of just cluttering the log files. Not only can the system administrator now fully examine accesses to the primary systems but they can now fully monitor all systems in the environment with reduced time requirements.

Even when represented visually, the amount of information available will often be overwhelming since it is in the details through which the activity of the users will be recognized. The activity of individual users whose activity is of interest must be pulled from the visual clutter of the remaining users. We must compare each user's activity against what is typical. This comparison is a necessary step in identifying the relevance and importance of the data for further investigation [8]. This identification of what is typical is dependent on an organization's unique usage of its computer and network infrastructure. The knowledge of the local system administrator is imperative in this analysis.

3. Previous Work

The prior work of Erbacher et al. [9] described the basic visualization environment used in this research. As we continue to analyze the data collected by the Hummer intrusion detection system [10] and other collected data we have realized the need for a more organized approach to analyzing this data. Most previous work on visualizing network data has been based on measuring performance or bandwidth characteristics. Very little prior work has dealt with visualizing network intrusion data, particularly real-time network intrusion data as is our ultimate goal.

Ultimately, we are attempting to visualize the actions of an enterprising hacker actively seeking to counter the attempts being made to identify the hacker's actions. This dynamism is an attribute not seen in other visualization tasks resulting in the need for novel solutions.

3.1. Intrusion Detection Systems

Little previous work has been done towards the use of visual analysis as an aid to intrusion detection. For instance, many have proposed use of a simple "odometer-like" or metered scale to indicate the estimated level of attack a system is enduring. This is embodied in the Hummer "perceived level of threat" [10] indicator. Earlier systems, such as DIDS [11, 12], provided graphical representations in the form of color to indicate when a system had experienced a sequence of suspicious events. While useful, these approaches do not provide adequate detail to do more than observe that attacks are in progress and do little to aid diagnosis. Dr. Frincke has performed preliminary investigations towards identifying likely models for depicting system state [13]. This expands on lessons learned from that system and its prototype.

3.2. Visualization systems

In contrast to intrusion detection, quite a bit of visualization research has been applied to network accesses. The principal body of work related to network intrusion is from the information exploration shoot-out, organized by Georges G. Grinstein and supported by the National Institute of Standards and Technology (NIST) [14]. In this project, researchers were given access to a data set consisting of network intrusions. The idea was to identify which researcher's techniques were effective at identifying the intrusions. The driving philosophy was that little work has been done to compare visualization techniques in a formal setting. Perceptual studies have been done to identify characteristics of the human visual system [15, 16] that should be used as a basis for the development of visualization techniques but little has been done to actually compare and contrast visualization techniques. There is no body of literature that identifies what visualization techniques definitively work better on a given data set.

Heydon et al. [17] discuss the application of visual languages to security specification. The use of visual metaphors reduces errors and assists in clearly identifying the correctness of UNIX file permissions.

Most previous work involving visualization related to networks has emphasized graphics that depict network performance and bandwidth usage [18, 19, 20, 21], even down to the router [20], individual packets [22], and individual e-mail messages [23]. The techniques developed for these purposes do not provide sufficient detail or handle sufficient numbers of nodes and attributes in combination for our needs. The work by Eick et al. [23] strictly deals with e-mail and subsequently resolves many fewer nodes and attributes than is needed for intrusion detection. Other work has been geared towards visualizing systems for program analysis and program development [24]. These environments typically deal with small numbers of processors that are working on a single task and thus have a common grounding. This research into network usage has not been applied to network intrusions. It does, however, provide a starting point.

Becker et al. [25] discuss the SeeNet environment that provides linkmaps for visually representing the amount of data being sent between two network nodes. It can identify when a node is overloaded, shows the network's behavior, how data moves from one location to another and its volume. This is important when a crisis occurs and usage increases dramatically, e.g., after a California earthquake. Understanding the consequences of events, such that telephone companies can be prepared for changing demands, is imperative.

Livelihood [26] is an environment for visualizing and measuring the web. By probing web accesses they gather statistics as to the number of hits web sites are receiving. This information is presented in statistical form as charts and graphs. An extension to the environment provides a more graphical representation. In this advanced form, the approximate location of network nodes is represented, showing geographical association between web sites. The visual representation of each node is then presented in such a way as to reveal the activity of the site. Each node can represent several parameters simultaneously.

Much effort has been done to formulate algorithms for efficiently laying out nodes onto the display [21, 23, 27]. The goal of most algorithms is to reduce the complexity of the resultant rendering, as seen in [27]. Other techniques use geographically based representations [20, 21, 25] or weighted techniques [23]. In our case, the node layout needs to provide a representation as to whether nodes are local or remote and the criticality of the system. The node placement must assist the system administrator in analyzing and understanding the log information and should not necessarily be geared towards what the visualization expert interprets as an efficient node layout.

4. Data Collection

This paper examines visualization techniques in the context of protecting large-scale networks. Examination of both commercial and research efforts to identify security violations consistently results in the observation that considerable quantities of data are generated—usually considered to be far too much to be evaluated effectively using current techniques [28, 29]. Some of this is due to the way that data-gathering choices are made [28]. It is our belief that refinements in the data gathering decision making process will not suffice: as networks grow larger, the amount of misuse-relevant data will also grow. Hence, better methods for analyzing the data are needed, rather than continued reliance on primarily textual techniques.

The majority of the data we will be analyzing consists of data collected by the Hummer intrusion detection system. This data consists of much of the information available through normal log files with additional statistics and other available system information. In addition, the log files for all systems under consideration are merged into a single Postgres database, providing for easy querying. The total volume of data collected through log files tends to be relatively small, approximately a megabyte per day per system. However, the total number of events that must be analyzed and interpreted in relation to one another is huge, in the thousands.

Additionally, we are incorporating information collected by a network sniffer program. A network sniffer allows the collection of network traffic data (i.e., individual network packets). This additional data is imperative for detecting systems illegally connected to the network and detecting the misuse of compromised machines that may no longer be providing complete or correct information to the system administrator through log files. Network sniffers can generate gigabytes of data per hour, resulting in a huge amount of data. It should be noted that once analysis has been performed, data sets not containing relevant information can be removed. Additionally, subsets of the data not relevant can be removed if it is deemed necessary to archive samples of network traffic.

Most modern UNIX shells provide the ability to save a history of the commands the user has executed. This is generally to assist the user in re-executing a sequence of commands. While we have found this useful in analyzing the activities of novice hackers, more sophisticated hackers will disable this ability. By using system based tools to collect this information, i.e., user space programs, for ALL users we provide another level of data that can be used to analyze the activity on the system for intrusions or misuse.

5. Visualization Techniques

The current problem is with too much information and inefficient analysis techniques. Currently, system administrators must select what information they collect and analyze on a daily basis. The information they generally collect is much less than the information that is actually available. Only when a problem is detected will additional information be collected. This leads to situations where misuse or intrusions may not be detected for quite some time, for example the misuse of CIA computers for IRC usage [30]. In general, log information is collected which reports the connection information of a system and usage that may be indicative of break ins. User level applications are not reported in this log information due to the volume of information that would be generated. Also, network traffic is also generally collected due to its volume. This information while useful would take so long to analyze that it would effectively be useless and is not collected.

The problem with analyzing log files results from the basis that reading textual information is inherently a perceptually serial process. Interpretation of graphical images, on the other hand, is perceptually a parallel process [15, 16]. Forcing the user to use textual information therefore slows the analysis process substantially in comparison to the use of graphical imagery. An additional advantage of imagery is that more “concepts” can be presented in a single image. Thus, rather than observing individual reports or report summaries, it is possible to observe a single image that embodies the same information. This will reduce the amount of mental context switching required by users, making system assessment both easier and more efficient.

Visually, we must provide informative and perceptually based techniques that allow the system administrator to examine the activities in the computing environment as a whole and quickly identify activities that require further investigation. Figure 1a shows the typical usage of a single system under normal usage. The system in the center is the system under analysis, which is accessing two systems over NFS mounts. It is being accessed by four systems. Privileged ftp’s are represented by long dashes, anonymous ftp’s by short dashes, and NFS mounts by solid lines. All lines are directed, showing the direction of access. The intensity of the node or line shows the duration since the last active access. The node under examination also shows the number of users through the protruding spikes. The gray border around the image is indicative of the time of day.

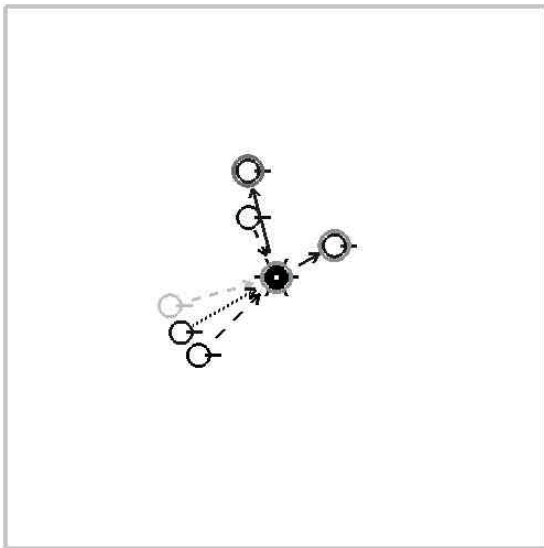


Figure 1a: Moderate volume activity analysis in the morning with time based representation of known nodes.




Figure 1b: Analysis of low volume activity in the middle of the night.

5.1. Attack Analysis

Analyzing the collected information and determining whether or not an attack or misuse is occurring requires that the intent or behavior of the individual be analyzed. Currently, when suspicious behavior is noted the individual’s activities are examined, most often after the fact. With visualization it is possible to examine an individual’s activity as it is occurring and determine immediately, before substantial harm has been done, that the individual’s activities are unacceptable. Even suspicious activity can be difficult to detect with standard log file based approaches that require the system administrator to peruse textual information.

Network traffic data can be incorporated into the display, allowing the user to quickly examine the data for particular types of traffic, such as illegal systems on the network, improper application usage, connections from unknown systems or users, etc. With this integration additional cases of misuse and intrusions can be detected very quickly. For example, the case of the personnel at the CIA running an illegal chat room could have been detected through the analysis of network traffic information that would have identified the characteristic IRC packets on the network that would have been a clear indication of misuse. Since the information is not being read textually but rather visually

through a graphical display, the gigabytes of information related to network traffic and user space applications can quickly be analyzed at intervals. It should be noted that the entire set of data need not be saved rather the system administrators can take snapshots over regular intervals, examine the collected data, and remove it. No loss of information will result.

Individually, single actions by an individual do not provide much context or basis for their motivation in their activity. Certain activities are clearly indicative of illegal system usage, however, these actions are most often identified in users who are inexperienced in subverting a system. These types of novice users are easily identified with conventional techniques. Our principal concern must be with the more experienced users who will attempt to hide their tracks or camouflage their actions. In these situations, even though the user may be attempting to cover their tracks or hiding their true intent, the overall actions when taken together will clearly indicate the overall motive of the individual. In this fashion, we are providing system administrators with tools that allow them to visually examine the activity on the computer systems under the administrators control as well as network usage in a merged environment.

5.2. Behavior Identification

In everyday life we must ascertain the intent and motivation of individuals on a daily basis. In a computing environment the same level of information we use socially isn't available. We must collect the information that is available and provides the information in a form such that the activities of the user can be examined. The behavior of an individual can be derived to a limited extent from the activities the user performs, when these activities are performed, the order in which they are performed, and how the presence of others affects their activities. At issue is the need to collect much information that system administrators currently allow to go unnoticed due to the clutter it adds to typical log file analysis.

Notice in figure 1b that there are two nodes being accessed by NFS mounts. These systems are the locales where our collected data is being stored. The third node is a user performing an anonymous FTP. In and of itself this wouldn't be much of an issue. However, we must take into account other aspects of the user's behavior. First, the user's node and directed line are colored red, indicative of the fact that we cannot do a reverse hostname lookup on the system. Second, the user is performing the FTP after midnight. While this may not be wholly unusual for a student, taking the situation in its entirety indicates the need for further examination of the user's activity and the circumstances coinciding with the user's activity. It is this ability to take multiple characteristics, through multiparametric visualization techniques, and integrate them to find a greater understanding that is the key to analyzing network and computer usage for intrusions and misuse. This becomes even more important when analyzing systems that are more heavily used and examining multiple systems, particularly when a user's actions must be deciphered across multiple systems.

A second example of interest is show in Figure 1a. The three nodes in the lower left are connections that were made in rapid succession from different IP addresses. Notice that these nodes aren't within the University's local network itself. Is this an indication of an attack? Had they been local to the University's network they would have been deemed to be students logging on immediately after class. Once removed from the University directly, seeing such sequences should raise a level of concern, particularly if they are even more removed from the University's network locale than is exemplified here. At most organizations such sequences will be seen when classes or meetings end or typical starting times for employees arrive. Had the nodes been telnet connections, port probes, or even the same type of connection then they would be of greater concern. However, visually, particularly when animated and showing history, this display clearly shows the actions of users on the system and how connections are being made. These types of activities are warning signs as to possible intrusions or misuse. The administrator's knowledge of local behavior is imperative to making sense of the data and understanding its meaning.

Figure 2a shows an example where a user after connecting to a local system from a remote node has immediately jumped to a second local node, effectively using the first node as an intermediary. What is the meaning of this activity? Does the meaning change when many users are behaving similarly, Figure 2b? In this case, the systems on the local network are configured to deny connections from outside of the local network, except for one system

designed to accept such incoming connections (e.g., a firewall). Consequently, the network configuration is enforcing this behavior. Had the network not been configured in such a way or a user was using a different system as an intermediary then the meaning of such behavior would be much more dismal. Knowledge of the local network and system configurations is imperative to understanding the meaning of certain behavioral patterns.




Figure 2a: Local behavior with a single user using a node as an intermediary.




Figure 2b: Local behavior with multiple users connecting through a single node as an intermediary.

6. Technique Comparison

There are other possible mechanisms for examining the volume of data that can be collected related to network and computer system usage, most notably data mining. However, they are severely limited and do not meet the needs of system administrators. In particular, data mining provides the ability to analyze data and identify patterns or absence of patterns within the data. Data mining, while a powerful tool for analyzing data, does not apply well to the scenario of intrusion and misuse detection. First, it doesn't contain the necessary dynamism that will allow it to deal with the proficient hacker who is attempting to avoid detection, especially once the hacker is aware of such a system. In other words, data mining may not be suitable for dealing with the human element present in security. Second, current data mining techniques require considerable time to execute. This delay provides substantial time for the intruder to cause significant damage to the system.

We showed examples in which knowledge of the local network and system configurations are necessary to determine the meaning of a set of actions, relating to multiple accesses in rapid succession and users jumping through intermediary systems. This is difficult to incorporate into a data mining system. The data mining system can be trained on a system administrator's local network, however, this requires substantial knowledge and ability on the part of the system administrator related to data mining. It also opens up the possibility of mistraining the data mining system. Our goal is to reduce the burden on the system administrator not complicate it.

Data mining can be used to assist in the intrusion detection cycle. It can aid in identifying users who may not be attempting to intrude on the system but may be beginning to experiment with the tools and techniques for breaking into systems. Second, it will be helpful in identifying casual users attempting to break-in to the systems, acting as a second level of defense against intruders.

Used together, visualization and data mining can create a powerful set of capabilities for detecting intrusions and misuse. This will allow system administrators to eradicate unacceptable activities before substantial harm and expense are incurred.

7. Conclusions

Computer and network security are becoming critical issues. The capabilities are not yet in place to allow system administrators to efficiently detect and counter act intrusions and misuse of the systems and networks under the administrators control. Only through the innovation of new technologies can we hope to be able to counter act the growing threat from hackers. Of the techniques available visualization appears well placed to take on the brunt of this task. Perusal of textual log files is totally inadequate. Data mining has promise but currently can not provide the necessary timeliness or handle the human element.

By providing sufficient attribute mappings within the visualization we can represent substantial characteristics as to the overall behavior of users within the environment. By analyzing user behavior as a whole we can gain insight into the user's intent and ultimate goals. It is only through the combination of attributes when taken together that the whole of the meaning of the user's activity can be discerned. By focusing analysis on the user's behavior we are reducing the number of false alarms and increasing the reliability of the systems administrator's analysis. Ultimately, the incorporation of visualization tools should prove to greatly improve the detection of intrusions and misuse before damage is incurred to the systems.

The visualization tools will also aid in reducing false alarms and identifying potential problems that would otherwise go undetected. These types of situations can be a great drain on a system administrator's time. Ultimately, this will become much more than just an early warning system for system administrators, rather it will become a filtering device allowing the system administrator to filter out unwanted details and identify real activity of concern.

References

1. Greg Farrell, "Police have few weapons against cyber-criminals. Problem stems from lack of funds, training," *USA Today*, pp. 5B, December 6, 2000.
2. Richard A. Clarke, "Convergence and Transition, Privacy and Security," Remarks at SafeNet 2000, Redmond, WA, December, 2000.
3. Jim Wolf, "U.S. Draws Attention to Information Warfare," Yahoo! News and Reuters, December 26, 2000.
4. Joe Mohen and Julia Glidden, "The Case for Internet Voting," *Communications of the ACM*, Vol. 44, No. 1, pp. 72-85, January 2001.
5. Deborah M. Phillips and Hans A. Von Spakovsky, "Gauging the Risks of Internet Elections," *Communications of the ACM*, Vol. 44, No. 1, pp. 72-85, January 2001.
6. Rebecca Gurley Bace, *Intrusion Detection*, Macmillan Technical Publishing, 2000.
7. Edward R. Tufte, *The Visual Display of Quantitative Information*, Graphics Press, 1983.
8. Edward R. Tufte, *Visual Explanations*, Graphics Press, 1997.
9. Robert F. Erbacher and Deborah Frincke, "Visualization in Detection of Intrusions and Misuse in Large Scale Networks," *Proceedings of the International Conference on Information Visualization '2000*, London, UK, July, 2000, pp. 294-299.
10. Polla, D., J. McConnell, T. Johnson, J. Marconi, D. Tobin, and D. Frincke, "A FrameWork for Cooperative Intrusion Detection," *21st National Information Systems Security Conference*, pp. 361-373, October 1998.
11. Snapp, S. et al., "DIDS (Distributed Intrusion Detection System) Motivation, Architecture and An Early Prototype," *National Information Systems Security Conference*, 1991.
12. C. Ko and D. Frincke and T. Goan et al., "Analysis of an Algorithm for Distributed Recognition and Accountability," *ACM conference on Computer and Communication Security*, Vol. 1, No. 1, 1993.
13. Vert, G., J. McConnell, and D. Frincke. "Towards a Mathematical Model for Intrusion," *21st National Information Systems Security Conference*, pp. 329-337, October 1998.
14. Georges Grinstein, "Workshop on Information Exploration Shootout Project and Benchmark Data Sets: Evaluating How Visualization does in Analyzing Real-World Data Analysis Problems," *Proceedings of the IEEE Visualization '97 Conference*, IEEE computer Society Press, Phoenix, AZ, pp. 511-513, 1997.
15. Markus Gross, *Visual Computing, The Integration of Computer Graphics, Visual Perception and Imaging*, Springer-Verlag, 1994.
16. William R. Hendee and Peter N.T. Wells, *The Perception of Visual information*, Springer-Verlag, 1994.

17. Allan Heydon, Mark W. Maimone, J. D. Tygar, Jeannette M. Wing, and Amy Moormann Zaremski, "Miro: Visual Specification of Security," *IEEE Transactions on Software Engineering*, Vol. 16, No. 10, pp. 1185-1197, October 1990.
18. Richard Becker, Stephen Eick, and Allan Wilks. "Graphical methods to analyze network data," In *IEEE International Conference on Communications ICC '93 Proceedings*, Geneva, Switzerland, pp. 946-95, May 1993.
19. Taosong He and Stephen G. Eick, "Constructing Interactive Visual Network Interfaces," *Bell Labs Technical Journal*, Vol. 3, No. 2, pp. 47-57, April-June 1998.
20. Kenneth Cox, Stephen Eick, and Taosong He, "3D geographic network displays," *ACM Sigmod Record*, Vol. 25, No. 4, pp. 50, December 1996.
21. Eleftherios E. Koutsofios, Stephen C. North, Russel Truscott, and Daniel A. Keim, "Visualizing Large-Scale Telecommunication Networks and Services," *Proceedings of the IEEE Visualization '97 Conference*, IEEE Computer Society Press, San Francisco, CA, pp. 457-461, 1999.
22. Deborah Estrin, Mark Handley, John Heidermann, Steven McCanne, Ya Xu, and Haobo Yu, "Network Visualization with Nam, the VINT Network Animator," *IEEE Computer*, Vol. 33, No. 11, pp. 63-68, November 2000.
23. Stephen G. Eick and Graham J. Wills, "Navigating Large Networks with Hierarchies," In *Visualization '93 Conference Proceedings*, San Jose, California, pp. 204-210, October 1993.
24. Robert F. Erbacher, "Visual Assistance for Concurrent Processing," University of Massachusetts at Lowell Doctoral Dissertation (1998 CS-3), Lowell, MA 01854.
25. Richard Becker, Stephen Eick, and Allan Wilks, "Visualizing Network Data," *Readings in Information Visualization: Using Vision To Think*, Stuard Card, Jock D. Mackinlay, and Ben Shneiderman, editors, Morgan Kaufman Publishers, pp. 215-227, 1999.
26. Tim Bray, "Measuring the Web," *Readings in Information Visualization: Using Vision To Think*, Stuard Card, Jock D. Mackinlay, and Ben Shneiderman, editors, Morgan Kaufman Publishers, pp. 469-492, 1999.
27. Giuseppe Di Battista, Peter Eades, Roberto Tamassia, and Ioannis G. Tollis, *Graph Drawing: Algorithms for the Visualization of Graphs*, Prentice-Hall, 1999.
28. Picciotto, J., "The Design of an Effective Auditing System," Technical report, 1987.
29. D. Zerkle et al. "A Data-Mining Analysis of RTID Alarms," *Recent Advances in Intrusion Detection*, Sept 1999.
30. Vernon Loeb, "Chat Room Causes Trouble for CIA Employees," *The Washington Post*, pp. A10, November 12, 2000.