

A Mission-Centric Visualization Tool for Cybersecurity Situation Awareness

Nichole Stockman
Four Eyes Lab
UC Santa Barbara
nstockman@cs.ucsb.edu

Tobias Höllerer
Four Eyes Lab
UC Santa Barbara
holl@cs.ucsb.edu

Kyriakos G. Vamvoudakis
CCDC*
UC Santa Barbara
kyriakos@ece.ucsb.edu

Richard Kemmerer
Computer Security Lab
UC Santa Barbara
kemm@cs.ucsb.edu

Laura Devendorf
School of Information
UC Berkeley
ldevendorf@ischool.berkeley.edu

João P. Hespanha
CCDC
UC Santa Barbara
hespanha@ece.ucsb.edu

ABSTRACT

We present a novel visualization tool that provides high-level situation awareness for cybersecurity scenarios. Our visualization tool is tailored to intrinsically higher-level information presentation centered around cyberdefense missions. We applied our visualizations to data logs from the 2011 UCSB International Capture The Flag competition, which was designed to mirror general cybersecurity scenarios that would be encountered by the military. In this paper, we present our system and user interface design and discuss some anomalies in the data that are highlighted through the use of our visualization tool.

Categories and Subject Descriptors

H.5.2 [Information Interfaces and Presentation]: User Interfaces

General Terms

Design, Security

Keywords

Cybersecurity, Situation Awareness, Cybermissions, Visualization

1. INTRODUCTION

Cybersecurity is a growing concern, particularly in the military domain where there are strict requirements for the time-critical completion of cybermissions. The goal of situation awareness (SA) [6] is a common aspect of many cybersecurity solutions, but the majority of these solutions promote SA for low-level data like network flows or error logs, requiring the analyst to perform higher-level analyses manually which is time-consuming and error-prone [2, p. 5]. Offloading some of these analyses from the human analyst to

the visualization tool is helpful and many tools have been developed that focus on event monitoring and attack correlation. However, mission-centric cybersecurity scenarios require something more; namely, a complete overview and understanding of the state of the missions and any potential threats to their completion.

Missions, which can be represented as finite state machines, consist of an end goal and the intermediate steps needed to complete that goal. Each step (ie state) in the mission requires certain services in order to execute and progress to the next state. Thus, a security analyst must be able to assimilate high-level data about the states of the mission and threats to any required services. This is the driving motivation behind our visualization tool and its design. We not only display attack and vulnerability analysis data, but we also show an overview of the entire mission and provide interactions which correlate relevant data. We demonstrate our visualization tool's capabilities using what is to our knowledge the first dataset that includes mission descriptions as well as attack logs and the statuses of any services required by the missions. This dataset is the result of a distributed, wide-area security exercise hosted each year by the Security Lab at the University of California, Santa Barbara. Recently, this competition, known as the International Capture The Flag competition (iCTF), has been designed and developed with Cyber-Situation Awareness in mind [16] [5]. The mission-centric structure of the 2010 and 2011 UCSB iCTFs, and the competitive (attack/defense) aspects of the game closely match practical cybersecurity mission scenarios, and enable us to evaluate our high-level visualizations for situation awareness using a controlled scenario (playback of the iCTF logs), demonstrating the benefits to analysts performing tasks such as identifying competent attackers rather than searching through network traces to determine if an attack even occurred.

We evaluate our system using a set of motivating questions divided into the three phases of situation awareness as described by Endsley: Perception, Comprehension, and Projection [6]. Respectively, these refer to being aware of current data, understanding enough to draw conclusions about one's status in regard to this data, and speculating on one's future status. A full list of questions is provided in section 5 with an analysis of how well our visualization tool aided

in answering them.

In this paper, we introduce a visualization tool for situation awareness in mission-centric cybersecurity scenarios. We begin by discussing some previous work in section 2 and then we give an overview of the visualization tool and its motivation in section 3. In section 3.1 we describe our dataset and the ways in which it complements as well as deviates from the general scenario/motivation for which our tool was designed. Next, in section 4, we discuss the visualization tool's design as well as our design process in more detail. Afterward, we present our results in section 5 and include a discussion of these results in section 6. Finally, we present our conclusions in section 7.

2. RELATED WORK

Many visualization tools exist in order to assist security analysts and network administrators in identifying security risks to online systems. However, the majority of these systems visualize low-level data. Some analysis tools like tcpdump and libpcap simply display textual readouts of network traffic capture data [15]. Wireshark additionally provides a graphical display and allows users to interactively browse the traffic running on a computer network [18]. However, these analysis tools still require time and effort from the user to correlate data and determine where and when an attack has occurred.

Higher levels of situation awareness can be attained by automating some of this work, for example through attack correlation. Tools such as RedSeal's Vulnerability Advisor [12] and SkyBox's Risk Exposure Analyzer [14] aid security analysts by doing some of this work and generating attack graphs for them. Further work in attack modeling and detection has been done by Abbot et al [1], and NAVIGATOR [4] generates attack graphs and models the effects of different types of attacks on the system being monitored. Rasmussen et al also developed NIMBLE [11] which aids analysts in responding to threats by providing defensible recommendations.

Attack correlation was also used in visualization tools designed for the 2011 VAST competition mini-challenge on cybersecurity situation awareness. Participants were asked to design situation awareness interfaces that would aid security analysts in identifying and preventing cyberattacks on a fictional freight company. Common themes to these solutions were: preprocessing data, displaying multiple correlated views of the data, and providing an interface with interactions that allowed users to select and view certain sections of data in more detail [8] [9] [10].

Interaction and correlation are very useful tools for an analyst exploring security data. Our visualization tool makes heavy use of interaction and highlights corresponding data throughout the tool relating to which element(s) the user has selected. However, providing interaction only helps so much if there is still a large amount of data for the analyst to search through, especially since the amount and complexity of network data today has simply become too much for a human analyst to correlate manually and achieve situation awareness [13, p. 15]. Erbacher et al note that data analysis is greatly aided by abstract data which the user can

comprehend and then drill down into for more detail [7]. Attack correlation, where low-level data is correlated and displayed as attacks with the option of drilling down to view the underlying network data or intrusion detection logs, is one method of achieving this. Our interactive visualization tool enables an analyst to link specific attacks to high-level mission requirements and resources, as demonstrated on the iCTF dataset, whose completeness in terms of data logs facilitates the analysis of such interdependencies.

Furthermore, although situation awareness visualizations have come a long way in helping analysts identify possible attacks and decide on defensive courses of action, to the best of our knowledge there are no visualization tools that relate this information to the user's missions. This is a critical concept particularly in military cybersecurity scenarios. In order to efficiently coordinate their attention and defensive actions, security professionals must have a way of prioritizing which services to defend and determining which attackers are the biggest threats. By using a mission-centric approach in which attacks are correlated with mission-critical resources over a common Timeline, we are able to build upon existing solutions for identifying attacks and focus our efforts on exploring the benefits of situation awareness visualization tools designed to help analysts make informed decisions.

3. TOOL OVERVIEW

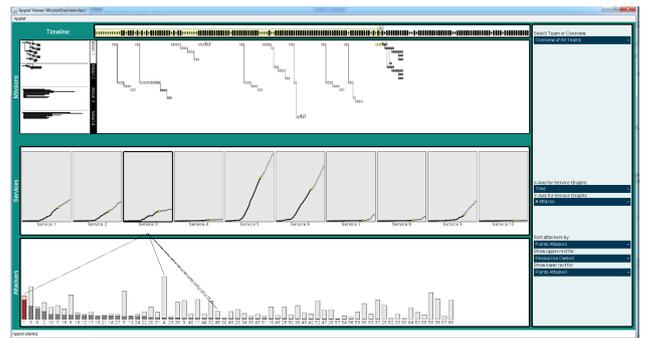


Figure 1: Our visualization tool displaying information from the 2011 UCSB iCTF cybersecurity competition.

Our visualization tool, depicted in figure 1, was designed specifically for mission-centric cybersecurity scenarios. Thus, we divided the screen real estate into three distinct views (with a fourth area that serves as a Control Panel and a fifth which serves as a Timeline). Naturally, one section (top third) is devoted to an overview of the missions themselves. It displays the various stages of the missions and shows the user which stage each mission is currently in as well as possible future stages and potential times-to-completion. The second view (middle row) displays information on the services that are required by the missions. Finally, the third view (bottom row) shows information on attackers.

Each section is useful in its own right; however it is the combination of and interaction among the views that truly promotes a higher level of situation awareness. By correlating key data, analysts can gain a deeper understanding

of the motivation behind attackers’ actions and what effects they will have on the defender’s (user’s) services and the success of the missions.

While our tool has been designed in terms of a general cybersecurity scenario, there are some aspects (such as the possible criteria for sorting attackers) that are customizable for the specific dataset/scenario under review. Next, we will briefly describe the dataset that we utilized in our evaluations.

3.1 The Competition and Dataset

A generic cybersecurity scenario involves a goal or mission that needs to be accomplished, the set of services needed to do so, and various attackers that threaten the mission’s success. The 2011 UCSB iCTF was designed with this in mind. [16] [5] We will include a brief summary of the competition here to provide the necessary background to understand some of our evaluations. In this section, we will describe the structure of the iCTF and highlight the ways in which it models a cybersecurity scenario as well as the ways in which it deviates from a real-world scenario.

3.1.1 The iCTF Competition

In order to apply the concepts of situation awareness, we opt to look at the competition from the perspective of a single (defending) team. Our goal is to visualize this team’s missions, the services required by their missions, and their attackers. The 2011 iCTF, which was a competition with 66 participating teams, thus provides a rich collection of data, consisting of 66 distinct but interlinked data sets.

The competition was a day-long event, and 9 hours of data were logged. The game designers kept track of four missions and which state each one was in at any time. Mission lengths varied between 6 and 48 minutes, and were cyclically iterated over the entire time period. Every 1 - 2 minutes, a time period referred to as a “tick”, a probability function would be used for each mission to determine which of its possible next states it would progress to. The designers would then look at which of the 10 services were required by the current state of each mission. From this information, different values would be assigned to each service to indicate its current level of criticality to the missions. These values were provided to the teams at every tick. Thus, although the teams had no visual representation of the missions, these values kept the teams informed of which services were currently critical to their missions’ success.

While these values could be set to anything that makes a service appear more attractive to an attacker, we will briefly describe the design used in the 2011 iCTF competition. (We will refer to these concepts when describing the modeling of an optimal attacker in section 4.4.1). Game success was measured by achieved points. Teams gained points by hacking into other teams’ services (ie “attacking” those services), thus “capturing the flag” for that service, and submitting this flag to the game coordinators along with some amount of their resources that they wished to convert to points.¹

¹The teams’ resources were actually digital money that they obtained through a separate set of challenges created specifically for the purpose of providing the teams with money.

Depending on which service the submitted flag belonged to, different values were used in calculating the conversion rate of resources to points. These values, which were provided for each service at each tick, were called Cut (C), Payoff (P), and Risk (R). They were partially determined by the number of missions that relied on that service at that point in time with the result that the more critical a service was to the missions, the more attractive it was to an attacker. The Cut represents a sort of processing fee applied to the conversion. The Payoff represents the conversion rate for the remaining resources. The Risk is used to calculate the probability that the conversion succeeded; thus, there was the possibility of losing resources if they used a flag from a service whose risk at the time was too high.

3.1.2 The Dataset

Throughout the competition, data for each team’s actions was recorded and stored in a SQL database. This includes information on the resources owned by the teams as well as when and through which services they attempted to convert their resources to points. The statuses of each team’s services (whether they had been attacked or were unreachable, etc) were also recorded. Additionally, we had access to the python files used to define the missions, which were originally defined as Petri nets. Using these, we were able to generate finite state machine representations, which our tool then used in order to determine how to graphically represent the mission (described further in Section 4.2).

4. VISUALIZATION DESIGN

Our tool is divided into five parts: the Timeline, the Missions View, the Services View, the Attackers View and the Control Panel. To facilitate higher levels of situation awareness, we display data from each of these parts concurrently and allow the user to explore correlations between them through various interactions. In this section, we will describe the functionality and design decisions related to each part.

4.1 Timeline



Figure 2: The Timeline represents the duration of the iCTF competition. A yellow highlight covers the time period for which data is being displayed.

The Timeline, shown in figure 2, represents the time period of interest; in this case, the duration of the competition. The user is able to auto-play through time or jump backwards and forwards to specific ticks by clicking on tick marks in the Timeline. The tick marks represent the one to two minute “ticks” into which the competition was divided. Data is encoded into this part of the visualization tool by enlarging and coloring certain tick marks. Light gray represents ticks where all of the defender’s services were down (i.e. unreachable by the competition organizers during that tick) and black ticks denote that there was at least one attack during that tick. This allows the user to quickly identify key temporal locations in the data to focus on. After However, we will use the term “resources” throughout the paper since it is more general.

observing certain anomalous situations, we decided to use red to highlight ticks where at least one of the services was down/unreachable and yet its flag was still captured, which indicates a certain level of sophistication on the part of the attacker. Lastly, when a service in the Services View is hovered over, the highlights for any ticks that correlate to that particular service are drawn longer and thicker in order to stand out. This can also be seen in figure 2.

One challenge in designing the Timeline was that it has a direct physical mapping to the Missions View as manifested by them extending over the same width, whereas the mapping to the Services View and the Attackers View is not as direct. In order to highlight this difference, we included a spacing that separates the Timeline and Missions View from the Services View and Attackers View. These decisions will be discussed further in the following sections. In general, we decided on this design with the goal of keeping the tool general and applicable to different cybersecurity scenarios.

4.2 Missions

This part of the visualization tool provides an overview of the defending team’s (in this case: four) missions, which it is the ultimate goal of a cybersecurity scenario to complete. Our design process for this view was highly iterative. A previous version of the mission tool was designed as a standalone project designed to visualize petri nets (the original format of the missions for the competition) in a Gantt chart style format, essentially generating linearized versions of non-deterministic finite automata. This representation makes it much easier to see and understand the relative lengths of time different paths would take, which can be an important factor to consider for time-sensitive missions. Figure 3 shows a screenshot from this program. Each row in the chart represents a path from the current state to the final state with back arrows to indicate where possible cycles exist. Furthermore, the services required by each state are represented as colored circles below the state. When a circle is hovered over, all other circles that don’t correspond to the same service disappear and allow the user to quickly determine which paths in the mission contain a state requiring that service. This allows the user to determine which missions would be most at risk if that service were compromised.

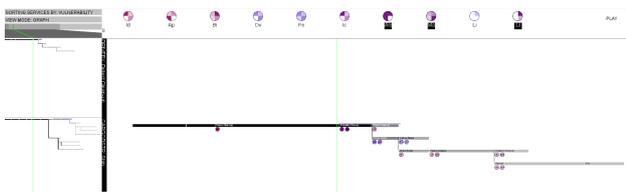


Figure 3: A precursor to the Missions View in our visualization tool showing missions represented as Gantt charts.

In developing the Missions View for our visualization tool, we wanted to further highlight the temporal component of each mission. As can be seen in figure 4, the past paths traversed through the Gantt chart are drawn for each execution of the mission, although previous executions are grayed out. The current iteration of the mission shows all possible

paths from the currently active state. Paths will disappear as choice points are passed and they become unreachable, resulting in a sparser graph as can be seen with the previous iterations of the mission. The rectangles representing the states in each iteration are shrunk down in order to line up with the ticks in the Timeline, thus preserving the temporal correlation. Finally, the states traversed in the currently active iteration are colored bright yellow to indicate current progress through the mission.

Additionally, any iteration of the mission can be hovered over to provide more information. In figure 4, a state in the third iteration of mission 1 is being hovered over. A highlight is drawn around that iteration of the mission and a pop-up displays information about the duration of that iteration as well as the hovered state.

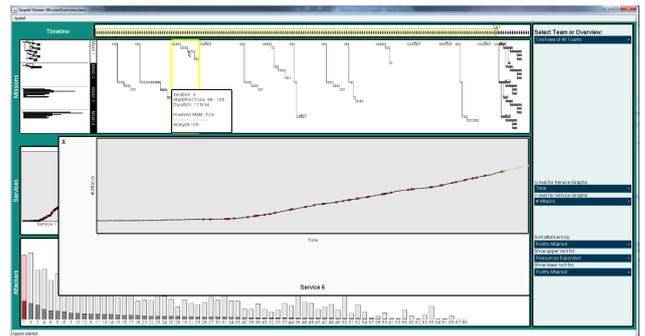


Figure 4: At tick 241, mission 1 is in its 15th iteration. The 3rd iteration is hovered over and thus highlighted. In this figure, Service 6 has also been zoomed in on.

4.3 Services

The middle row of the main screen of our tool is dedicated to the main resources to be utilized and protected over the course of the missions. For our evaluation scenario, we depict the ten services here that each team in the iCTF competition was tasked to host and defend. The Services View, which can be seen in the upper half of figure 5 and the middle of figure 1 is thus divided into ten rectangular areas, one for each service in which data over time relating to that service are graphed. The user can select which data are displayed using a drop-down menu in the Control Panel.

Additionally, the user can shift+click on any service and see an enlarged view in a pop-up as shown in figure 4. The width of this pop-up is set to match the width of the Timeline, thus providing the user with a more direct correlation between the service they select, the Timeline, and the Missions View as well (since it is already correlated with the Timeline). After some discussion with faculty in the Security Lab at UCSB, we realized that the main usefulness of the Services View lies in its ability to help security analysts decide where they most need to focus their attention. We allowed each rectangle to be minimized to further aid the user in focusing on only a subset of the services, thus giving users some control over screen estate management. This is also shown in figure 5 where services 3 and 7 are minimized. The red coloring for each service is a result of hovering over the optimal attacker and will be further discussed in section

4.4.1. Additionally, when a service is hovered over, the ticks when it was attacked are enlarged on the Timeline as can be seen in figure 2.

4.4 Attackers

Finally, the Attackers View, shown in the lower half of figure 5, is where any known or inferred information about attackers is displayed. In a typical cybersecurity scenario, this data would consist of attackers' IP addresses, the attacked IPs and ports, and possible collusion of the attackers. The success rate of attacks from each IP or geographical location might also be displayed along with attack frequency. In our evaluation, working with the more abstracted iCTF dataset, the information we display for attackers includes: a unique ID, their attack frequency and success rate, the amount of resources they own and the amount they have used, and also the severity of their attacks, which is determined by their efficiency in converting their resources to points. Each time a new attack occurs, information about it is added to the view as a stacked bar and sorted according to a user-selected criterion. Due to the high-level data in the iCTF dataset, when an attack occurs we are able to determine which team generated it and either add a new attacker to the display or update data for a previously seen attacker. This allows us to provide a display that shows an object (in this case a stacked bar) for each attacker rather than displaying multiple graphs, which the user must then correlate by hand to determine when attacks have occurred and where they originated.

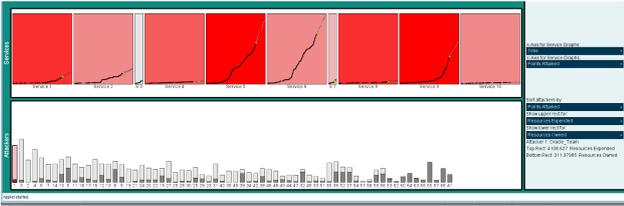


Figure 5: Attackers are displayed as stacked bars labeled with an ID number. Here, the bars represent resources owned and resources expended. The service graphs are colored red according to the frequency of the optimal attacker's projected attacks over the next 5 ticks.

In addition to being able to choose which data to view for the identified attackers and how to sort them, the user is able to hover over attackers causing the ticks during which they attacked to be enlarged on the Timeline in the same way they are enlarged when the user hovers over a service (shown in figure 2). This allows the user to see more detailed temporal data relating to each attacker's activity. One attacker in particular is colored red rather than gray. This is actually a simulated attacker we generated and introduced for comparison purposes and to enable look-ahead situation awareness. This attacker follows an 'optimal' attacking strategy which we will describe further in the next section.

4.4.1 Optimal Attacker

Highlighting the optimal attacker's strategy is a way to inform the user which services would suffer the most if a sophisticated attacker were to strike at any given moment.

This is done by calculating which services the optimal attacker would target, and when, over the next 5 ticks. A clever attacker will optimally allocate their resources u_k^i at every tick k and through every service i during the competition such that the total number of points at the end of the game is maximized subject to the estimated dynamics and constraints (e.g. an upper bound on the resources available till the end of the game). A moving horizon optimization scheme is an ideal approach since the attacker continuously extends their horizon as time progresses, allowing them to incorporate new data at any point in time. If the horizon is small ($N \leq 5$ ticks) the attacker can make very good predictions by using the history of existing data and then allocating their resources to obtain the maximum number of points. For easier accessibility, we will first define an optimization scheme as if all the data was known to the attacker a priori and then we will describe the moving horizon optimization with estimated dynamics for the data getting into the system. The problem, for 248 time ticks and 10 services, is described as:

$$\begin{aligned} & \text{maximize} && \sum_{k=1}^{248} \sum_{i=1}^{10} \mathbb{E}[\rho_k^i p_k^i u_k^i] \\ & \text{subject to} && \sum_{k=1}^{248} \sum_{i=1}^{10} u_k^i \leq U \\ & && \text{w.r.t. } u_k^i \in [0, \infty), \forall i \in \{1, 2, \dots, 10\}, k \in \{1, \dots, 248\}, \end{aligned}$$

where U is an upper bound on the available resources and for each service $i \in \{1, 2, \dots, 10\}$, the $\rho_1^i, \rho_2^i, \dots, \rho_{248}^i$ are independent and identically distributed (i.i.d.) Bernoulli random variables with

$$P(\rho_k^i = 0) = \min\left\{\frac{R_k^i}{30} u_k + 0.4, 1\right\}$$

where $p_k^i \equiv P_k^i(1 - C_k^i) \in [0, 1]$. Because of the i.i.d. assumption, and since it never makes sense to choose a value for u_k^i for which $\rho_k^i = 0$ with probability one, we can expand the expected value in the optimization criterion and obtain

$$\begin{aligned} & \text{maximize} && \sum_{k=1}^{248} \sum_{i=1}^{10} p_k^i \left(0.6 - \frac{R_k^i}{30} u_k\right) u_k^i \\ & \text{subject to} && \sum_{k=1}^{248} \sum_{i=1}^{10} u_k^i \leq U \\ & && \text{w.r.t. } u_k^i \in \left[0, \frac{0.6}{\frac{R_k^i}{30}}\right], \forall i \in \{1, 2, \dots, 10\}, k \in \{1, 2, \dots, 248\}. \end{aligned}$$

which is a concave maximization problem (convex minimization) with linear constraints, a duality gap of zero, and which can be solved relatively fast through a Matlab-based convex optimization solver such as CVX [3]. For space reasons, more details on our optimization algorithm are presented in [17].

Once the optimal attacker's actions have been determined, corresponding highlights are added upon mouse-over of their stacked bar in the Attackers View. The background of each service's graph is colored a shade of red corresponding to the frequency at which it will be attacked by them over the next 5 ticks (figure 5). Since the optimal attacker always focuses on the most critical services, this coloring provides the user with a sense of which services being targeted would

cause the most damage to their mission's completion in the near future, thus allowing them to make more informed defensive decisions.

When comparing the attack trends of the optimal attacker to the actions of the actual attackers, we found that they matched quite well. Figure 6 shows the resources the optimal attacker converted to points through each service whereas figure 7 shows the same thing for all the real attackers. The most resources were converted through services 2, 5, 6, and 9 in both plots. However, services 3, 4, 7, 8, 9, and 10 were targeted substantially less by the actual attackers than expected. In figure 1, since service 3 is hovered over, we can see that, as of tick 190, it had only been exploited by 4 teams. Through further exploration using our visualization tool, we found that services 4, 7, 8 and 10 were likewise attacked by fewer than 10 teams overall. Service 9 was attacked by numerous teams, but not until much later in the game. This indicates that the vulnerabilities in these services were difficult for teams to find and exploit. For the optimal attacker, we assume it knows the vulnerabilities of every service and can attack any one at any time. Had teams been able to exploit all services from the beginning of the game, the plots may have been much closer.

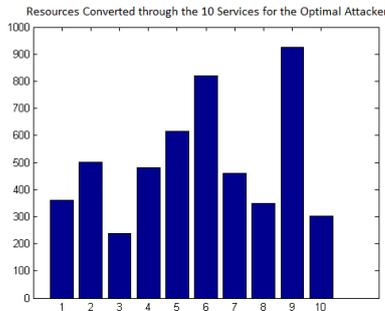


Figure 6: The resources the optimal attacker converted to points through the 10 services

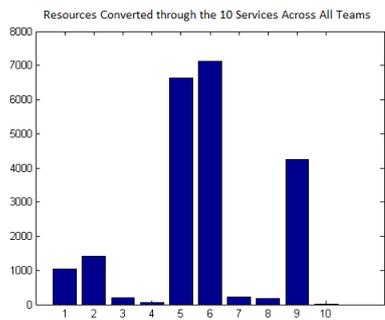


Figure 7: The resources converted to points by all teams (attackers) through the 10 services

5. RESULTS / LESSONS LEARNED

The goal of our visualization tool is to help security analysts achieve situation awareness regarding their missions, services and potential threats. To evaluate the effectiveness of our visualization tool in achieving this result, we now look at a set of motivating questions. In this section, we will address the main questions in each of the three categories of

SA as described by Endsley - Perception, Comprehension, and Projection [6] - and discuss examples of how they can be answered.

5.1 Perception

Example questions that address a user's perception of the cybersecurity situation include:

- What is the current state of our missions?
- Who is currently attacking us?
- Which services are they targeting?

These types of questions are used to establish the current situation before the user makes any conclusions or predictions. A first glance at the visualization provides the relevant answers. Figure 8 shows the visualization screen at tick 148 for team Enoflag. In the Missions View, we can see which path has been taken through the mission during its current iteration and that it is one step away from completion. We can also see that services 8, 9 and 10 are required for this step to be completed, because their outlines are highlighted in the Services View. Lastly, the attackers' bar heights show that there is one attacker who has far more resources at their disposal than any of the others, yet because the attackers are sorted by points attained, we are also aware that that attacker has made relatively fewer points and, therefore, is likely a less efficient attacker and worth less of our attention.

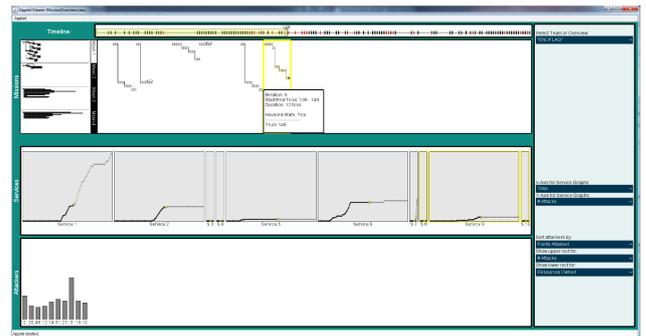


Figure 8: The visualization tool showing data from the perspective of team Enoflag at tick 148.

It is also easy to spot that there were many times when this team had a service that was down and yet was attacked (shown as a red tick line on the Timeline), indicating a certain level of sophistication on the part of the corresponding attacker and thus providing additional information to the user when evaluating which attackers pose the highest threat.

5.2 Comprehension

Example questions that directly facilitate a user's comprehension of the cybersecurity situation include:

- Which services have been targeted the most?
- Which attackers have attacked us the most?
- Which attackers have attacked us most effectively?

The customizable graphs displayed for each service allow the user to explore whatever data they feel is relevant, controlled

by simple drop-down menus. For example, by graphing the number of attacks for each service as in figures 1 and 8, it quickly becomes clear that services 5 and 6 have been targeted much more often than the others. However, it is interesting to note that even though more attacks have targeted service 5 overall (figure 1), more points have actually been attained through service 6 (figure 5). This particular situation was a surprise to the game designers and shows that the attackers who were converting their resources to points using service 6 were doing so more efficiently which may indicate that they had a better understanding of the game dynamics and could be more dangerous attackers. It may also indicate that service 6 is simply more attractive for attackers in terms of reward per effort and may therefore warrant more defense.

We can further explore data about these attackers in the Attackers View. By setting the stacked bars to represent the number of attacks from an attacker or the amount of resources they own, the user can easily scan the Attackers View and pick out the attackers with the highest potential threat. Sorting the attackers by points attained also conveys an understanding of which attackers have executed the most efficient attacks and may mean that the user wants to try blocking these teams specifically. Note that information in this section of our tool may be more incomplete and conjectural in a real-world cyberscenario. Using the iCTF dataset, we have the luxury of fairly complete attacker information, something that would have to be provided through external threat detection algorithms in a real-world application of our software. The computer security group that this tool was developed in close collaboration with has several such attacker identification efforts under development.

5.3 Projection

Example questions that drive a user's exploration of the possible futures of the current cybersecurity situation include:

- Which services will it be most effective to concentrate defense on?
- Which attackers pose the most serious threat to the completion of the missions?

Our visualization tool's initial state is to show an overview of information. An analyst is then able to narrow things down and focus on areas of interest. For example, in deciding which services will most need to be defended, we might start by setting the service graphs' y-axes to be the number of attacks that have targeted that service over time. The user can also narrow down the number of services they are interested in keeping a watch on. For example, figure 8 shows a situation where services 3, 4, 7, 8 and 10 have been minimized by the user, allowing them to focus on the remaining services. However, in order to still maintain an awareness of the status of the minimized services, interactions with them are preserved despite their minimized status. For example, hovering over service 7 highlights any ticks on the Timeline when it was targeted by attacks. Figure 2 shows the Timeline at the end of the competition for team Enoflag when service 7 is hovered over. The ticks during which Enoflag's service 7 was attacked are enlarged and we can see that it was never attacked until quite late in the game. Once it began to be attacked, around tick 200, the user would have

been able to maximize it again to keep a closer eye on it and potentially shift their defense efforts to include it as well.

The Optimal Attacker is also a great asset for Projection. By hovering over its bar in the Attackers View, the user is provided with color coding on the services' graphs indicating which ones this attacker would most likely target over the next 5 ticks, as seen in figure 5. While there is certainly no guarantee that this is what an attacker will do, it shows the user which services are the most critical to their missions for 5 - 10 minutes into the future. This is highly valuable information since it is impossible to know exactly what attackers will do and so it is often useful to concentrate one's defenses on the most critical or vulnerable services in order to mitigate damage that might compromise the mission.

6. DISCUSSION

The design of our visualization tool went through several iterations as we worked to make it both intuitive and informative. We originally allowed the user to choose between displaying aggregate or instantaneous data in both the Services and Attackers views. However, early feedback showed that this was more confusing than helpful because the type of data could be inconsistent across views. Instead, we adopted an overview-then-zoom structure where the default mode always shows aggregate data from the start of the game to the currently selected tick, but which provides users with the ability to zoom-in on certain items and see the data for the current tick only. This keeps the visualization tool in the same state overall and maintains consistency while still allowing more fine-grained data to be shown.

Our tool provides a first look into what situation awareness visualization tools can offer in the presence of high-level cybersecurity objectives, and we hope our experiences will help inform future work in the area. In particular, it is a first step in designing a SA visualization tool that relates all data to the operator's *missions*, which are critical for many cyberdefense scenarios.

Also of particular interest, is the implementation of an optimal attacker, supporting the determination of which services will be most critical in the near future. Since predictions of attackers' actions always embody some amount of uncertainty, it is highly important to inform the user about known data concerning which services could most threaten the mission's success if they were attacked.

6.1 Limitations and Future Work

One important caveat in judging the specific visualization tool successes we presented in our evaluation is that they are all based on a dataset that, while it has several great advantages over artificially created simulation data, still is an abstracted version of actual cyberdefense missions. The structure of the competition allowed us to determine when attacks were made and which teams were responsible for them. In general, determining when an attack has been made is a difficult and uncertain task involving correlating large amounts of low-level data like network logs and security alerts. Furthermore, it can be even more difficult to determine whether two attacks were performed by the same entity. Thus, improving methods of attack correlation and automated defensible recommendations is still an existing

challenge. However, solving these issues was not the focus of our work. As attack correlation and network monitoring methods improve, situation awareness models will need to adapt to data that will become less focused on packet traces and more related to the operator's mission and how to defend against threats. Our work makes a first step in exploring the possibilities of situation awareness in the presence of this data.

We have received feedback requesting the ability to select subsets of the attackers or services and to be able to view data for them alone. This would provide another way to "zoom-in" on parts of the data. We are currently implementing this functionality, but it was not completed in time for this submission.

We also have further plans for the Missions View display. While it currently provides the user with information about the mission's states on mouse-over (in a pop-up as well as by highlighting the services required by that state), we also plan to provide overview information about the mission that is always visible. This includes visual indications about which are the most endangered branches of the current mission continuation options considering which services have been most targeted by attacks or are likely to be targeted heavily in the near future. We also have plans to include more interactions with the mission such as highlighting all states that require a particular service when that service is moused-over in the Services View or when it is required by a state that is being moused-over.

Lastly, our greatest concentration of future work will be on supporting and providing analysis tools for "what if" scenarios based on past data and any known future data (such as the structure of the mission and which states follow from the currently active state). To this end we are developing optimization schemes for the defender's possible actions such as taking a service offline when not needed or extending the duration of a state that would be unable to progress if a certain service were compromised.

7. CONCLUSION

In this paper we presented a visualization tool for mission-centric cybersecurity situation awareness. Through interaction with our visualization tool, we were able to discover interesting elements of a realistic example dataset that would not otherwise have been readily apparent. Additionally, based on feedback from domain experts in the field of cybersecurity, we discussed decisions a security analyst might make when using our visualization tool at various points in time over the period of our iCTF data. By reviewing the actual development of the situation according to how the iCTF played out, we then verified these decisions by noting that the services which we deemed most in need of defense were in fact the ones most targeted by attacks. We hope that this work will provide a starting point for developing mission-centric visualization tools for situation awareness in order to aid security analysts in making informed defensive decisions to ensure their missions' success.

8. ACKNOWLEDGMENTS

This material is based upon work supported by the U.S. Army Research Laboratory and the U.S. Army Research

Office under MURI grant No. W911NF-09-1-0553

9. REFERENCES

- [1] J. Abbott, J. Bell, A. Clark, O. D. Vel, and G. Mohay. Automated recognition of event scenarios for digital forensics. In *Proceedings of the 2006 ACM symposium on Applied Computing*, pages 293 – 300, 2006.
- [2] P. Barford, M. Dacier, T. G. Dietterich, M. Fredrikson, J. Griffin, S. Jajodia, S. Jha, J. Li, P. Liu, P. Ning, X. Ou, D. Song, L. Strater, V. Swarup, G. Tadda, C. Wang, and J. Yen. *Cyber Situational Awareness*. Springer US, first edition, 2010.
- [3] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [4] M. Chu, K. Ingols, R. Lippmann, S. Webster, and S. Boyer. Visualizing attack graphs, reachability, and trust relationships with navigator. In *Proceedings of the Seventh International Symposium on Visualization for Cyber Security*, pages 22 – 33, 20110.
- [5] A. Doupé, M. Egele, B. Caillat, G. Stringhini, G. Yakin, A. Zand, L. Cavedon, and G. Vigna. Hit 'em where it hurts: A live security exercise on cyber situational awareness. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC 2011)*, Orlando, FL, December 2011.
- [6] M. Endsley. *Theoretical Underpinnings of Situation Awareness: A Critical Review*, chapter 1, pages 3 – 32. L. Erlbaum Assoc., 2000.
- [7] R. F. Erbacher, D. A. Frincke, P. C. Wong, S. Moody, and G. Fink. A multi-phase network situational awareness cognitive task analysis. In *Information Visualization*, pages 204 – 219, 2010.
- [8] N. A. Giacobe and S. Xu. Geovisual analytics for cyber security: Adopting the geoviz toolkit. In *the 2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 315 – 316, 2011.
- [9] L. Harrison, W. Dou, A. Lu, W. Ribarsky, and X. Wang. Guiding security analysis through visualization. In *in the 2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 317 – 318, 2011.
- [10] W. M. Lamagna. An integrated visualization on network events. In *in the 2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 319 – 321, 2011.
- [11] J. Rasmussen, K. Ehrlich, S. Ross, S. Kirk, D. Gruen, and J. Patterson. Nimble cybersecurity incident management through visualization and defensible recommendations. In *Proceedings of VizSec*, 2010.
- [12] Redseal systems: Vulnerability advisor. <http://www.redseal.net>, March 2012.
- [13] H. Shiravi, A. Shiravi, and A. A. Ghorbani. A survey of visualization systems for network security. *Visualization and Computer Graphics, IEEE Transactions on*, 2011.
- [14] Skybox security: Risk exposure analyzer. <http://www.skyboxsecurity.com>, March 2012.
- [15] tcpdump & libpcap. <http://www.tcpdump.org/>, March 2012.
- [16] The ucsb ictf. <http://ictf.cs.ucsb.edu/>, July 2012.
- [17] K. G. Vamvoudakis and J. P. Hespanha. Optimal

attacks for the ictf game.

<http://www.ece.ucsb.edu/~hespanha/techrep.html>,
July 2012.

[18] Wireshark. <http://www.wireshark.org/>, March 2012.