

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YY) 21-10-16		2. REPORT TYPE Conference Proceedings		3. DATES COVERED (From - To) 10/2015 – 03/2016	
4. TITLE AND SUBTITLE Coordinated Displays to Assist Cyber Defenders			5a. CONTRACT NUMBER FA8650-14-D-6501-0009		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Alex Vieane ¹ , Gregory Funke ² , Vincent Mancuso ³ , Eric Greenlee ⁴ , Gregory Dye ² , Brett Borghetti ⁵ , Brent Miller ² , Lauren Menke ⁶ , Rebecca Brown ⁶			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER H0HJ (53290813)		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) ¹ Colorado State University, Fort Collins, CO; ² Air Force Research Laboratory, Wright-Patterson AFB, OH; ³ MIT Lincoln Laboratory, Lexington, MA; ⁴ Texas Tech University, Lubbock, TX; ⁵ Air Force Institute of Technology, Wright-Patterson AFB, OH; ⁶ Ball Aerospace & Technologies Corporation, Fairborn, OH			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Materiel Command Air Force Research Laboratory 711 th Human Performance Wing Airman Systems Directorate Warfighter Interface Division Applied Neuroscience Branch Wright-Patterson Air Force Base, OH 45433			10. SPONSORING/MONITORING AGENCY ACRONYM(S) 711 HPW/RHCP/RHCPA		
			11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.					
13. SUPPLEMENTARY NOTES 88ABW Cleared 03/18/2016; 88ABW-2016-1322. Human Factors and Ergonomics Society 2016 Annual Meeting (HFES) 19 – 23 September 2016.					
14. Cyber network analysts must gather evidence from multiple sources and ultimately decide whether or not suspicious activity represents a threat to network security. Information relevant to this task is usually presented in an uncoordinated fashion, meaning analysts must manually correlate data across multiple databases. The current experiment examined whether analyst performance efficiency would be improved by coordinated displays, i.e., displays that automatically link relevant information across databases. We found that coordinated displays nearly doubled performance efficiency, in contrast to the standard uncoordinated displays, and coordinated displays resulted in a modest increase in threat detections. These results demonstrate that the benefits of coordinated displays are significant enough to recommend their inclusion in future cyber defense software.					
15. SUBJECT TERMS Cyber network analysts, cyber defense, coordinated displays					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT: SAR	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON Gregory Funke
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER

COORDINATED DISPLAYS TO ASSIST CYBER DEFENDERS

Alex Vieane¹, Gregory Funke², Vincent Mancuso³, Eric Greenlee⁴, Gregory Dye², Brett Borghetti⁵, Brent Miller², Lauren Menke⁶, Rebecca Brown⁶

¹Colorado State University, Fort Collins, CO; ²Air Force Research Laboratory, Wright-Patterson AFB, OH; ³MIT Lincoln Laboratory, Lexington, MA; ⁴Texas Tech University, Lubbock, TX; ⁵Air Force Institute of Technology, Wright-Patterson AFB, OH; ⁶Ball Aerospace & Technologies Corporation, Fairborn, OH

Cyber network analysts must gather evidence from multiple sources and ultimately decide whether or not suspicious activity represents a threat to network security. Information relevant to this task is usually presented in an uncoordinated fashion, meaning analysts must manually correlate data across multiple databases. The current experiment examined whether analyst performance efficiency would be improved by coordinated displays, i.e., displays that automatically link relevant information across databases. We found that coordinated displays nearly doubled performance efficiency, in contrast to the standard uncoordinated displays, and coordinated displays resulted in a modest increase in threat detections. These results demonstrate that the benefits of coordinated displays are significant enough to recommend their inclusion in future cyber defense software.

Effective cyber defense is crucial to the success and security of modern commercial, industrial, and governmental organizations. Dependence on cyber systems continues to increase as network traffic rises and computerized networks are extended to interconnect a wider variety of functional assets (Maybury, 2015). This increased reliance on cyber networks has amplified the potential degree of harm that may be inflicted by adversarial cyber attacks (e.g., malware, worms, viruses), insider exploitation (e.g., phishing), and other threats to cyber security. These threats have the potential to impact a wide variety of networked resources including life sustaining utilities, such as power and water, and mission-critical military assets (Maybury, 2015). As noted within the Air Force's "Cyber Vision 2025," the scope, capability, and utilization of cyber networks are expected to continue increasing (Maybury, 2015). Unfortunately, malware development is expected to surge as well, resulting in an estimated 1000% increase in unique, malicious software by the year 2025. In order to counter the dangers of ever-evolving cyber threats, cyber defense systems must be optimized to minimize the risk to human and technological assets.

In modern cyber defense, intrusion detection is the first line of protection against immediate cyber threats. Intrusion detection "is the process of monitoring the events occurring in a computer system or network and analyzing them for signs of possible incidents, which are violations or imminent threats of violation of computer security policies, acceptable use policies, or standard security practices" (Scarfone & Mell, 2007, p. 2-1). This intrusion detection process is typically initiated by an algorithmic, automated Intrusion Detection System (IDS), which inspects all network events and compares them to a database of known "signatures," i.e., profiles of malicious activity. When the IDS detects potential suspicious activity, such as the occurrence of a network event that is *similar* to a known attack signature, the system generates an alert which is then presented to a human *computer network defense analyst*, or more succinctly, a *network analyst*, who must evaluate the veracity of that alert.

To accomplish this evaluation, a network analyst will broadly follow a standard pattern of activity (Dye, in press).

First, an analyst will consult an IDS display and select an alert to investigate further. Though alerts generally include some information about the nature of a potential threat, the information may be incomplete or the analyst may be unfamiliar with the specific details of the threat. In such a case, the analyst may then consult a signatures database, which includes more detailed information about the nature of a threat and evidence indicating its presence on a computer or network. Next, the analyst must interrogate potential sources of evidence, such as network packets, firewall logs, network diagrams, etc., to determine if sufficient evidence exists to support the IDS's assertion of malicious activity. If adequate evidence is present, the alert information is passed to escalation analysts for further investigation (D'Amico & Whitley, 2007); if not, the alert is typically marked as "closed" (in some fashion) and further investigation is terminated.

Recently, a number of studies have been conducted to investigate potential human factors challenges inherent in a network analyst's job, revealing that these operators face numerous task-related and work-related challenges (e.g., Champion, Rajivan, Cooke, & Jariwala, 2012; Mancuso et al., 2015). First and foremost among these challenges is the sheer volume of alerts which line operators must investigate. Generally speaking, IDS systems are liberally biased in the identification of suspicious network events in order to minimize the possibility of missing an attack. As a consequence, however, most IDS alerts are false alarms (D'Amico & Whitley, 2007). Common types of false alarms include instances where the IDS misidentifies normal, benign activity as malicious; detects only partial evidence matching a known signature; or correctly recognizes that malicious activity is present, but the system is not vulnerable to that activity (e.g., an attack exploits a vulnerability that has been patched, or an attack targets a closed port). The net result is a constant flood of alerts that must be investigated by network analysts (Champion et al., 2012), but the signal-to-noise ratio of these alerts is extremely low, perhaps as low as .01% (though it should be noted that it is exceedingly difficult to impossible to assess this ratio in practice; Dye, in press).

Another challenge for line operators is the task of collecting relevant information from multiple network sensors in order to make an accurate decision about an IDS alert. As mentioned previously, an analyst must inspect data from multiple sources to confirm the assertion that malicious activity was present (Dye, in press). Each of these data sources is often presented by a separate computer program, in separate display windows, tabs, or even on separate computer monitors. This requires the analyst to sequentially investigate each data source to locate the information that is relevant to a given alert. This investigation is complicated by the fact that different data sources are usually uncoordinated. Typically, timestamps serve as the common link between data sources, meaning that network analysts will note the time at which an alert was generated and search through multiple uncoordinated sources for sensor data that was recorded in a similar timeframe.

This is a tedious and potentially error prone process. In speaking with one of the current authors, interviewed cyber analysts have confirmed that they find the task frustrating and time consuming. Anecdotally, these subject matter experts indicated that the cyber community has a term for this process of manually searching, coordinating, and investigating multiple data sources: they call it “pivoting,” because the operator is required to mentally “pivot” between data sources, and in some cases, they must also physically “pivot” between different workstations equipped with specialized software. Given the huge volume of alerts analysts must interrogate, performance efficiency is key to maintaining successful cyber defenses.

However, the need for manual search and synchronization could be eliminated by enabling an analyst’s workstation with software to automatically coordinate related data across sources, a concept akin to “coordinated views” (e.g., Andrienko & Andrienko, 2007), though software providing coordinated views also typically includes some form of data visualization. A simple way to coordinate multiple databases would be to link them using timestamps, so that when an analyst selects an alert to investigate, the network data sources are automatically searched to display the temporally relevant information to the analyst for inspection. Given that analysts often use timestamps in their manual search, it seems likely that an automatic, timestamp-based process of coordination would expedite alert assessment in cyber intrusion detection tasks.

To test this possibility, we conducted the current study using a validated synthetic cyber task environment to simulate a network analyst’s task (i.e., evaluating IDS alerts for evidence of malicious activity). We predicted that the use of *coordinated displays* (i.e., timestamp-linked displays) would result in superior performance efficiency, indexed by time to complete the task, compared to *uncoordinated displays* (i.e., displays requiring manual search and synchronization). While coordinated displays were predicted to elicit superior performance efficiency, they were not expected to affect performance efficacy (i.e., accuracy in identifying threats and non-threats) because all alert-related information was present regardless of display condition (coordinated or uncoordinated).

METHOD

Participants & Experimental Design

In this experiment, 46 people (19 men, 27 women) were recruited from local universities, available Air Force personnel, and the local community. They ranged in age from 18 to 35 ($M = 22.13$, $SD = 4.66$). Participants received a single payment of \$30 as compensation for their time. All participants were cyber novices and had no previous experience in cyber defense.

Experimental Design

This study featured a single experimental factor, display condition, and a control factor, alert list. Both were between-participants factors. Display condition had two levels, coordinated and uncoordinated displays. The alert list factor, which had two levels, determined which of two sets of IDS alerts participants engaged during the experimental task (please see below for a full description of the lists). Dependent variables assessed in this experiment included time to complete the experiment and alert judgment accuracy.

Air Force Cyber Intruder Alert Testbed (CIAT)

Participants in this experiment took on the role of network analysts charged with defending a hypothetical Air Force network from malicious cyber attacks. The synthetic task environment (STE) employed in this experiment was the Air Force’s Cyber Intruder Alert Testbed (CIAT; Funke et al., in press), presented in Figure 1.




Figure 1. Example of the CIAT interface. Represented in the figure are 1) the intrusion detection system (IDS), 2) the query and signatures database, 3) the packet capture software, 4) the network list, and 5) the participant response buttons (i.e., “Not a Threat” and “Threat”). Though these disparate components appear together in the figure to conserve space, during the experiment, each of the enumerated elements existed on separate “tabs” in the display, with the exception of the response buttons, which appeared below the IDS. As a result, participants had to dynamically alternate between tabs to complete the task. This was done to emulate real-world constraints of network analysts who must frequently switch between different pieces of software (and even workstations) necessary for them to accomplish their work. It should be noted that once an alert in the IDS was selected by a participant, it was highlighted in blue.

This STE was designed to broadly emulate key functions of Enterprise-level cyber defense platforms, such as Hewlett-Packard's ArcSight, AlienVault's Unified Security Management (USM), and IBM's Security Network Protection (XGS). Specifically CIAT includes a simulated intrusion detection system, signatures database, packet capture software, and network list. We will now briefly consider each of these task components with regard to their functionality in CIAT.

Intrusion detection system (IDS). Alerts for investigation were displayed for participants on the IDS tab. To complete the experiment, participants had to evaluate each of the alerts presented and mark them as "not a threat" or a "threat" to the hypothetical Air Force network they were defending. As described above, participants made this judgment by evaluating evidence gathered from the other elements of the task, and they made their response by clicking on the appropriate button (i.e., "not a threat" or "threat"). Though somewhat artificial in nature, these responses broadly correspond to similar judgments made by network analysts to either mark an alert as closed (non-threats) or forward the alert details to an escalation analyst (threats). All signatures utilized in this experiment were drawn from Baumrucker and colleagues (2003) and are representative of the kinds of alerts that network analysts encounter while performing their duties (the alerts are in fact "real" alerts, though the signature content is a bit out of date).

Each alert in the IDS tab possessed four characteristics: 1) a numeric and visual representation of the severity of the alert, 2) the time the alert was generated, 3) the name and number of the signature associated with the alert, and 4) an "action field." Alert severity is a categorical representation of the potential harm an attack matching the signature could cause to the computer or network affected. Typically, IDSs represent this information with a numeric value (usually on a scale of 1-5 or 1-10, where higher values indicate greater harm) and with redundant color coding. In this experiment, all alerts were coded as 4 (orange) or 5 (red).

The action field of each alert initially was blank. When participants marked an alert as "not a threat" or as a "threat," this judgment was indicated in the action field so that participants could maintain awareness of alerts they had already addressed.

At the start of the experiment, 45 alerts were present in the IDS tab. Of these, 40 of the alerts were designed to provide insight into network analyst decision making and the effects of coordinated displays; the remaining alerts were designated "catch alerts," designed to identify participants who were not fully engaged in the task.

Twenty unique signatures were utilized to generate the 40 non-catch alerts (i.e., each signature appeared in the IDS list twice). Each of these signatures required participants to verify the presence of 4 to 6 pieces of evidence from the packet capture software or network list in order to render a threat judgment. To be a threat, all elements of a signature had to be present in the packet capture data and/or the network log; otherwise, the alert was to be assigned a status of "not a threat." Non-threat alerts included in this experiment lacked a

single piece of evidence required to meet their associated signature requirements.

In all conditions, 10 alerts met all signature requirements and therefore represented "threats." However, the specific threat signatures presented to participants were determined by the alert list condition to which they were assigned. In the first list condition, 10 non-catch signatures (from the 20 signature list) were presented as threats; in the second list condition, the other 10 non-catch signatures were threats. As mentioned above, all signatures were presented twice in the IDS list. Therefore, participants in both list conditions were presented with 1 threat and 1 non-threat instance of 10 signatures (differentiated based on their assigned list condition), and 2 non-threat instances of each of the remaining 10 signatures.

The IDS list also included five "catch" alerts, designed to identify participants who may not have been fully engaged in the task, who responded in a haphazard fashion, or who may not have fully understood the task even after satisfactorily completing training (e.g., Oppenheimer, Meyvis, & Davidenko, 2009). Each catch alert utilized a unique signature that was not repeated in the IDS alert list. The same five alert signatures were employed in both list conditions. These alerts were specifically designed to be exceptionally easy to verify, requiring only a single piece of evidence from the packet capture tab to evaluate. In all cases, sufficient evidence was present in the packet capture data to judge the alerts as "threats." This approach was adopted to anticipate and counter a potential participant bias to mark alerts as non-threats without careful consideration (e.g., to avoid working hard, finish quickly, etc.) since the majority of alerts featured in this experiment fell into that category. Though they were not informed of the requirement, participants had to correctly identify 4 of the 5 catch alerts during task performance as a "threat" or their data would be excluded from the experiment.

Signatures database. The signatures database tab had two primary fields: a query window and a response window. Participants could input a signature number in the query window and the signature description would populate the response window. In the coordinated displays condition, clicking on an alert in the IDS tab would cause the signature description to automatically populate the response window. In the uncoordinated displays condition, participants had to manually input the signature number to receive its description.

Packet capture software. The packet capture tab was the primary source of evidence participants could consult to determine if an alert represented a threat. Each packet had seven data fields, corresponding to: 1) the number of the packet, determined by its serial position in the packet list; 2) the time the packet was generated; 3) the source internet protocol (IP) address of the packet; 4) the destination IP address of the packet; 5) the protocol type; 6) the packet length (i.e., its size); and 7) any additional information associated with the packet. Evidence associated with an IDS alert signature could be found in any of the packet data fields. For example, evidence of one signature might include repeated packets sent from the same source IP address, while a different signature may require packets to be of a particular protocol type.

In addition to packets associated with each IDS alert in the experiment, the packet capture list also included distracter packets unrelated to any specific alert. These packets were included to create a more representative list, such as might be encountered by network analysts. Distracter packets differed from IDS-relevant packets in that they included source and destination IP addresses or protocol types that were not part of any included IDS alert signature.

In the coordinated displays condition, when participants clicked on an alert in the IDS tab, any packets timestamped within a 30 second window of the alert timestamp were automatically highlighted and the first highlighted packet was centered in the packet tab window. Highlighted packets could include both relevant and irrelevant packets, so participants still had to interrogate individual packets for signature evidence in this condition. However, the process was significantly truncated by comparison to that in the uncoordinated condition, which required participants to manually navigate through the packet tab and initially identify potentially relevant packets by comparing packet timestamps and the alert timestamp.

Network list. The network list tab included information about the computers participants were responsible for on the Air Force network. Entries on the list included two pieces of information about each computer: 1) the IP address of the computer, and 2) any additional, potentially relevant information about that computer, such as membership in a subnet or potential vulnerabilities (due to missing patches, etc.).

In the coordinated displays condition, when participants clicked on an alert in the IDS tab, any potentially relevant computers, based on factors such as source or destination IP address indicated in the alert signature, in the network list were automatically highlighted. In the uncoordinated displays condition, participants had to manually search the computer list for relevant information.

Procedure

Upon arrival in the laboratory, participants were assigned at random to either the coordinated or uncoordinated displays condition, and to either the first or second IDS alert list. Participants were then told that they would be taking on the role of an Air Force network analyst defending a critical Air Force network against malicious cyber attacks. Next, participants completed a computer based training course on the CIAT task, which presented information about the different task tabs, how to evaluate IDS alerts to make threat judgments, and examples of how to interact with the task. Training was condition specific (i.e., coordinated vs. uncoordinated displays), so that participants understood how to complete the task in their assigned condition. The computer based training was self-paced and was typically completed in approximately 20 minutes.

Participants then completed a CIAT practice trial in the condition to which they had been assigned. This trial featured 3 IDS alerts. Alert signatures included in the practice trial were designed to be comparable to those of the experimental trial (i.e., each required 4 to 6 pieces of evidence in order to

render a threat judgment); however, practice trial signatures did not appear again in the experimental trial.

For the first alert, the researcher demonstrated how to collect evidence to make a threat judgement while talking aloud. In completing the second alert, participants were able to practice the task. Participants were required to talk aloud as they did so, and the researcher was available to answer any questions or intervene when necessary to ensure understanding. Participants were required to complete the final alert on their own while still talking aloud. After the practice trial was over, the researcher was available to address any questions participants had before starting the experimental trial.

Participants were then told that the experimental trial would feature a larger IDS alert list and that they had to address each alert, judging them as either a threat or non-threat, to finish the experiment. Further, they were told that they were free to investigate the alerts in any order they wished and at their own pace.

Participants typically completed the entire experiment within 2-3 hours.

RESULTS

Catch trials

As described previously, five “catch” trials were included in the experiment to permit detection and exclusion of participants who were perhaps not fully engaged in the task, who were responding haphazardly, or who may not have fully understood the task even after satisfactorily completing training. Participants were required to correctly categorize 4 of the 5 catch items as a “threat” for their data to be included in the final experimental set. Of the 46 participants in the current sample, two participants, one participant assigned to the uncoordinated condition and one participant assigned to the coordinated condition, failed to meet this inclusion criterion. As a result, these two participants’ data were excluded from all succeeding data analyses (i.e., subsequent analyses proceeded with the remaining 44 participants).

List effects

Participants in this experiment were assigned at random to one of two lists of alert stimuli. To determine if list assignment influenced task performance, the time to complete the experiment and the total number of correct signal judgments (calculated as number of correct rejections + number of threats correctly detected) were compared for the two list conditions using an independent samples *t*-test and a Mann-Whitney *U*-test, respectively.

The results of the list analyses indicated that stimulus list did not significantly influence time to complete the experiment, $t(42) = .45, p > .05$, nor did it influence the number of correct signal judgments, $U = 212.50, p > .05$. As these analyses indicated that alert list did not influence task performance in this experiment, the factor was dropped from consideration in all subsequent analyses.

Time to complete the trial

Total time to complete the alert list in the experiment was analyzed using an independent samples *t*-test. The analysis indicated that participants in the uncoordinated condition took significantly longer ($M = 86.21$ minutes, $SE = 4.62$ minutes) to complete the alerts than did participants in the coordinated condition ($M = 43.61$ minutes, $SE = 3.08$ minutes), $t(42) = 7.67$, $p < .001$, Cohen's $d = 2.31$. In other words, participants in the coordinated displays condition finished in approximately half the time of participants in the uncoordinated condition.

Responses to alerts

Correct rejections. As mentioned previously, 30 of the 45 alerts participants evaluated in this experiment lacked sufficient evidence to conclude that those alerts were generated by a legitimate cyber attack. When participants encountered such alerts, the correct response was for them to be marked as "not a threat" (i.e., to correctly reject those items as threats). To examine if there were differences in correct rejection rates between the uncoordinated and coordinated conditions, a Mann-Whitney *U*-test was computed. The results of the analysis of correct rejections indicated there was no statistically significant difference between conditions in correct rejections, $U = 231.00$, $p > .05$. Participants in both conditions correctly rejected approximately the same number of alerts (the median in both conditions was 27.00).

Correct threat detections. Also as described previously, 10 of the 45 alerts included in the experiment represented actual cyber attacks, i.e., sufficient evidence existed supporting the conclusion that a cyber attack had taken place. When participants detected these critical signals, the correct response was for them to be marked as a "threat." To examine if there were differences in detection rates between conditions, an additional Mann-Whitney *U*-test was calculated. The results of this test indicated that participants in the uncoordinated condition detected significantly fewer threats (median = 8.50) than participants in the coordinated condition (median = 9.50), $U = 137.00$, $p = .01$.

DISCUSSION

The goal for this experiment was to examine the effectiveness of coordinated displays to assist network analysts in performing their primary duties. We initially hypothesized that access to coordinated displays would reduce task completion times relative to an uncoordinated display condition. In addition, we hypothesized that access to coordinated displays would not influence correct rejections and threat detections, as the same evidence required to make those decisions was present in both task conditions. The results of the experiment indicated that, as hypothesized, coordinated displays reduced task completion times. However, coordinated displays were also found to modestly improve threat detections in this experiment.

To some extent it is not particularly surprising that access to coordinated displays reduced task completion times, since it drastically reduced the amount of time participants required to correlate timestamps across evidentiary sources in this experiment. However, the magnitude of reduction, coupled with the increase in correct threat detections, suggests that coordinated displays may also have facilitated task comprehension, allowing participants to make correct judgments about potential malicious activity more frequently and efficiently. Though the increase in correct threat detections in this experiment was relatively small, given the number of alerts that network analysts must routinely examine, even modest differences may be meaningful.

Overall, the improvements in performance observed in this experiment suggest that coordinated displays may provide a powerful tool for improving analyst performance. Developers of cyber defense software should seriously consider including such functionality into their future products to allow cyber defenders to realize these potential benefits.

Next steps for this research should include testing the benefits of coordinated displays with participants who have previous experience in cyber defense. Though it is likely that coordinated displays will also facilitate their performance, the magnitude of the effect is likely to be somewhat reduced due to greater familiarity with the task.

Acknowledgements. This project was supported by grant no. F4FGA05076J003 from the Air Force Office of Scientific Research (Benjamin Knott, Program Officer).

REFERENCES

- Adrienko, G., & Adrienko, N. (2007, July). *Coordinated multiple views: A critical view*. Paper presented at the Fifth International Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV 2007), Zurich, Switzerland.
- Baumrucker, C.T., Burton, J., Dentler, S., Dubrawsky, I., Osipov, V., & Sweeney, M. (2003). *CISCO Guide to Secure Intrusion Detection Systems*. Syngress Publishing.
- Champion, M. A., Rajivan, P., Cooke, N. J., & Jariwala, S. (2012, March). *Team-based cyber defense analysis*. Paper presented at the 2012 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), New Orleans, LA.
- D'Amico & Whitley (2007). The real work of computer network defense analysts: The analysis roles and processes that transform network data in to security situation awareness. In J.R. Goodall, G. Conti, & K.-L. Ma (Eds.), *VizSEC 2007: Proceedings of the workshop on visualization for computer security* (pp. 19-37). Heidelberg, Germany: Springer-Verlag.
- Dye, G. (In press). *Using imprint to guide experimental design of simulated task environments* (Technical Report No. AFIT-ENG-MS-15-J-052). Wright-Patterson Air Force Base, OH: The Air Force Institute of Technology.
- Funke, G., Dye, G., Borghetti, B., Mancuso, V., Greenlee, E., Miller, B., Menke, L., Brown, R., & Vieane, A. (in press). *Development and validation of the Air Force Cyber Intruder Alert Testbed (CIAT)*. Proceedings of the 7th International Conference on Applied Human Factors and Ergonomics.
- Mancuso, V. F., Greenlee, E.T., Funke, G., Dukes, A., Menke, L., Brown, R., & Miller, B. (2015). Augmenting cyber defender performance and workload through sonified displays. *Procedia Manufacturing*, 3, 5214-5221.
- Maybury, M. (2015). Toward the assured cyberspace advantage: Air Force cyber vision 2025. *IEEE Security & Privacy*, 13, 49-56.
- Oppenheimer, D.M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45, 867-872.
- Scarfone, K., & Mell, P. (2007). *Guide to intrusion detection and prevention systems (IDPS): Recommendations of the National Institute of Standards and Technology* (Special Publication 800-94). National Institute of Standards and Technology, Technology Administration, U.S. Department of Commerce.