12TH ICCRTS
ADAPTING C2 TO THE 21ST CENTURY

Impact of Network Performance on the C2 Mission: Visualization for Network Awareness

Topic: Cognitive and Social Issues

Brian P. McKenna, James W. Gualtieri, & William C. Elm

Point of Contact: Brian P. McKenna

Resilient Cognitive Solutions, LLC
113 South 26th St., Suite 200
Pittsburgh, PA 15203
412-499-3064 x29 (voice) 412-499-3077 (fax)
bmckenna@resilientcognitivesolutions.com

Abstract

The transfer of intelligence related data is occurring over larger and larger computer networks. In these networks, data are exposed to choke points and dropped nodes, where they can be delayed or lost, impacting the larger intelligence mission. As the command center is the end user of this information, it must also be responsible for making decisions about how to control this process. An example of a network visualization that enables more effective, real-time network monitoring is described, highlighting the cognitive support provided to the intelligence operations center command team. This visualization allows operators to monitor over 20 times the number of nodes which could be monitored in the previous system, while also providing explicit *Observability* of network information that was previously hidden in individual drill down screens for each node (showing over 20,000 parameters at one time) and revealing the relationships between network nodes. Cognitive Systems Engineering provided the necessary work domain insights and associated representational design techniques that enabled the development of an innovative solution to the long-standing networking monitoring need. The underlying, generalizable design principles used in this visualization are discussed as a means of enabling their use by other C2 decision support designers.

Introduction

Successful Command and Control decision-making is dependent on having the right information at the right time. Sun Tzu (1981) argues that the person who knows the enemy and knows himself will win every battle. Following this maxim, much effort is spent on the intelligence collection and analysis process. The goal of this process is to

collect enough information about the enemy so that the proper preparations can be made, all the weaknesses of the enemy can be exploited, and all strengths of the enemy can be countered. Knowing what the enemy is going to do before it is done is a decisive advantage that must be exploited at every opportunity; this is a must for the command center.

Recently, modern technology has greatly increased the reach and speed of intelligence collection. Data can be collected in one location and transmitted via computer networks to a location anywhere else in the world. A command center is able to get information from across the world in very little time. Whereas in the past, it might take hours or days for information to reach a commander, information can now be transmitted almost instantaneously to those who need it. Decisions can be made while the information is new, maximizing the impact that information can have. This is also the impetus for assembling the Global Information Grid (GIG) – a globally interconnected, end-to-end set of information capabilities (Daly & Tolk, 2003). By quickly sharing information, decision-makers around the world can collaborate, sharing information and arriving at the best conclusion on a course of action.

While the benefits of transmitting data via a computer network are clear, the use of computer networks also leaves the intelligence process (and thus the entire decision-making process) wholly dependent on the network. Due to the importance of intelligence collection, it becomes imperative to monitor the transmission of data from the original source to command center. Data must arrive at the command center for it to inform and influence the decision-making process. Any problem within the network will impact the larger intelligence mission, and therefore, decisions in general. Being transmitted

through a computer network exposes data to choke points and dropped nodes within the network; the data can easily be delayed or lost, greatly impacting the decision-making it is intended to support. Therefore, the command center, which is dependent on the transmission of the information, must also be responsible for making decision related to controlling the information flow based on the current state of the network world; this makes network monitoring a command and control process.

A key to ensuring successful network operations is providing visualizations that offer a thorough understanding of how the network behaves (Becker, Eick, & Wilks, 1995). This paper will describe a visualization that has been created to assist a command center with monitoring one such intelligence information network. Further, the principles behind designing the visualization will be explored, as they can be extrapolated to any command and control visualization. Finally, lessons learned from dealing in such a difficult domain will be detailed in an effort to help others who may face similar issues.

Display Description

The network for which the visualization was built works in a store and forward fashion. Messages come into a system, where they are stored until they are ready to be processed. After processing, they are stored again, until they can be sent to the next system. Problems can occur for a number of reasons, including corrupt messages and broken processes. Whenever a problem occurs in one system, it inevitably affects other systems, as the queues in which messages are stored, with nowhere to send messages to, become backlogged.

Understanding the need for advanced decision support for network monitoring, a new visualization was created for a customer. The previous version of the display (Figure 1) being used allowed the analysts to monitor around 150 systems at one time on the main screen, but contained very little information. No information was provided on the main display about what was causing errors or even where in the subset of information the error is occurring. The status of the sensor was color-coded based on the most serious condition found in the pop-up menu, which contained the number of messages in each queue as well as the status of each process and hard drive.

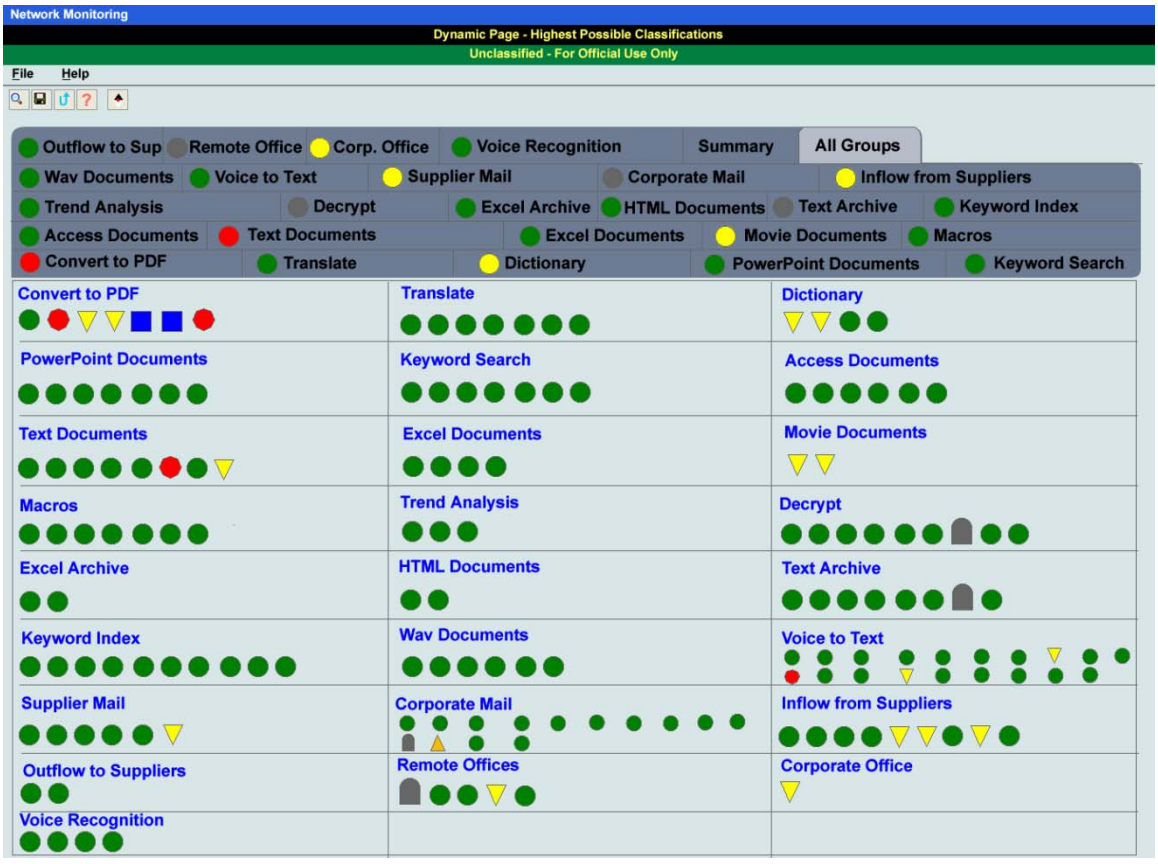

Figure 1. The previous version of the network monitoring display showed only a minimal amount of information – allowing it to show approximately 150-200 pieces of information at one time.

Despite the paucity of information provided, the main display was nearing its upper limit on the number of data indicators that could be placed on the screen at one time. However,

analysts foresaw the need to monitor many more sensors at one time (into the thousands), and realized that all forms of scaling the current display design would be insufficient. Additionally, the display provides only vague indicators of existing problems. It groups alarms, making it impossible to determine, without seeking more information on a different screen, how many problems exist within that system, what is the nature of the problem(s), or where within the system the problem(s) exist.

The new display sought to alleviate the burden placed on the user by the rigors of the task. In order to do so, the developed display introduced an entirely new concept for monitoring the network (Figure 2). Rather than a generic indicator for each system that rolled up information from the previous screen, the display took the entire set of data available on the individual system pop-ups and presented it on a single display.

As opposed to the main display of the previous system, which could present only a minimal amount of information on 150 systems, the new display makes available the continual monitoring of approximately 4000 systems on one screen. This enables operators to monitor over 20 times the number of systems than was previously possible. Additionally, whereas the previous system limited the amount of information available on any one screen (the main screen held around 150 data elements, while each pop-up contained about 50 data elements), the new visualization combined all of these elements on the same display. It is estimated that the number of data elements available to the user in the new visualization presents upwards of 20,000 data elements at one time.

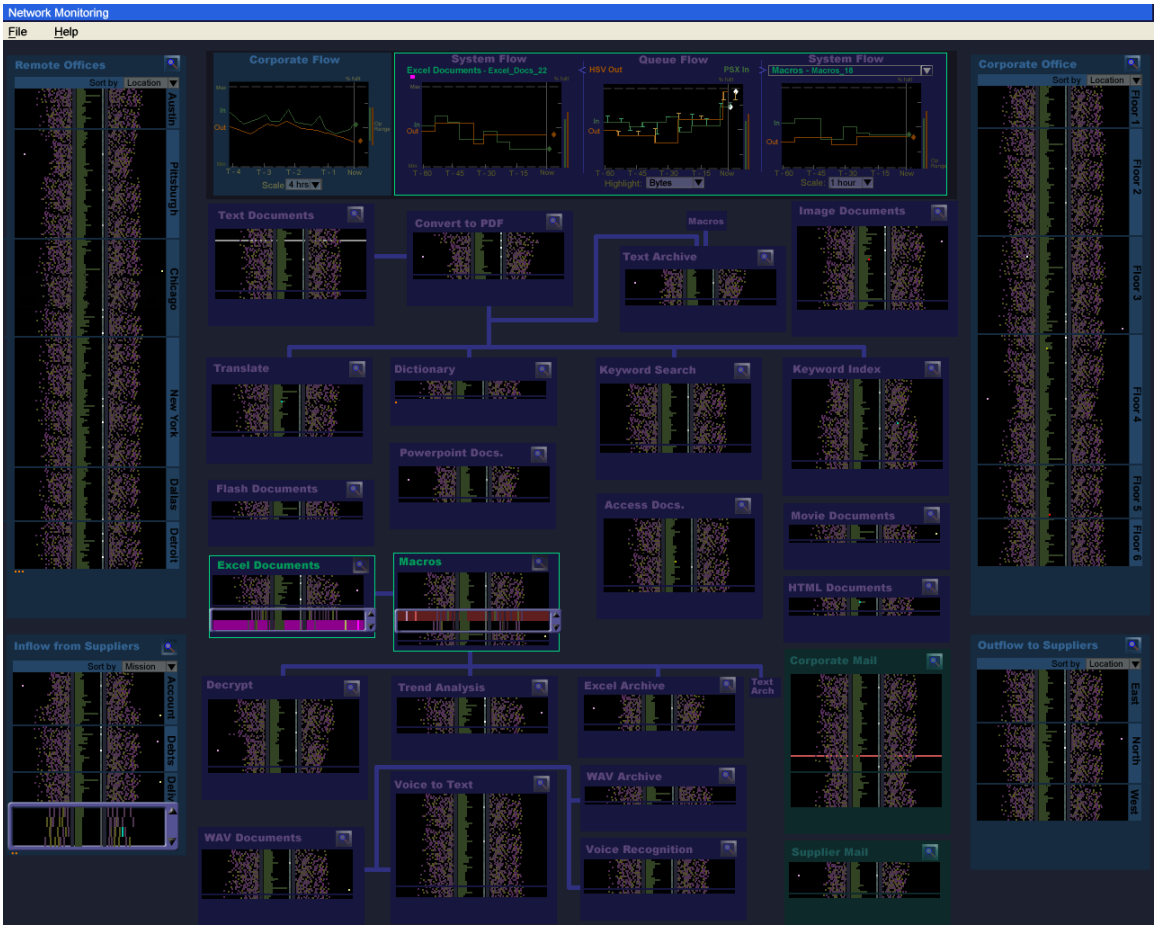

Figure 2. The redesign of the network monitoring display includes all information for a system that had previously been hidden within a separate pop-up menu. Operators no longer have to seek out a problem; their attention is drawn to systems where problems are occurring.

The visualization also provides a look up at the missions that are being affected by the data flow. A pop-up (Figure 3) is provided to operators that show, not only individual intelligence questions that are being asked, but how dependent those questions are on the selected system. Therefore, when a system encounters problems, the command center will know exactly what will be impacted, and how great that impact will be. Data related to time-critical, high priority Intelligence Needs can be rerouted through alternate systems – systems that only support the lower priority Intelligence Needs. In this way, the Command Center can ensure that it gets the information it needs, when it needs it, in order to make combat effective decisions.

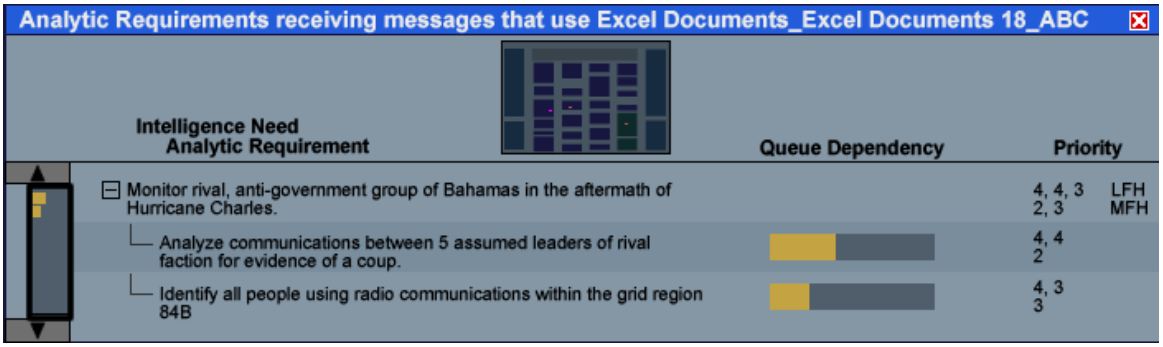

Figure 3. This pop-up provides operators with the missions that rely on a selected system for information flow. Additionally, it provides how dependent the mission is on that system. (Note: Missions are notional and are only intended as an example)

CSE principles important for the development of the display

The design of this display follows several design principles from Cognitive Systems Engineering in order to support the decision-making capabilities within the Command Center. Cognitive Systems Engineering (CSE) is the study of "how humans can cope with and master the complexity of processes and technological environments" (Hollnagel & Woods, 2005) and the application of those means. These principles will be described here in the context of how they impact the network monitoring display. Additionally, the impact that each principle can have on system design in general is presented.

Directed Attention

The role of the network monitoring display is to allow operators to identify where problems in the system exist or are developing. In the most basic sense, it is an alarm system; it alerts the command center to problems that are developing in the network that will impact decision-making. Alarm systems are intended to serve as Directed Attention mechanisms. They are added to the command center system to draw the attention of an

operator to a problem area, allowing the operator to decide whether or not the problem should be dealt with at that time.

In order to draw the attention of the operators to problems, the display identifies message back-ups by increasing the brightness of a message queue (represented as a single dot within a row) as it becomes fuller. Additionally, each queue is on a normalized scale, so that the queues that are becoming increasingly full move towards the outer edge of the scale, causing separation from the rest of the message queue indicators. Processes and hard drives that are problematic also become more salient on the screen (although, they have a fixed location within a display element). With all the data on the single display, every instance of a problem is identifiable in the same view. And, since the anomalous data is highly salient on the screen, the operator is easily able to identify where a problem has developed.

This is in direct opposition to the previous display, where problems were buried within a number of pop-up screens. The only indication a problem existed on the main display was through a group alarm for that system. However, group alarms are too vague to be effective; they do not help the operator decide where to focus their attention. With the group alarm, there is no indication of which problem might be the most important. The user is forced to go through each system's pop-up one-by-one. Deciding where to spend the most effort becomes very difficult.

By placing all the information on a single screen, the new network monitoring display avoids placing the user in a keyhole. The operators, seeing all data, no longer have to guess where to look at a given time. Additionally, the display no longer requires the user to be looking at the right portion of the data set at the right time to discover where a

problem is occurring – the display draws the operator's attention to the correct place. Operators in the command center are able to identify problems in the network faster when all the data are presented at once, ensuring decision-making can go on as needed.

Alarm systems should assist the operators in determining where to place their attention at time when their attentional capacity is most severely strained. At any given time, operators may face time pressure, competing goals, and a high risk of failure (Woods & Hollnagel, 2006), and the alarm systems must support operators as they try to work through the problem at hand. The alarm system must be constantly available to be viewed by the operator, but it must not force the operator to look at a certain place.

Providing Directed Attention does not mean that the system tells the operator where to pay attention. Directing Attention involves drawing the user's attention to where a problem exists (through attention grabbing encoding) and providing some insight into what the problem may be, allowing the operator to decide the best place to focus at a time. As many command and control centers employ alarm systems, the principles of Directed Attention can play an important role in reshaping the way that these are used.

Observability

The goal of a system design is to present information on the screen so that operators will have an in depth understanding of the fundamental principles of the domain. One way this is achieved is through providing Observability of the domain. Observability is the ability to form insights into a process based on feedback received. Rather than being simply an interface, the display is a window into the domain. Operators can see sequences and evolution over time, future activities and contingencies, and the patterns

and relationships in a process. A display that provides Observability allows operators to have a more accurate understanding of what is occurring in the domain, allowing them to respond more quickly and correctly to events, even those that were unanticipated by the system design.

Observability of the domain is provided by the network visualization in a variety of ways. A key feature of computer networks (and networks in general) is that systems are related to each other in a known way. The physical connections between systems means that data flow between networks can be known and predicted. When a problem occurs in one system, it will knowingly impact a host of other systems. In the previous visualization, no Observability into these relationships was provided. The operators were required to store this information in their heads. They had to remember the names of all systems in the network, as well as the systems to which each of these systems connected. With thousands of systems in the network, operators are extremely limited in the number of connections that they can recall. One reason these relationships were not preserved is that maintaining the relationships between systems has been a notoriously difficult problem for designers. Large networks generally encounter three main problems (Eick, 1996):

- Close proximity of physical locations can lead to operators relating two nodes that have no physical relationship,
- Exact positioning of the nodes can be difficult, with erroneous positioning leading to false conclusions, and
- The glut of information on the display causes clutter.

In the new visualization, each of these problems has been mitigated by a unique approach to representing networks. Rather than trying to represent the physical layout of the

network, on a map, for example, the network systems are grouped into functions. While each system has a different physical location, there are generally only a few functions that need to be performed - in this case, there are roughly 25 functions. This means that systems do not have to be positioned on a map; there is no opportunity to incorrectly place the system, nor is it possible to draw false conclusions about the relationships between systems. All systems grouped together share a function. By grouping systems in this manner, the first two difficulties are overcome.

The final problem is solved by only showing related systems dynamically. The physical connections between systems are only evident when a system is selected, at which point the operator is able to see all systems that feed data to that system as well as all systems that receive data from that system. This removes much of the clutter on the screen, which, in general, is very static. The freed up display space can then be used to show the dynamic data for which monitoring is necessary. The operators are no longer required to remember a large set of relationships between systems; the relationships are available at the click of the mouse.

Another form of Observability provided within the display is in the comparison of connected queues over time. Problems within the system rarely develop instantaneously - queues slowly become full over time. Additionally, a full queue is not necessarily indicative a problem. As with any computer process, it may take a long time to process extremely large files. While this may cause a back-up in one system, there may be no cause for concern, since as soon as the large message is processed, the remaining queues can be forwarded to be processed as well.

This information cannot be gleaned from a single static queue indicator. Determining if a problem exists or not, and, then, what might be done to solve the problem requires that the data over time be trended for the user. In the previous system, the data were available, but only as data values. The operators had to do the math in their head to decide if there was a problem.

The new visualization eases this cognitive work by providing trend graphs of these data. By explicitly representing the trends of connected queues over time (Figure 4), the operators are able to determine whether or not a problem truly exists, or if a high queue is merely a happenstance of the current situation. The operators can tell if a problem is getting worse, has leveled off, or is already improving. Providing this Observability allows operators to more effectively combat problems in the domain, as they are able to differentiate those backlogs that are truly a problem from those that just appear to be problems.

Figure 4. The trend graphs provide operators with the ability to easily see the history of queues so that it can be easier to make network diagnoses.

Having the needed data visible on the screen does not mean that the Observability has been provided. Observability entails presenting the data on the screen in a manner where it becomes clear what the situation is. Observability helps to reduce the cognitive workload of an operator, whether through limiting the conversions that the operator needs to do in order to make sense of the data presented, or by presenting information on the

screen that previously had been kept “in the operator’s head”. Information known to be useful for decision-making (such as the relationships between systems in a network) should be represented to the user, allowing more cognitive work to be spent deciding the best manner to use the information provided, rather then figuring out what the data means.

Data in Context

A key to making decisions is the placement of needed data in context. In order to make decisions, we must be able to extract meaning from what we perceive in the world (von Uexkull, 1934). Woods (1994) expands on this and argues that most technology viewpoints emphasize the ability to provide more and more data, but focus little on how to encode the data so that they provide meaning to the operator. In many C2 visualizations, operators complain of data overload; the data that is provided produces no meaning for the operator. The previous network monitoring display was heading towards the same trap. The problem is generally not with the amount of data presented to the operator; rather, the problem is that the data are not presented in any meaningful context to make decision making simple.

What operators need to make decisions is not data, but information. Presenting the data so that they are meaningful to the operator turns the data into information (Woods, 1994). One prerequisite is to know what relationships are informative in what contexts in the field of practice. In the case of network monitoring, the number of messages in the queue is not important for the operator, although these are the data reported. For the sake of ensuring messages can flow between queues, the more important relationship is the

number of messages in the context of the number of messages that a queue can hold. Because of this, all data values for a system can be put on a common frame of reference. This realization dramatically increased the number of data elements that could be put on the screen at one time. No longer were individual message counts being presented to the user. Instead, the user is shown a percentage of queue fill. Any queue that is getting too full moves away from those that are not (as well as becomes brighter). By finding the right context of the data, and representing this context, it becomes easy to see those queues that may be problematic.

Data in context, essentially, goes hand-in-hand with Observability. Finding the correct context for the data is a necessary step in creating good Observability. Knowing what the data should mean to the operator helps identify how to clearly display that information, so that the operator has a good understanding of the current situation.

Lessons learned

While the display is remarkably effective at showing the relationships in the domain, there were several issues that this revolutionary design brought forward. First, with so many data elements on a single display, the selection of an individual element becomes difficult. Each display element is 2 pixels wide by 2 pixels high; this is nearly an impossible target to hit with a mouse click. Therefore, to aid the selection, magnifiers were added to the display. However, this only increased the height of each display element (to 10 pixels high). While selection is possible, it still is not easy. The display is not as usable as others, requiring additional non-intuitive steps to execute a selection; however, with the added functionality and the new understanding created by the display,

this was seen as a small price to pay to greatly increase the Operators ability to monitor the network. While usability is important for making a good display, having a usable display is not sufficient for good operations.

Part of placing data in context entails finding the correct frames of reference for the data. However, with large amounts of data on the display, clutter becomes an issue; when a data set can be viewed in multiple frames of reference, it can be difficult to show all of the frames of reference at one time. Allowing the user to choose the frame of reference at a given time can help provide a work around for this issue.

In the network monitoring display, systems can be related in a number of ways. Mainly, systems are related by function. However, certain systems have no main function – they instead hold several relationships that may be of importance. Systems may exist at the same location, send to the same function, be of a similar type, or work towards the same mission. Which relationship is of importance to the operator at any one time depends on the situation. If, for example, equipment fails at a location, the systems at that location will all encounter problems simultaneously. However, problems may be related along any of the identified frames of reference; providing only one frame of reference could cause an operator to miss related problems if the problem has not developed in the chosen frame of reference.

The way around this is to allow the operators to dynamically change the frame of reference as they see fit. While the correct frame of reference may not be available as the problem occurs, the operator can easily shift the frame of reference to see which frame of reference, if any, is of importance. Other solutions to a similar problem try to force the operator to use one frame of reference, or may try to choose the frame of reference at

which the operator should be looking based on some metric (Bibighaus & Spink, 2006). While the metric may predict the correct frame of reference a majority of the time, forcing the user to work with a frame of reference in those times when it is incorrect severely hampers the operator, and can lead to false conclusions being drawn. If the metric is of value, it should be provided to the operators so that they may choose to use it as a piece of information in deciding the display to look at, but it should not force the decision.

Finally, the group for which the display is designed has multiple color anomalous operators (Red-Green protanopes). Usually, in order to overcome the possibility of color anomalous users, display elements are double coded; the original display had redundant coding of shape and color for each system indicator. While double coding suffices with a limited amount of data, it takes away available encoding possibilities as the information on the display increases. Providing a single coding mechanism in these situations is a far better solution.

Creating this single data coding mechanism to support both standard color and color anomalous operators again required a unique solution. Salience management is a key factor in designing for users. The more important and dynamic the information is, the more easily visible the information should be on the screen. Salient objects to a standard color observer are often not salient to anomalous color observers, and vice versa, however. Therefore, a unique color set was created for the color anomalous observers that preserved the salience map created for the standard color observer. Figure 4 shows a side-by-side comparison of the two color schemes used as well as a look at how they appear to a color anomalous protanope. Although the two displays are quite different in

color, the salience of elements on the screen is maintained. Equivalent events will be equally noticeable to all types of viewers.

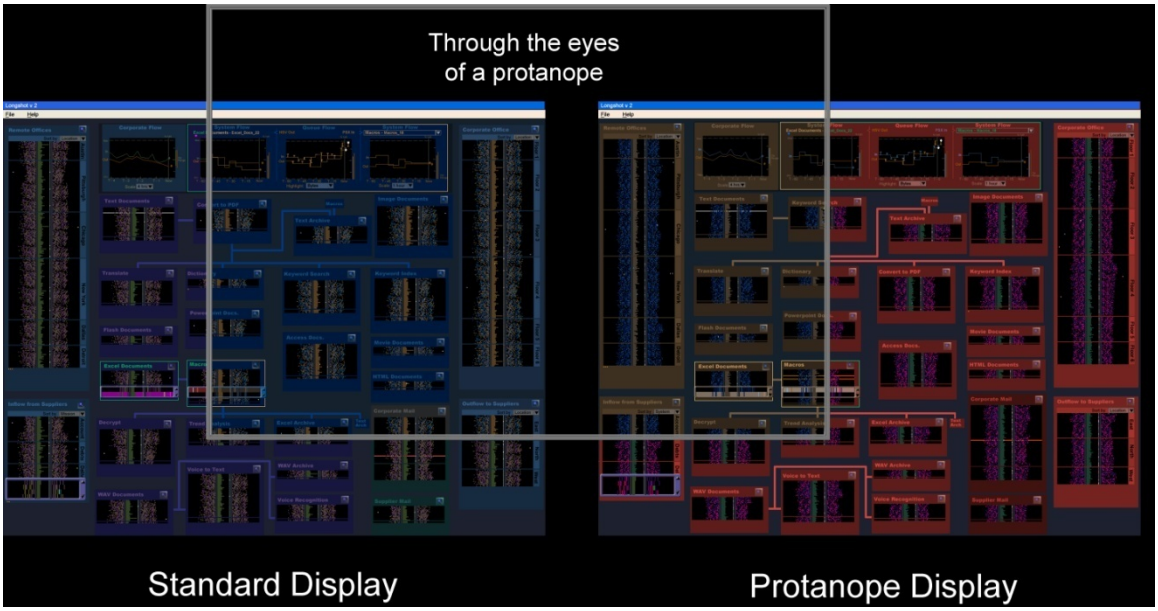

Figure 5. A side-by-side comparison of the standard display and the display built specifically for a color anomalous protanope shows the changes that must be made to the coding in order to maintain the relative salience of display elements.

Options such as these can be stored as properties attached to a log-in of an operator. The display does not work any differently, just the hues used. Note that this is different from allowing the user to choose how the information is represented on the screen. The displays were specifically tailored to control the salience of data elements. Standard color observers who choose to use the display encoded for the color anomalous observer would not be able to detect system problems as easily as they would on the display designed for them.

As displays are placed more and more on powerful computers and servers, the ability to create dynamic displays to suite all sorts of users is increased. While there had been standard solutions set when displays were not as flexible, these standard solutions should not constrain innovative ways of solving problems now that this dynamism is possible.

Whether allowing the operator to change the frame of reference or providing a new color scheme for a color anomalous protanope, displays can now be designed to overcome any number of issues that, in the past, required sub-optimal solutions (such as double coding).

Conclusion

Decision-making in the command center is only as good as the information that it receives through the intelligence collection process. This process has been greatly improved by advancements in technology, which allow a command center to receive data from across the globe in a short amount of time. And, while technology has enhanced the capabilities of the intelligence collection process, it also creates difficulties when it falters. Without information from the field reaching the command center, decision-making becomes very difficult. Due to its impact on the decision-making process within the command center, the command center itself must be responsible for monitoring the network's status, performance, and ability to transmit collected information.

A visualization is offered as a solution to the network monitoring problem for one such intelligence gathering process. The visualization has increased both the number of systems that can be monitored at one time, as well as the number of data elements that are on display on the screen. The design was built in line with Cognitive Systems Engineering principles in order to fully support the operators as they monitored the network. By supporting the redirection of operator attention, providing Observability of important domain information, and finding the correct context for the data, operators are better at identifying troubled spots within the network, leading to quicker solutions. Operators are able to keep the intelligence process running as effectively as possible,

ensuring the network is working as intended, so that commanders can make sure they get the information they need, when they need it.

References

Becker, R.A., Eick, S. G., & Wilks, A.R. (1995). Visualizing Network Data. *IEEE Transactions on visualization and computer graphics, 1(1)*. 16-28.

Bibighaus, D.L. & Spink B. (2006). Situational Awareness for the Information Operations Innovation Network. *The Next Wave (Available from the National Security Agency, Ft. George G. Meade, Suite 6541, MD 20755-6541),* 36-39.

Daly, J.J. & Tolk, A. (2003). Modeling and Simulation Integration with Network-Centric Command and Control Architectures. *Fall Simulation Interoperability Workshop*. Orlando, FL. 1-10.

Eick, S.T. (1996). Aspects of Network Visualization. *IEEE Computer Graphics and Applications, 16(2),* 69-72.

Hollnagel, E. & Woods, D. D. (2005). *Joint Cognitive Systems: Foundations of Cognitive Systems Engineering*. Boca Raton, FL: CRC Press / Taylor & Francis.

Tzu, S. (1981). *The Art of War*. London, UK: Hodder & Stoughton.

von Uexkull, J. (1934). A stroll through the worlds of animals and men. In K. Lashley (Ed.), *Instinctive behavior*. International Universities Press, New York.

Woods, D.D. (1994). *Visualizing Function: The theory and practice ƒ representation design in the computer medium*. Unpublished Manuscript, The Ohio State University, Columbus, OH.

Woods, D.D. & Hollnagel, E. (2006). *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. Boca Raton, FL: CRC Press / Taylor & Francis.