

High-Resolution Displays Enhancing Geo-Temporal Data Visualizations

John Booker¹, Timothy Buennemeyer², Andrew Sabri¹, and Chris North¹

¹Center for Human-Computer Interaction, Department of Computer Science and

²Bradley Department of Electrical and Computer Engineering

Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061 USA

{jobooker, timb, plastk, north}@vt.edu

ABSTRACT

Intelligence analysis is difficult due to the volume and complexity of the data, as well as the subtlety of the connections that have to be made in order to identify threats. Information visualization techniques are necessary for human interpretation of this data, but many use a single monitor that lack the resolution and area needed to effectively display it. We developed an application specifically for a high resolution tiled display, taking advantage of its increased area and pixel count to show a detailed and thorough geospatial and timeline view of terrorist activity. These displays were well suited for intelligence analysis visualization due to their ability to provide access to numerous details from multiple views, and allowed users to maintain the context of data. This work showcased the ability of high resolution visualizations to assist intelligence analysts.

CR Categories and Subject Descriptors: H.5.2 [Information Interfaces and Presentation]: User Interfaces – Screen design,

Keywords: Geo-temporal views, high resolution tiled displays, intelligence analysis, visualizations

1. INTRODUCTION

Intelligence analysis is a difficult task that involves mining complex datasets for clues. Distinguishing between evidence of suspicious activity and background noise is complicated, as is making connections between events, people, and other evidence. While algorithms to parse, extract, and make connections between data are important to this field, at some point the data must be interpreted by human users. Information visualization allows people to employ their innate visual strengths to gain insights about data and is necessary to effectively understand both raw data and the filtered output of algorithms.

While there are many barriers to effective intelligence analysis, based on related research and our own experience we focused our efforts on the following problems:

1. Viewing geospatial and temporal data simultaneously is difficult. However, it is desirable to do so in order to identify “geo-temporal” relationships that might show movement trends or where a terrorist plot might culminate into an attack.
2. Intelligence data is often displayed with multiple views

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. ACMSE 2007, March 23-24, 2007, Winston-Salem, N. Carolina, USA Copyright 2007 ACM 978-1-59593-629-5/07/0003...\$5.00

overlaid with numerous data points, connections, and details. This data can overwhelm the display space of a single monitor.

3. Maintaining context within large datasets is challenging because it is mentally demanding to remember what was seen previously, where it was, and its potential relationship to current information.

We developed a Gigapixel Intelligence Analysis Navigation Tool (GIANT) (Fig. 1) to implement solutions to these problems. We used node-linking techniques similar to the type often employed to represent social networks, and a geospatial visualization that placed nodes overtop of a map in their appropriate positions (Fig. 2). The geo-temporal problem was addressed by mapping activity age to the opacity of the edges between nodes. For example, a phone call that occurred several days ago would be represented by a faintly transparent edge between two nodes, whereas an opaque edge indicated current activity. We also included a separate timeline view, but the edge opacity allowed a basic understanding of event ordering without requiring users to switch their focus to a different view.

We addressed the remaining problems by using a high resolution tiled display. It allowed us to show greater amounts of data at once, while still presenting multiple views. Because we had sufficient area for the data, we were able to avoid traditional information visualization navigation strategies such as panning and zooming, overview + detail, and focus + context, and their associated problems with maintaining context. By designing to take advantage of the increased pixel count, we were able to implement many features not possible with a single monitor solution.

In our research, we sought to better understand how to integrate time and geospatial data simultaneously in a single view, while maintaining perspective about the data. Experts in the information visualization and human-computer interaction field evaluated different methods of conveying this type of data. We used rapid prototyping techniques paired with a formal user evaluation to understand effective ways to aid analysts in finding threatening events and persons. These methods allowed us to explore the benefits and challenges associated with using high resolution tiled displays for intelligence visualization.

During our system evaluation we found that combining the geospatial and timeline views offered a more efficient means to search than using isolated views. Users also located specific nodes faster with the timeline view because node names were shown as an index. Geospatial searching improved with location-based nodes because users relied on their familiarity with map geography. Lastly, large tiled displays overcame many smaller display shortcomings, yet still required filtering and other methods to reduce the amount of data shown.

In this paper we discuss related work and known problems with visualizing intelligence information. We next describe our system



Fig. 1. Interaction with GIANT.

design and evaluation. In closing, we present our preliminary results, conclusion, and future work.

2. RELATED WORK

Intelligence analysis was a demanding problem for a number of reasons. Elm et al. [26] identified intelligence information as troublesome because it consisted of large, complex, and highly heterogeneous data. He also acknowledged the difficult goals of performing inferential analysis and developing hypotheses from incomplete or seemingly disparate data. “The Sign of the Crescent” [12] case study was an excellent representation of these challenges.

Badalamente et al. [22] gave a good overview of ten tools that intelligence analysts could use. Many applications dealt with backend database integration, but analysts particularly needed dynamic visualization tools that could handle a changing dataset as new intelligence reports arrived over time. There was a clear need for visualization tools that could display events in the order they occurred.

With so much data to map to the visualization, simply plotting events along a single axis was insufficient. Havre et al. [23] developed a technique called “Theme River,” in which the themes of a document collection were mapped to smaller paths within the overall timeline. This was a good method to show an overview of the characteristics of a timeline, without requiring users to understand each data point. However, this concept still used a single axis and was not well suited for showing the activities of multiple people.

“Lifelines” by Plaisant et al. [6] was a good representation of multiple events within a single person’s lifetime. Effective at showing various aspects of an individual’s medical history, this idea could be adapted to a criminal record or to terrorist intelligence data. However, like Theme River this mapping was not immediately extendable to multiple individuals in an overview setting.

Kapler and Wright [25] acknowledged that using multiple views for spatiotemporal data had shortcomings that could be potentially overcome with one view showing both space and time. Their solution that combined the two ideas was called “GeoTime,” a spatial map with time mapped to the 3rd dimension.

Example tools that specifically analyze intelligence data were CrimeLink Explorer and the related Spatio Temporal Visualizer (STV) for law enforcement, both from the University of Arizona, [15] [24]. CrimeLink noted the importance of performing computational link analysis to identify important data, while the



Fig. 2. GIANT’s social network graph.

STV approached the problem from the visualization angle, offering spatial and temporal views of the data. Both systems acknowledged the problem of information overload, and STV demonstrated the advantages of seeing data in geospatial and temporal views.

Van Ham and van Wijk [9] demonstrated that displaying all nodes and edges, while integrating detail, was not a viable option for large node-link graph structures. Their method for “Interactive Visualization of Small World Graphs” showed several innovative aggregation techniques for managing densely clustered graphs while retaining the overall structure.

With the volume and diversity of the data involved, it was important to have the proper canvas to display the data. A tiled display has increased physical size, which helped maintain context [8] [17], while an increased pixel count allowed the presentation of additional details [21]. This was important for maps, particularly in military intelligence settings [18], as well as large datasets such as telecommunication networks [3]. In addition to benefits, tiled displays come with their own set of interface layout problems [20], new physical navigation issues [1], and a host of difficulties managing the increased information and screen content [11]. With this in mind, we set out to design and evaluate a system that would make the most of the large display, and deal with any issues it presented.

3. DESIGN

Our research effort proceeded along two parallel tracks. First, we sought to design a database capable of handling text-based intelligence reports. The second track involved design and implementation of the visualization specifically for high resolution, large screen displays.

3.1 Data Organization

We deployed a MySQL database [19] to store and organize the data. This relational database format provided a reliable backend for our visualization application stored the following tables: person, event, establishment, link, node, report, and location.

Data was organized around individual nodes, which could be a person, establishment, or event (Fig. 3). A report entry associated with each node contained the full text report as well as a summarized version. The link table contained two nodes and a connection to join them. Connections between nodes were comprised of communication, travel, money, business and personal relations, and crime. The link table also contained a date field so that interactions between nodes could be filtered by time.

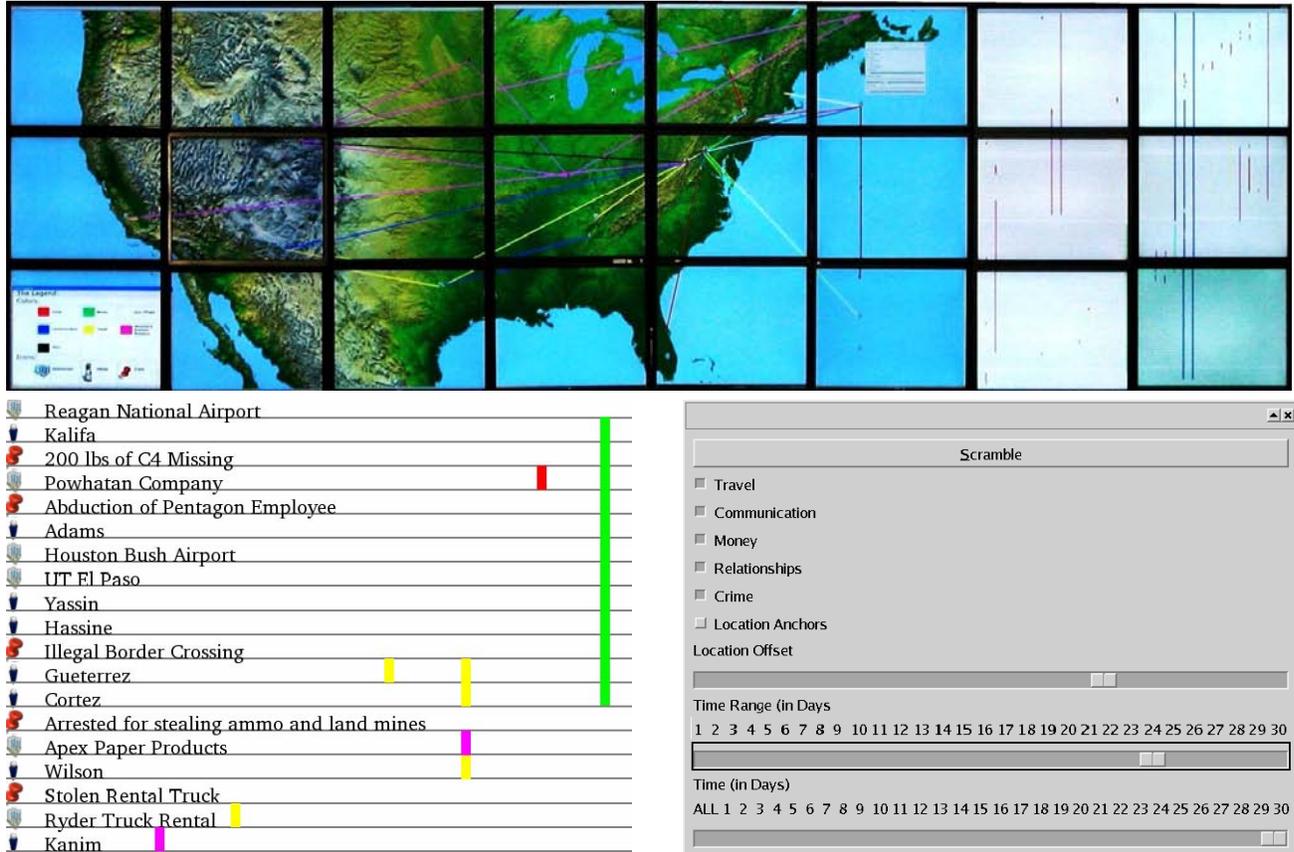


Fig. 4. a) GIANT display, b) Timeline view, c) Moveable user interface.

The method used to create links was to only connect nodes if they were both present in the same report. For example, two crimes committed by the same person would not be linked to each other. Instead as each report was entered, a link would be made back to the suspect's node. The connection between the crimes themselves would only be seen once the data was visualized.

3.2 Application Features

GIANT (Fig. 4a) was developed in C++ with the OpenGL and QT 4.0 libraries on a Linux platform. It included a geospatial view that contained nodes mapped to their actual latitude and longitude coordinates on a mercator map projection, a timeline view (Fig. 4b), and an interface specifically designed for a tiled display.

The moveable interface (Fig. 4c) gave users access to available application features regardless of where their focus was. The first slider controlled the distance of each node's offset from its true location in the geospatial view. In (Fig. 5a) the offset is zero, and the nodes are stacked on top of each other, showing concentrations of activity. In (Fig. 5b) the nodes are offset, separating the individual nodes to show details about activity in the busy areas. The interface also contained a pair of sliding time controls, which filtered out events based on when they occurred.

Nodes, which consisted of persons, establishments, or events, were represented as person, building, and thumbtack icons respectively in the geospatial view. The activities that occurred in the intelligence reports were mapped to the edges, with the link's opacity determined by the difference between the dates of the report and the time selected in the interface's time slider (with the exception of future events not shown at all). In (Fig. 5c) GIANT

shows a number of solid edges that have occurred on the selected date, and a number of relatively transparent edges which occurred prior to that time. The activity type was mapped to the edge color, communication edges were blue, money transactions were green, travel links were yellow, business and personal relations were purple, and crimes were red. Black edges signified that two people nodes (real and alias name) were found to have been the same entity. White edges anchored offset nodes back to their original location. All of these edges could be filtered, using the checkbox controls on the interface.

In the timeline view nodes were displayed as persistent horizontal lines, and edges were vertical lines that connected nodes at the time a report occurred. The color and opacity of the edges had the same meaning as they did in the geospatial view. Brushing and linking was implemented, so that when items were chosen in either view they would become selected in the other view.

The high resolution of the display allowed for text labels as well as brief summaries to be presented next to all of the vertices.

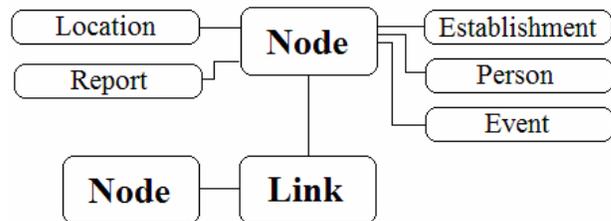


Fig. 3. Database table layout.

Users could get further information such as the location of the person, establishment or event by placing the mouse over the graphic. Clicking on a node would cause the icon to grow and become highlighted. The adjacent edges would become wider, distinguishing them from surrounding edges (Fig. 5d). The larger screen area allowed for their associated report summary to be displayed alongside these edges.

3.3 System Architecture

Since the architecture of a system could have constrained quality, the decisions made during design and early implementation had considerable impacts on the resulting solution (Fig. 6). Distributing the computing workload increased the overall system's complexity and software development effort but added a measure of redundancy.

Development was done on a standard Linux workstation, and the code was designed to run on a 12-node Linux cluster with 24 LCD monitors (two monitors per server). The display was arranged in an 8x3 tiled monitor configuration with a physical size of approximately 3x1 meters and a resolution of 10240x3072 pixels (approximately 34 million pixels). The cluster used Distributed Multihead X Project (DMX) [7] to provide the user with a single desktop interface and Chromium [4] to distribute the real-time rendering of the GIANT's OpenGL code. A MySQL database [19] server was setup external to the server cluster and operated as the backend to the software. This supporting architecture satisfied the performance, quality of service, data storage, reliability, and rapid access time requirements needed to sustain our visualization.

4. REQUIREMENTS ANALYSIS AND INITIAL FEEDBACK

We used an iterative feedback process through which experienced information visualization designers identified problems as well as possible enhancements that were then addressed with later revisions.

4.1 Initial Evaluation

Five participants, computer science graduate students, evaluated our prototype in order for us to gain a better understanding of how we could optimize and clarify our visualization. The prototype emphasized the design layout, so the geographical map with timeline integration could be conceptualized. It used highly developed image maps to represent link changes between different dates.

A series of questions were posed to the test participants to gain necessary design feedback and useful feature suggestions for incorporation into the application. Users described enhancements

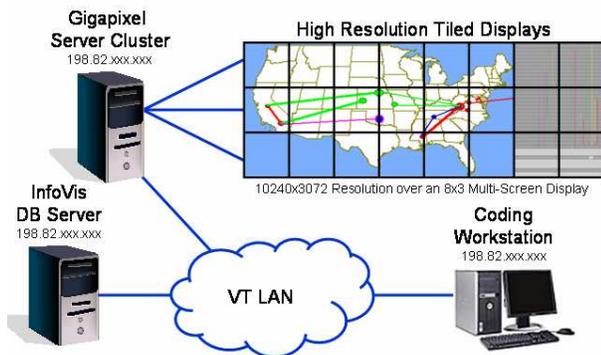


Fig. 6. Implemented system architecture.

to better convey the meaning of the icons and improve the connection aggregation. One resulting innovation supported by our evaluation was to fade the edges to indicate time with opacity. With each evaluation cycle, we further incorporated feedback to remedy user cited issues. The prototype allowed our participants to directly engage in the visualization process. With these issues in mind, we transitioned from prototype visualization and commenced implementation of the fully featured application.

4.2 Prototype Methods

With the full visualization implementation, we considered the appropriate methods to assess the application. We looked at heuristic evaluation [13] [14] and cognitive walkthrough [5] for a way ahead. Heuristic evaluation was considered simple to perform, however, it was only likely to report obvious cosmetic problems while more serious system issues could be easily overlooked [16].

We obtained user feedback by using a modified cognitive walkthrough because that method was designed to test the learnability of user interfaces, so it could be used at the beginning of the interface design and repeated as the features evolved. Also, it was renowned for uncovering serious system flaws [5].

To determine the usefulness of our visualization, we evaluated its capabilities against "The Sign of the Crescent," a case study from the Joint Military Intelligence College [12]. This study was used to train intelligence analysts to critically examine field reports to determine terrorist plots. Five system engineers experienced in the information visualization field examined the application following a series of expert questions. The first set of tasks was more subjective and thus evaluated using a Likert Scaling method. The remaining tasks were direct manipulations of the GIANT tool. The participants read a series of questions, which directed them to search for specific information using certain view aspects of GIANT. Usability issues were further investigated with a series of timed questions.

We primarily focused on usability and evaluated user searches with isolated and combined views. Participants were prompted by the questions to locate nodes, links and summary data using a spatial search of the geographic map or by searching the timeline index. These tasks were designed to investigate usability and to improve the effectiveness of showing all the data in geospatial and timeline views, and its effect on the overall usability of the GIANT.

5. PRELIMINARY RESULTS

While there was already a sizeable knowledge base about tiled displays, GIANT was one of the first applications designed to take advantage of the increased space and resolution for intelligence analysis. We discovered many interface issues specifically related to large displays during the development process. The initial user feedback further verified some of these issues, as well as exposing other areas for work, which are related here.

The traditional visualization mantra of "overview first, zoom and filter, and details on demand" [2] still applied to GIANT, and the increased pixel count allowed us to display a larger overview and more data concurrently, even with multiple views. While simply showing one view in even more detail would be tempting, users appreciated the multiple views. Each of these views had sufficient room to show an overview and details, without being forced to yield screen real estate to the other at any time.. This would not have been possible on a single 1280x1024 monitor, where we would be forced to lose either 96% (23 of the 24 screens) of the detail or overview.

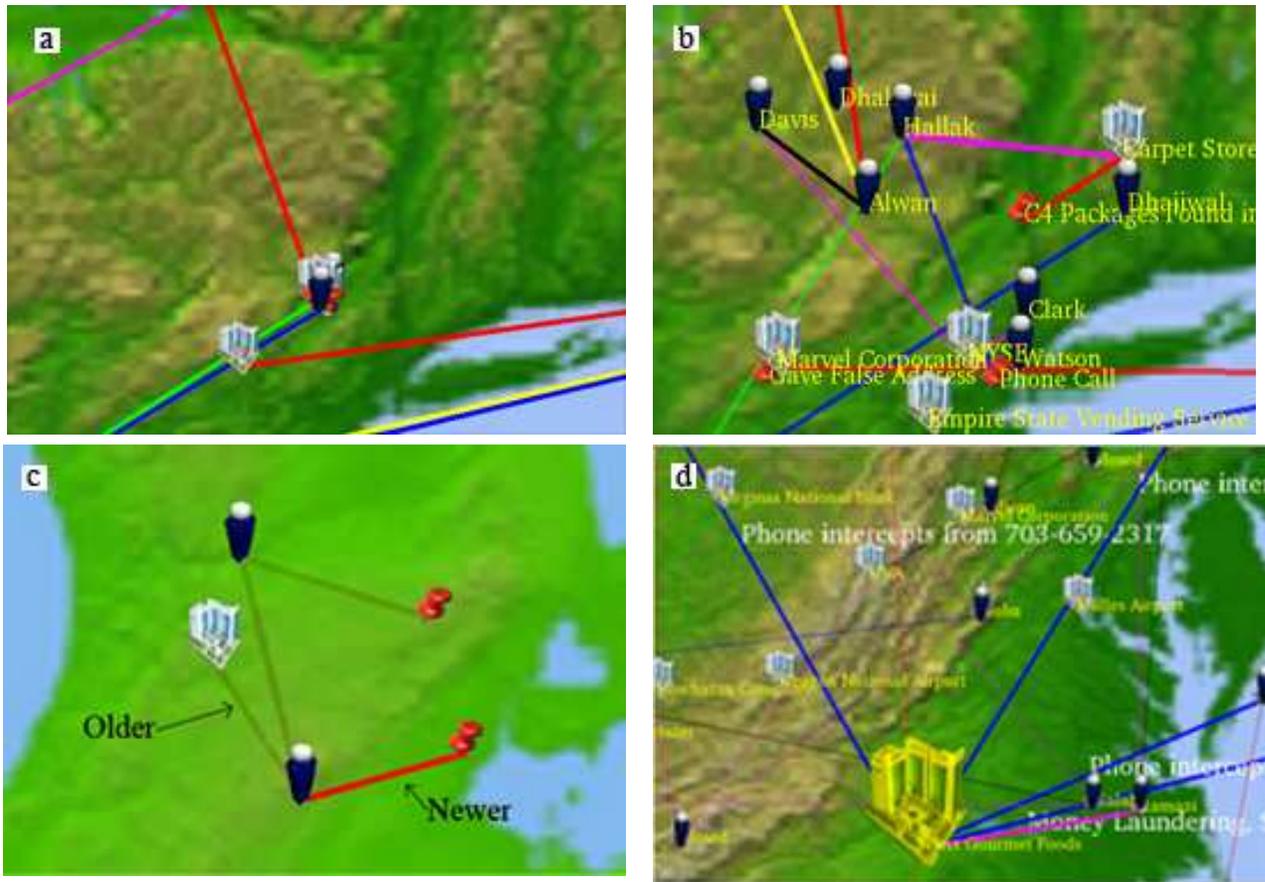


Fig. 5. a) Aggregated nodes, b) Dispersed nodes, c) Edge opacity, d) Selected node.

Further support of the benefits of having a consistent presence of two views comes from the observation the use of the timeline view as an index into the geospatial view, due to its alphabetical listing of names. In smaller visualizations, search tools or an index window may be explicitly accessed, but the constant high resolution alternative view provided a quick way to scan the list of names and still identify the desired person. Linked views then allowed them to select it to be highlighted in both views. This feature is not groundbreaking, but its desirability combined with our displays unique support for it warrants the further exploration of applications like GIANT. Other domains that are heavily dependant on multiple views for information exploration will likely show even stronger benefits from the use of a large high resolution display.

Using high resolutions to show all the data allows for faster access to details and less loss of context. Even though nodes occluded each other in highly active areas, we allowed examination of individual words without the use of a zoom. Instead we dispersed the nodes around their original location, and anchored them back home with a white line. Since the map was stationary, users were able to exploit their spatial memory [10] to maintain context about where earlier clues were located.

The interface design changed considerably when working with a screen space that is physically large. Controls had to be near the user's focus and appropriately sized. A typical slider configuration for a single monitor display was the full width of the screen, positioned at the bottom. In our large display, this would translate to a three meter slider positioned half a meter below the data that users examined. For this reason we used direct

manipulation when possible and a floating window that contained other filters. These features were necessary in GIANT to give users the ability to quickly interact with and analyze the data so that they could make intelligence-related knowledge discovery.

6. CONCLUSION AND FUTURE WORK

By designing and developing a baseline solution for the cited challenges discussed in this paper, a few key techniques arose that contribute to effective visualization of intelligence data.

The first technique provided a solution for the difficult visualization problem of displaying a full range of time on a two dimensional coordinate plane, such as a geographical map. User evaluations revealed that modifying the transparency of edges effectively indicated age and intuitively identified time differences in the data.

Using a high resolution display avoids some of the problems that arise from having multiple views of a large dataset on a single 1280x1024 screen. Both views and extra data used up valuable pixels, yet with our tiled display we kept both views constantly visible, avoided panning and zooming, and still had space leftover for more details.

Even with the extra pixels, there are still design issues with large displays that we developed solutions for. In the cases where information was too densely clustered, we implemented a novel way of examining individual data points without disturbing the context. To deal with interacting with such a physically large display, we developed a mobile and easily accessible user interface. Features like those above are part of the key to maximizing the benefit from developing for these large displays.

Tiled displays, such as the one presented in this paper, show great potential for intelligence analysis. Analysts need visualizations that provide multiple views of large datasets, allow rapid access to details, and preserve the context of data. We demonstrated that the GIANT system satisfies these requirements, thus applications designed for large high-resolution displays are well suited for solving intelligence analysis problems.

This research compilation, application development effort, and noted visualization contributions have moved forward our understanding of the role high resolution large displays can play in intelligence analysis. We have a better comprehension of the benefits of these displays, and the methods that should be used with them. GIANT provided innovative means to show nodes and connections that allowed users to gain further insight into intelligence data. The residual value from our research was high because of the design, architecture, and potential real-world applicability of this system. Despite the accomplishments achieved thus far, there are numerous visualization research questions remaining to be resolved with future work.

Our future work could combine our visualization methods with parsing algorithms or fuzzy logic search tools. Understanding which aspects of intelligence analysis are best suited for visualization, which are best suited to machine processing, and how to best integrate the two is the logical next step. Features are needed that allow analysts to annotate data and store their own hypotheses. Lastly, the system could be extended to better support data entry from the field in order to test how analysts work with a dynamic, growing dataset. GIANT's practicality and effectiveness in providing data connections, trends, and user insights into intelligence data strongly suggests this design be a basis for future research.

7. REFERENCES

- [1] A. Bezerianos, R. Balakrishnan, "Interaction and Visualization Techniques for Very Large Scale High Resolution Displays", University of Toronto Technical Report DGP-TR-2004-002, July 21 2004.
- [2] B. Shneiderman, "The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations", Proceedings of IEEE Visual Languages (Boulder, CO) IEEE Computer Society, pp. 336-343, 1996.
- [3] B. Wei, et al., "Visualization Research with Large Displays: Analysis of Communication Networks and Services", Computer Graphics and Applications, Vol. 20(4), pp. 50-54, 2000.
- [4] Chromium Documentation. Version 1.8. Available at <http://chromium.sourceforge.net/doc/index.html/>, March 23 2005.
- [5] C. Lewis, P. Polson, "Theory-Based Design for Easily Learned Interfaces", Human-Computer Interaction. Vol. 5(2) pp. 191-220, 1990.
- [6] C. Plaisant, R. Mushlin, A. Snyder, J. Li, D. Heller, B. Shneiderman, "LifeLines: Using Visualization to Enhance Navigation and Analysis of Patient Records", Revised version in 1998 American Medical Informatic Association Fall Symposium, AMIA, Bethesda MD pp. 76-80, Nov. 9-11, 1998.
- [7] Distributed Multihead X Project. Available at <http://dmx.sourceforge.net/>
- [8] D. Tan, et al, "With Similar Visual Angles, Larger Displays Improve Spatial Performance", Proceedings of CHI, Ft. Lauderdale, FL, ACM, pp. 217-224, 2003.
- [9] F. Van Ham, J. Van Wijk, "Interactive Visualization of Small World Graphs", IEEE Symposium on Information Visualization, Austin, Texas, Oct. 10-12-2004.
- [10] G. Robertson, "Data Mountain: Using Spatial Memory for Document Management", Proceedings of the 11th Annual ACM Symposium on UIST 98, San Francisco, CA, ACM, pp. 153-162, 1998.
- [11] G. Robertson, "The Large-Display User Experience", Special Issue of Computer Graphics and Applications, Vol. 25(4), 2005.
- [12] J. Hughes, D.A. Schum, "Evidence Marshaling and Argument Construction: Case Study no. 4, The Sign of the Crescent (analysis)", Manuscript developed for exclusive use by the Joint Military Intelligence College, January 2003.
- [13] J. Nielson, R. Mack, "Usability Inspection Methods", New York Wiley, 1994.
- [14] J. Nielson, R. Molich, "Heuristic Evaluation of User Interfaces", Proceedings of ACM Conference on Human Factors in Computing, New York, pp. 249-256, 1990.
- [15] J. Schroeder, J. Xu, H. Chen, "CrimeLink Explorer: Using Domain Knowledge to Facilitate Automated Crime Association Analysis", Proceedings of ISI, in LNCS Vol. 2665 pp. 168-180, 2003.
- [16] K. Hornbaek, E. Frokjaer, "Usability Inspection by Metaphors of Human Thinking Compared to Heuristic Evaluation", International Journal of Human-Computer Interaction. Vol. 17(3) pp. 357-374, 2004.
- [17] Lauren Shupp, Robert Ball, Beth Yost, John Booker, Chris North, "Evaluation of Viewport Size and Curvature of Large, High-Resolution Displays", Graphics Interface 2006.
- [18] M. Miller, "Requirements for High Resolution Displays in Military Systems", High-Resolution Displays & Project Systems, SPIE, San Jose, CA. pp. 90-95, 1992.
- [19] MySQL. Available at <http://dev.mysql.com/downloads/index.html/>
- [20] P. Baudish, "High-Density Cursor: A Visualization Technique that Helps Users Keep Track of Fast-Moving Mouse Cursors" INTERACT 2003, ACM, pp. 236-243.
- [21] R. Ball, M. Varghese, A. Sabri, E.D. Cox, C. Fierer, M. Peterson, B. Carstensen, C. North, "Evaluation of the Benefits of Tiled Displays for Navigating Maps", IASTED International Conference on Human-Computer Interaction, pp. 66-71, November 2005.
- [22] R.V. Badalamente, F.L. Greitzer, "Top Ten Needs for Intelligence Analysis Tool Development", Proceedings of the First International Conference on Intelligence Analysis Methods and Tools. McLean, VA, 2005.
- [23] S. Havre, B. Hetzler, L. Nowell, "Theme River: Visualizing Theme Changes over Time", Infovis, IEEE, 115, 2000.
- [24] T. Buetow, et al, "A Spatio Temporal Visualizer for Law Enforcement", Proceedings of ISI, in LNCS Vol. 2665, pp. 181-194, 2003.
- [25] T. Kapler, W. Wright, "GeoTime Information Visualization", Proceedings of the IEEE Symposium on Information Visualization, INFOVIS, IEEE Computer Society, Washington, DC, pp. 25-32, October 10-12 2004.
- [26] W.C. Elm, M.J. Cook, F.L. Greitzer, R.R Hoffman, B. Moon, S.G. Hutchins, "Designing Support for Intelligence Analysis", Proceedings of the Human Factors and Ergonomics Society", New Orleans, LA, Sept. 20-24, 2004.