Visualizing Cyber Security: Usable Workspaces

Glenn A. Fink*, Christopher L. North**, Alex Endert**, and Stuart Rose*

* Pacific Northwest National Laboratory ** Virginia Polytechnic Institute and State University

ABSTRACT

The goal of cyber security visualization is to help analysts increase the safety and soundness of our digital infrastructures by providing effective tools and workspaces. Visualization researchers must make visual tools more usable and compelling than the text-based tools that currently dominate cyber analysts' tool chests. A cyber analytics work environment should enable multiple, simultaneous investigations and information foraging, as well as provide a solution space for organizing data. We describe our study of cyber-security professionals and visualizations in a large, high-resolution display work environment and the analytic tasks this environment can support. We articulate a set of design principles for usable cyber analytic workspaces that our studies have brought to light. Finally, we present prototypes designed to meet our guidelines and a usability evaluation of the environment.

KEYWORDS: Cyber analytics; cyber security; visualization; usability; large, high-resolution displays; cognitive task analysis.

INDEX TERMS: H.5.2 [User Interfaces]: Interaction styles, Prototyping, Screen design, User-Centered Design

1 INTRODUCTION

Cyber analysts who defend our computer infrastructures use primitive, command-line tools that are ineffective at the high volume and velocity of the data they must process. They have resisted using visualizations, partly because no visualization has yet met their complex needs. We believe this is because their tasks, work environments, and requirements have not yet been studied sufficiently. We need more user-centered design in the solutions we offer. Large displays have been valuable in other applications with massive data [1], and we suspect that they can be helpful in this application, too.

Cyber analytics is a new science of analysis for understanding the behavior of computers and computer networks from the data they generate—discerning the story hidden inside massive cyber data. Many job descriptions include cyber analytic tasks, such as system administration, cyber security, and design and maintenance of computer infrastructures. In this paper, we concentrate on cyber analysis for securing enterprises and large infrastructures of related organizations. We found the behaviors of the cyber analysts we studied were distinct from behaviors of analysts in other domains such as intelligence analysis.

For cyber security professionals, a usable workspace should support multiple, simultaneous, open-ended investigations. Separate tasks that arise from distinct tip-off points may eventually connect, implying the need for an overview of all active tasks and their relationships. Analysts desire to find connections that point to the sources of threats to the system they are defending. Using the analogy of information foraging [2], cyber analysts are tracking big game. Individual clues are only valuable if they support other clues that point to the same root cause. To acquire multiple, complementary information items, cyber analysts rapidly switch between analytic inquiries, multi-tasking and refining or broadening queries as they investigate potential leads.

Analysts need tools that interoperate. Their tools (and queries) are highly specialized, and they spend much of their time joining data tables and translating information between tools. Often the story they seek is hidden within complex correlations that no single view adequately reveals.

This paper contains the results of our study of cyber security analysts, our proposed solutions to selected problems, and users' reactions to our solutions. We identify lessons learned and present a set of design principles for usable cyber analytics workspaces. Our goals were to identify sensemaking processes in cyber analytic work and critical usability issues in cyber analytics workspaces, and to elicit cyber analysts' ideas about how large, high-resolution displays can help them work more effectively. We observed cyber analysts using large displays with a sample problem. Then we created mockups that identify effective uses of large displays in cyber analytics. We sought feedback on our prototypes from the users who participated in our study.

2 RELATED WORK

Many visualizations for cyber security data are special-purpose representations of a particular kind of data. Few tools support interoperability with other applications and utilities. Most tools are neither fast enough nor flexible enough for cyber analysts. There are visualizations [3][4][5] for packet-headers, network flows, system log files, IDS alerts, etc. Optimizing a tool for one type of data separates the tool from the context of an overall investigation. Unless a visualization tool fits into the broad context of the overall investigation it will limit its utility to the analyst. For instance, a cyber analyst might find some interesting alerts via a Snort IDS alert visualization and then wish to investigate the network flows, packets, host log entries, application logs, etc., that are related to these alerts. Most visualizations do not support this kind of rapid, open-ended foraging activity.

Snap-together visualization [6] provides a flexible visualization system of coordinated views that links legacy tools at the database (relation-query) level. Thus a user-defined series of visualizations can be driven from direct interaction with any visualization in the series. Snap would provide the power and flexibility of database and command-line tools, but it has not been applied to cyber analytics. We believe that a coordinated visualization system that interprets interactions and coordinates views relationally is an essential building block of a usable cyber analytic workspace.

Displays of the size we used in our study are not in broad use, so it is not surprising that many visualizations are not optimized for this environment. Mitigations like tabbed windows that help conserve space in small displays are counterproductive in a large display. To make use of these displays, we must develop new window and display management techniques.

Glenn Fink <Glenn.Fink@pnl.gov>, Chris North <north@cs.vt.edu>, Alex Endert <aendert@cs.vt.edu>, Stuart Rose <Stuart.Rose@pnl.gov>

3 CYBER ANALYST ETHNOGRAPHIC STUDY

We interviewed eight cyber analysts at a major government laboratory to find out how large, high-resolution displays can help solve important problems in analysis. Three interviewees were strategic analysts whose job included understanding threats to the organization, four were tactical defenders whose job was primarily to protect their machines and networks from attacks in real-time [7], and one was a developer with experience in tactical cyber analytics. Strategic analysts are accustomed to accessing feeds of compressed network data from multiple installations via a SQL database. Their objectives include understanding the adversaries and broadly characterizing the threat. In contrast, tactical analysts access a wide variety of information from numerous sources using many different tools. Tactical analysts' objectives include maintaining situational awareness and rapid remediation of security problems.

We used a large, tiled display made of eight thirty-inch panels arranged in two stacked rows (Figure 1). The total display area was nearly 33 megapixels in volume and was four to six times larger than the displays the analysts were accustomed to using.

Figure 1: Our 33 MPixel display setup

To provide a framework for talking about displays and visualizations, we presented analysts with generic visualizations of Netflow and Snort alert data using SpotFire (http://spotfire.tibco.com) on the large screen. We also presented a cyber security scenario derived from the upcoming 2009 VAST challenge data set [8] to several of the analysts in each category. We used these exhibits as media for conducting semi-structured interviews.

The analysts told us about their duties, the tools they use, and how they would use a large display. Some were enthusiastic about the display; others were openly skeptical. Some liked the idea of visualizations, while others thought of visualizations as a waste of time. One analyst was particularly critical of the visualizations he knew about. While describing how visualizations, "get in the way of the data," and are "good for people who want to be spoon-fed," he casually noticed an interesting feature in a scatterplot visualization. He said, "that's interesting," drew closer and said, "Wow, that's very interesting!" Then he went back to Excel and the command line to quantify exactly why it was interesting. This feature turned out to be the solution to the problem he had been working at solving for two hours.

3.1 Why do cyber analysts dislike visualizations?

Many of the talented analysts we interviewed, prefer the command line because of its unparalleled flexibility and expressive power. While high-end graphics workstations with speedy processors, and large amounts of online storage make visualization of cyber data viable, not all cyber analysts embrace visualization.

Visualizations of cyber data frequently do not interoperate efficiently with other applications and utilities. Cyber analysts report that visualizations waste time because they require so much effort to import and export data with other tools. The visualizations that frustrate them display particular types of data in specific formats, and were rigidly designed to support preconceived workflows rather than open-ended investigation. Many visual tools were designed to be a monolithic collection of all possible functionality (for example, [9][10][11]). Rather than text-based input and output, many tools use proprietary data formats that further limit interoperability. In contrast, many of the text-based tools were built on the highly interoperable UNIX model of "small is beautiful" and "do one thing well." Analysts can build complex pipelines that connect the output of one tool to the input of the next. While visualizations provide useful information to analysts, they avoid using them because of the frustration of non-interoperability.

The cyber defenders we interviewed distrust visualizations that hide or smooth the underlying data. Access to the source data is critical, but even with details on demand, many seemed more comfortable looking at the actual data line by line. Another problem is that they want to be able to filter, join, and transform the data without losing or altering the original. Visualizations seldom allow flexible manipulation of data, and they can give a feeling of distance and lack of control. One user feared he would irrevocably alter the source data by manipulating the visualization.

Cyber defenders distrust automated reasoning about their data in general because they are accustomed to poor performance of intrusion-detection systems (IDS). The number of false positives emitted by fielded IDS is truly staggering [12], and cyber attackers specialize in adapting their methods to produce false negatives [13]. The defenders we interviewed rely on their own experience, domain knowledge, and hunches over any automated warning system. Cyber security is essentially a human-on-human adversarial game played out by automated avatars. Human cyber attackers succeed by learning to outwit defensive measures, and they often don't follow the rules. For instance, pornographers use their customers as free labor to beat the defenses of sites that use CAPTCHAs (Completely Automated Public Turing test to tell Computers and Humans Apart) [14]. Cyber defenders have learned the hard way that automated detection technologies are unable to defeat determined and creative human attackers. Thus, visualizations that simplify the real, messy data are suspect. It is the simplification, not visualization itself that is untrustworthy, but the two are often conflated.

Cyber domain experts from this and previous studies [15] often consider visualizations to be gentle training for aspiring defenders who have not yet gained sufficient knowledge to handle raw data. Experts consider the ability to read and manipulate massive streams of textual cyber data as a hallmark of their expertise. Although these individuals are often very talented, unfortunately, as one of our interview subjects admitted, "We usually don't find the bad guys until after the fact."

When the analysts we interviewed would find the information they needed in a visualization, they would cross correlate it with other data manually. After noticing an important feature in a visualization of the VAST 2009 challenge data set, one analyst immediately switched to Excel to see exactly when the suspicious events occurred. The visualization was showing a join of the network activity and building access tables by person, but time was not shown. The analyst had selected several outliers, but the visualization (prototyped in SpotFire) had no way to take the selected data items and visually compare the periods of time when they

occurred. So the user switched to a different tool. This lack of flexibility makes visualizations an interesting curiosity, but useless for practical, open-ended investigations.

3.2 Huge volume and velocity of cyber data

Another reason cyber defenders do not use visualizations is that they typically cannot handle the amount of data that they must analyze. For example, the U.S. Department of Energy (DOE) is a widely dispersed enterprise with over 100 sites and approximately 500,000 computers. This enterprise yields an estimated 500 million events per day from sources distributed across North America. In the near future the number of daily log events generated is expected to approach several billion. Contrast this need with the abilities of visualizations that strain to produce pictures of a million items [16] and are unable to do this in real time with streaming data. DOE analysts need a national perspective within seconds to minutes of an event to maintain situational awareness. Even if current visualizations could keep up, there would not be enough pixels on the typical screen to handle this much data.

Strategic analysts we interviewed reported having to analyze about 500,000 new NetFlow records daily in their individual assignments. They would also analyze trends by comparing current traffic with entries in a huge historical database. Tactical analysts dealt with a much wider variety of data types including packet-header data, host log files, vulnerability scanner reports, and external information from numerous sources. Thus, it was much more difficult to quantify an average amount of data that tactical analysts dealt with daily. One analyst reported typically getting to only 25 to 30 percent of the investigations he wanted to achieve in a given day. Analysts expect the data volume to increase by several orders of magnitude in the future.

One tactical analyst described how he begins his daily routine by examining a collection of blogs and websites to find out the new vulnerabilities that developed during the night. Being in North America, he found it helpful to examine the reports from Europe because new attacks often hit them first since they are ahead in the day. From these reports he identifies a list of approximately 50 threats that he needs to examine that day. He prioritizes this list to determine the top 10 that must be addressed. Of these, he is usually able to tackle the top 3-5 during the day.

Slow database access, update, and retrieval causes slow refresh on visualizations, leading to user frustration. Frustrated users said things like, "[It's] quicker to query!" and "I'd rather use grep." Partly this could be alleviated by better database management; e.g., via query previews and incremental return of results. But slow databases increase the need for a work environment that supports multi-tasking. As the user waits for a query to finish, she may be browsing the Internet looking for more information on a vulnerability, or perhaps researching compatibility issues others had installing patches. Then, when the query finishes, the user should be alerted to draw her attention back to that task.

The most commonly used analytical tool was Microsoft Excel. Many analysts used Excel PivotTables to summarize large data sets. Because Excel 2003 PivotTables were limited to 65,536 rows, analysts examined their data in small chunks. Several upgraded to Excel 2007 that allows a million rows in a PivotTable. Other analysts preferred to use databases to analyze data. In either case, the true limiting factor was the speed of access.

One cyber security analyst used Microsoft Excel's conditional formatting to display the numbers of successful connections made by pairs of IP addresses as a colored heat map. He used the map to identify unique pairings of IP addresses and connection patterns indicative of botnet activity. The data-size limitations of Excel and the encumbrances it presents to import and manipulate the data are a price he is willing to pay to create visual patterns that he can quickly recognize.

Visualization designers must consider carefully how much data to store, how long to store it, and how to provide timely access to data that is needed. Perhaps an adaptive strategy that predictively fetches data associated with features like the ones the user has most recently investigated would help. Unfortunately, the huge volume and velocity of the data cyber analysts must deal with causes them to have to investigate in real time or risk losing the data forever. For example, when analysts identify a communication pattern they need to investigate more closely, they must immediately remotely log in to the affected machine to examine various transient details such as current processes, performance statistics, and file system status, which are not logged [5]. Thus, they are likely to miss transient indicators of the problems they are investigating.

3.3 Diversity of data sources challenges tools

Cyber analysts draw information from a variety of internal and external sources to create a context to help them interpret individual events and sequences of events. External sources include news stories, announcements from vendors, official bulletins, vulnerability clearinghouses, and social media like blogs, wikis, Twitter, Facebook, etc. Internal sources include network flows, packet traces, host and application log files, logs from firewalls, IDSs, special host-based monitoring software, and reports from vulnerability scanners, etc. In large hierarchies of organizations, all defenders generally have access to external data sources, but defenders outside a particular site may not have access to data sources internal to that site. Strategic defenders who are responsible for multiple sites may not be interested in any more than high-level network flow records and IDS alert data since they would be overwhelmed if the internal data from all their sites were available. Additionally, internal data are often very sensitive and may be proprietary to the originating organizations.

Pulling information from online free-form text sources to use as a query for a visualization is another task that analysts expressed interest in. For example, an analyst might ask, "show me any of my database servers that made web responses on port 1311 to a host listed as 'bad' in an official bulletin." Such queries are easy to frame verbally but they require a great deal of tool flexibility.

One reason cyber analytics is so difficult is that separate data sources are difficult to join in an absolute time sequence. First, sensors on separate machines may not be time synchronized. Second, some alerts are time-stamped when an event is triggered; others receive a timestamp only *after* a series of events has occurred. Finally, some alerts are logged without any timestamp, making them extremely difficult to join with other data sources.

3.4 Need to have direct access to the data

To the cyber analysts we interviewed, visualizations hide what is going on with the data. They want to know why the visualization shows what it shows. Being able to drill down and get as much detail as possible when needed is a critical requirement.

When using visualizations in our study, analysts would typically investigate spikes (or other irregularities) to determine what caused the feature. Unfortunately, in the interest of efficiency, many visualizations do not store all the data or may over-aggregate, effectively smoothing out "noisy" data. But the adversary seeks to hide in the noise, and over-aggregation contributes to his camouflage.

Most cyber analysts have been exposed to poorly designed visualizations that prejudice them against all visualization. Commonly cited examples are the simplistic visual charts that accompany many of IDS software systems. These charts over-aggregate the data and reveal only the dominant patterns in IP traffic. Sometimes that is useful for seeing major spikes in traffic, but the devil they seek is quite literally hidden in the omitted details.

Such charts often provide extremely limited interactivity so analysts cannot drill down to investigate.

Other cited examples of poorly designed visualizations are those that are generic data visualization tools and not designed for the specific nuances of cyber security. For example, viewing textual values of cyber data is extremely important in certain cases. We observed cyber analysts scanning rows and columns of data in textual spreadsheet format for specific sets of values. Cyber analysts are very skilled at recognizing specific IP address octet values, specific IP port numbers, or specific countries of origin. For example, in cyber security there is a big difference between the meaning of port 80 and port 81, but in a generic scatterplot of packet header data where port number is visually encoded or mapped to an axis, these two ports may be visually indistinguishable. Displaying values when mousing over a dot in the plot would require the analyst to hover over every dot in the vicinity to check its actual value. Simply visually scanning a column of text values or using grep is actually much faster. Embedding text values directly into the visualization could be a potential solution on large displays that offer more space.

3.5 Quest for a query

As we observed cyber analysts in action, we began to refer to their general approach as the "quest for a query." When they investigate an incident, they proceed through a complex analytic process of data foraging and sensemaking to identify the suspicious phenomenon. They explore the data in a variety of ways, essentially looking for a descriptive "query" that returns only the data that concerns the phenomenon that they are investigating. The query must have acceptable degrees of precision and recall so that they can associate the query with a named phenomenon like "users who accessed their computers during odd hours."

Thus, an important product of their analysis is the set of data that represents the occurrence of a phenomenon. However, an even more important product is the "query" itself, the process they used to find the data. In some cases, the query is directly identified at the end of their process, perhaps as a final SQL query. But more often, the query is a record of the *process* an analyst must undertake to obtain the same results in another instance of the same problem. That is, the long sequence of interactions they performed to analyze (process, filter, sort, visualize, reorganize, etc.) the data essentially forms the query. The process is the product. This observation emphasizes the importance of capturing the interactive process and reformulating it into query space that can be reused and shared.

We believe this query-seeking behavior is related to specifying signatures that can be used to automatically find suspect activity or to filter out known safe activity. Finding reusable queries that become part of their domain knowledge is a key goal of cyber analysts that makes them more effective. Once a rule is developed for a specific threat, it can protect the analyst's network from that threat forever. However, it can also be reused to assist in developing new rules for similar threats, such as with virus derivatives, thereby enabling the analyst to rapidly adapt to the constantly changing cyber battlefield. Maintaining rule sets is a common problem for analysts. Rule sets grow large and individual rules become out-of-date quickly. Analysts may not be motivated to share rules because they represent the hard-earned expertise that is an analyst's livelihood.

Visualizations could become more effective tools for cyber analysts if they took advantage of this query-seeking behavior to automatically generate queries based on the features the analyst spent the most time investigating. Some of the cyber analysts we talked to were SQL experts, not by desire or interest, but by necessity. We believe we could better serve these users by providing tools that help frame queries through natural interactions with a visualization rather than via manual SQL statements [17].

A problem with command-line queries is that they force the analyst to formalize their hypotheses too soon. At the beginning of an investigation when there is much uncertainty, analysts are frequently unsure of what to query and need more exploratory means. Visualization tools can support a form of 'incremental formalism' [18][19] that gives analysts the opportunity to begin with informal hypotheses and gradually increase the rigor of their query until the security threat is clearly identified. An example of incremental analysis is a tool called ProSPECT that enables analysts to arrange data sources, marshal data that is relevant to the problem at hand, and create and analyze multiple competing hypotheses that are supported or refuted by the data [20][21].

3.6 Long sequences of activities

Frequently, the quest for just the right query (with acceptable levels of observed precision and recall) takes analysts through a long series of views of the data. But when this process takes hours, or even days to accomplish, they easily forget the steps that helped them arrive at their conclusion. Without a clear recollection of the steps, it is hard for the analyst to report on or share their process with others.

For example, analysts that used SQL queries frequently iterated through many versions of a query while refining it. For version control, they would label resulting table views incrementally, helping them to backtrack and remember how they got to their final query. Unfortunately, the views do not keep a history of the queries used to create them, and very few analysts write notes about their queries, so other forms of process tracking are necessary. One analyst pointed out that when he reuses a query he created in the past, he frequently forgets why he made the query as he did. By looking at a query history, he could reconstruct his thought process and remember the reasoning.

Similarly, the analysts that use Excel frequently save versions of a dataset as they try different pivot tables. However, they often forget which strategies they have already tried and cannot easily return to previous results. Worse, most of the analysts used multiple tools, and tracking processes across them was difficult.

Recording query development history would allow an analyst to learn from his/her own mistakes and to help others avoid wasting effort on unproductive paths. A large display such as the one we used could allow presentation of these steps by taking snapshots as windows changed. But producing a record that spans multiple activities, times, and tools is difficult. Command-line users frequently use their command history to repeat variations on previous actions, but no analogous method exists that spans the many tools available to cyber analysts.

3.7 Many Windows and Multi-tasking

Each investigation typically involves many windows, and analysts typically multi-task among several open investigations. The typical analyst workstation we saw in our study had one or two moderate-resolution displays with 20 to 40 windows open at a time representing multiple active investigations. This meant that more windows were covered or minimized than were visible at any given time. One analyst cited a typical scenario by saying, "I've pulled up an Excel file and I'll look, and say 'I don't understand what I'm looking at,' and that's because it's another case I started two days ago, and it's just the wrong window or tab."

Most of them had a dual screen system, which they complained were not nearly large enough. They typically used one screen as

their primary workspace for analysis tools, and the other for reference or awareness tools. But the layering of windows affords no spatial memory of where a window was located. With a large display, after a study, we would conduct the post interview and black out the screen. We found it was common for analysts to point to areas of the black screen and say things such as, "when I was working with the prox data, over here…" demonstrating to us that they had naturally organized their work space and remembered where they had put different data sources.

Some windows were awareness tools such as security alert websites that they frequently referred to; others were tools that they commonly used during most analyses, such as DNS lookup. These reference windows were left open for easy occasional access. Other windows were more transient with data-analysis tools pertinent to the current investigation. Some analysts said they had to constantly flip through several reference alert windows or tabs throughout the day to avoid missing alerts. Another analyst described a situation in which he had to open many windows to simultaneously log in to a many remote hosts to perform administration tasks and monitor performance.

Analysts had difficulty organizing these windows effectively on their dual-screen workstations. They wanted to be able to organize the windows for each analysis into a separate group, or "case." Some applications used tabs to reduce the number of windows, but each tab actually belonged to a different task—thus the window aggregation was by application rather than by investigation. One analyst said he frequently had over 20 Excel files open at a time, for different tasks, and each time he needed to switch he would have to cycle through all of them because he could not keep track of them via the task bar. Window management could be greatly enhanced with space a large display offers.

3.8 Low-level thinking and tasks

A frequent strategy employed by analysts was to first identify what is normal and then use that definition to highlight the abnormal. We observed one analyst sorting through a large number of alerts to manually filter out normal activity because he "did not want to miss something by using a filter" that might remove too many items at once. This highlights both the importance of showing the user the actual data and the difficulty of distinguishing what is important to the analyst's task and what is not. However, the care this analyst used in examining the accuracy of the filters is a cautionary tale for designers of visualizations that attempt to automatically generate queries. The goal is noble, but accomplishing it in a way that will be acceptable to users will be difficult.

We noticed that several of the analysts tended to work at a very low level, examining the records one-by-one. This was surprising given the large amount of data that they needed to investigate. Even with visualization tools, where we thought the advantage would be to view large amounts of data in parallel, they would frequently use the tool to focus in on one piece of data at a time and sequentially iterate through a dataset. Perhaps this indicates that they are accustomed to tasks in which they are looking for a needle in a haystack.

Mandiant Highlighter (http://www.mandiant.com) was cited as a helpful tool for tasks like this. It focuses on textual representation of log files, but offers helpful interactive features such as highlighting records based on a selected field value, and easily filtering selected records.

We found that analysts also tended to follow a mental cookbook approach. When presented with a new problem, their tendency was to proceed by iterating through a list of common potential targets. Each analyst had his or her own personal list, built up over years of experience. For example, in network traffic, they would look for large uploads, late night traffic, or traffic occurring at extremely regular intervals. However, in some cases, this cookbook process took precedence over higher-level sensemaking to the detriment of their analysis. We believe that if they paused to think more abstractly about the problem that they could think of more fruitful scenarios to investigate. It is possible that their tools and the nature of the data tend to lead them into low-level thinking and a more rote style of work.

Cyber data and tasks are highly structured, making them a tempting target for automation. One analyst said, "a monkey could do it", referring to a repetitive 'look-up' task that he had to perform. For example, if a vulnerability scan returned a suspect IP address, he would then have to go through several different tools in different windows to get information about the IP, such as the host name, its location in the network or building, its OS version and update status, its owner, and the owner's phone number. Our observations of this analyst's activities indicate that he was saying that the steps of this activity were very repetitive and required only very minimal domain knowledge, not open-ended analysis. However, some activities that humans consider very simple are quite difficult to automate (CAPTCHAs [14] are an excellent example).

Unfortunately, analysts are very busy with many such rote, low-level tasks. But if more rote tasks could be automated or reduced, analysts might be freed to pursue higher-level tasks and sensemaking activity. From our study, these higher-level thoughts are often the key piece in solving their task, as they link collections of lower-level findings.

3.9 Beyond 'Yet another packet-header visualization'

Numerous visualizations of packet header data, network flows, and IDS alerts have been proposed and implemented. Several have been released as commercial products. But special-purpose visualizations that work for only one sort of data and do not visualize correlations among different types of data are not adequate for real investigation. Fink showed that correlation of data types was more effective at improving analyst performance than visualization alone, but that visualizing correlated data was significantly more effective than either visualization or correlation alone [22]. Analysts perform standard types of correlation in the course of their normal work, such as correlating network flows to process activity. But many unexpected types of correlation may arise during investigation of new types of problems.

Having 20 to 40 windows open at any given time representing a variety of tools underscores the need for correlating data from one tool to the other. To be effective, visualizations must show users how several sources of information are related. Still more useful would be an approach where users specify the data that needs to be correlated, and the visualization tool presents the correlated data in a comprehensible way. Feedback into the visualization tool is essential to promoting a true dialogue with the data [17].

To support this dialogue, visualization designers must perform careful requirements analysis to understand the true problems hidden beneath the surface problems analysts say they are trying to solve. Most analysts will agree that a packet-header visualization should be helpful. Yet, most analysts do not use packet visualizations in their work. Although these visualizations seem useful at the surface level, the cognitive gaps in cyber analytics occur at a deeper level. Visualization designers must uncover the deeper problems that analysts do not realize they have.

Visual analytics (http://nvac.pnl.gov/agenda.stm) presents an opportunity to take information visualization beyond merely efficient display of all the data to become a means for analysts to

actually work with the data. Analysts should be able to reorganize the visualization to create a result. Ultimately, analysts should be able to go beyond information foraging to modeling solutions.

4 Analysts solving VAST09 cyber security challenge

In addition to the ethnographic interviews, we observed four cyber analysts solving a scenario generated for the VAST 2009 contest [8]. The particular portion of the contest data consists of a collection of building-access and network traffic data—exemplifying how cyber security analysts must correlate different kinds of data to solve problems.

4.1 Study Design

The study consisted of cyber security analysts being given the VAST 2009 challenge dataset consisting of network traffic and physical access information (prox records) for each employee in a fictitious embassy. Records were kept for entering the building, and entering and exiting the classified section (which did not have computer or network access). The remainder of the office was divided up into offices, with two employees per office (Figure 2). The challenge was to use network traffic data, prox access records, and physical office locations to determine whether or not there was a malicious insider exfiltrating information from the embassy . The challenge was designed so that no single source of data was sufficient by itself to solve the mystery.

Each analyst used the large, high-resolution workspace setup as depicted in Figure , curved to provide the optimal setup for a single user [23]. We provided the analysts with a standard set of tools running on Windows XP such as Microsoft Office, a general visualization tool (Spotfire), and the other standard tools native to the operating system. The user interface was a regular wireless keyboard and mouse. We added a mouse pad the armrest of the chair to increase the freedom of the user to move without being tied to a desk. Analysts were allotted two hours to come up with a hypothesis based on their exploration of the data.

4.2 Analysts' processes

Each analyst had a unique approach to solving the challenge. Some used a set of pre-determined queries for specific IP values, ports, etc. Others generated complex pivot tables in Excel to show connections within the data. In some cases, they searched for protocol information on the Internet to understand the data presented. These "cookbook" activities appeared to be a way to orient themselves to the data. As they exhausted their cookbook searches, they began to test hypotheses specific to the scenario.

Most of them heavily relied on viewing the raw data in Excel for the majority of their analysis. Given the relatively small size of the dataset, Excel was able to handle the number of rows easily, and searching did not take a very long time.

Throughout the individual studies, it was apparent the background of the analyst plays a large role in both their tool preference. Some were very familiar with the features of Excel and could produce quick charts and graphs based on their pivot tables. Some analysts who were accustomed to databases rather than Excel found it difficult to adjust to the "find" feature of Excel.

Excel, and most database engines, provide an environment that empowers the user to search, visualize, and edit information. However, we quickly discovered that these tools did not naturally encourage use of the large display area. For instance, Excel was designed to conserve window usage by opening tabs for each active file. While this may be helpful for small displays, this hindered analysts from spreading the various windows out to compare two or more sets of data. We do not recommend this space-conserving behavior because it limits the user's ability to make use of more screen real estate.

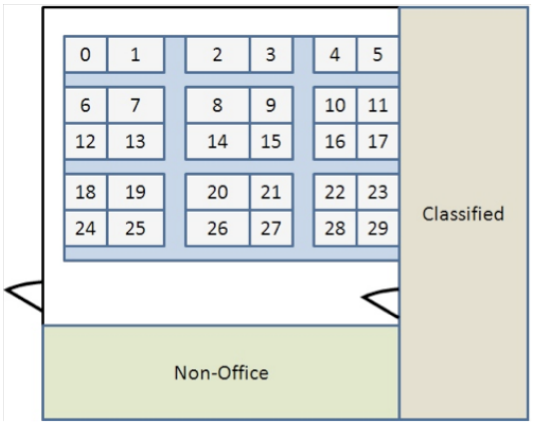

Figure 2: Layout of the office space used in the VAST 2009 dataset.

While most analysts used Excel to orient themselves, one analyst used Spotfire to obtain a visual overview of the data and, "get to know the data." He wanted to "see what the usual routine of a single employee is," and match that visually to what everyone else was doing on a certain day, or compare the visual pattern to what that employee does during the course of the week. He proceeded to conduct the study almost entirely within the visualization, and was among the analysts who attained most accurate hypotheses in the study. He would perform tasks such as coloring prox records a certain color, network traffic another color, visually query the information with the use of filters, and display them in a large scatterplot.

The text-oriented analysts used queries to develop a "normal pattern" against which suspicious abnormalities would stand out. In comparison, analysts could pick out abnormalities with less effort and greater speed visually than with textual queries. In addition, it was easier for visualization users to share their findings and processes with others.

We encouraged analysts to think aloud as they worked, but periodically we had to ask what they were doing, how they arrived at a certain point, or why they chose to explore the current subject further. They were often unable to easily provide us with a history of how they arrived at a certain subset of data, or how that related back to the overall challenge. Reportedly, analysts have this problem in their regular jobs, as well. One analyst explained how he would keep around multiple versions of the same file and use sequential filenames to keep an overview of the entire task.

After the time to solve the problem expired, we would point the text-oriented analysts towards the visualization and show them how solving the task visually might have been beneficial, regardless of their reported hypothesis. At times, this was all we had to do, and certain characteristics and patterns would strike the analysts as interesting, engaging them again in the scenario. Often after becoming more familiar with the visualization, they were able to track down and visually investigate the data by applying filters and certain encodings to the data.

The analysts gave positive feedback about Spotfire. They were surprised that it was easy to learn how to use and very effective for many tasks. They liked the simple yet powerful interaction of the dynamic queries (especially the text search filter for partial IP addresses) and changing the scatterplot axes. In addition, they mentioned the ease of importing data and quickly visualizing it as

an extremely nice feature. However, they found the zooming technique difficult and clumsy, which was further exaggerated when the visualization was enlarged to cover the entire display.

4.2.1 Usefulness of Large, High-Resolution Displays

The goal of information visualization is to compactly present all the data in a single screen, if possible, using strategies such as aggregation to compress the representation down to a manageable size. This compaction naturally causes some features of the detailed data to disappear within the overview. But with large, high-resolution displays, much more space is available, and less aggregation is needed.

Large, high-resolution displays can alleviate some of the constraints of information visualization, allowing analysts to arrange their workspace to reflect their understanding of the data. The extra space can form a "solution space" that allows the analyst to organize multiple lines of inquiry simultaneously. The display flexibly conforms to the mental models of the users. This enables us to go beyond information visualization and begin to understand how analysts use space itself as a problem-solving medium.

4.3 Analysts' solutions

The solutions the analysts provided at the end of their two-hour study varied. Mainly, there was a difference in the amount of detail that the analyst deemed sufficient to warrant further investigation. For example, the specificity of their hypotheses would vary from "a section of the office" that should be placed under further review, to "this IP, which they are sending information to, looks suspicious". Their caution was, in part, driven by their usual reporting system, where the level of confidence plays a large role.

Even with their varying techniques (textual versus visual) and different hypotheses, there was a distinct (and unique) process each analyst went through to get to their hypothesis. A key landmark in their analysis happened when they chose to synchronize and merge the prox data with the network traffic data. Doing so enabled them to see where each employee with respect to his workstation. This often led to comments such as, "How is this employee's machine sending email while the employee is in the classified section?"

When working within the visualization, the investigation went smoothly and quickly, with the analysts finding a view that displayed the events that should be investigated. However, getting them to use the visualization in the first place was difficult. With the text-oriented approach, the majority of the investigation was spent running formalized queries on the raw network data. Partly this was because network data was an information source they were used to and prox data was not. Only when they began to use the visualization were the analysts able to gain a broad overview of the problem. Then, equipped with a map of the building, most were able to come to their correct hypotheses.

5 Large Display Analytic Workspace Prototypes

Using the insight gained from the interviews with the cyber analysts, we derived the following set of design principles for usable workspaces for cyber analysts:

1. Provide history and traceability for investigations
2. Support multiple, simultaneous investigation cases
3. Design visualization tools to be flexible and interoperable to support a broad spectrum of the analytic process
4. Enable the user to interact with the data by direct manipulation of the visual space whenever possible
5. Consider the inherent differences that large displays will have on space, rendering, and ergonomics

We created a set of visual prototypes guided by these principles and used them to elicit further feedback from cyber analysts. Although the tasks of a cyber analyst vary greatly depending on their particular job and task, the general challenges these mock-ups address span many situations.

5.1 Provide history & traceability for investigations

Cyber analysts we observed performing their tasks expressed difficulty staying oriented throughout an investigation. Following their "hunches" may guide them down different paths of analysis, but they often stated that getting back to where they decided to pursue one of these branches could be very difficult. Analysts said that this lack of traceability made it difficult to report their strategies, interactions, and findings. Analysts reported that both productive and unproductive hunches are important to formally document in the written report. The final report is the most common way cyber analysts "share with their supervisors and fellow analysts." The number and quality of reports may contribute to analysts' performance evaluations, so analysts are motivated to present them clearly. Maintaining temporal and logical orientation over the lifetime of an investigation would help analysts clearly state their conclusions and how they arrived at them.

When we suggested that tools should allow analysts to take notes on their analytical process as they solved the problem, some responded that "training themselves to take notes" might be additional work at the time, but could save hours of work writing the final report. Merely providing a history of activities in the context of each investigation would improve traceability and the thus the quality of the reports. Such a history would also save time.

Figure 3: History trees, a mockup visual workspace showing history and allowing workflow traceability through the use of key frames, resulting in the final report on the far right

Figure 3 shows a PowerPoint prototype of a history tree workspace that can provide orientation and traceability over the life of an investigation. History trees provide a means for easily retracing their steps when it comes time to produce a report. Using the additional space of a large display, analysts can "fork off" instances of their tools, and pursue branching hunches in parallel. Windows along each branch of the history tree are running instances of their tools or windows enabling the user to easily backtrack to an earlier state and remain oriented to the entire task.

The larger windows are "key frames" that mark a state in the investigation that an analyst deems particularly important. They might be branch points where the analyst can create a new instance of the tool he is using to pursue a new hunch. The size of each history window is proportional to the age of its most recent use or to the frequency it is consulted. Seldom used displays slowly become smaller unless the user refers to them by hovering the mouse over them, clicking, or resizing them. They never disappear until the user deletes them or the branch they live on, so the user can easily review his thought process and regain orientation quickly when switching among branches.

Although a usual investigation will be much more complex than the one pictured in Figure 3, we received positive feedback on the value of history trees. This approach flexibly enables users to visually represent their thought processes as any number of these paths, write their thoughts using "sticky notes" appended to points along the paths, and to freely interact with any part of this visual workspace. One analyst said that "an integrated workspace/work flow tool that is self-documenting as much as possible but provides space for the analyst notes and thoughts" would be ideal. This workspace would provide a historical record so that an analyst could re-visit the work and re-assess his recommendation when new information became available. Analysts said that the history tree would be beneficial when training other cyber analysts by showing both the complexity of the cases and the thought process that led to the final report. One analyst also commented that this layout would be extremely helpful to the legal and human resources departments to be able to trace back through the analyst's process and see how they arrived at each conclusion.

5.2 Support Multiple Simultaneous Investigations

Most of the cyber analysts we interviewed handled multiple cases at once. Only one handled investigations serially. Many analysts relied on tools such as Big Brother (www.bb4.com) for a quick overview of the state of each machine they were monitoring. Typical views provided by this tool show the status of a set of machines as a set of green, yellow, orange, or red glyphs. Given a limited amount of space, this works fine.

Unfortunately, these overview tools were not designed to show more information when given more visual space. Enlarging the window should provide more information, but instead it just makes the simple glyphs larger. We observed that regular usage of this software requires a user to click through a series of options and windows to get to the detailed data. At times, the events happening in real-time that are causing these glyphs to change color are not logged. An analyst must be quick to react and sometimes must remotely log on to the machine to investigate the problem. This demands both mental workload and additional display space.

When an analyst pulls up detailed information to de-aggregate the glyph, he must organize and manage a new set of tools, windows, and tasks on the machine of interest. Investigations into why a glyph reports a "bad" state can be very complex and may require multiple tools. Analysts must observe and investigate many machines at once, thus, they build up multiple cases that they are working on simultaneously. Often, information from one case applies to another. The ability to have all the collected information visible is important, as it allows the analyst to make semantic connections between the gathered information more easily. On a dual-monitor workstation, the many layers of overlapping windows and a task bar overflowing with minimized windows cause organizational difficulties. Window layering and minimization are interactions required by a lack of display space. There is not enough space to group a set of tool windows into "cases" so an analyst can gain a quick overview of all the material pertaining to a single investigation.

Figure 4 shows Cases, a PowerPoint prototype enabling users to organize their workspace into multiple cases each with a set of tools of their choice and an aggregated view of the overall state of each case for quick reference. The Cases prototype encapsulates all of the visualizations, terminal windows, and analyst's notes into a case area. It encodes the aggregated state of each case as the background color of the case area. Inside the case area are the tools and displays the analyst can use to make a more accurate judgment on what is actually happening. These cases may be minimized, rearranged, and shared with remote collaborators. For instance, once an analyst has the information necessary to know what to watch for on a given server, the case can be minimized, or aggregated, to a smaller graph representing the specified activity. This is more beneficial than minimizing a group of windows to the task bar, as it keeps them coupled based on the case. Likewise, rearranging the cases on the screen space is also beneficial, as it allows for persistence of the information and freedom to organize the cases however a user sees fit.

The ability to collaborate, both locally and remotely, is one that the interviewed analysts greatly desired. Barrett [24] noted that communicating system state and other context information among information technology workers was a typical source of problems despite access to a variety of communication means. The same ethnographic study noted that gathering all the relevant informa-

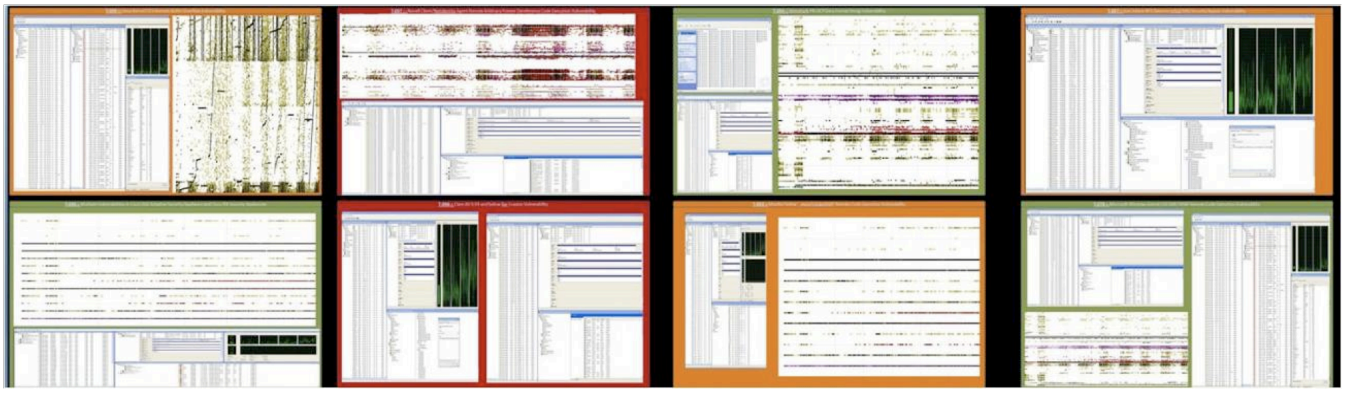

Figure 4: Cases, a workspace showing a collection of cases containing their relevant information within the aggregated status (color)

tion together to share with a collaborator was also difficult. Collaboration is crucial because diverse expertise allows each analyst to approach the problem from a slightly different angle. One analyst said he often reached a point where he would stop and ask someone else for her insight on the situation. Rather than a piecemeal transaction, sending an entire case would allow analysts to provide the collaborator(s) complete insight into their discoveries up to that point. The case includes all the visualizations, history, and notes that the analyst has used to arrive at his conclusions. This complete set of information and insight will accelerate the time needed to reach common ground among the collaborators, allowing them to spend more time solving the task instead of getting up to speed on the problem.

Analysts believed that this form of collaboration support would be mainly beneficial to network administrators, as they are the ones who would most commonly monitor individual cases of machine status. However, analysts may be investigating instances of a potential security issue at multiple sites, and we believe using a Cases approach would help organize these investigations as well. Strategic analysts could share their cases with tactical analysts as living, dynamic reports rather than as a single final report that often lacks critical information about the analyst's sensemaking process.

5.3 Support a Broad Spectrum of the Analysis Process through Flexible, Interoperable Tools

Cyber analysts use many text-based tools for analysis, but no single tool can support the entire spectrum of the analysis process. Command line tools can be linked output-to-input in pipeline fashion. Visualization tools need to be capable of similar flexibility and interoperability. Unfortunately, many visual tools are monolithic [9][10][11], attempting to be the solution for every data processing need. This generally implies that all the data must be imported into the monolithic tool before analysis can occur. This is not suitable for streaming data from multiple, diverse sources. Another common approach (usually found in the research community) is to focus a visualization on a single type of data such as packet-header data or network flows. This approach yields stovepiped visualizations that cannot communicate with other tools automatically.

The analysts voiced their frustration with stand-alone visualization tools, as they are very specific in what they can show. In addition, the tools often did not allow for integration into other tools, forcing the analyst to import and export data to manually link to other tools. It is critical for adoption that tools be built to interoperate easily with the tools cyber analysts already use. While many of the analysts we interviewed were quite open to new technologies, many of them used specialized tools that performed niche functions specific to their work. Interoperability must extend to these legacy tools to create a usable workspace.

Cyber analysts need visual tools that interoperate with the same ease as their text-based tools. Since much of the data cyber analysts use is imported into databases, visualizations should support the relational model of data organization. Visualizations should represent tuples from database relations and provide: unique identifiers for tuples (primary keys), well-defined data extraction capabilities (queries), and explicit representation of relationships (joins) [6]. The relational model goes beyond merely providing a network of filters or a set of constraints connecting visualizations– it enables tools to be joined flexibly at the data relation level.

Cyber problems continually change character as attackers adapt their methods. Neither monolithic nor stovepiped visualizations

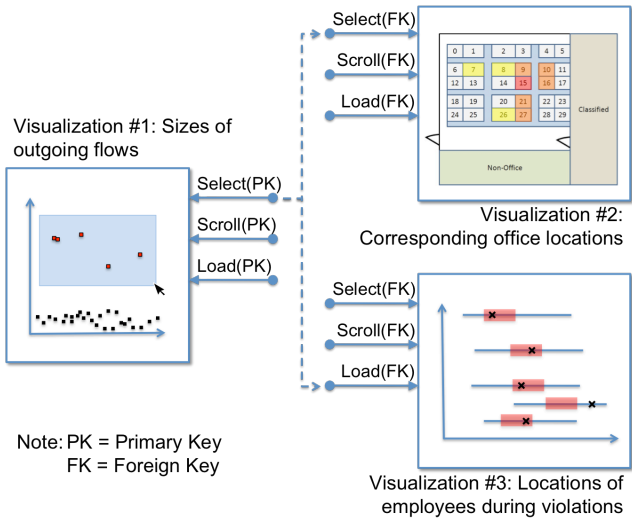

Figure 5: SNAP-linked visualization where selecting large outflows in visualization #1 selects corresponding offices in visualization #2 and loads employee locations into visualization #3

will solve the problem once and for all. The changing cyber battlefield requires highly configurable, interoperable tools that can be flexibly arranged to discover new patterns.

5.4 Enable Interaction via Direct Manipulation

Many visualizations analysts used, such as Excel charts and graphs and other visualization packages, accepted interaction only via the data. The visualization was a passive object that showed data that had previously been manipulated using other means. Since they do all their data manipulation in text-based tools, visualizations are peripheral to the analysis process. Often visualizations were only used when we suggested that an analyst try them. We believe that visualizations will only become intrinsic to the cyber analytic process when they are strongly tied to the underlying data and when interactions with the visualizations are direct dialogue with the data.

This direct tie implies that visualizations should connect using the relational model discussed in the last section. Selecting tuples in one visualization may cause another visualization to load corresponding data in a different view. For example, using the VAST 2009 Challenge data set, a user might visualize the traffic and identify the largest outgoing flows. Using Snap-like connections, selecting these flows in one visualization could highlight the offices where the flows originated in a map view. Another visualization that displayed all the in and out prox records could show where those workers were when the large out-flows took place (Figure 5).

Direct manipulation would mitigate the reticence to use visualizations we witnessed in cyber analysts. Coupled with flexibility and interoperability, direct manipulation will help researchers meet the challenge of making tools so useful and compelling that analysts will choose to use them over text-based interfaces.

5.5 Consider the Impact of Large Displays

When designing a usable workspace, one should consider the inherent differences that large displays will have on space, navigation, and ergonomics.

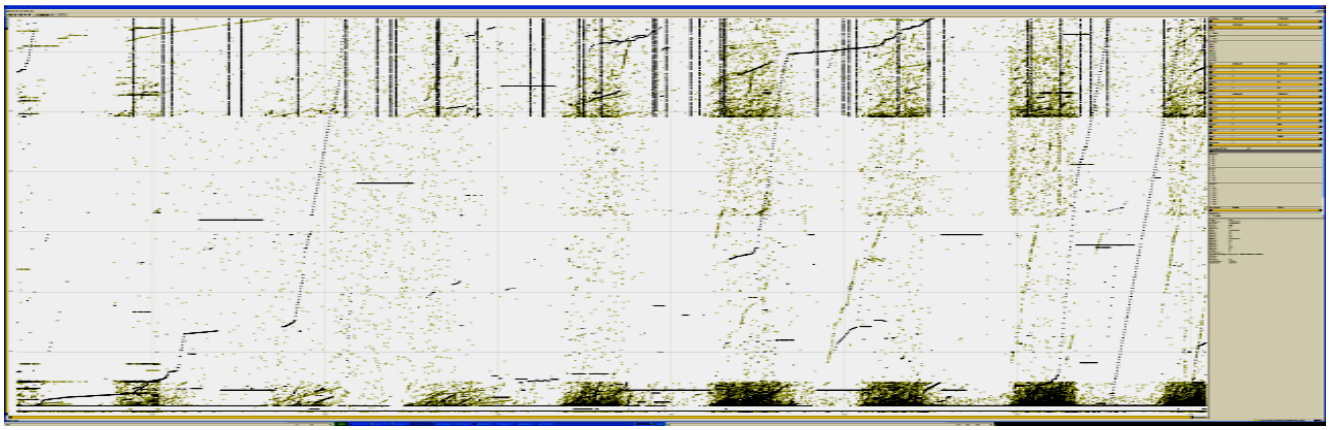

Figure 6: Scatterplot in Spotfire displaying more than 533,000 datapoints without aggregation. Analysts were able to use even simple visualizations to identify recognizable features in the data.

Space: Large, high-resolution displays give a tremendous advantage to users because they can display so much information. For example, over 533,000 snort alerts may be displayed legibly without panning, zooming, or aggregation (Figure 6). The large display allows the visualization to show much more data without aggregation (unlike an enlarged projection of a regular-sized screen).

Cyber analysts pointed out that often the actual threat resides within the noise, so "removing statistically insignificant data due to space constraints will often remove the threat". On a smaller zoomed-out visualization, these would not be recognized because the algorithm is forced to smooth the data, hiding some of these details from the analyst. By displaying more data, we are also able to show the user a longer window of time into the data. For instance, when displaying snort alerts on a normal-sized display, the sheer number of alerts will quickly fill up a visualization requiring aggregation or scrolling. With a large enough display, analysts can see days, or even weeks, of information, enabling them to recognize patterns of events over a longer timeframe.

Navigation: With a visualization on a large display a user can physically navigate the within the visualization, moving his body away from the display to get an overview and coming in close for details. Physical navigation has been shown to improve a user's performance and awareness [25]. Designers might take advantage of physical navigation by using low-contrast colors or thin lines for details that can only be seen close-up. This way the details do not distract from the larger patterns visible only in the overview. When moving closer to the display, one can analyze the full details of the data, often recognizing detailed patterns and sequences within the noise. But these details virtually disappear when the user has stepped back for an overview.

The degree of freedom a user has to physically navigate depends on the configuration of the display (Figure 7). We used two configurations: open (nearly flat) for groups of standing analysts and closed (horse-shoe shaped) for an individual analyst who is usually sitting. Standing analysts have more freedom to physically navigate, but sitting analysts in the closed configuration need only rotate in their chairs to navigate.

Physical navigation is limited, though, by legacy input devices such as the mouse and keyboard. Users preferred not to be tethered, but wireless keyboards were difficult to carry and use. We experimented with using a rolling keyboard stand (for standing users), but at overview distances, the font was too small to allow typing. Users seemed to have the best results when seated in the center of a wrap-around display. This kept them near enough to see the text but still allowed them to back up a little to get an overview. Using tethered keyboards encouraged standing users to come close the screen to type. Touch screens and motion capture are other possible interaction approaches we hope to experiment with in the future.

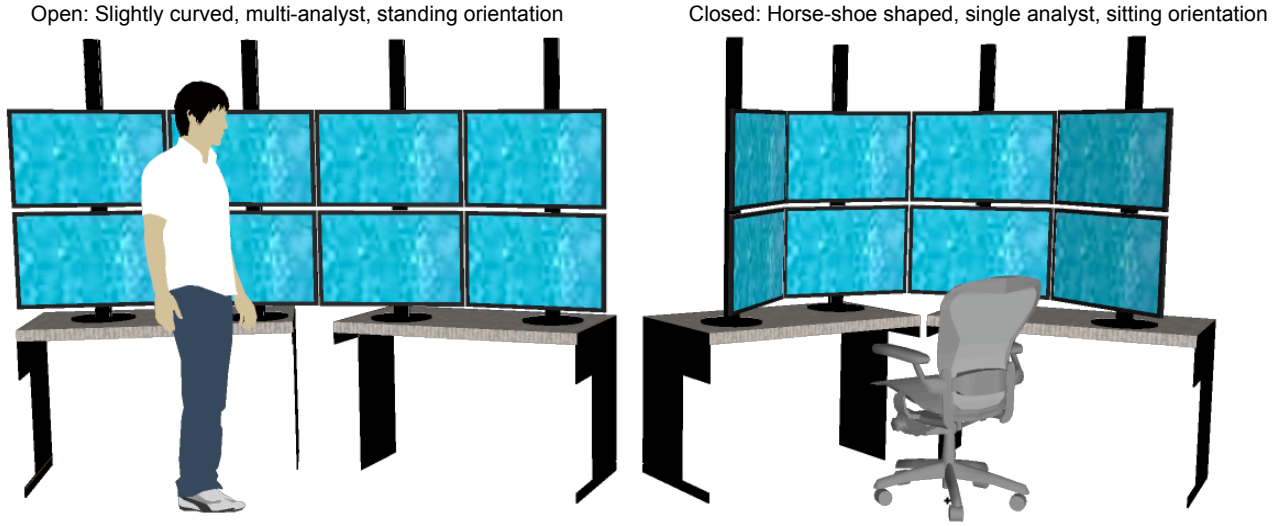

Figure 7: Alternative display setups

Navigation difficulties also stemmed from the design of existing window managers. When a modal dialogue box appears five feet away from the place the user is looking, it takes some time to figure out why the machine seems to have frozen. Users also found it difficult to precisely click on objects that are relatively far away, and losing the mouse cursor was a common malady. One analyst declined to participate in a follow-on study because, "with the mouse and keyboard interface the displays are too cumbersome to navigate." The impact of navigation differences is an important factor to consider in workstation design. We believe window manager and input device improvements could go a long way toward making large displays more pleasant to use.

Ergonomics: Designers of large-display workspaces must consider the ergonomic impacts of their designs. We had an ergonomics professional evaluate our large display and make recommendations. Her recommendations were based on our descriptions of the tasks rather than actually being able to watch analysts using the display for their work. Thus, some findings may be overly general. The interview subjects' impressions often disagreed with the professional opinion of the ergonomics specialist, but the analysts only used the display for two hours at a time, not over many days. We found the professional review instructive and we plan to include ergonomics in future, *in-situ* studies where we can be obtain longitudinal findings.

Our display is made of eight 30-inch panels (four columns by two rows) with each column mounted on a two-monitor desktop stand. The desk is a standard, 30-inch, height and the monitor space extends vertically from 30 to 64 inches above the floor. We showed the ergonomics specialist both configurations of the display.

In the open configuration for a group of standing analysts our ergonomics advisor suggested that the height of the top row was fine, but the bottom row may be too low for effective work. In general, monitors should be placed at just below eye level and about 20 inches away from the user. Subjectively, we estimated the comfortable distance between the users' eyes and the monitors would be 40 inches, which is too far away to see normal-sized text clearly. While using touch-panel displays might help users stay close to the display so they could see the type, users might be uncomfortable standing so close to the huge display space. After a two-hour period of use, some users felt almost snow-blind from the brightness of the displays. For prolonged work, close to the display, the monitors became uncomfortably hot, requiring us to add ventilation to the lab. We found we could reduce both these ill effects by using a black screen background and by decreasing the intensity of the monitors to the lowest comfortable setting.

The specialist believed that the large display would be overwhelming for up-close work. With a group of analysts working simultaneously, she suggested we would also want to have space between the individual workspaces to avoid confusion over who owned what display space. The specialist suggested that the bezels surrounding each monitor would help people group their tasks and could provide the visual space between analysts as they worked. The bezels did influence users to arrange their windows so that they did not cross monitor boundaries, but the overwhelming user consensus was that the bezels were annoying and restrictive.

For the sitting, closed configuration, the ergonomic specialist suggested that the upper row displays were too high for use as primary workspaces and should be used only for reference windows. The horseshoe configuration encourages the user to remain 20 to 30 inches from the displays but again, the ergonomic specialist believed that many people would find the display overwhelming and mentally taxing. In contrast, several analysts we interviewed commented favorably on having a wall of screens surrounding them, but this was based on only a two-hour trial.

With the closed arrangement, the specialist was concerned that repeated head turning could cause neck injuries. She suggested that we should find a way to encourage users to turn their bodies instead. In several cases, we used a gyro-mouse attached to the back of the user's chair to cause the cursor to move across the screen as he or she turned in the chair. Analysts who used this arrangement found it so completely natural that several didn't notice that this wasn't a normal behavior. One said, "that's the way it should be." This "chair mouse" would help encourage users to turn their bodies and avoid repetitive-stress neck injuries.

6 Conclusion

From our study, we gained insight into the complex realm of cyber analysis and made progress toward providing usable workspaces for analysts. To confirm what we learned, we observed analysts on a large, high-resolution workspace solving the VAST2009 challenge. Then, we prototyped visualizations that address the concerns we received from the interviews and conducted follow-up interviews to receive feedback on our designs. Finally, we conducted an ergonomic review of the setup to estimate the long-term effects of large displays on analysts.

Cyber analytics is a clear fit for both visualizations and large, high-resolution displays. We learned that indeed the additional display space and resolution is beneficial. However, the tools, input devices, and window managers currently used in this field do not make good use of the added display space. Worse, regardless of display space, the available visualizations do not fit the needs of the analysts in the total context of their investigation. For this and other reasons we have discussed, visualizations are currently not widely used among cyber analysts.

Cyber analytics is a relatively new science, and the community relies heavily on custom-made command-line tools to support investigations. The ability to conduct investigation visually helped analysts to raise their level of awareness beyond manipulating data to thinking more holistically about problems. We believe visualizations will help cyber analysts to both identify problems, and to work visually towards finding solutions. Through our prototypes, we proposed several ways that large displays could help cyber analysts by making visualizations more useful and usable.

Our recommendations form an agenda for cyber analytics research. The following solutions would greatly enhance the performance of cyber analysts:

1. A way to provide rich linkages among multiple visualization tools that better support the entire process of analysis.
2. Tools that help frame queries built from natural interactions with the data rather than via SQL statements.
3. A means of keeping a visual history of the manipulation steps analysts took to achieve a particular representation.
4. Input devices, controls, and window managers that work well for large displays.

Cyber analytics is an emerging discipline with a distinctively different approach from other analytics domains. Researchers should respond to the unique needs of the analyst community with usable tools and workspaces. We encourage more ethnographic studies that will help researchers understand the entire spectrum of cyber analysts' investigative work.

In the end, our challenge is to help analysts increase the safety and soundness of our digital infrastructures by providing tools and workspaces that are more effective than those that are currently available. We believe large, high-resolution displays with interoperable, flexible, and compelling visualization tools are core components of a usable workspace for cyber analysts.

ACKNOWLEDGEMENTS

The authors are grateful to the Information and Infrastructure Integrity Initiative (I4) and the National Visual Analytics Capability (NVAC) located at Pacific Northwest National Laboratory for their support of this work. We also thank the analysts and their managers who provided access to so much useful information.

REFERENCES

[1] Ball, R., C. North, et al. (2007). Move to improve: promoting physical navigation to increase user performance with large displays. *Proceedings of the SIGCHI conference on Human factors in computing systems*. San Jose, California, USA, ACM.

[2] Pirolli, P. and S. Card, Information foraging in information access environments, in Proceedings of the SIGCHI conference on Human factors in computing systems. 1995, ACM Press/Addison-Wesley Publishing Co.: Denver, Colorado, United States.

[3] Conti, G.: Security Data Visualization. No Starch Press, SanFrancisco, CA, USA (2007)

[4] Gregory Conti, Kulsoom Abdullah, Julian Grizzard, John Stasko, John A. Copeland, MustaqueAhamad, Henry L. Owen, Chris Lee, "Countering Security Information Overload through Alert and Packet Visualization," *IEEE Computer Graphics and Applications*, vol. 26, no. 2, pp. 60-70, March/April, 2006.

[5] Fink, G. A., Duggirala, V., Correa, R., and North C. "Bridging the Host-Network Divide: Survey, Taxonomy, and Solution." In *Proceedings of the 20th Large Installation System Administration Conference* (LISA '06). Washington, DC pp. 247-262, USENIX, 2006.

[6] North, C., "A User Interface for Coordinating Visualizations based on Relational Schemata: Snap-Together Visualization," University of Maryland Computer Science Dept. Doctoral Dissertation, (May 2000).

[7] Fink, G. A., D. McKinnon, S. Clements, and D.Frincke. "Tensions in collaborative cyber security and how they affect incident detection and response" chapter in *Collaborative Computer Security and Trust Management*. IGI Global, to appear.

[8] VAST challenge dataset, http://hcil.cs.umd.edu/localphp/hcil/vast/index.php

[9] Roth, S., Lucas, P., Senn, J., Gomberg, C., Burks, M., Stroffolino, P., Kolojejchick, J., Dunmire, C., "Visage: a user interface environment for exploring information", Proc. Information Visualization, IEEE, pp. 3-12, (October 1996).

[10] Livny, M., Ramakrishnan, R., Beyer, K., Chen, G., Donjerkovic, D., Lawande, S., Myllymaki, J., Wenger, K., "DEVise: integrated querying and visual exploration of large datasets", *Proc. ACM SIGMOD'97*, pp. 301-312, (1997).

[11] Risch, J.S., et al. The STARLIGHT information visualization system. In Information Visualization, 1997. Proceedings., 1997 IEEE Conference on. 1997.

[12] Axelsson, S. The Base-Rate Fallacy and the Difficulty of Intrusion Detection in ACM Transactions on Information and System Security (TISSEC) archive. 2000: ACM PressNew York, NY, USA.

[13] T. H. Ptacek and T. N. Newsham, "Insertion, Evasion and Denial of Service: Eluding Network Intrusion Detection", Secure Networks, Inc., Jan. 1998. http://www.aciri.org/vern/Ptacek-Newsham-Evasion-98.ps

[14] Luis von Ahn, Manuel Blum, Nicholas J. Hopper, and John Langford, "CAPTCHA: Using Hard AI Problems for Security" chapter in *Advances in Cryptology — EUROCRYPT 2003*. Lecture Notes in Computer Science, Springer Berlin / Heidelberg. 2003

[15] Fink, G., R. Correa, and C. North. "System Administrators and their Security Awareness Tools." 2005. http://people.cs.vt.edu/~finkga/Research%20Defense/System%20Admins.html.

[16] Keim, D.A., et al., Pushing the Limit in Visual Data Exploration: Techniques and Applications, in KI 2003: Advances in Artificial Intelligence. 2003. p. 37-51.

[17] Endert A, North CL, Andrews CP, and Fink GA, "The Visualization Pipeline is Broken", Pacific Northwest National Laboratories report number PNNL-SA-65619, 2009.

[18] Frank M. Shipman, I. and C. C. Marshall (1999). "Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems." Comput. Supported Coop. Work **8**(4): 333-352.

[19] Hsieh, H. and F. M. Shipman (2002). Manipulating structured information in a visual workspace. *Proceedings of the 15th annual ACM symposium on User interface software and technology*. Paris, France, ACM.

[20] Sanfilippo, A., *et al*. "Building a Human Information Discourse Interface to Uncover Scenario Content" in*Proceedings of the 2005 International Conference on Intelligence Analysis*. 2-6 May, 2005 McLean, VA, https://analysis.mitre.org.

[21] Sanfilippo, A., *et al*. A Layered Dempster-Shafer Approach to Scenario Construction and Analysis. in Intelligence and Security Informatics, 2007 IEEE. 2007.

[22] Fink, G. *Visual Correlation of Network Traffic and Host Processes for Computer Security*. (Ph.D. dissertation, Virginia Polytechnic Institute and State University, 2006), pp. 105-112.

[23] Shupp, L., R. Ball, et al. (2006). Evaluation of viewport size and curvature of large, high-resolution displays. *Proceedings of Graphics Interface 2006*. Quebec, Canada, Canadian Information Processing Society.

[24] Barrett, R., *et al*., Field studies of computer system administrators: analysis of system management tools and practices, in *Proceedings of the 2004 ACM conference on Computer supported cooperative work*. 2004, ACM: Chicago, Illinois, USA.

[25] Ball, R. and C. North (2005). Effects of tiled high-resolution display on basic visualization and navigation tasks. *CHI '05 extended abstracts on Human factors in computing systems*. Portland, OR, USA, ACM.